

Fabricating Translucent Materials using Continuous Pigment Mixtures

Marios Papas^{1,2}
Philip Jackson³

Christian Regg^{1,2}
Wojciech Matusik^{1,4}

Wojciech Jarosz¹
Steve Marschner⁵

Bernd Bickel¹
Markus Gross^{1,2}

¹Disney Research Zürich

²ETH Zürich

³Walt Disney Imagineering

⁴MIT CSAIL

⁵Cornell University



Figure 1: A side by side comparison of real translucent materials (right) next to their silicone replicas (left) fabricated using our method.

Abstract

We present a method for practical physical reproduction and design of homogeneous materials with desired subsurface scattering. Our process uses a collection of different pigments that can be suspended in a clear base material. Our goal is to determine pigment concentrations that best reproduce the appearance and subsurface scattering of a given target material. In order to achieve this task we first fabricate a collection of material samples composed of known mixtures of the available pigments with the base material. We then acquire their reflectance profiles using a custom-built measurement device. We use the same device to measure the reflectance profile of a target material. Based on the database of mappings from pigment concentrations to reflectance profiles, we use an optimization process to compute the concentration of pigments to best replicate the target material appearance. We demonstrate the practicality of our method by reproducing a variety of different translucent materials. We also present a tool that allows the user to explore the range of achievable appearances for a given set of pigments.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

Keywords: subsurface scattering, material design, fabrication

Links:  DL  PDF

ACM Reference Format

Papas, M., Regg, C., Jarosz, W., Bickel, B., Jackson, P., Matusik, W., Marschner, S., Gross, M. 2013. Fabricating Translucent Materials using Continuous Pigment Mixtures. *ACM Trans. Graph.* 32, 4, Article 146 (July 2013), 12 pages. DOI = 10.1145/2461912.2461974 <http://doi.acm.org/10.1145/2461912.2461974>.

Copyright Notice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Copyright © ACM 0730-0301/13/07-ART146 \$15.00.
DOI: <http://doi.acm.org/10.1145/2461912.2461974>

1 Introduction

Most of the materials in our manmade environment are colored by dyes, pigments, or other colorants suspended in a scattering medium. Paints, plastics, papers, textiles, stained glass, ceramic glazes, candy—nearly all surfaces that are not metallic or completely transparent fall under this description. Many natural materials are also well approximated as colored scattering media—skin, leaves, flowers, foods—in which the colorants are naturally occurring. Because of the ubiquity of colored scattering materials, the technology of predicting and controlling their color is very mature, as epitomized by systems that automatically mix paints to match a given sample.

But color is not the only attribute of a colored scattering medium; pigmented media are, by their very nature, translucent. Some materials are so dense (wall paint, for instance) that the translucency can be ignored at macroscopic scales, but for others it is subtly (“opaque” plastic, skin) or obviously (translucent plastics, stained glass) part of the appearance. Translucency is a more complex phenomenon than diffuse color, and currently the appearance of such materials is normally controlled by trial and error.

The goal of this paper is to create the fundamental technology of controlling translucency as precisely as color can already be controlled, including accurately predicting the appearance of translucent materials, automatically matching existing real or virtual materials, and synthetically adjusting mixtures with feedback about translucent appearance. Over the last dozen years, the field of computer graphics has developed an increasingly mature understanding of how to simulate [Jensen et al. 2001; dEon and Irving 2011], measure [Hawkins et al. 2005; Dorsey et al. 2008; Weyrich et al. 2009a], and manipulate [Xu et al. 2007; Song et al. 2009] the appearance of translucent objects in rendered scenes, and we believe this technology is becoming mature enough to be applied to the more demanding application of manipulating materials in the real world.

The resulting methods are directly useful in design applications involving pigmented translucent materials, such as industrial design

of consumer products where appearance is important, or design of dental materials, prostheses, or animatronics that should match a given person's body—appearance in all these examples is currently matched manually by trial and error. Together with the ability to accurately render translucent materials, this will make translucent materials an integral part of realistic product pre-visualization, letting the user see on screen the exact appearance the manufactured product will have.

In addition, as multi-material 3D printers become sufficiently capable, the same fundamental technology of controlling the appearance of scattering materials can ultimately enable 3D printing of nearly arbitrary materials with precise control over appearance. Our work complements existing work on spatial mixtures of 3D printed materials [Dong et al. 2010; Hasan et al. 2010] by examining how to control the properties of individual materials. Printers that can control both material properties and spatial arrangement precisely will enable *physical* appearance prototyping as well as the direct manufacture of products whose appearance is important.

In this research we have developed a system to generate a pigment mixture to match the subsurface scattering of a homogeneous target material. Our system is analogous to computerized paint matching systems that measure the spectral reflectance of a target material and then reproduce it by combining pigments in a scattering or clear base material. We use a similar concept, but instead of treating the material as having only diffuse surface reflection, we model it as a volume with corresponding subsurface scattering effects. This approach extends our system beyond just color matching to cover translucent materials with significant subsurface scattering in a way that can be easily integrated into existing systems. Figure 1 shows a sample of different materials that can be reproduced using our system. Our methods are applicable to any process where pigments are used to control the appearance of a material.

Our process for matching subsurface scattering of a given target material has the following steps. First, we manufacture a collection of samples with different concentrations of pigments mixed into a base material. We use a custom-built measurement system (Section 5) to acquire multi-spectral reflection profiles for each material sample. This allows us to establish a mapping from pigment concentrations to reflection profiles that capture subsurface scattering properties. Using the same measurement device we acquire multi-spectral reflection profiles for a given target material. We design an iterative optimization procedure that computes the pigment concentrations required to reproduce the target material accurately (Section 6). We demonstrate the whole process by reproducing a number of organic and inorganic materials (Sections 9 and 10). In addition, we describe a tool that enables interactive exploration of the range of translucent appearances that is achievable with a given set of pigments.

2 Related Work

Computerized Color Matching Systems. Many commercial systems for color matching have been developed, and they have been successfully used in many industries (e.g., house paint, automotive, etc.). They typically use a spectrophotometer to determine material reflection as a function of wavelength. Then they either determine the closest material that is already in the database or they determine the combination of base pigments to obtain the best match. The most common systems include Sher-Color™ by Sherwin-Williams [Sherman and Simone 1989] and ChromaVision® by DuPont [Kelly 1987]. While some of these systems try to match sheen as well, the translucency of a material is never considered.

BSSRDF Models and Measurement. The first practical model for representing subsurface scattering, in computer graphics, was

developed by Jensen et al. [2001]. They also used an image-based measurement system to estimate model parameters for a number of different materials. Hawkins et al. [2005] developed a method to estimate the scattering phase function of smoke. Donner et al. [2009] proposed a homogeneous BSSRDF model that accounts for both single and multiple scattering. However, their model is based on simulations only and no extensive validation based on measurements has been performed. Similarly, d'Eon and Irving [2011] introduced a BSSRDF model that accurately decouples single and multiple scattering. Their model works well for highly absorbing materials and very thin layers. Munoz et al. [2011] estimate an approximation of the reflectance characteristics of translucent materials, using a single photograph as input. Weyrich et al. [2006] used a contact-based measurement device to estimate subsurface scattering parameters of human skin. We use an improved version of this device to estimate multiple scattering in arbitrary materials. Narasimhan et al. [2006] developed a method to estimate scattering parameters by diluting low concentrations of a material in a clear solvent. They did not address the problem of computing mixtures to match other materials. Although our method is limited to moderately to highly scattering materials, it is practical and efficient and it applies to solids as well as liquids.

Editing. There has been some recent work on authoring and editing the appearance of translucent materials [Xu et al. 2007; Song et al. 2009]. All of this work has however focused on editing and prescribing the subsurface scattering appearance in simulation, where inaccuracies in the BSSRDF model can be completely ignored, and not for physical fabrication.

