Affordable Spectral Measurements of Translucent Materials

TOMÁŠ ISER, Charles University, Czech Republic
TOBIAS RITTIG, Charles University, Czech Republic
EMILIE NOGUÉ, Imperial College London, United Kingdom
THOMAS KLAUS NINDEL, Charles University, Czech Republic and Berufsakademie Sachsen, Germany
ALEXANDER WILKIE, Charles University, Czech Republic

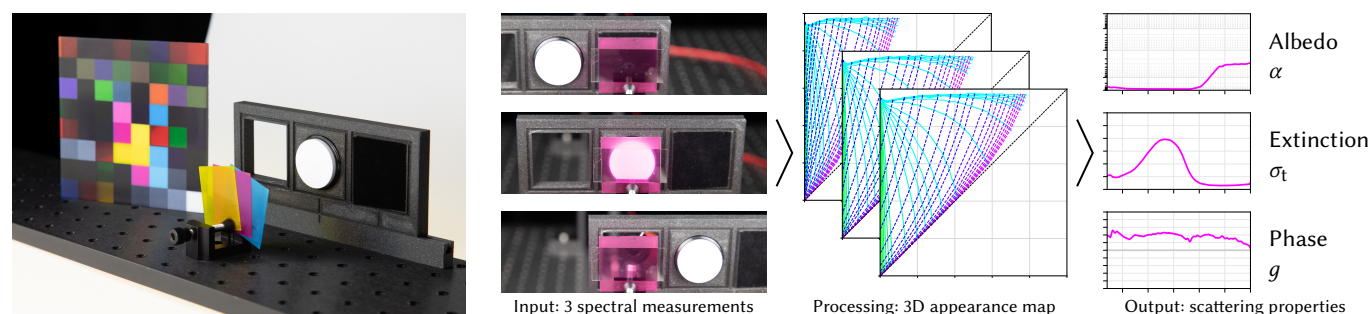| Input: 3 spectral measurements | Processing: 3D appearance map | Output: scattering properties |

Fig. 1. We present a method for measuring bulk optical properties of translucent materials, such as the base inks of a full-color 3D printer. Starting from a thin sample, we measure three spectral intensities: on a black background, white background, and of a collimated light beam. After a fitting step through our precomputed 3D appearance map, we obtain the material's spectrally-varying albedo, extinction coefficient, and phase function anisotropy.

We present a spectral measurement approach for the bulk optical properties of translucent materials using only low-cost components. We focus on the translucent inks used in full-color 3D printing, and develop a technique with a high spectral resolution, which is important for accurate color reproduction. We enable this by developing a new acquisition technique for the three unknown material parameters, namely, the absorption and scattering coefficients, and its phase function anisotropy factor, that only requires three point measurements with a spectrometer. In essence, our technique is based on us finding a three-dimensional *appearance map*, computed using Monte Carlo rendering, that allows the conversion between the three observables and the material parameters. Our measurement setup works without laboratory equipment or expensive optical components. We validate our results on a 3D printed color checker with various ink combinations. Our work paves a path for more accurate appearance modeling and fabrication even for low-budget environments or affordable embedding into other devices.

CCS Concepts: • **Computing methodologies** → **Rendering**; *Reflectance modeling*; • **Applied computing** → **Computer-aided manufacturing**; *Physics*.

Authors' addresses: Tomáš Iser, tomas@cgg.mff.cuni.cz, Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 118 00, Czech Republic; Tobias Rittig, tobias@cgg.mff.cuni.cz, Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 118 00, Czech Republic; Emilie Nogué, e.nogue@imperial.ac.uk, Imperial College London, Exhibition Rd, South Kensington, London, SW7 2BX, United Kingdom; Thomas Klaus Nindel, thomas@cgg.mff.cuni.cz, Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 118 00, Czech Republic and Berufsakademie Sachsen, Hans-Grundig-Strasse 25, Dresden, 01307, Germany; Alexander Wilkie, wilkie@cgg.mff.cuni.cz, Charles University, Faculty of Mathematics and Physics, Malostranské nám. 25, Prague, 118 00, Czech Republic.

Additional Key Words and Phrases: spectral measurements, optical properties, translucent materials, light scattering, appearance modeling, predictive rendering, 3D printing, computational fabrication

1 INTRODUCTION

The world around us is filled with translucent materials which partially absorb or scatter light that passes through them. Examples include many liquid-state materials (water, wine, milk), aerosols (fog, clouds, smoke), biological matter (tissue, skin, blood), and many solids (plastic, wood, porcelain). Characterizing and measuring the optical properties of such materials is therefore an important topic not only for appearance predictions in computer graphics, but also for various other fields, including biology and medicine [Jacques 2013; Tuchin 1993] or atmospheric sciences [Emde et al. 2016].

Its importance also shows in full-color 3D printing, where recent advances enabled fabricating customized objects with desired appearances. The appearance reproduction is achieved by depositing base printing inks layer-by-layer, forming a high-resolution optically heterogeneous grid. Each voxel of this grid is translucent, and finding their adequate arrangement is important to realize the desired color mixing. Hence, for accurate printouts without texture blurring or color inaccuracies, it is key to know the inks' optical properties [Elek et al. 2017; Nindel et al. 2021; Sumin et al. 2019].

Measuring these optical properties is our main goal. More importantly, we focus on estimating them *spectrally*, allowing for accurate predictions of the colors when the base inks are *spatially mixed*

within the 3D printouts. As we show in our supplemental document, spectral simulations of translucent materials eliminate color inaccuracies caused by metameric ambiguities and attenuation non-linearity, as opposed to three-channel RGB simulations.

The optical properties are estimated by designing and solving an appropriate *inverse problem*, which consists of first measuring light interacting with the material particles in bulk, and then numerically finding which optical properties must have led to such results. For simplicity (unlike recently studied non-exponential media [Bitterli et al. 2018]), we model each material as *homogeneous*, with molecules distributed in an *uncorrelated, uniform* fashion. Under this assumption, the material is characterized by the wavelength (λ) dependent *extinction coefficient* $\sigma_t(\lambda) \left[\text{mm}^{-1}\right]$, the *single-scattering albedo* $\alpha(\lambda)$, and the *phase function* $p(\lambda, \theta)$. The extinction coefficient $\sigma_t(\lambda)$ characterizes the *exponential attenuation* (known as Beer's law [Jarosz 2008]) of the light intensity $I(\ell)$ after traveling through the medium for a distance ℓ with the initial intensity I_0:

$$I(\ell) = I_0 \cdot \exp\left(-\ell \cdot \sigma_t\right). \tag{1}$$

The attenuation is caused by absorption and out-scattering of the photons, parametrized by the *absorption and scattering coefficients* $\sigma_a(\lambda)$ and $\sigma_s(\lambda)$. It holds that $\sigma_t(\lambda) = \sigma_a(\lambda) + \sigma_s(\lambda)$, and $\alpha(\lambda) = \sigma_s(\lambda)/\sigma_t(\lambda)$. The phase function $p(\lambda, \theta)$ characterizes the angular θ distribution of the light scattering. We assume a simple single-parameter phase function of Henyey and Greenstein [1941], which approximates the more accurate Mie scattering model, but only requires a single parameter $g \in [-1, 1]$; $g < 0$, $g = 0$, and $g > 0$ for dominantly backward, isotropic, and dominantly forward scattering, respectively. We acquire the parameters $(\alpha(\lambda), \sigma_t(\lambda), g(\lambda))$ for the given materials with a high spectral resolution of approximately 1000 wavelengths within the visible spectrum.

Such measurements using the existing methods (Sec. 2.2) require complicated calibrations and/or expensive laboratory-grade equipment such as gonio-photometers, rotating platforms, integrating spheres, time-resolved detectors, or lasers. Although there are alternative approaches with emphasis on simplicity and affordability, they rely on acquiring one-dimensional spatial profiles, which is impossible with a spectrometer that can only measure a single point. Hence, these methods only use RGB cameras, and using them for spectral measurements would require expensive hyperspectral photography. Furthermore, many of these setups suffer from ambiguities that make them unable to estimate the phase function.

Our contribution. To enable high-resolution spectral measurements, our core question is how to resolve the three unknown material parameters $(\alpha(\lambda), \sigma_t(\lambda), g(\lambda))$ as quickly as possible per each wavelength, and with a minimum number of point measurements with a spectrometer. This is not a simple question, as we need to account for various factors, including the similarity relations [Wyman et al. 1989; Zhao et al. 2014]. After studying the existing approaches and their behavior, we found and proved that there exists a simpler and easy-to-visualize relation between material appearance and its optical properties (Sec. 3), which we call the *appearance map*. The core of our approach is how light intensity changes in three different settings: first, a given material sample is diffusely illuminated and placed against a black background, and a white background; then,

it is observed with a collimated illumination. These three combinations are not only easy to simulate using a Monte Carlo renderer, but also easy to capture in real life with a simple and affordable physical setup with single spectrometer (Sec. 4), which allows spectrally resolving the whole parameter triplet (α, σ_t, g).

While we believe our method to be applicable in many fields, we remain concise and demonstrate the results on the inks for full-color 3D printing. We measure the base CMYKW (cyan, magenta, yellow, black, white) inks and validate our method's accuracy by predicting the spectral appearance of their various mixtures within a halftoned color checker (Sec. 5). This implies the promised applicability of our method to the aforementioned printing optimization pipelines.

2 RELATED WORK

We begin by introducing relevant *prediction models* for simulating light behavior in translucent materials (Sec. 2.1). Using these models inversely is then the core of all *measurement methods* (Sec. 2.2), which show strategies for capturing materials and inversely fitting their matching optical properties. Lastly, we relate our work to the context of translucency appearance and 3D printing (Sec. 2.3).

2.1 Prediction models

We briefly look at prediction models that have been extensively used by methods for estimating the bulk optical properties in the past. We refer to Frisvad et al. [2020] for a complete survey.