Appearance Fabrication. Computer graphics researchers have recently addressed the problem of using optimization methods to derive materials with desired appearance properties [Fuchs et al. 2008; Weyrich et al. 2009b; Matusik et al. 2009; Hullin et al. 2011]. In particular, two recent methods have addressed computational design and fabrication of materials with desired subsurface scattering [Hasan et al. 2010; Dong et al. 2010]. These approaches express the output as a material composed of discrete voxels or layers with known subsurface scattering properties. However, the materials for the layers/voxels are fixed. These methods can approximate materials with heterogeneous subsurface scattering as well. However, they also have some disadvantages, in particular, materials with homogeneous subsurface scattering cannot be accurately reproduced since these methods rely on either voxel dithering or layering of discrete materials. Our proposed approach explores an orthogonal research direction that can be naturally combined with the layered-material methods. Instead of manufacturing a compound multi-layer material as in previous approaches, we fabricate materials using continuous pigment mixtures within a single homogeneous layer. Therefore, by default all our materials have homogeneous subsurface scattering properties.

3 Method Overview

At a high level, the goal of our system is to provide a recipe of how to mix pigments with a base material (clear silicone in our system) in order to reproduce a measured target material. Our process can be divided into five main stages (illustrated in Figure 2): database and target measurement, pigment parameter estimation, concentration estimation, and, finally, fabrication.

In a pre-process we fabricate and measure various mixtures of the available pigments with base material (Database Measurement). This collection of measurements becomes our appearance database which we then use to estimate a *global* set of pigment parameters for




Figure 2: Overview of our system.

predicting the subsurface scattering appearance of a silicone mixture (Pigment Parameter Estimation).

To create a silicone replica, we first measure the diffuse reflectance and translucency of a target material (Target Measurement). We then perform an optimization that will estimate pigment concentrations for the target (Mixture Optimization), assuming global pigment parameters (Pigment Parameter Estimation). This process is refined by an iterative local pigment parameter estimation step, with the goal of enabling better results in regions of the domain where the forward model, and consequently the global set of pigment parameters, is no longer a good approximation.

In Section 4 we describe our chosen forward appearance model, which balances our needs for both accuracy and performance. We use a single, custom-built measurement device, which we describe in Section 5, to perform both the database and target measurements. In Section 6, we describe both the pigment parameter estimation (Section 6.3) and mixture optimization stages (Section 6.4) used for computing a pigment concentration recipe that will match a measured target appearance. In Section 7 we describe an improvement to this method (Local Pigment Parameter Estimation) which accounts for inaccuracies in the forward model approximation. With this system in place, we enable users not only to replicate target materials, but also to edit and pre-visualize (Section 8) the appearance of a desired material before fabrication. The output of the appearance matching and appearance editing process is a recipe, which we fabricate by mixing pigments with silicone (Section 9). We compare the fabricated results of our system against a set of target translucent materials in Section 10.

4 The Forward Appearance Model

Before we can fabricate a replica, we need a computational model to predict the appearance of scattering materials which we can later invert to derive the scattering parameters necessary for fabrication. According to the theory of scattering media, a homogeneous material can be described by a phase function and two parameters, the absorption coefficient σ_a and the scattering coefficient σ_s , or equivalently by the extinction coefficient $\sigma_t = \sigma_a + \sigma_s$ and albedo $\alpha = \sigma_s/\sigma_t$. If any two of these parameters are known the other two can be computed. In highly scattering materials, the flow of light can be well modeled with a diffusion equation, which leads to approximate analytical models that are useful to describe translucent materials. In such materials, one can replace the scattering coefficient with the reduced scattering coefficient σ'_s , and then treat the phase function as isotropic (see Jensen and Buhler [2002] and references therein

for details). In this paper we will always use the reduced parameters σ'_s , σ'_t , and α' but will omit the customary primes for notational convenience.

Scattering Profile. From the many available diffusion models, we chose quantized diffusion [dEon and Irving 2011] to analytically express the subsurface reflectance profile since it remains relatively accurate even in the case of moderately absorbing materials. For the purposes of our task, the chosen diffusion model simply returns an analytic reflectance profile between two surface points x and y :

$$\frac{dL_r^\lambda(x)}{d\Phi_i^\lambda(y)} = R_d^{QD}(\alpha^\lambda, \sigma_t^\lambda, d; r), \quad (1)$$

as a function of their distance $r = \|x - y\|$, thickness d , and the reduced albedo α and reduced extinction coefficient σ_t per wavelength band λ . This also depends on η . In this paper we always use a measured silicone index of refraction value of 1.41. The internal details of the diffusion model are largely unimportant for the rest of our pipeline, hence R_d can be treated as a black box.

Diffuse Reflectance. We also require a computational model for the diffuse reflectance ρ . We evaluated the accuracy of various analytical and numerical models for diffuse reflectance, but found that none matched Monte Carlo simulations well enough (see the supplemental document for our detailed analysis). We therefore created a dense tabulation of diffuse reflectance values from brute force Monte Carlo [Wang et al. 1995] simulations:

$$\rho(\alpha, \sigma_t, d) = 2\pi \int_0^\infty \frac{R_d^{MC}(\alpha, \sigma_t, d; r)}{\pi} r dr. \quad (2)$$

For our tabulation we parametrize ρ according to reduced albedo (α) and optical thickness ($\sigma_t d$, where d is the measured thickness of the sample) and store the values in a 2D table for lookup and interpolation.

Please note that our reflectance term assumes that the multiple scattering process (including internal reflection) is isotropic, and that ρ is the reflectance for incoming light inside the surface to outgoing light inside the surface.

5 Measurement Setup

Using our fabrication process, we would like to reproduce a wide range of organic and inorganic materials that exhibit subsurface scattering. To allow measuring of target materials, that cannot be diluted or modified in other ways, we use a non-invasive contact-based measurement device inspired by, but extended from, the design proposed by Weyrich et al. [2006]. We use this device to measure both the database silicone samples as well as the target materials we want to replicate.

Our setup incorporates a single housing to perform two distinct types of measurements. The first type of measurement is used to extract the **diffuse reflectance**, $\bar{\rho}^\lambda$, of the target material, and the second type is used to measure the bulk **scattering profile**, \bar{R}_d^λ , of the target material. In Figure 3 we illustrate a cross-section view of our enclosed device, which we fabricated using a 3D printer. The device has a height of 11 cm and a diameter of 25 cm. The sample is placed at the bottom of the measurement device, in contact with the base opening circle (A), and we place a monochromatic QImaging Retiga-2000R camera into the circular opening (B) at the top of the device. This non-invasive method only requires that a small circular patch of the target material is fairly flat and can be




Figure 3: Top: An illustration of the setup used to measure the appearance of a target sample. The measurement device performs a diffuse reflectance measurement and as well as a bulk scattering profile measurement without any moving parts. Bottom Left: An example greyscale HDR image of a diffusion profile measurement. Bottom Right: An example greyscale HDR image of a diffuse reflectance measurement.

brought in contact with our measurement device for a few seconds or minutes, depending on how dark the material is. All measurements are performed for five distinct wavelength-bands (indexed by λ) using two identical sets of five color LEDs, one for reflectance and one for profiles. This allows us to capture reflectance and profile measurements in sequence without having to move the sample or reconfigure the measurement device.

The camera used for our measurements was an actively cooled, monochromatic QImaging Retiga-2000R which captures 12-bit RAW images. After extensive testing we found that these RAW files have a linear radiance response and variations due to vignetting are negligible for the regions we measure. To create HDR images we sum up the unclipped pixels from each image and divide by the total exposure time. This approach effectively applies low weights to photographs with low exposure times, which is desirable since extremely low exposure photographs are prone to noise. For these measurements, an absolute radiance value is not required.

Diffuse Reflectance. The diffuse reflectance measurement is performed using an array of five fiber optic cables located in the top right edge of the device (C), aiming towards the center of the sample (A) as seen in Figure 3. The other ends of these fiber optic cables are mounted to an LED holder with five color LEDs. We place a 5 mm diameter opal glass diffuser at the outgoing end of the fiber optic cable to ensure a constant angular intensity distribution on the sample. Figure 3 shows a greyscale HDR diffuse reflectance capture (bottom-right) with one LED turned on.

In reality, we cannot easily observe $\bar{\rho}^\lambda$ directly. Instead, we observe values of reflected radiance, \bar{L}_s^λ , off the sample. However, we design the geometric configuration so that we can derive $\bar{\rho}^\lambda$ from our observed measurement under some reasonable assumptions. In particular, if we assume that the material is a homogeneous medium with a smooth Fresnel boundary and that single scattering is negligible, we have the following general expression for the observed

radiance:

$$\bar{L}_s^\lambda(\bar{\omega}_o) = 2\pi \int_0^\infty \int_\Omega \bar{L}_i^\lambda(\bar{\omega}_i) (\bar{\omega}_i \cdot \bar{\mathbf{n}}) \frac{S_d^\lambda(r, \bar{\omega}_i, \bar{\omega}_o)}{\pi} r d\bar{\omega}_i dr, \quad (3)$$

where $\bar{\omega}_i$ and $\bar{\omega}_o$ are the incident and outgoing directions respectively, $\bar{\mathbf{n}}$ is the surface normal, and $S_d^\lambda(r, \bar{\omega}_i, \bar{\omega}_o) = F_i(\bar{\omega}_o) \bar{R}_d^\lambda(r) F_i(\bar{\omega}_i)$ is the BSSRDF with Fresnel reshaping.

In our case, we observe the sample from directly above, $\bar{\omega}_o = 0^\circ$, and we illuminate the sample from a single direction $\bar{\omega}_i$ at 45° to normal incidence (to avoid imaging direct specular reflection of the light). We measure the per-wavelength-band radiance of the sample, $\bar{L}_s^\lambda(0^\circ)$, by averaging an approximately 1 cm^2 square patch centered at (A). Assuming that the incident direction is constant at contributing regions, the observed radiance therefore simplifies to:

$$\bar{L}_s^\lambda(0^\circ) \approx \frac{I_i^\lambda}{t^2} \cos(45^\circ) F_i(0^\circ) F_i(45^\circ) 2\pi \int_0^\infty \frac{\bar{R}_d^\lambda(r)}{\pi} r dr. \quad (4)$$

Where I_i^λ is the intensity of the light source, located at distance t . To estimate the intensity we also perform a one-time measurement, $\bar{L}_c^\lambda(0^\circ)$, for a grey diffuse calibration target placed at (A). By assuming that the calibration target is perfectly Lambertian we have:

$$I_i^\lambda = \frac{t^2 \bar{L}_c^\lambda(0^\circ)}{\bar{\rho}_c^\lambda(0^\circ, 45^\circ) \cos(45^\circ)}, \quad (5)$$

where $\bar{\rho}_c^\lambda(0^\circ, 45^\circ)$ is the reported reflectance of the diffuse calibration target. Though we allocate a small warmup time for stability within a measurement, we cannot assume that the intensity of the LEDs will be consistent across measurements over multiple days. To correct for fluctuations in LED intensity, we additionally use four small reflectance standard patches (D) which are always visible in our measurements. We then scale I_i^λ by the average radiance ratio of these four patches in the material sample measurement and the calibration target measurement.

By combining Equations (4) and (5) and rearranging terms we obtain a simple expression for use as our model reflectance $\bar{\rho}$ (Equation 2):

$$\bar{\rho}^\lambda = 2\pi \int_0^\infty \frac{\bar{R}_d^\lambda(r)}{\pi} r dr \approx \frac{\bar{\rho}_c^\lambda(0^\circ, 45^\circ) \bar{L}_s^\lambda(0^\circ)}{F_i(0^\circ) F_i(45^\circ) \bar{L}_c^\lambda(0^\circ)}. \quad (6)$$

Bulk Scattering Profile. Our contact measurement device also contains a second set of five LEDs and fiber optic cables originating from the left side (E).¹ These are in contact with the material sample and illuminate it at a location (F) which is not directly visible by the camera. The horizontal distance between the center of the 1 mm diameter fiber optic cable end and the first measurable location on the sample is 0.9 mm.

Light from the fiber optic cables propagates through the scattering material and into the field of view of the camera. We show an example greyscale HDR capture with one LED turned on in Figure 3 (bottom-left). We capture five such HDR images, one for each LED. We extract the horizontal scanline, vertically aligned with the currently active light source, and use this as the bulk, per-wavelength-band, diffusion profile measurement, \bar{R}_d^λ . The measured profile has

¹We measured the spectral distributions of all LEDs using a Photo Research SpectraScan PR 730 spectrometer to ensure there is negligible variability between corresponding LEDs across the two sets.

an inherent arbitrary scale factor, due to the unknown intensity of the LEDs. Additionally by assuming that for our profile measurements the Fresnel transmission terms are not spatially varying, then the Fresnel terms can also be folded into this scale factor. We will later on show how knowledge of this unknown factor is not necessary for our purposes.

6 Mapping Measurements to Pigments

Once a target material is measured, our goal is to reproduce the material by mixing one or more pigments into the base material. To compute a recipe, or a vector containing a concentration for each available pigment, our system must be able to predict the appearance (both the reflectance and the scattering profile) that will result from any given set of concentrations. This mapping from recipe to appearance is then inverted in an optimization process to determine the best recipe for matching the appearance of a target material.

Our key task is to derive a mapping from concentrations c_p for $p = 1, \dots, n_p$ to the observed appearance characteristics ρ^λ and R_d^λ for $\lambda = 1, \dots, n_\lambda$. We build this model from a number of example materials that are fabricated and measured ahead of time, forming what we call a *database* of appearance measurements (Section 6.1). Our approach leverages the approximate forward model from the previous section to intelligently interpolate among the database samples and to provide starting points for nonlinear optimization of pigment concentrations.

Given the measurements of \bar{R}_d^λ and $\bar{\rho}^\lambda$ for any particular material, the parameters of that material (α, σ_t) can be estimated by fitting to the forward model, seeking a match in the (relative) diffusion profile and the (absolute) total reflectance:

$$(\alpha^\lambda, \sigma_t^\lambda) = \underset{\alpha, \sigma_t}{\operatorname{argmin}} F(R_d(\alpha, \sigma_t, d), \rho(\alpha, \sigma_t, d), \bar{R}_d^\lambda, \bar{\rho}^\lambda) \quad (7)$$

$$F(R_d, \rho, R'_d, \rho') = \left[E(R_d, R'_d) + (\rho - \rho')^2 \right]$$

where E is a profile difference measure we will describe in Section 6.2.

In principle, this fitting approach could be used to determine the material parameters of each of the training examples, from which the properties of each individual pigment could be derived. The parameters of a target material could then be determined in a second fit and used to find the pigment concentrations required. However, the diffusion approximation is not accurate enough to directly achieve a visual match using this simple approach, i.e. the model parameters are not linearly related with pigment concentrations. In the following subsections we describe our approach for finding a set of global pigment parameters, linearly related to pigment concentrations, that best predicts appearance.

6.1 Measurement Database Selection

The first step in estimating the pigment parameters is designing the input set, which we call the **Measurement Database**. Our goal is to use the methodology and machinery described in Section 5 to acquire the per-wavelength-band diffuse reflectance measurement, $\bar{\rho}^\lambda$, the per-wavelength-band bulk scattering profile, \bar{R}_d^λ , for a set of samples with known pigment concentrations, and then estimate α^λ and σ_t^λ for each pigment and for the base silicone using the methods of the previous section.