The following models have been developed for a simplified situation of diffusely illuminated, infinitely wide, homogeneous layers of translucent materials. Kubelka and Munk [1931] show the computation of the total reflectance and transmission, a model popular for modeling textiles, paints, and 2D printing. It was later generalized to support refractive boundaries, non-diffuse illumination, and 3D radiative transfer [Nobbs 1985; Sandoval and Kim 2015; Yang and Hersch 2008]. More accurate is the *adding-doubling model* [Prahl 1995; van de Hulst 1980], which generalizes the problem to multi-layer material stacks (stacked in one dimension), including correctly handling refractive boundaries. It is also useful for fast rendering of layered materials [Jakob et al. 2014; Zeltner and Jakob 2018]. While these models are fast, they lack flexibility and simplicity in defining the complete 3D geometry, including the light sources and sensors, and samples of finite dimensions, which are important in our method to accurately compute the appearance map.

The *diffusion* approximation [D'Eon and Irving 2011; Haskell et al. 1994; Jensen et al. 2001] is another often used model. It is based on the assumption that scattering events are more frequent than absorption. This makes it inaccurate for low-albedo materials, excluding the application for example to inks of dark colors.

In the end, we decided to rely on *Monte Carlo* methods that solve the radiative transfer by probabilistically sampling the space of possible light paths through the scene, and calculating their radiance contribution [Kajiya 1986]. This is the most flexible solution that supports various scenes including translucent materials [Elek et al. 2017; Novák et al. 2018], and offers accurately simulating a wide range of material properties within our method's geometry. We used a modified fast and flexible implementation of Nimier-David et al. [2019] and Jakob et al. [2022].

2.2 Measurement methods

The simplest measurement can be performed by shining a collimated beam of intensity I_0 through a material sample of thickness ℓ, measuring its attenuated intensity I, and simply inverting Eq. (1):

$$\sigma_{\mathrm{t}} = \ell^{-1} \left(\ln I_0 - \ln I \right), \qquad (2)$$

but this approach gives no information on light scattering, necessitating a more sophisticated approach.

A common spectral measurement method is based on capturing the total *hemispherical reflectance and transmittance* of a material sample [Pickering et al. 1992, 1993; Prahl et al. 1993] using two integrating spheres, or a Coblentz hemisphere [Schröder et al. 2015]. Two material parameters (either $\alpha, \sigma_{\mathrm{t}}$; or α, g; assuming that the third parameter is known, e.g., from Eq. (2)) are then fitted using *inverse adding doubling* [Prahl 2011]. Our method does not require integrating spheres and instead uses significantly less expensive and easily switchable reflective and absorptive backgrounds.

Another group of methods relies on fitting the optical properties to *one-dimensional brightness curves*. They illuminate a sample in a specific way, and then capture how the brightness changes spatially over a certain region. These methods are affordable, because they use a simple RGB camera to take a photograph, from which the brightness curve is extracted. The main downside is that one cannot use a spectrometer as it can only measure a single point of interest, and hyperspectral cameras or various spectral filters would be an expensive alternative. Examples of such methods include the one of Narasimhan et al. [2006], which uses dilution to reach such low concentrations of the material that single scattering dominates over multiple scattering. Other methods combine fitting the one-dimensional profile to the diffusion-based prediction model with also measuring the total diffuse reflectance [Jensen et al. 2001; Weyrich et al. 2006]. Papas et al. [2013] used a similar approach with a custom-made container with five LEDs of different spectra, and used a combination of Monte Carlo and quantized diffusion [D'Eon and Irving 2011] as their prediction models.

Elek et al. [2021] presented an alternative method acquiring a *lateral scattering profile* on a step-edge black and white background. Similarly to us, they demonstrate their results on 3D printing, fit the properties to a Monte Carlo simulated dataset, and use contrasting backgrounds. However, they have exactly the same disadvantage as above, requiring a camera capture. A similar idea of using different backgrounds, including a mirror, was also opened in a short concurrent work by Pranovich et al. [2021]. However, they only used a simplistic light transport model, did not estimate phase functions, and did not analyze the appearance map like we did.

With the exception of the hemispherical reflectance and transmittance measurements, none of the methods above are capable of *measuring the phase function*. Elek et al. [2021] hinted that it might be possible with their profile, but our own experiments concluded that the discriminability is not high enough. In our supplemental document, we discuss similarity relations that are the general culprit: there exist equivalent optical properties that give the same material appearance under given conditions.

Methods that are primarily aimed at measuring the phase function are based on measurements of a collimated beam from various angles [Gkioulekas et al. 2013; Leyre et al. 2014], or assume that $(\alpha, \sigma_{\mathrm{t}})$

is already known [Minetomo et al. 2018]. While Gkioulekas et al. [2013] accurately match various phase function shapes of reference materials, not limited to the simple Henyey-Greenstein model, they require rotation platforms, accurate calibrations, and bright collimated illuminants. Our method can only estimate a single-parameter phase function, but is affordable and less complicated.

2.3 Translucency appearance and color 3D printing

While our measurements aim to be objective and give physically meaningful parameters, it is important to note that the human visual system perceives translucency and translucent objects in a bigger context, combining parameters such as lighting direction or object's shape, especially its edges [Fleming and Bülthoff 2005; Xiao et al. 2014, 2020]. This has led to the important question of how to define and measure translucency in a both physically and perceptually meaningful way [Urban et al. 2019]. Especially in full-color 3D printing, the complex heterogeneous light scattering inside the printouts poses challenges such as texture blurring or inaccurate color reproduction. Methods that aim to counteract these problems and control the printout's appearance can be roughly categorized into two classes: "top-down" phenomenological approaches [Brunton et al. 2018, 2015; Chen and Urban 2021; Urban et al. 2019], and "bottom-up" simulation-based methods [Elek et al. 2017; Nindel et al. 2021; Rittig et al. 2021; Sumin et al. 2019]. The latter works use Monte Carlo simulations to predict the appearance of a given heterogeneous 3D grid of solidified droplets of various base inks, which is then used inside an optimization loop. So far, these approaches relied on RGB measurements, which suffered from the issues discussed in our supplemental document. Our spectral results overcome these issues and are directly applicable and pluggable into the existing pipelines.

3 METHOD

Our method is based on constructing a mapping, or a relation, between optical properties of a given material sample, and light intensities in simple geometrical configurations that are easy to capture with a spectrometer. We have to ensure that the mapping is one-to-one between the optical properties and different measurements, at least for physically meaningful settings. We call the mapping the *appearance map* because it contains every possible material appearance in the given setting. In the following sections, we define and study a three-dimensional and an attendant two-dimensional appearance maps using Monte Carlo simulations, and we explain the actual measurement procedure later in Sec. 4.

Separation of surface and volume light transport. Solid and liquid translucent materials typically have a well-defined boundary where the material starts and ends. For example, a glass of water has an exterior air-to-glass boundary, and an interior glass-to-water boundary. Because light interacts not only with the medium itself, but also its boundary, it is beneficial to treat surface and internal optical properties independently. Our work focuses on volume properties of translucent materials. Hence, throughout this work, we eliminate the influence of the surface as much as possible by assuming that it is a simple, perfectly *smooth dielectric surface layer* governed by Fresnel equations. This assumption is valid both in theory and

(a) Appearance map (b) Monochromatic visualization (c) Spectral visualization
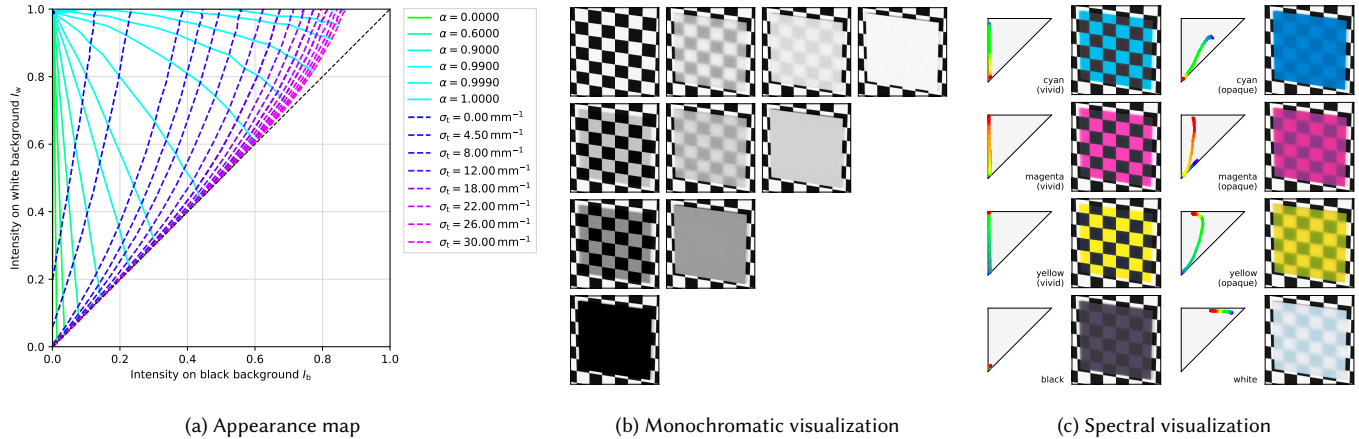
Fig. 2. Our two-dimensional appearance map relates the appearance of a material sample to its optical properties, here assuming the Henyey-Greenstein phase function $g = 0.4$, sample thickness $\ell = 0.5\,\text{mm}$, and refractive index $n = 1.0$. The black and white backgrounds are assumed to be perfectly diffuse with reflectances 1% and 99%. The contours in (a) are of constant α and σ_t, respectively. Refer to Sec. 3.1 for a detailed analysis. In (b), we show rendered examples with their layout roughly corresponding to their position in the triangle in (a): the left edge contains purely absorbing, non-scattering materials; the top edge purely scattering, non-absorbing materials; and the diagonal perfectly opaque materials. In (c), we expand the visualization spectrally for 8 examples, with the colored curves in the small triangles representing per-wavelength coordinates in the appearance map, from shorter (blue) to longer (red) wavelengths.