The main challenge, when the input set of pigments contains highly absorbing entries, is the design of database samples that will not

violate the assumptions of diffusion theory which will be used to estimate their parameters. The two main assumptions for diffusion theory to hold are $\sigma_a^\lambda \ll \sigma_t^\lambda$ and that the multiple anisotropic scattering in the material can be well approximated using an approximately equivalent isotropic material with reduced scattering parameters.

We use a total of 6 pigments. With the exception of white and yellow pigments, the remaining (red, green, blue, and black) pigments are highly absorbing. We created for each pigment a set of database entries, which we call a dilution set. Each such dilution set consists of fabricated silicone samples with varying concentrations of that pigment, always mixed with white pigment at a concentration of 0.05%. For each dilution set, we used the minimum amount of pigment such that both the color and the profile differentiate enough from the appearance of white pigment at 0.05% concentration. The maximum concentration was chosen such that the smallest measurable profile is at least 3mm long. We also include a dilution set with varying concentrations of only white pigment. This ensures that the scattering parameters of white pigment can be distinguished from the parameters of the base material. The physical size of our database samples is $10 \times 10 \times (2-4)$ cm, achieving a minimum optical thickness of about 10 along their minimum dimension. A set of plots showing the concentrations used for our dilution sets, along with the forward model fits, can be found in our supplemental document.

6.2 Fitting the forward model to a single material

The simplest fitting operation is to fit the forward model to the measured appearance data \bar{R}_d^λ and $\bar{\rho}^\lambda$ for a single wavelength band of a single material.

Initial Guess. To find a starting point for the optimization, we use our diffuse reflectance lookup table to find α such that $\rho(\alpha, \sigma_t, d) = \bar{\rho}^\lambda$ initially assuming that the sample is semi-infinite ($d = \infty$). To obtain an initial guess for σ_t , we perform an asymptotic simplification of the quantized diffusion model, valid for $r \gg 1/\sigma_t$:

$$R_d(\alpha, \sigma_t, d; r) \approx k \frac{e^{-r\sqrt{\sigma_a/D}}}{r}, \quad (8)$$

where D is the diffusion coefficient and k is a constant. As shown in Figure 4, this asymptotic approximation states that, for large enough r , we can expect a plot of $\log(rR_d(r))$ against r to be a straight line with slope $-\sqrt{\sigma_a/D}$. Hence, by fitting a line to $\log(r\bar{R}_d^\lambda)$, we obtain $\sqrt{\sigma_a/D}$, from which we compute σ_t using the currently estimated value of α . We repeat these two steps (α and σ_t estimation), but for the following iterations we no longer assume a semi-infinite sample but instead we use the measured thickness of the sample, d . This process usually converges after 3–5 iterations.

Non-linear Optimization. Starting from these estimated values for σ_t and α , we use the Levenberg-Marquardt algorithm to compute the minimum of (7). To compute the difference between two scattering profiles we use the metric:

$$E(R_d, R'_d) = \frac{1}{r_1 - r_0} \int_{r_0}^{r_1} \left[(R_d(r)/\mu)^{\frac{1}{3}} - (R'_d(r)/\mu')^{\frac{1}{3}} \right]^2 dr \quad (9)$$

where we divide the profiles by their mean values (μ and μ') to account for the unknown intensity of the light source in the diffusion profile measurement. The interval $[r_0, r_1]$ is a range of distances over which the model is expected to fit well. This range is determined by shrinking the interval until a line fits within a given tolerance, and can be manually overridden to avoid any glitches in the measured profiles. An example from our measurements highlighting this range is shown in Figure 4.

As a final step in the single-material fitting process, we summarize the residual error of each database sample using a confidence:

$$z_m^\lambda = \min \left(\frac{\mu_d^\lambda}{d_m^\lambda - d_{75\%}^\lambda + \mu_d^\lambda}, 1 \right), \quad (10)$$

where d_m^λ is the residual (the minimum value of (7)) in wavelength band λ for the m^{th} sample, $d_{75\%}^\lambda$ is the 75th percentile residual over the whole database for this wavelength band, and μ_d^λ is the median error for this wavelength band over the entire database. We use this confidence later as a weight in fitting pigment parameters.

This process of fitting to a single profile produces material parameters that correspond to the observed appearance, but because the model is only approximate, the best-fit parameters may not be close to the true parameters of the material. This is particularly problematic in the case of anisotropically-scattering materials for which the forward model is less accurate. To obtain more meaningful results we subsequently fit larger collections of samples at once, as described in the next section.

6.3 Global Pigment Parameter Estimation

Once we have separately estimated the parameters of all the database samples (each of which consists of a known mixture of one or more pigments with the base material), we have material parameters for each sample, which describe that sample's appearance. However these pigment parameters are not linearly related to pigment concentrations as radiative transport theory predicts (see Figure 5). As a result, interpolating between the independently fit parameters can lead to poor prediction results.

To get more reliable predictions, we instead use the results of independent fitting to initialize a larger fitting problem that finds material parameters for each pigment that are globally consistent with all samples in the database, under the radiative transport theory assumption of a linear relationship between pigment concentrations and the parameters of the mixture. This linear relationship can be succinctly expressed using a matrix Σ , which contains the properties of all samples in all wavelength bands:

$$\Sigma = \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_{n_p} \end{bmatrix} \quad (11)$$

where

$$\sigma_p = [\sigma_{s,p}^1 \quad \sigma_{t,p}^1 \quad \cdots \quad \sigma_{s,p}^{n_\lambda} \quad \sigma_{t,p}^{n_\lambda}] \quad (12)$$

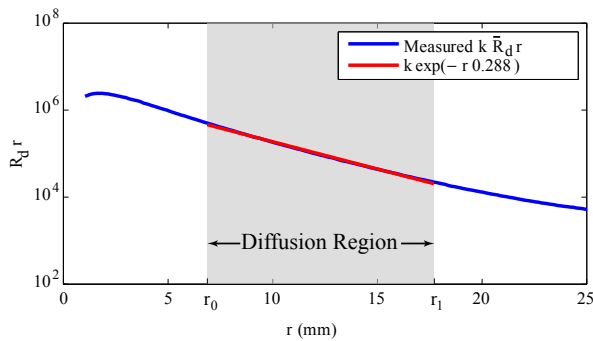


Figure 4: A log plot of a measured bulk scattering profile as a function of distance. Only the grey highlighted region is used for fitting with the diffusion model.




Figure 5: Per-wavelength-band scattering coefficients for the blue pigment dilution set as a function of concentrations. Note how σ_s is non-linear with respect to concentration.

are the material parameters of the p^{th} pigment, and a matrix C , which contains the known concentrations of all pigments in all database samples; entry c_{mp} is the concentration of pigment p in sample m . The $n_p \times 2n_\lambda$ matrix Σ has a row for each pigment (including the base material) and a column for each parameter in each wavelength band. The $n_m \times n_p$ matrix C has a row for each material in the database and a column for each pigment. With these definitions, the matrix $M = C\Sigma$ contains the material parameters of every material in the database.

To find globally consistent material parameters for the pigments, we fit the same objective function (7) independently for each wavelength band, except summed over all materials:

$$\sigma_{\text{global}}^\lambda = \underset{\Sigma^\lambda}{\operatorname{argmin}} \sum_{m=1}^{n_m} F(R_d(\alpha_m^\lambda, \sigma_{t,m}^\lambda, d), \rho(\alpha_m^\lambda, \sigma_{t,m}^\lambda, d), \bar{R}_d^\lambda, \bar{\rho}^\lambda) \quad (13)$$

where

$$\sigma_{s,m}^\lambda = \sum_{p=1}^{n_p} c_{mp} \sigma_{s,p}^\lambda, \quad \sigma_{t,m}^\lambda = \sum_{p=1}^{n_p} c_{mp} \sigma_{t,p}^\lambda, \quad \text{and} \quad \alpha_m^\lambda = \frac{\sigma_{s,m}^\lambda}{\sigma_{t,m}^\lambda}.$$

To ensure convergence to the global minimum, in this initial phase we estimate the properties of the pigments one at a time, using the one- and two-pigment dilution sets described in Section 6.1. We begin with the white dilution set, optimizing (13), summing only over the materials in that set, for the properties of the base material and the white pigment. For each color dilution set, we then similarly optimize for the properties of the color pigment, holding the white and base materials fixed. We initialize these optimizations by fitting a line to the scattering parameters (from the previous step) of all entries in a dilution set.

6.4 Mixture Optimization

Once we have parameters for each of the available pigments, we can compute a recipe to match a target material using the same tools. Given the measured diffuse reflectance $\bar{\rho}^\lambda$ and scattering profile \bar{R}_d^λ for the target material, we first use the fitting process of Section 6.2 to estimate the $2n_\lambda$ -vector of scattering parameters $\hat{\Sigma} = [\hat{\sigma}_s^1 \cdots \hat{\sigma}_t^{n_\lambda}]^T$ for the target mixture. Then we solve the linear system $\mathbf{c}_i^T \Sigma = \hat{\Sigma}$ to get a p -vector of pigment concentrations $\mathbf{c}_i = [\hat{c}_1 \cdots \hat{c}_{n_p}]^T$. Using \mathbf{c}_i as an initial guess, we optimize the predicted appearance to the target:

$$\mathbf{c} = \underset{c_1, \dots, c_{n_p}}{\operatorname{argmin}} \sum_{\lambda=1}^{n_\lambda} F(R_d(\alpha^\lambda, \sigma_t^\lambda, d), \rho(\alpha^\lambda, \sigma_t^\lambda, d), \hat{R}_d^\lambda, \hat{\rho}^\lambda) \quad (14)$$

where α^λ and σ_t^λ are defined by

$$\sigma_s^\lambda = \sum_{p=1}^{n_p} c_p \sigma_{s,p}^\lambda, \quad \sigma_t^\lambda = \sum_{p=1}^{n_p} c_p \sigma_{t,p}^\lambda.$$

The resulting vector \mathbf{c} is the recipe to replicate the appearance of the target material using the given pigment set.

7 Local Pigment Parameter Estimation

The mapping of measurements to pigment concentrations as it is described until now assumes that the forward model can globally fit the entire database with a single set of pigment parameters. However, even though our forward model is relatively accurate, we cannot expect it to work well over the entire parameter range when using a single global set of pigment parameters. In particular, the forward model will be less accurate for low optical thickness, low albedo, or anisotropically scattering materials. This is the case, for instance, for Silicone Mixture 2 shown in Figure 9 which by design is within the gamut of our pigments, but is not replicated accurately.

To overcome inaccuracies in the forward model, we introduce a local refinement strategy to find a set of pigment parameters that locally fits the samples in the database that are most similar to the target. To accomplish this, we apply a higher weight on the neighbors, with respect to the pigment concentration, when estimating the “effective” pigment scattering parameters for finding the recipe.

This can be performed in an iterative procedure which interleaves the Parameter Estimation and Mixture Optimization stages with the difference that row weights are used in the Parameter Estimation stage to bias the error to be lower for neighboring mixtures already in the database. In practice, we found that the dot product of the normalized pigment concentration vectors between the currently predicted pigment concentrations and a database entry (excluding base silicone concentration) provides consistent results.

The procedure for the local pigment parameter estimation parallels the global optimization in Section 6.3; however, the objective function (13) is replaced by:

$$\sigma_{\text{local}}^\lambda(\mathbf{c}) = \underset{\tilde{\Sigma}^\lambda}{\operatorname{argmin}} \sum_{m=1}^{n_m} \left[(w^\lambda(\mathbf{c}, \mathbf{c}_m) + k_{\text{reg}}) F(R_d(\tilde{\alpha}_m^\lambda, \tilde{\sigma}_{t,m}^\lambda, d), \rho(\tilde{\alpha}_m^\lambda, \tilde{\sigma}_{t,m}^\lambda, d), \tilde{R}_d^\lambda, \tilde{\rho}^\lambda) \right], \quad (15)$$

where the material parameters $\tilde{\alpha}$ and $\tilde{\sigma}_t$ are the local ones (derived from the optimization variable $\tilde{\Sigma}^\lambda$). The k_{reg} parameter regularizes the problem so that even when some pigments are not used by the nearby samples, their parameters stay close to the global parameters. The regularization parameter k_{reg} is set just high enough (around 10^{-4} relative to a unit maximum) to stabilize the optimization, while still low enough not to affect the quality of the local fit.

The weights w^λ are set to

$$w^\lambda(\mathbf{c}, \mathbf{c}_m) = z_m^\lambda D_{\text{mixture}}(\mathbf{c}, \mathbf{c}_m) \quad (16)$$

where

$$D_{\text{mixture}}(\mathbf{c}_1, \mathbf{c}_2) = \text{normalize}(\mathbf{c}_1) \cdot \text{normalize}(\mathbf{c}_2). \quad (17)$$

These weights cause the pigment estimation stage to find pigment parameters that fit well to database materials similar in composition to the target mixture, resulting in better prediction of appearance for the optimized recipe. After a new recipe has been found, we update the weights, and we repeatedly solve re-weighted systems until convergence or for a maximum of 5 iterations.




Figure 6: Top: Screenshot from the color selection phase of our editor. The variations along the CIELab primaries are shown above and below the current pigment recipe. Bottom: A screenshot from the translucency selection phase. The user can navigate along the T, S, R coordinate system. The blacked out images show directions in translucency that are limited by our pigments’ gamut.

8 Appearance Editing

In addition to matching the appearance of measured targets, we also developed an authoring tool which allows the user to pre-visualize recipes before fabrication and edit the desired color and translucency, while staying within the gamut imposed by the minimum and maximum possible pigment concentrations. Our approach provides an intuitive editing workflow by decoupling edits to the color and the translucency of a desired material.

Overview. Our editing process starts with an initial recipe and sample thickness provided by the user. We then provide intuitive browsing control to fine tune the color followed by translucency. We took inspiration for our system from the “variations” control interface in Adobe Photoshop. Screenshots from our editor are shown in Figure 6.

For editing color, our interface displays the diffuse reflectance color of the current recipe and six color-variations, estimated by moving in both the positive and negative directions of the three CIELab primaries. When the user selects a variation, the current recipe gets updated and the process repeats until the user has reached the desired color.

For editing translucency, our interface renders a synthetically fabricated recipe under step edge illumination to visualize the reflectance profile (similar to the captured photographs in Figure 9). We again show the current recipe and translucency-variations which can be selected to modify the shape of the diffusion profile without affecting

the overall diffuse reflectance.

We detail the steps necessary to accomplish this below.

8.1 Color Reproduction

To display the CIELab color predictions of a recipe, we require a mapping, $(L, a, b) = P_{\text{Lab}}(\mathbf{c}, d)$, to convert pigment concentrations \mathbf{c} to CIELab color values, given the depth d of the sample. We use globally linear estimates of pigment scattering parameters from our database, and then linearly mix these parameters according to \mathbf{c} to get the resulting scattering and absorption coefficients of this mixture. These coefficients are then converted to a 5D reflectance vector using the interpolation method described in Section 6. We then convert our 5D spectral color values to spectral reflectance distributions by training a set of “eigen-spectral reflectance functions” from spectral reflectance measurements of our database [Park et al. 2007]. Once we have an approximate spectral distribution for the color, we convert to CIEXYZ and subsequently to CIELab and sRGB. The CIELab is used for computing the variation distances and sRGB for displaying the predicted reflectance values to the user.