3.1 Material appearance against diffuse backgrounds

A trivial property of translucent materials is that their appearance depends on the background behind them. For example, in Fig. 3, red wine appears red over a white background, but is black over a black background, while a strongly scattering milk remains white in both conditions. Our crucial observation is that there exists a *structured mapping* between the material's optical properties α, σ_t and the
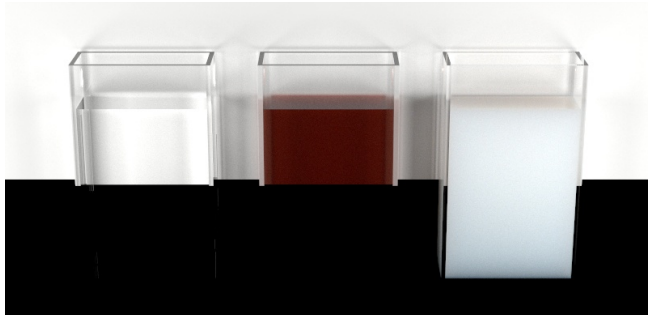

Fig. 3. Three translucent materials in glass cuvettes placed against ideal black and white diffuse backgrounds. Water (left) is virtually non-absorbing and non-scattering, so one can see the backgrounds clearly without any color cast. Red wine (middle) is similar, but strongly absorbs light of wavelengths other than red, giving it a red color cast. Milk (right) is a strongly scattering material, so most of the light is scattered before reaching the backgrounds, which makes it appear white even against the black background.

observed light intensities I_b, I_w against wavelength-independent diffuse black and white backgrounds. We call this mapping the *two-dimensional appearance map* $A_2(I_b, I_w) \rightarrow (\alpha, \sigma_t)$, and we show its example diagram in Fig. 2. With this mapping, one can simply determine the material properties (α, σ_t) from the contours by measuring the intensities (I_b, I_w) on the horizontal and vertical axes, per each wavelength independently, since the diagram is wavelength-independent (assuming the black and white backgrounds are both wavelength-independent reflectors). Generating the mapping was done in the inverse way: we ran Monte Carlo simulations on pairs of (α, σ_t) and acquired the corresponding (I_b, I_w), which led to the contours. The simulations used the same geometry as in Fig. 7a.

Basic properties. As shown in Fig. 2, this notion of appearance map is only valid in the upper left triangle, as anything below the diagonal would break the principle of energy conservation. The diagonal $I_b = I_w$ represents perfectly opaque materials, as they appear the same regardless of the background. The materials along the diagram's left edge $I_b = 0$ represent non-scattering absorptive materials with $\alpha = 0$, and the intensities along that edge trivially follow Eq. (1): $I_w \propto \exp(-\ell \cdot \sigma_t)$. Likewise, materials towards the top edge $I_w \rightarrow 1$ represent non-absorbing scattering materials with $\alpha \rightarrow 1$. The bottom-left corner $I_b = I_w = 0$ represents a perfectly absorbing material: $\sigma_t \rightarrow \infty, \alpha = 0$. The top-right corner $I_b = I_w = 1$ represents an ideal diffuse reflector: $\sigma_t \rightarrow \infty, \alpha = 1$. The top-left corner represents a perfectly transparent material: $\sigma_t = 0$.

Uniqueness. We conclude that there is a unique pair (α, σ_t) for each pair (I_b, I_w), hence the appearance map is *one-to-one*, from the fact that the α and σ_t-contours in Fig. 2 always intersect at exactly one particular point. Very importantly, the individual contours are also not self-intersecting, with the following exceptions. The exceptions are the α-contours collapsing at $I_b = 0, I_w = 1$, and the

in practice (left column continued)

in practice: liquids can be placed in a smooth glass container, and solid samples' rough surface can be sanded, polished, or optically smoothed by adding a thin layer of liquid and a microscope glass slide on top (Figs. 8 and 9, and Elek et al. [2021]; Pickering et al. [1993]; Prahl et al. [1993]).

σ_t-contours collapsing along the diagonal, where $\sigma_\text{t} \rightarrow \infty$. This essentially means that there is a large σ_t gradient and uncertainty for highly-absorbing materials, because after a certain threshold, the material absorbs so much light that any further difference becomes negligible. The ambiguity likewise happens for a low α value, where the detected backscattered light intensity gets too low to discriminate the exact α. Also note that the α-contours show exponential spacing, with most of the map covered by $\alpha > 0.9$.

Triangle deformations. The appearance map contours' shapes also depend on other parameters (Fig. 4), e.g., the sample thickness ℓ, the phase function parameter g, or the material refractive index n,
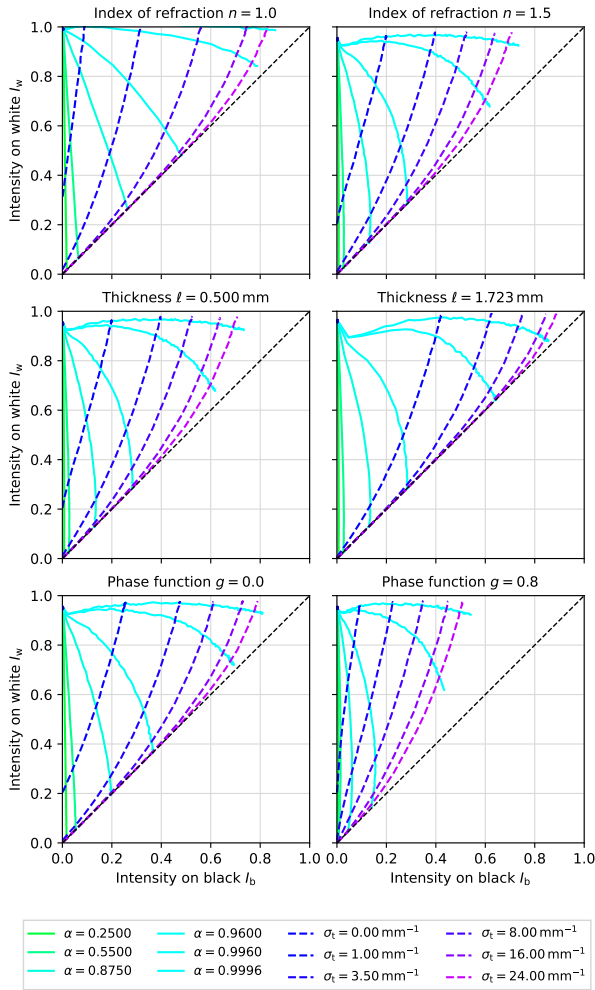

Fig. 5. The g-contours of the three-dimensional appearance map seen from two different angles. The three axes are $(I_\text{b}, I_\text{w}, \log_{10} I_\text{c})$, the contours correspond to $g \in \{0.0, 0.2, 0.4, 0.6, 0.8\}$. Note how the contours are sufficiently spaced and non-intersecting, except for their collapse to a single line at $I_\text{b} = 0$, where no scattering occurs and the phase function plays no role.

which we discuss in detail later (Secs. 3.2 and 3.3). The contours also stretch if the assumed background materials are not ideally diffuse with reflectances 0% for the black, and 100% for the white background, which is impossible to achieve in real measurements. In the measurement procedure (Sec. 4), the true background reflectances need to be modeled in the simulations to prevent a bias.

3.2 Phase function

While the two-dimensional appearance map A_2 allows estimating $(\alpha, \sigma_\text{t})$, it assumes that we have chosen a known phase function with a fixed anisotropy g. While the map's contours change with various g (Fig. 4 bottom), there are infinitely many possibilities to which g was the correct one. To pin down this unknown g, we need to include a third observable to match the number of observables to the number of unknowns. We need to account for the similarity relations [Wyman et al. 1989; Zhao et al. 2014] (see our supplemental document), such that this third observable indeed allows for the disambiguation of g, while keeping the setup still affordable.

Three-dimensional appearance map. To disambiguate the phase function anisotropy factor g, we introduce a collimated beam placed at the backside of the material sample (Fig. 7b). The combination of front-illuminating diffuse lights and a back-illuminating collimated beam gives more specific information regarding backward and forward scattering. Together, this allows "breaking out of" the assumptions in the similarity relations. Formally, we extend the previous two-dimensional map A_2 into a *three-dimensional appearance map* $A_3(I_\text{b}, I_\text{w}, I_\text{c}) \rightarrow (\alpha, \sigma_\text{t}, g)$, i.e., A_2 is a partial map of A_3, where I_c is an attenuated intensity of a collimated beam passing through the medium (Fig. 7b). We show an example diagram of A_3 in Fig. 5.

Unscattered transmission. We generate the contours similarly to the two-dimensional case. Even though we used a Monte Carlo simulation also for computing I_c, we discovered that computing it analytically using Eq. (1) has an accuracy indistinguishable from Monte Carlo. This analytical intensity, referred to as *unscattered transmittance* in literature [Prahl 2011], seemingly differs from the actual measured intensity for highly-scattering materials ($\alpha \gg 0$) because of light in-scattering to the non-zero solid angle captured by a real sensor. However, as we compare in Fig. 6, the differences

Fig. 4. Visualizing how the appearance map changes with different parameters n, ℓ, g. For example, increasing the refractive index n shrinks the map, as according to Fresnel equations, some energy is simply reflected away from the geometry (Sec. 3.3). Increasing the sample thickness ℓ shifts the σ_t-contours towards the diagonal, as more light is being absorbed in the thick sample. Changing the phase function anisotropy g shrinks the whole map, which we discuss in Sec. 3.2 and show in three dimensions in Fig. 5.
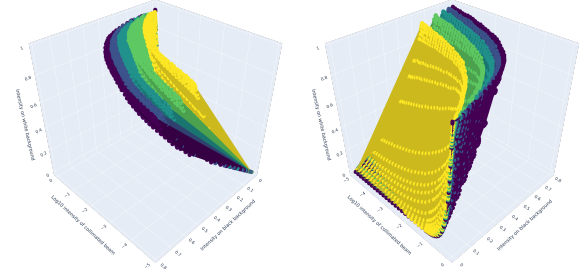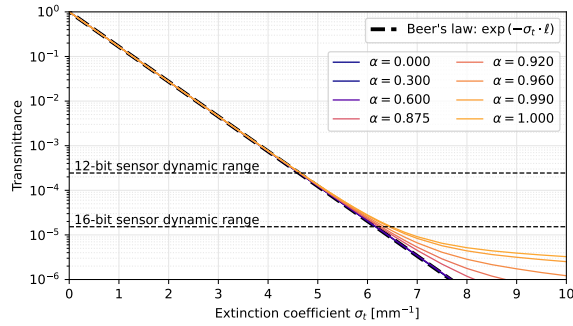
Fig. 6. Transmittance observed with a 4 mm aperture detecting a 1 mm diameter collimated beam passing through a $\ell = 1.8$ mm thick sample of varying (α, σ_t) and a dominantly forward-scattering Henyey-Greenstein phase function $g = 0.8$. Notice how the transmittance of low-scattering materials with a small α follows the Beer's law, Eq. (1). For highly scattering materials, the in-scattering from the beam is superimposed on the detected transmittance, which causes the true measured intensity to be higher than Beer's law. However, this arguably only affects measurements whose dynamic range falls below what an ideal 16-bit sensor could capture. For lower g and thinner samples, this becomes even more negligible.

are not noticeable in our setup, even for a dominantly forward-scattering phase function and a relatively thick sample, though still thin enough compared to the inverse of the extinction coefficient.