8.2 Color Editing

Our goal is to find variations of our recipes with a given distance in CIELab. For this, we first compute the partial derivatives $J_{\mathbf{c} \rightarrow \text{Lab}}$ of $P_{\text{Lab}}(\mathbf{c})$ with respect to pigment concentrations \mathbf{c} , using finite differences. This matrix represents the predicted color change with respect to pigment concentration changes. In order to find the pigment concentration changes that allows movement along the CIELab primaries, we first perform a singular value decomposition on $J_{\mathbf{c} \rightarrow \text{Lab}}$: $[U, \Sigma, V] = \text{svd}(J_{\mathbf{c} \rightarrow \text{Lab}})$ and compute the pseudo inverse of $J_{\mathbf{c} \rightarrow \text{Lab}}$ as $V_{\text{Lab}} = V\Sigma^{-1}U^T$.

Each column of the resulting 6×3 matrix V_{Lab} is now the derivative of the pigment vector with respect to the CIELab primaries. The difference in pigment concentrations, $\Delta_{\mathbf{c}}$, to achieve a desired change in CIELab, Δ_{Lab} , is given by $\Delta_{\mathbf{c}} = V_{\text{Lab}}\Delta_{\text{Lab}}$.

8.3 Translucency Editing

Given the scattering and absorption coefficients for each wavelength band, we render a column image using the diffusion model described in Section 6. Each pixel representing the radiance of our 5 wavelength bands is transformed into CIELab using the process described in Section 8.1.

For navigating the translucency space without affecting the color, we reuse a part of the 6×6 matrix V from the singular value decomposition of the color derivatives matrix $J_{\mathbf{c} \rightarrow \text{Lab}}$. We split V into two submatrices $[V_c V_t] = V$, where V_c are the first 3 columns of V and V_t the remaining 3 columns. While the pigment vectors in V_c affect the color of the final recipe, the pigment vectors in V_t do not and can be used to influence just the translucency.

Unfortunately, the pigment vectors in V_t do not correspond to meaningful directions. One intuitive direction for controlling translucency is increasing the average reduced extinction coefficient $\langle \sigma_t^\lambda \rangle$, over all wavelength bands, by scaling the current pigment concentrations with the desired $s = \langle \sigma_t^\lambda \rangle$. For a semi-infinite slab, this directly corresponds to scaling all profiles by s , but for finite-slabs this can inadvertently affect the diffuse reflectance color. We can find the closest direction T in the space defined by V_t by solving the linear system $V_t T = C$. Now, given T , we find the orthonormal directions S and R that should not change color, nor the average reduced extinction coefficient $\langle \sigma_t^\lambda \rangle$. These vectors T , S , and R are used to define the variations in translucency that are presented to the user.

When editing translucency, it is often impossible to move in some directions, most often the S and R directions, because of the constraints on how much pigment can be used. For steps that exceed the limits, we black out that direction to indicate to the user that they have reached the edge of the gamut.

9 Fabrication

Fabricating the generated recipes with a high degree of precision is critical to the correct evaluation of our method. The main challenges in fabrication are: ensuring that the correct amount of each pigment is added, avoiding air and other impurities, and finally ensuring that there is at least one side on the sample which appears near-specular. In this section we provide a brief description of the materials and machinery used for fabrication and then we describe the hierarchical dilution process we employ for improving the concentration accuracy, reducing waste and streamlining the fabrication process.

Silicone and Pigments. We use SortaClear40 Translucent, Addition Cure silicone rubber by Smooth-On. This silicone rubber cures at room temperature with a shrinkage of less than 0.01%. The two-component silicone rubber requires a catalyst, mixed in with a ratio 1:10, to activate the curing process. The curing is roughly 24 hours. The silicone is completely cured after 7 days.

For manipulating the appearance of the silicone we use silicone pigments which are used for coloring silicone rubber. The pigment colors we use are: White (Pantone White C), Yellow (RAL 1018), Red (Pantone Red C), Green (Pantone 3292), Blue (Pantone 2757C), Black (Pantone Black C). These pigments are mainly absorbing except for White, Yellow and Red, for which scattering is significant.


Hierarchical Dilution. For fabricating the generated recipes as faithful as possible, we employ a hierarchical dilution scheme.

For each pigment we initially produce 1 kg master batches of 5% pigment concentration. This is the maximum ratio of pigment that still allows the silicone to cure. To achieve a target concentration, we then dilute some quantity of this master batch, in an iterative fashion, with base silicone. We start with some of the 5% mixture and repeatedly mix the current dilution with an equal amount of base silicone to halve the pigment concentration. Once the concentration is roughly twice the target concentration for our recipe, we mix with the exact ratio of base silicone needed to achieve the desired target concentration.


This dilution tree structure is automatically generated by a simple script using a bottom-up approach. Intermediate dilutions that are needed by many target recipes are merged and created only once. We ignore concentrations of pigments with absolute concentration less than 10^{-8} of the total sample weight, or less than $3\mu\text{g}$ for a typical 300g sample. Our process also accounts for catalyst that will be added at the very end for curing the final samples.

We use a digital balance with accuracy of $\pm 0.05\text{g}$ when mixing dilutions. To ensure high accuracy, we always use a minimum of 10g of both dilution and base silicon when mixing. Though limiting the minimum weight increases the number of steps, we obtain a relative accuracy lower bound of 1% at each dilution step. This relative error is reduced with each dilution step by the mixing ratio.

Mix Preparation Process. During the fabrication process, preserving homogeneity is very important. Once all ingredients are added, we stir the mixture for several minutes until it is homogeneous. Since this process accumulates air, the mixture is placed in



Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl
Figure 7: Synthetic target and replica comparisons for the global method. For each sample, on the left we render the target under a shadow edge illumination scenario and on the right the replica. Both renderings use a QD fit to ground-truth path traced results.



Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl Tar Repl
Figure 8: Synthetic target and replica comparisons when using a simpler optimization that fits on parameters directly. Fitting on parameters directly does not perform as well as using our appearance distance term (cf. Figure 7).

a vacuum chamber for about 10 minutes to further improve homogeneity. We then pour the mixture into a mold, constructed using acrylic plates. We use acrylic plates so the sample has smooth, near-specular surfaces for better compliance with our model assumptions and more angularly uniform reflectance for our $45^\circ/0^\circ$ measurement setup, due to the lack of micro-surface roughness. Finally when the samples cure, they are measured using the method described in Section 5.

10 Results

We validated our forward model, measurement setup, and optimization procedures first in simulation and also by fabricating physical replicas using silicone. We describe these here in turn.

10.1 Reproduction of Simulated Examples

To ensure that our method can properly measure and replicate real-world translucent materials, we first validated our entire pipeline in simulation by fitting on synthetic images generated using brute-force volumetric path tracing with the Mitsuba renderer [Jakob 2010]. We designed an analogous database of measurements and tested our method on synthetic targets. We created 6 synthetic pigments with prescribed RGB material parameters and isotropic phase functions. The database samples and targets are created by linearly blending the pigment parameters according to their concentrations. For our synthetic tests, we recreate the geometry of our physical setup: the camera and light sources are placed at the same positions and orientations, all the database samples had a thickness of 30 scene units and our diffusion profile light source was modeled as a disk light touching the sample.

Our path traced measurements (which mimic our real measurements) omit the beginning of the profile, so to enable rendered validations in a realistic shadow edge configuration (which requires access to the entire profile) we performed an unconstrained fit of our forward model to these path traced measurements. Given the fitted profiles we use the same method as in Section 8 to render the shadow edge illumination seen in Figure 7.² The differences between the targets and the replicas in this synthetic test are barely noticeable. In fact, over this set of 14 synthetic targets the average relative reflectance

²We confirm in our supplemental material that these fitted profiles are indistinguishable from the 1D Monte Carlo profiles.

error was 0.10% and the mean relative reduced mean free path error was 6.44%. This confirms that our measurement and optimization procedures are able to accurately replicate unknown translucent materials.

In Figure 8 we perform a similar synthetic evaluation, but using a simpler optimization strategy and error metric to demonstrate that our more involved approach is necessary to obtain accurate results. The only difference compared to Figure 7 is that the replicas are obtained by minimizing the MSE between the estimated/measured parameters for the target (reduced scattering coefficient and absorption coefficient) and the linearly blended global parameters of pigments of the tested recipe. This can be modeled as a simple linear system, similar to our initialization step in the Mixture Optimization section. The results show that, fitting on parameters directly does not perform as well as using our appearance distance term. This is due to two reasons: a) the material parameters recovered by a single target measurement are highly under-constrained; and b) our appearance distance term can leverage “appearance metamerism” (several recipes and pigment parameters can lead to indistinguishable translucent appearance) to tolerate ill conditioning in the mapping from optical properties to appearance.

10.2 Reproduction of Real-World Examples

In addition to validating on synthetic measurements, we also demonstrate the quality of our pipeline on real-world examples. We measured and replicated several homogeneous real-world materials such as strawberry yogurt drink, fabric softener, full-fat milk, low-fat milk, soap, and white chocolate. We computed the corresponding mixing recipes using our iterative optimization procedure and fabricated the samples in various shapes. Each replica has the same shape and thickness, d , as its corresponding target. The computation time for estimating a recipe is about 2 minutes with the global method and 30 minutes for the local method. The local method training database also includes the replicas fabricated using the global method.

For our real-world examples, we capture photographs of the original targets and the replicas using a Nikon D800 camera. We construct HDR images from RAW images with 6 different exposure times. The RAW LDR images are processed with DCRAW and white balanced using a white reference card. We recover and account for the camera response curve using the technique of Mitsunaga and Nayar [1999]. This results in white balanced linear sRGB images which are then converted to HDRs using the aforementioned




Figure 9: Measured target and fabricated replica HDR pairs for a variety of materials. The mixing recipes for the fabricated replicas of each target were generated using the global method (left) and the local method (right). All samples are partially illuminated with an illuminant with a color temperature of 4700K. A horizontal blocker is placed between the samples and the light source to better emphasize their translucency.

method. At the very end, we tonemap the HDR image to an 8-bit sRGB image by applying the sRGB gamma curve. We compare the target material against the fabricated replicas under two different illumination configurations.

In Figure 1 we showcase side-by-side visual comparisons of various liquid and solid translucent targets next to their replicas, in their natural form under top and side illumination.

Figure 9 shows side-by-side shadow edge comparisons with an illuminant that approximates a color temperature of 4700K. For each triplet, the target is in the middle, and we show the results of the global and local methods to the left and right of the target respectively. Visually, even our global method is able to match both the translucency and reflectance of many of these materials quite well. For the global method, the appearance distance (defined in Equation (7)) averaged over all 9 targets is 0.00644, with a standard deviation of 0.01142. Our local refinement usually performs better, and sometimes significantly better, than predicting the appearance of the replicas using the global parameters (Table 1). In fact, the mean distance after local refinement decreases to 0.00276, with a standard deviation of 0.00213. The improvement is especially noticeable for the dark Silicone Mixture 2 seen in Figure 9, and some improvement is also visible for the Blue Fabric Softener and Strawberry Yogurt Drink samples.

We provide plots of the simulated, predicted, and measured diffuse reflectance and scattering profiles for our targets in the supplemental material. In general, we observed a good prediction for materials with high reflectance and optical thickness. With a decrease of the reflectance and optical thickness, we notice an increase in matching error. This is not surprising, since diffusion theory is not a very good approximation for very dark materials with low optical thickness.

To isolate the performance of our optimization, we tested our pipeline with materials that are in theory exactly reproducible by our system. To ensure that our target material is within the gamut of reproducible materials, we fabricated target samples out of silicone mixed with pigments that are available to our fabrication process. We then measured these samples and used the diffuse reflectance and scattering profiles as input to our optimization process (these

Table 1: Global estimate of reduced scattering and absorption coefficients for pigments and base silicone used in our fabrication process. Units are in $1/\mu\text{m}$.

	Reduced Scattering Coefficient (σ_s)					Absorption Coefficient (σ_a)				
	blue	cyan	green	orange	red	blue	cyan	green	orange	red
White	848.9	847.6	783.6	669.4	718.3	0.0	0.0	0.0	0.0	0.0
Yellow	92.6	74.0	78.8	64.0	55.8	30.1	1.8	0.5	0.2	0.1
Red	0.0	0.4	0.3	620.3	96.7	515.4	839.8	999.6	163.9	9.2
Green	7.1	0.0	10.0	0.1	0.0	19.5	9.4	15.8	164.0	264.4
Blue	92.3	9.1	0.0	0.1	0.2	41.6	90.2	181.1	1080.9	1045.9
Black	25.3	21.8	22.9	25.9	26.6	321.6	331.4	330.1	323.7	325.4
Base	0.078	0.053	0.061	0.054	0.021	0.002	0.002	0.001	0.001	0.001

samples are not included in our measurement database). Table 2 shows the ground truth pigment concentrations of two samples and the estimated concentrations using our optimization process. In Figure 10, we compare the profiles of the original measured sample, the predicted appearance, and the measured fabricated replica. Note that the appearance of the replica closely matches the target, even though the pigment concentrations are different. This is an example of an “appearance metamer” where we cannot perceive the appearance difference between two different recipes. Such a replica would never be created when using a simpler parameter optimization as shown in Figure 8.

11 Discussion & Future Work

We have presented a method to physically replicate homogeneous translucent materials by mixing pigments into a substrate. By using local fits to a diffusion model, we intelligently interpolate between measurements of known materials to predict the appearance that will result from novel mixtures. The method only considers homogeneous materials, and its accuracy drops off for highly forward scattering and highly absorbing materials, with low optical thickness, where diffusion theory is less accurate. We believe that there is room for improvement on the reproduction accuracy of our method.

In Figure 11, left, we provide spectral reflectance measurements for a set of target materials that are outside of the gamut of our method. The three target reflectance measurements exhibit a significant positive gradient in the region between 540 and 580 nm. On the plot on the right we can see that the spectral reflectance values of our pigments do not exhibit a significantly positive gradient in that region, indicating that these targets are outside of the gamut of our pigments.

There are various avenues for improving the reflectance gamut of our pigments. One such option would be to add additional pigments with useful spectral features that are not present in our current set. Currently our system computes the reflectance error of the true spectral distribution function of our targets, another option would be

Table 2: Concentration percentages prediction generated using our local method for the 9 targets used. The last 2 rows show the actual concentration of the fabricated silicone mixtures used as targets.

	white	yellow	red	green	blue	black
silicone mixture 1	0.20590%	0.20897%	0.00345%	0.00140%	0.00735%	0.00000%
blue fabric softener	0.16384%	0.00000%	0.00037%	0.00170%	0.01407%	0.00000%
low-fat milk	0.13572%	0.01291%	0.00000%	0.00012%	0.00000%	0.00000%
red soap	0.41249%	0.19489%	0.01169%	0.00247%	0.00000%	0.00000%
strawberry yogurt drink	0.25791%	0.06423%	0.00467%	0.00000%	0.00062%	0.00000%
white chocolate	0.32361%	0.24976%	0.00261%	0.00284%	0.00000%	0.00000%
full-fat milk	0.35230%	0.02702%	0.00000%	0.00000%	0.00000%	0.00012%
mocca yogurt drink	0.21805%	0.37069%	0.00953%	0.00543%	0.00124%	0.00000%
silicone mixture 2	0.00000%	0.19234%	0.00142%	0.00000%	0.00547%	0.00702%
silicone mixture 1 (actual)	0.20000%	0.20000%	0.00300%	0.00000%	0.00750%	0.00000%
silicone mixture 2 (actual)	0.00000%	0.20000%	0.00300%	0.00000%	0.00750%	0.00000%




Figure 10: We evaluate the appearance prediction ability of our system, using the local method, by comparing the target measurements (a) with the model prediction (b) and finally the fabricated replica measurements (c). The first row shows comparisons of our system’s prediction ability for Silicone Mixture 1, whereas the second row shows comparisons for Silicone Mixture 2. Both target samples are in gamut and their recipes can be found in Table 2. Note that even though Silicone Sample 2 has low optical thickness and reflectance values, the fabricated replica is a good match.

to minimize the perceptual error. To improve the measurement gamut of our method, we could (with minimal changes to our method) use a spectrally homogeneous illuminant and a multi-spectral camera instead of a greyscale camera combined with a discrete set of LEDs.