Uniqueness. We show that the parameter g can be uniquely found in the A_3-map. It follows from Fig. 5, as the g-contours are not self-intersecting and they only collapse into a single line at $I_b = 0$, corresponding to a non-scattering material with $\alpha = 0$; and at $I_w = 0$, a strongly absorbing material. It is reasonable, as one cannot measure a scattering phase function of a non-scattering material, or of a material that absorbs all the light it was illuminated with.

3.3 Index of refraction

The material's index of refraction n is the last factor that noticeably influences the appearance map. Our Monte Carlo simulations follow the Fresnel equations governing the reflection and transmission of light from the material boundary. As some incoming light gets completely reflected away, it is expected that the contours shrink towards lower values of I_w for increasing n (see Fig. 4, top). However, it may seem unintuitive that the peak of the top α-contours is not in the top-left corner, but rather the contours form a hill that raises with an increasing σ_t. This is a result of two factors: first, highly-scattering materials with a high α and σ_t back-scatter a lot of light before it even reaches the background boundary; second, our Monte Carlo simulations show that there are significant multiple reflection bounces occurring between the white background and the material boundary, which also explains why the top-left corner for $n = 1.5$ is higher than one would expect from the Fresnel equations alone.

4 MEASUREMENT

We design a simple, affordable physical setup and measurement procedure that follows from our proposed appearance maps A_2 and A_3 (Sec. 3). The end goal is to place a material sample into the

setup (Sec. 4.1), spectrally measure the intensities I_b, I_w, and I_c with a spectrometer (Sec. 4.2), and then interpolate within the appearance map to find the material's corresponding optical properties (α, σ_t, g) independently for each wavelength (Sec. 4.3).

For an affordable setup, we eliminated expensive and specialized equipment as much as possible: we used a simple pocket-size spectrometer, household illuminants, inexpensive electronics, a limited number of optical components, and a custom-made background holder. The calibrations that we require are simple and not time-consuming. In principle, our method is also usable for various samples, including liquids in glass cuvettes.

4.1 Measurement setup overview

We propose the following setup for capturing the intensities I_b, I_w, and I_c, illustrated in Fig. 7. The main components are the *material sample*, two *reflective diffuse backgrounds* (black and white), the *detector part*, and the *illuminants* (diffuse for I_w and I_b, and collimated for I_c). For detailed technical specifications of all the components mentioned in this section, please refer to the supplemental: this lists the exact supplier and part numbers for each component.

Our setup supports quickly switching between three different configurations: behind the sample, there can be a black background, a white background, or a collimated light beam (Fig. 7a, Fig. 7b). Using an inexpensive FDM 3D printer, we printed a custom sliding background holder to change between those three configurations in

(a) Diffuse illumination setup for measuring I_b, I_w

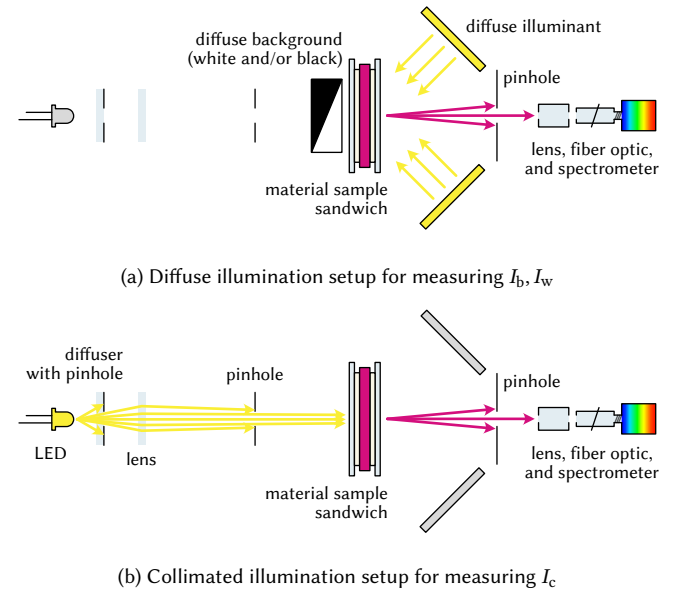(b) Collimated illumination setup for measuring I_c

Fig. 7. Diagram of our measurement setup, in two different configurations. The first configuration is used to measure I_b and I_w and is composed of the sample over a black or white diffuse background, illuminated by two diffuse illuminants. The reflected light is filtered through a pinhole, collected by a lens, and focused at an optical fiber connected to a spectrometer. The second configuration is used for the phase function estimation: the diffuse illuminants are covered, the background is removed, and the sample is backlit by a beam that was collimated from an LED.

a matter of seconds. The printable STL file for this holder is attached in the supplement. The holder, visible in Fig. 9, is split into three parts: a rectangular hole for collimated measurements, a Spectralon holder for the white background, and a rectangular space for gluing the black background.

In the end, measuring all three I_b, I_w, and I_c can be performed within a minute, with the only manual activity being sliding the background holder twice, and covering the diffuse illuminants for the collimated measurement. Since the intensities are measured spectrally with a spectrometer, the wavelength dependency is implicitly incorporated into each of the three measurements.

Material sample. The sample of the material needs to conform to certain requirements. First of all, samples of materials with a high σ_t have to be thin, otherwise we would detect zero transmitted light I_c. In our case, this was critical for the white ink that approaches $\sigma_t = 25\,\mathrm{mm}^{-1}$ in the blue spectral region. Then, as mentioned in Sec. 3.1, we assume the sample to have a smooth surface with a known refractive index n. For liquids, we recommend using a rectangular glass cuvette. Its geometry has to be modeled in the appearance map simulation to compensate for the additional layers of material. For moderately rough solid samples, such as the 3D printed samples, we used a combination of manual polishing and then layering both sides with a small quantity of index-matched immersion oil and a thin borosilicate-glass microscope cover slide (Fig. 8). The glass slides adhere to the sample with the help of the oil. The sample and the slides can then be washed with soap and reused.

Reflective diffuse backgrounds. Ideally, the backgrounds would be perfectly diffuse materials with a uniform spectral response, with a 100% reflectance (for the white one), and a 0% reflectance (for the black one) over the whole visible range. As such materials do not physically exist, we use the closest commercially available alternatives. We use a 1-inch Spectralon disk with a 99% diffuse reflectance as the white material, but one could also use a much cheaper 92% high-reflectance PTFE sheet. For the absorbing black material, we use a black flocked paper with a reflectance around 2%, which unfortunately substantially increases for $\lambda > 650\,\mathrm{nm}$. Alternatively, a black aluminum foil with a more uniform reflectance around 5% could be used, but it is noticeably more glossy than the flocked paper. Another option would be a larger cavity lined with black material, analogous to a beam dump in laser experiments.

Detector. The light that has transmitted or reflected from the sample first passes through a pinhole, which spatially filters only a small region of interest on the sample. This beam is then collected by a small camera lens mounted on an optical post and centered on the optical path. Finally, the light is focused on the entry of an optical fiber, which delivers it to the spectrometer.

Diffuse illumination. For the diffuse illumination for measuring I_b and I_w, we use two off-the-shelf 28.3×14.1 cm floodlight LED panels. They are placed at a distance of approximately 30 cm with an orientation of $\pm 45°$ with respect to the sample surface.

Collimated illumination. The illumination for measuring I_c was created by collimating a simple, bright, warm-white LED with a small beam angle. Its light first passes through a diffuser immediately
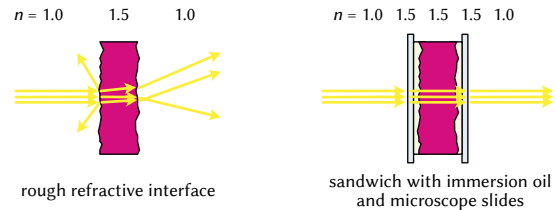

Fig. 8. Solid materials, such as 3D printed samples, still have a somewhat rough surface even after moderate polishing. Such a rough refractive interface scatters light away from the measurement geometry, which is difficult to calibrate. To compensate, we "sandwich" solid samples between two borosilicate glass microscope cover slides ($0.14 - 0.17$ mm thick), and a thin layer of index-matched immersion oil. Since all these materials have a very similar index of refraction $n \approx 1.5$, the only significant reflections and refractions occur at the smooth outer layer.

followed by a pinhole, which approximates a point light source. The light is then gathered by a lens that collimates it. We further restrict the beam diameter with a bigger pinhole in order to limit the in-scattering intensity within the sample.

4.2 Measurement procedure

Measuring the material consists of a trivial illuminants calibration, and then measuring the sample against the three backgrounds.

Calibrations. We begin by turning all illuminants on and letting them stabilize and warm up for several minutes. We proceed by measuring the three backgrounds without any sample. This directly gives us the spectrum of the collimated LED, and indirectly gives us the spectrum of the diffuse light (by dividing the measured value by the white background's uniform reflectance of 99%), and the reflectance of the black background (by dividing the measurement by the diffuse light spectrum acquired in the previous step). In case the illuminants are not perfectly stable, it is useful to perform these calibrations during the measurement of every sample again.