There are many other directions for continued development of our method. Improved transport models and/or better interpolation with larger databases can extend the gamut towards darker and less optically dense materials, ultimately to cover the whole range from crystal clear to opaque. The use of perceptual metrics (beginning with better color difference measures) for appearance matches will be important to achieving consistent results under a range of illumination conditions.

By itself, the ability to control the appearance of homogeneous materials has implications for many industries where pigmented materials are used, including plastics, food, prosthetics, and even dentistry, where the critical matching of appearance between natural teeth and artificial resins is still done painstakingly by eye. Translucent appearance matching is also important for fabricating more complex materials. Previous work on fabricating translucent appearance has been limited to spatial combinations of fixed materials, and our new method, with its ability to continuously tune material parameters, opens the possibility for more powerful new methods for inhomogeneous materials that optimize both the spatial mixture of materials and also the properties of the individual materials themselves. Ultimately these technologies will lead to future machinery that can automatically replicate the appearance of almost any material.

References

- DEON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics* 30, 4, 56:1–56:14.
- DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Transactions on Graphics* 29, 4 (July), 62:1–62:10.
- DONNER, C., LAWRENCE, J., RAMAMOORTHY, R., HACHISUKA, T., JENSEN, H. W., AND NAYAR, S. 2009. An empirical bssrdf model. *ACM Transactions on Graphics* 28, 3 (July), 30:1–30:10.
- DORSEY, J., RUSHMEIER, H., AND SILLION, F. 2008. *Digital Modeling of Material Appearance*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- FUCHS, M., RASKAR, R., SEIDEL, H.-P., AND LENSCH, H. P. A. 2008. Towards passive 6d reflectance field displays. *ACM Transactions on Graphics* 27, 3 (Aug.), 58:1–58:8.
- HASAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Transactions on Graphics* 29, 4 (July), 61:1–61:10.
- HAWKINS, T., EINARSSON, P., AND DEBEVEC, P. 2005. Acquisition of time-varying participating media. *ACM Transactions on Graphics* 24, 3 (Aug.), 812–815.
- HULLIN, M. B., LENSCH, H. P. A., RASKAR, R., SEIDEL, H.-P., AND IHRKE, I. 2011. Dynamic display of BRDFs. In *Computer Graphics Forum (Proc. EUROGRAPHICS)*, 475–483.
- JAKOB, W., 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. *ACM Transactions on Graphics* 21, 3 (July), 576–581.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, 511–518.
- KELLY, R. J., 1987. Process for matching color of paint to a colored surface. U.S. Patent Number 4692481. Filed Sep 27, 1984.
- MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H. P. A., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing spatially-varying reflectance. *ACM Transactions on Graphics* 28, 5 (Dec.), 128:1–128:9.

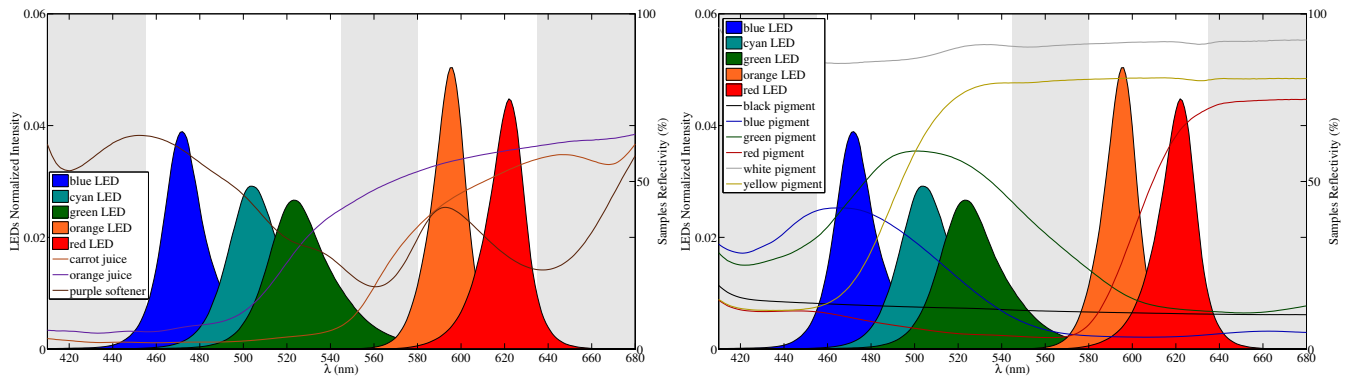


Figure 11: On the left we provide the spectral reflectances of 3 target materials that are not suitable for our method. On the right we provide spectral reflectance measurements of our pigments mixed with white. Both figures have the spectral intensities of our LEDs overlayed. The grayed out regions indicate spectral bands that are not sampled well with our LEDs. We observe that the 3 targets exhibit a significantly positive gradient in the region between 540 and 580 nm, which is not captured by our setup. In addition we can see on the right that our current set of pigments does not have a member that exhibits a significant positive gradient in this region, indicating that these targets are outside of the gamut of our pigments.

MITSUNAGA, T., AND NAYAR, S. 1999. Radiometric self calibration. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, 374–380.

MUNOZ, A., ECHEVARRIA, J. I., SERON, F., LOPEZ-MORENO, J., GLENCROSS, M., AND GUTIERREZ, D. 2011. BSSRDF estimation from single images. *Computer Graphics Forum* 30, 455–464.

NARASIMHAN, S. G., GUPTA, M., DONNER, C., RAMAMOORTHY, R., NAYAR, S. K., AND JENSEN, H. W. 2006. Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics* 25, 3 (July), 1003–1012.

PARK, J.-I., LEE, M.-H., GROSSBERG, M. D., AND NAYAR, S. K. 2007. Multispectral imaging using multiplexed illumination. In *ICCV*, 1–8.

SHERMAN, C. J., AND SIMONE, K. S., 1989. Process for manufacturing paints. U.S. Patent Number 4887217. Filed Jan 4, 1985.

SONG, Y., TONG, X., PELLACINI, F., AND PEERS, P. 2009. Subedit: a representation for editing measured heterogeneous subsurface scattering. *ACM Trans. Graph.* 28, 3 (July), 31:1–31:10.

WANG, L., JACQUES, S. L., AND ZHENG, L. 1995. MCML—Monte Carlo modeling of light transport in multi-layered tissues. *Computer methods and programs in biomedicine* 47, 2 (July), 131–146.

WEYRICH, T., MATUSIK, W., PFISTER, H., BICKEL, B., DONNER, C., TU, C., MCANDLESS, J., LEE, J., NGAN, A., JENSEN, H. W., AND GROSS, M. 2006. Analysis of human faces using a measurement-based skin reflectance model. *ACM Transactions on Graphics* 25, 3 (July), 1013–1024.

WEYRICH, T., LAWRENCE, J., LENSCH, H. P., RUSINKIEWICZ, S., AND ZICKLER, T. 2009. Principles of appearance acquisition and representation. *Foundations and Trends in Computer Graphics and Vision* 4, 2, 75–191.

WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics* 28, 3 (July), 32:1–32:6.

XU, K., GAO, Y., LI, Y., JU, T., AND HU, S.-M. 2007. Real-time homogenous translucent material editing. *Computer Graphics Forum* 26, 3 (Sept.), 545–552.