Sample measurement. Each sample is first prepared according to Sec. 4.1: in the case of the 3D printed inks, each sample is polished and then sandwiched with immersion oil and microscope
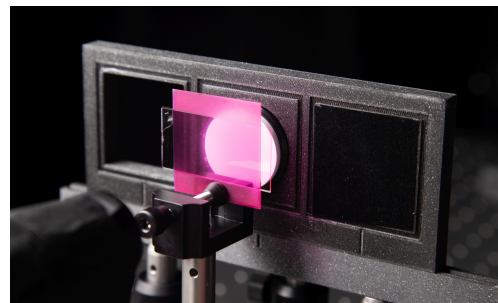

Fig. 9. Measuring a sample against a white background on our custom 3D printed background holder. Note how the appearance of the rough sample changes under the microscope slide, which optically smoothens the rough interface (see also Fig. 8).

glass slides. The sample is then inserted in a sample holder and centered within the optical path. The measurements can be done subsequently in any order by simply sliding the background holder. When measuring the collimated intensity I_c, we cover the diffuse illuminants by cardboards to make sure the spectrometer does not measure any parasitic diffuse backscattering from the sample.

Automation. While we only performed the described measurements manually, we believe that parts of the process could be easily automated, for example by using a motorized digitally controlled slider for the background holder. Such a setup could then be embedded into a self-contained device, to which a sample is inserted, and it then performs all three measurements fully automatically.

4.3 Fitting the optical properties

Once the relative intensities (I_b, I_w), or (I_b, I_w, I_c), are measured, the actual optical properties (α, σ_t), or (α, σ_t, g), respectively, have to be fitted via the appearance maps A_2, or A_3.

Datasets. We begin by constructing the appropriate appearance map for the given material refractive index n and sample thickness ℓ. As explained in Sec. 3, this is done using Monte Carlo simulations: first, we densely sample the space of possible parameters (α, σ_t, g), and then we use Monte Carlo rendering to get the corresponding monochromatic intensities (I_b, I_w, I_c) on a single-pixel sensor, according to the geometry in Fig. 7. Note that following the argumentation in Sec. 3.2 and Fig. 6, we can also compute I_c analytically from Eq. (1) without Monte Carlo. The range of parameters that we simulated was inspired by Elek et al. [2021], with low values of σ_t and high values of α simulated with finer steps due to the exponential behavior of light attenuation and scattering. We used a modified Monte Carlo implementation of Nimier-David et al. [2019] and Jakob et al. [2022], which already contained a volumetric path tracer for the diffuse measurements, but we needed to implement a simple volumetric light tracer for the collimated measurements, as such a simulation is impossible with a path tracer.

Fitting. As the dataset points do not form a regular grid (it is impossible to know the spacing of the intensities beforehand), we use a multi-dimensional linear interpolation of an irregular grid. Specifically, we used a SciPy [Virtanen et al. 2020] implementation based on constructing the interpolants through a triangulation of the dataset, and then performing linear barycentric interpolation on each triangle. We experimented with non-linear interpolations, but we were unable to find a solution that would be more robust than the linear one. Note that the interpolation is wavelength-independent (Fig. 2b), which means that spectral measurements are fitted separately per each wavelength (Fig. 2c). For example, our spectrometer had a resolution of about 1000 wavelength bins, so we performed 1000 independent interpolations to fit the spectrum. This operation is very fast, only taking a few second for a full spectrum.

Ensuring robustness. The triangulation and linear interpolation are not stable in certain regions of the appearance map, especially towards the diagonal, where all extinction coefficients collapse into a single line. Hence, naïvely interpolating the optical parameters

from the exact measured values will not be robust, since any additive noise can cause instability of the σ_t estimate along the diagonal.

Our solution is based on *random sampling*: we randomly sample a small ε-neighborhood around the actual measured values (I_b, I_w, I_c), we interpolate the optical properties for all of these samples, and then we choose the *median* result. We found that using 250 to 1000 random samples in an $\varepsilon \in [0.01, 0.04]$ neighborhood creates sufficiently robust results. Furthermore, we always combine data from at least two *different thicknesses* of the same material. For this, we simply combine the randomly sampled and interpolated values together, and choose the median of them. This also proves to be a useful consistency check, since both thicknesses should ideally result in identical optical parameter estimates. In Fig. 10, we show the results separately for each thickness, and for the combined estimate. It is also useful to visualize the σ_t estimates based on Eq. (2) to verify that the fitting was indeed robust.

5 RESULTS

We now demonstrate and validate our method on 3D printing inks from the Stratasys PolyJet material family. These inks have a wide range of properties, from almost non-scattering color primaries, to a strongly-scattering white material. This variety serves as a good stress test of our approach, and it also results in a wide gamut of appearances realizable by combining these inks together. For example, by combining a transparent yellow with a scattering white, we get an opaque yellow (Fig. 13).

In Sec. 5.1, we apply our measurement procedure to estimate the spectral optical properties of the cyan, magenta, yellow (VeroVivid), black (BlackPlus), and white (PureWhite) inks. We also discuss the repeatability of our method, its performance, sources of error, and some practical remarks on processing these samples. In Sec. 5.2, we perform an indirect validation of our results. We 3D print a color checker that contains various mixtures of the base materials, and we compare their real appearance and spectral reflectances to Monte Carlo simulations of the same virtual mixtures, based on the properties estimated in Sec. 5.1.

5.1 Measuring the base inks

We 3D printed two samples of each of the five inks with the dimensions 40×40 mm, and thicknesses 0.4 mm and 0.8 mm. As the printed objects have a rough surface, we polished all samples with sandpapers of grit sizes 400 to 2000, which also reduced the sample thicknesses by about 0.02 mm on average. In theory, one could continue polishing the samples to an almost perfect finish, but that becomes laborious with diminishing returns, so we instead sandwiched the samples in immersion oil and glass slides as in Fig. 8.

After applying our method, we obtain the results in Fig. 10, which include the spectrally resolved single-scattering albedo α, extinction coefficient σ_t, and phase function anisotropy g. The plots feature estimates from both sample thicknesses individually, and a robust combined estimate based on the procedure described in Sec. 4.3. The σ_t plots also contain estimates obtained from Beer's law (Eq. 2), which offers another consistency check.

From these measurements, we can trivially conclude that the color primaries and black are virtually non-scattering materials,
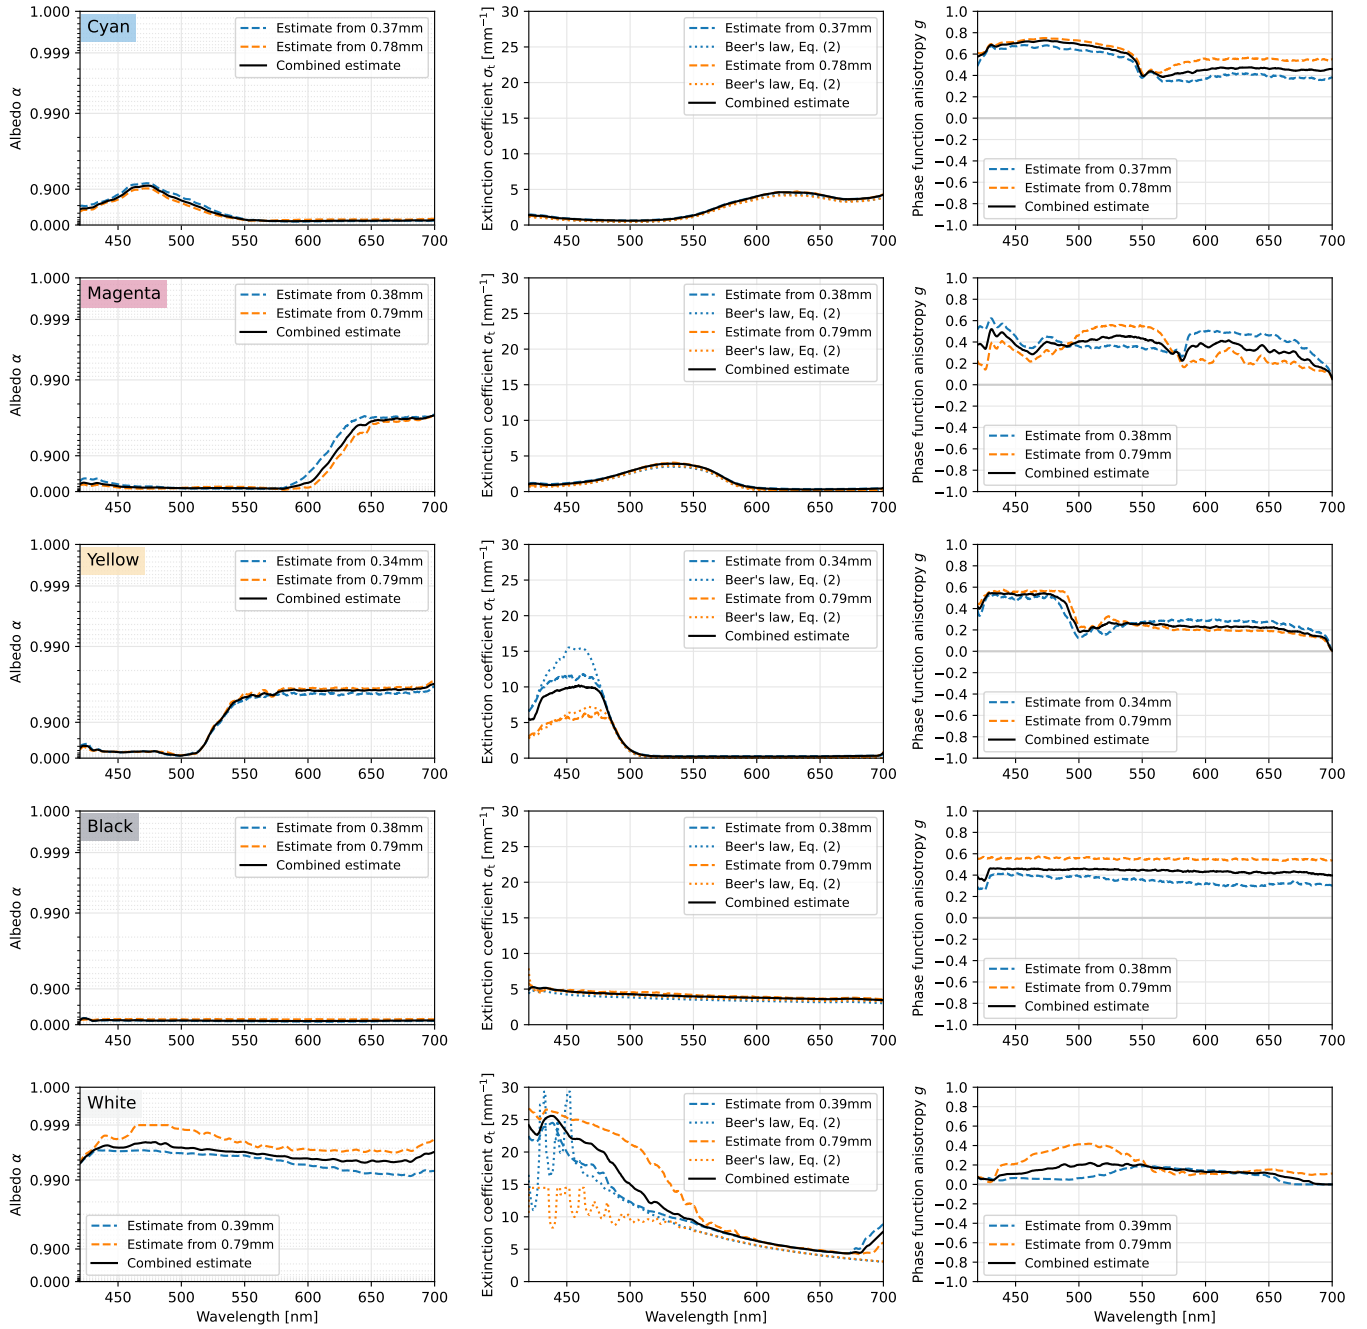
Fig. 10. Fitting the three optical properties (α, σ_t, g) of the base 3D printing materials, specifically, the Stratasys materials VeroVivid (cyan, magenta, yellow), BlackPlus, and PureWhite. The thick black curves show a combined estimate from both thicknesses to ensure robustness (Sec. 4.3), while the dashed curves are based on only one thickness. In the σ_t plots, we also visualize the corresponding Beer's law estimates as a consistency check. Note that the α-axis is in an exponential scale, because changes in low α values have a very small effect on the material appearance compared to values above 0.9, as shown by Elek et al. [2017], who numerically fitted a mapping that linearizes the effect, which we use in the axis scaling.

while the white material is strongly scattering. This is in contrast to the older family of Stratasys materials, which had significantly more scattering color primaries, as measured in RGB channels by Elek et al. [2021].

Sources of error. From Fig. 10, we can see that both the α and σ_t estimates are consistent throughout the materials, with a few exceptions. In the σ_t estimates, we can see inconsistencies in the blue wavelengths (around 400 to 500 nm). We suspect it may be caused by *fluorescence*, a wavelength-shifting effect of absorbing photons of lower wavelengths and re-emitting them at higher wavelengths. Our experiments with violet lasers indicated that wavelength shifting indeed occurs in these materials. This effect is not accounted for in the simulations, it is generally hard to measure, and none of the related methods could measure it either, so it remains as a challenging future work. This may also affect the white albedo, as the thicker sample would contain more fluorescent particles and appear more scattering than the thinner sample. We also observe some inconsistencies in the anisotropy g, however, one has to keep in mind that the estimates do not make physical sense for those wavelengths and materials, where scattering events seldomly occur, i.e., $\alpha(\lambda) \approx 0$ or $\sigma_t(\lambda) \approx 0$.

Repeatability. We assess the repeatability of our method by measuring the cyan ink several times, including different rotations of the sample. This experiment answers whether the samples are sufficiently homogeneous and rotationally invariant. From our results in Fig. 11, we conclude that the results are repeatable enough within the expectations from an affordable measurement method. For curious readers, Fig. 11 also includes a few failed measurements, where the sample was measured closer to its edge, where its thickness changes due to the manual polishing. It indicates the importance of accurately measuring the sample thickness at exactly the point that ends up being in the optical path of the measurement setup, for example using a micrometer screw gauge.

Time performance. The performance of our method can be split into two parts: precomputation of the appearance map, which is only done once per sample thickness, and the actual fitting process per material sample. The precomputation is using Monte Carlo, ideally with a very high sample count to avoid noise (for the final results, we used 640,000 samples per each datapoint). On a single machine, it takes less than a day, and it could be further improved by a significant factor by reusing the Monte Carlo path space samples for various α values. In contrast, the actual fitting process per each material sample is very fast, taking roughly two seconds for the whole spectrum of about 1000 wavelength bins.

5.2 Validation

We validate the measurements of the base inks by 3D printing a custom *color checker* and verifying its appearance. It is a 5 mm thick grid of $8 \times 8 = 64$ different mixtures of the base inks (Figs. 1 and 12). The mixtures correspond to 64 colors, each color being $15 \times 15 \times 5$ mm large and consisting of a different ratio of the base inks. Because 3D printers cannot physically mix the inks, the colors are printed in a 3D halftoned fashion: each color is made of voxels, and each voxel is randomly assigned a base ink such that the overall ink ratio is correct (similarly to mixture halftoning of Nindel et al. [2021]).

Using the optical properties measured in Sec. 5.1, we run a Monte Carlo simulation of this color checker to simulate each square's resulting reflectance. We then take the physically printed color
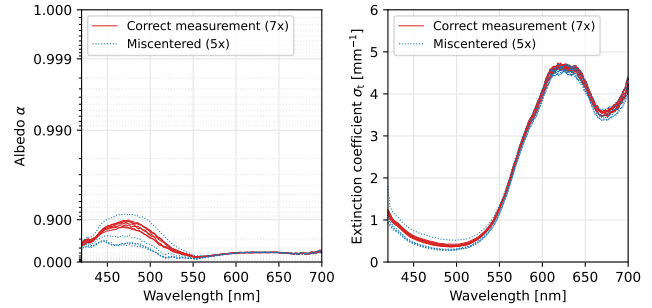
Fig. 11. Assessing the repeatability of our method by measuring the cyan ink several times, including various rotations. Poorly centered invalid measurements are also included. Refer to the text in Sec. 5.1 for discussion.

checker, and we measure each square with a real spectrometer. The goal is to compare the simulated reflectances to the real measurements, which we show for 20 of the squares in Fig. 13, visualized both with spectral curves, and with the corresponding sRGB colors.

Apart from visualizing the results obtained directly from the simulation, we also show these simulated curves multiplied by *correcting factors*, which were found individually for each square, minimizing the color error. The necessity for these factors reveals that there is a reasonable match in the wavelength dependency (in relative sense), but a mismatch in the overall intensity. The factors attempt to partially correct for geometry misalignments, e.g., because the color checker was manually polished and did not have a perfectly flat surface, and because it was manipulated by hand between every measurement. There are also additional sources of error such as the aforementioned fluorescence, which is unfortunately a wavelength-dependent effect, and the simplified phase function model, which only approximates the real light scattering. We also note that for
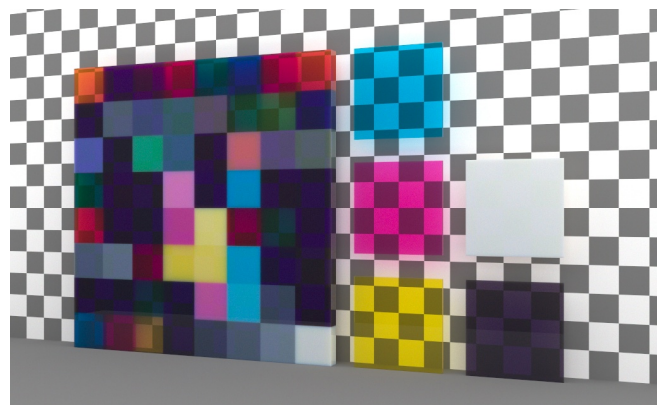

Fig. 12. Monte Carlo rendering of a demo scene containing our $120 \times 120 \times 5$ mm large color checker, and five 0.8 mm thick samples of the CMYKW base inks that the color checker is mixed from. The base inks are also located in the bottom row of the checker, leftmost four (CMYK) and the rightmost (W) square. We used a background that shows how some mixtures are significantly more translucent than others. A photograph of this color checker is in Fig. 1, and some of its spectral reflectances are shown in Fig. 13.
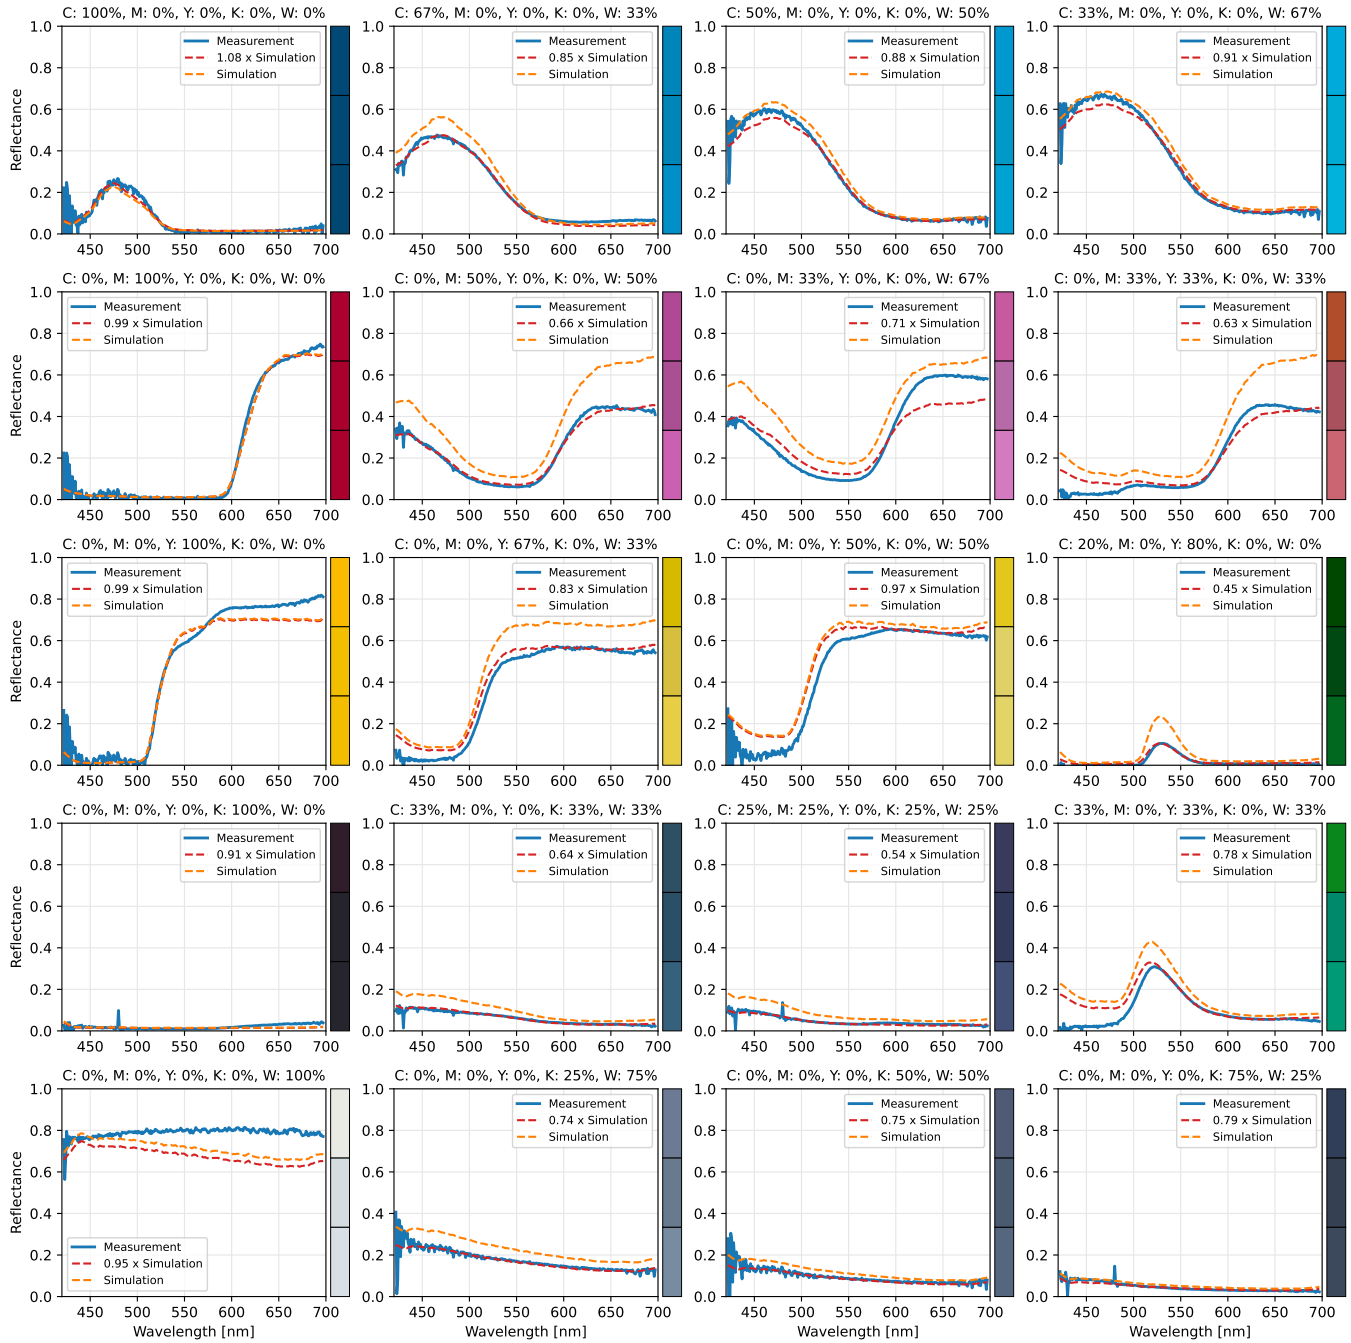
Fig. 13. A comparison between spectral reflectance measurements of a 3D printed color checker (Figs. 1 and 12) that mixes the five base inks (CMYKW) in various ratios, and a corresponding Monte Carlo simulation based on the estimated optical properties from Sec. 5.1, Fig. 10. Each plot also includes a curve with a correcting multiplicative factor (see Sec. 5.2 for discussion). The title of each plot describes the ink ratios corresponding to that measurement. The colored rectangles on the right visualize the (potentially clipped) sRGB colors corresponding to the reflectance spectra, assuming a D65 illuminant. Each top rectangle is the measured color, the middle one is the simulation with a correcting factor applied, and the bottom one is the pure simulation.

this simulation and rendering, our simulation software required

fixing a certain phase function, so we used $g = 0.4$, and the (α, σ_t) were refitted to that assumption.

Results with the base inks. The most critical result in Fig. 13 is that the predictions for the base inks themselves (left column, 100% ratios) are very accurate. Because we measured these inks on 0.4 and 0.8 mm thick samples, but the color checker is 5 mm thick, which is a significantly larger volume, the matching results are an important evidence of correctness. The most noticeable discrepancies are in the white and yellow inks, which we know to be fluorescent, which would explain why the real measurements always had a higher intensity in the longer wavelengths than the non-fluorescent simulations. The slight discrepancy in the black ink could be described by the fact that the absorbing black background used for the measurements in Sec. 5.1 is not perfect, and its reflectance increases around $\lambda > 650$ nm, causing a slight red shift.

Results with the mixtures. The highest accuracy was obtained in the mixtures of the cyan and white ink (top row). That is mainly because the cyan ink absorbs significantly above $\lambda > 550$ nm, so the error in the white ink, which is the highest in that range, is not being propagated to the final color. With all the other mixtures, we have achieved results corresponding to reasonably similar colors, although the exact color hues and saturation were mispredicted, mainly for magenta and green.

Rendering. Since the main use case of our method is Monte Carlo rendering and Monte Carlo based 3D printing optimization pipelines, we demonstrate such result in Fig. 12, a rendering of a demo scene containing the 120×120×5 mm large color checker, and 0.8 mm thick samples of the base inks. The scene has a checkered background to visualize the translucency of some of the material mixtures.

6 CONCLUSION

We have presented a simple and affordable technique for estimating spectral bulk optical properties of translucent materials in a way that is sufficient for predictive rendering purposes. Our results are directly pluggable into existing rendering pipelines, and 3D printing optimization pipelines based on Monte Carlo simulations. We demonstrated our method on a set of 3D printing materials from the Stratasys PolyJet family, and showed rendering of their various mixtures in a 3D printed object.

The key discovery of our paper was that the bulk optical properties can be estimated from only three simple point measurements, which is simpler than the existing solutions. A byproduct of this is our two-dimensional appearance map, which we believe could also be useful for artists as an intuitive tool for visualizing or editing translucent appearance, because all relevant possible combinations are encoded in the triangle in a perceptually meaningful way.

For the future, we see two additional measurement problems that would benefit from a simple and more affordable solution. First, the acquisition of bi-spectral (fluorescence) optical properties. Second, measuring the index of refraction, which is currently a necessary input to our method. Both problems currently require expensive equipment such as a monochromator or an ellipsometer, which are hardly found outside of specialized laboratories.

REFERENCES

Benedikt Bitterli, Srinath Ravichandran, Thomas Müller, Magnus Wrenninge, Jan Novák, Steve Marschner, and Wojciech Jarosz. 2018. A radiative transfer framework for non-exponential media. *ACM Transactions on Graphics* 37, 6 (2018), 225:1–225:17. https://doi.org/10.1145/3272127.3275103

Alan Brunton, Can Ates Arikan, Tejas Madan Tanksale, and Philipp Urban. 2018. 3D Printing Spatially Varying Color and Translucency. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 37, 4 (July 2018), 157:1–157:13. https://doi.org/10.1145/3197517.3201349

Alan Brunton, Can Ates Arikan, and Philipp Urban. 2015. Pushing the Limits of 3D Color Printing: Error Diffusion with Translucent Materials. *ACM Trans. Graph.* 35, 1 (Dec. 2015), 4:1–4:13. https://doi.org/10.1145/2832905

Danwu Chen and Philipp Urban. 2021. Deep learning models for optically characterizing 3D printers. *Optics Express* 29, 2 (Jan. 2021), 615–631. https://doi.org/10.1364/OE.410796 Publisher: Optica Publishing Group.

Eugene D'Eon and Geoffrey Irving. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics* 30, 4 (July 2011), 56:1–56:14. https://doi.org/10.1145/2010324.1964951

Oskar Elek, Denis Sumin, Ran Zhang, Tim Weyrich, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, and Jaroslav Křivánek. 2017. Scattering-aware texture reproduction for 3D printing. *ACM Transactions on Graphics* 36, 6 (Nov. 2017), 1–15. https://doi.org/10.1145/3130800.3130890

Oskar Elek, Ran Zhang, Denis Sumin, Karol Myszkowski, Bernd Bickel, Alexander Wilkie, Jaroslav Křivánek, and Tim Weyrich. 2021. Robust and practical measurement of volume transport parameters in solid photo-polymer materials for 3D printing. *Optics Express* 29, 5 (March 2021), 7568–7588. https://doi.org/10.1364/OE.406095 Publisher: Optical Society of America.

Claudia Emde, Robert Buras-Schnell, Arve Kylling, Bernhard Mayer, Josef Gasteiger, Ulrich Hamann, Jonas Kylling, Bettina Richter, Christian Pause, Timothy Dowling, and Luca Bugliaro. 2016. The libRadtran software package for radiative transfer calculations (version 2.0.1). *Geoscientific Model Development* 9, 5 (May 2016), 1647–1672. https://doi.org/10.5194/gmd-9-1647-2016 Publisher: Copernicus GmbH.

Roland W Fleming and Heinrich H Bülthoff. 2005. Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception (TAP)* 2, 3 (2005), 346–382. https://doi.org/10.1145/1077399.1077409 Publisher: ACM New York, NY, USA.

J. R. Frisvad, S. A. Jensen, J. S. Madsen, A. Correia, L. Yang, S. K. S. Gregersen, Y. Meuret, and P.-E. Hansen. 2020. Survey of Models for Acquiring the Optical Properties of Translucent Materials. *Computer Graphics Forum* 39, 2 (2020), 729–755. https://doi.org/10.1111/cgf.14023

Ioannis Gkioulekas, Shuang Zhao, Kavita Bala, Todd Zickler, and Anat Levin. 2013. Inverse Volume Rendering with Material Dictionaries. *ACM Trans. Graph.* 32, 6 (Nov. 2013), 162:1–162:13. https://doi.org/10.1145/2508363.2508377

Richard C. Haskell, Lars O. Svaasand, Tsong-Tseh Tsay, Ti-Chen Feng, Bruce J. Tromberg, and Matthew S. McAdams. 1994. Boundary conditions for the diffusion equation in radiative transfer. *Journal of the Optical Society of America A* 11, 10 (Oct. 1994), 2727. https://doi.org/10.1364/JOSAA.11.002727

L. G. Henyey and J. L. Greenstein. 1941. Diffuse radiation in the Galaxy. *Astrophysical Journal* 93 (Jan. 1941), 70–83. https://doi.org/10.1086/144246

Steven L. Jacques. 2013. Optical properties of biological tissues: a review. *Physics in Medicine & Biology* 58, 11 (May 2013), R37. https://doi.org/10.1088/0031-9155/58/11/R37 Publisher: IOP Publishing.

Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. 2014. A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics* 33, 4 (July 2014), 1–14. https://doi.org/10.1145/2601097.2601139

Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, and Delio Vicini. 2022. DR.JIT: a just-in-time compiler for differentiable rendering. *ACM Transactions on Graphics* 41, 4 (July 2022), 1–19. https://doi.org/10.1145/3528223.3530099

Wojciech Jarosz. 2008. *Efficient Monte Carlo Methods for Light Transport in Scattering Media.* PhD Thesis. UC San Diego.

Henrik Wann Jensen, Stephen R. Marschner, Marc Levoy, and Pat Hanrahan. 2001. A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques - SIGGRAPH '01.* ACM Press, 511–518. https://doi.org/10.1145/383259.383319

James T. Kajiya. 1986. The rendering equation. *ACM SIGGRAPH Computer Graphics* 20, 4 (Aug. 1986), 143–150. https://doi.org/10.1145/15886.15902

Paul Kubelka and Franz Munk. 1931. Ein Beitrag zur Optik der Farbanstriche. *Zeitschrift für Technische Physik* 12 (1931), 593–601.

Sven Leyre, Youri Meuret, Guy Durinck, Johan Hofkens, Geert Deconinck, and Peter Hanselaer. 2014. Estimation of the effective phase function of bulk diffusing materials with the inverse adding-doubling method. *Applied Optics* 53, 10 (April 2014), 2117. https://doi.org/10.1364/AO.53.002117

Yuki Minetomo, Hiroyuki Kubo, Takuya Funatomi, Mikio Shinya, and Yasuhiro Mukaigawa. 2018. Acquiring non-parametric scattering phase function from a single image. *Computational Visual Media* 4, 4 (Dec. 2018), 323–331. https://doi.org/10.1007/s41095-018-0122-z

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* 25, 3 (July 2006), 1003–1012. https://doi.org/10.1145/1141911.1141986

Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: a retargetable forward and inverse renderer. *ACM Trans. Graph* 38, 6 (2019), 203:1–203:17. https://doi.org/10.1145/3355089.3356498

Thomas Klaus Nindel, Tomáš Iser, Tobias Rittig, Alexander Wilkie, and Jaroslav Křivánek. 2021. A Gradient-Based Framework for 3D Print Appearance Optimization. *ACM Trans. Graph.* 40, 4 (July 2021). https://doi.org/10.1145/3450626.3459844 Place: New York, NY, USA Publisher: Association for Computing Machinery.

James H Nobbs. 1985. Kubelka—Munk Theory and the Prediction of Reflectance. *Review of Progress in Coloration and Related Topics* 15, 1 (1985), 66–75. https://doi.org/10.1111/j.1478-4408.1985.tb03737.x

Jan Novák, Iliyan Georgiev, Johannes Hanika, Jaroslav Křivánek, and Wojciech Jarosz. 2018. Monte Carlo Methods for Physically Based Volume Rendering. In *ACM SIGGRAPH 2018 Courses (SIGGRAPH '18).* ACM, New York, NY, USA, 14:1–14:1. https://doi.org/10.1145/3214834.3214880

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* 32, 4 (July 2013), 1–12. https://doi.org/10.1145/2461912.2461974

John W. Pickering, Christian J. M. Moes, H. J. C. M. Sterenborg, Scott A. Prahl, and Martin J. C. van Gemert. 1992. Two integrating spheres with an intervening scattering sample. *JOSA A* 9, 4 (April 1992), 621–631. https://doi.org/10.1364/JOSAA.9.000621 Publisher: Optical Society of America.

John W. Pickering, Scott A. Prahl, Niek van Wieringen, Johan F. Beek, Henricus J. C. M. Sterenborg, and Martin J. C. van Gemert. 1993. Double-integrating-sphere system for measuring the optical properties of tissue. *Applied Optics* 32, 4 (Feb. 1993), 399–410. https://doi.org/10.1364/AO.32.000399 Publisher: Optical Society of America.

Scott Prahl. 2011. Everything I think you should know about Inverse Adding-Doubling. (March 2011), 74. https://omlc.org/software/iad/manual.pdf

Scott A. Prahl. 1995. The Adding-Doubling Method. In *Optical-Thermal Response of Laser-Irradiated Tissue*, Ashley J. Welch and Martin J. C. Van Gemert (Eds.). Springer US, Boston, MA, 101–129. https://doi.org/10.1007/978-1-4757-6092-7_5

Scott A. Prahl, Martin J. C. van Gemert, and Ashley J. Welch. 1993. Determining the optical properties of turbid media by using the adding–doubling method. *Applied Optics* 32, 4 (Feb. 1993), 559. https://doi.org/10.1364/AO.32.000559

Alina Pranovich, Sasan Gooran, Jeppe Revall Frisvad, and Daniel Nyström. 2021. Optical properties and appearance of fused deposition modelling filaments. *Advances in Printing and Media Technology - Printing in the Digital Era* (2021), 134–140. Publisher: International Association of Research Organizations for the Information, Media and Graphic Arts Industries.

Tobias Rittig, Denis Sumin, Vahid Babaei, Piotr Didyk, Alexey Voloboy, Alexander Wilkie, Bernd Bickel, Karol Myszkowski, Tim Weyrich, and Jaroslav Křivánek. 2021. Neural Acceleration of Scattering-Aware Color 3D Printing. *Computer Graphics Forum* 40, 2 (2021), 205–219. https://doi.org/10.1111/cgf.142626

Christopher Sandoval and Arnold D. Kim. 2015. Extending generalized Kubelka–Munk to three-dimensional radiative transfer. *Applied Optics* 54, 23 (Aug. 2015), 7045. https://doi.org/10.1364/AO.54.007045

Sven Schröder, Alexander von Finck, and Angela Duparré. 2015. Standardization of light scattering measurements. *Advanced Optical Technologies* 4, 5-6 (Oct. 2015), 361–375. https://doi.org/10.1515/aot-2015-0041 Publisher: De Gruyter.

Denis Sumin, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Křivánek, Karol Myszkowski, and Tim Weyrich. 2019. Geometry-aware scattering compensation for 3D printing. *ACM Transactions on Graphics* 38, 4 (2019), 111:1–111:14. https://doi.org/10.1145/3306346.3322992

Valery V. Tuchin. 1993. Laser Light Scattering in Biomedical Diagnostics and Therapy. *Journal of Laser Applications* 5, 2 (Oct. 1993), 43–60. https://doi.org/10.2351/1.4745330 Publisher: Laser Institute of America.

Philipp Urban, Tejas Madan Tanksale, Alan Brunton, Bui Minh Vu, and Shigeki Nakauchi. 2019. Redefining A in RGBA: Towards a Standard for Graphical 3D Printing. *ACM Trans. Graph.* 38, 3 (June 2019), 21:1–21:14. https://doi.org/10.1145/3319910

H. C. van de Hulst. 1980. *Multiple light scattering: tables, formulas, and applications.* Vol. 1. Academic Press, New York.

Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. https://doi.org/10.1038/s41592-019-0686-2

Tim Weyrich, Wojciech Matusik, Hanspeter Pfister, Bernd Bickel, Craig Donner, Chien Tu, Janet McAndless, Jinho Lee, Addy Ngan, Henrik Wann Jensen, and Markus Gross. 2006. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics* 25, 3 (July 2006), 1013–1024. https://doi.org/10.1145/1141911.1141987

Douglas R Wyman, Michael S Patterson, and Brian C Wilson. 1989. Similarity relations for anisotropic scattering in Monte Carlo simulations of deeply penetrating neutral particles. *J. Comput. Phys.* 81, 1 (March 1989), 137–150. https://doi.org/10.1016/0021-9991(89)90067-3

Bei Xiao, Bruce Walter, Ioannis Gkioulekas, Todd Zickler, Edward Adelson, and Kavita Bala. 2014. Looking against the light: how perception of translucency depends on lighting direction. *Journal of Vision* 14, 3 (March 2014), 17. https://doi.org/10.1167/14.3.17

Bei Xiao, Shuang Zhao, Ioannis Gkioulekas, Wenyan Bi, and Kavita Bala. 2020. Effect of geometric sharpness on translucent material perception. *Journal of vision* 20, 7 (2020), 10–10. https://doi.org/10.1167/jov.20.7.10 Publisher: The Association for Research in Vision and Ophthalmology.

Li Yang and Roger David Hersch. 2008. Kubelka-Munk Model for Imperfectly Diffuse Light Distribution in Paper. *Journal of Imaging Science and Technology* 52 (May 2008). https://doi.org/10.2352/J.ImagingSci.Technol.(2008)52:3(030201)

Tizian Zeltner and Wenzel Jakob. 2018. The layer laboratory: a calculus for additive and subtractive composition of anisotropic surface reflectance. *ACM Transactions on Graphics* 37, 4 (Aug. 2018), 1–14. https://doi.org/10.1145/3197517.3201321

Shuang Zhao, Ravi Ramamoorthi, and Kavita Bala. 2014. High-order similarity relations in radiative transfer. *ACM Transactions on Graphics* 33, 4 (July 2014), 104:1–104:12. https://doi.org/10.1145/2601097.2601104