Physical Reproduction of Materials with Specified Subsurface Scattering

Miloš Hašan*[1] Martin Fuchs[2] Wojciech Matusik[3] Hanspeter Pfister[1] Szymon Rusinkiewicz[2]

[1]Harvard University [2]Princeton University [3]Disney Research

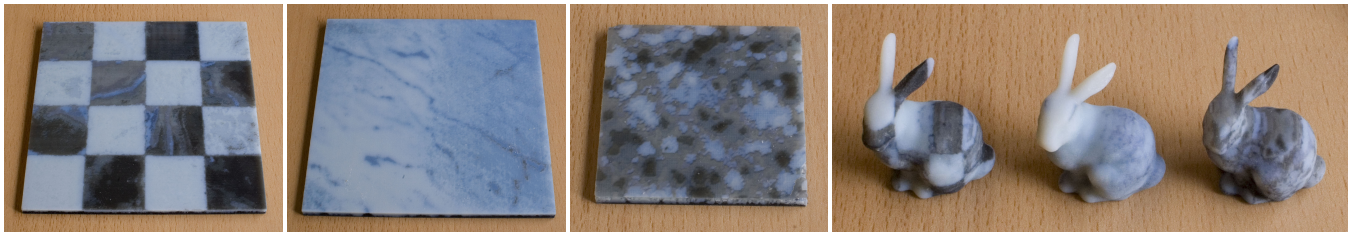

Figure 1: *Left: photographs of slabs fabricated using a multi-material 3D printer. We use a goal-driven optimization approach to find a volumetric arrangement of base material layers that best approximates an input heterogeneous subsurface scattering function. Right: the layering can also be applied to fabricate full 3D shapes.*

Abstract

We investigate a complete pipeline for measuring, modeling, and fabricating objects with specified subsurface scattering behaviors. The process starts with measuring the scattering properties of a given set of base materials, determining their radial reflection and transmission profiles. We describe a mathematical model that predicts the profiles of different stackings of base materials, at arbitrary thicknesses. In an inverse process, we can then specify a desired reflection profile and compute a layered composite material that best approximates it. Our algorithm efficiently searches the space of possible combinations of base materials, pruning unsatisfactory states imposed by physical constraints. We validate our process by producing both homogeneous and heterogeneous composites fabricated using a multi-material 3D printer. We demonstrate reproductions that have scattering properties approximating complex materials.

Keywords: fabrication, translucency, BSSRDF, scattering

1 Introduction

Today, many manufacturing processes are automated and under algorithmic control. 3D printers, milling machines, etching machines, laser-cutters, engravers, embroidery machines, and looms are just a few examples of computer-controlled output devices that have become commodities and are accessible to a wide population. Although these devices are currently used for rapid prototyping, we envision that they hold the ultimate promise of personalized mass-production of consumer goods.

*{milos,hpfister}@seas.harvard.edu, {mfuchs,smr}@cs.princeton.edu, matusik@disneyresearch.com

Reaching such an outcome will require re-imagining existing design and fabrication systems. Although the sophistication and capabilities of output devices have progressed, the pipeline of device characterization, simulation, and goal-driven fabrication has remained underdeveloped: taking full advantage of devices beyond simple 2D printers and monitors remains tedious and fraught with guesswork. We suggest that there are three major difficulties: (1) characterizing device capabilities and limitations; (2) accurately simulating fabricated appearance; and (3) optimization-based design, in which the user specifies the desired *output* instead of individual *inputs* to the fabrication device. Recent work has demonstrated the power of this approach in domains such as printing of spatially-varying reflectance [Matusik et al. 2009].

This paper addresses the fabrication of scattering objects. We are motivated by the recent availability of multi-material 3D printers such as the OBJET Connex series. Unlike previous technologies, these printers facilitate the creation of 3D objects with multiple constituent materials, including several in which subsurface light transport plays an important role. This suggests layering materials to print not only opaque objects, in which light immediately reflects from the surface, but also translucent ones, in which light penetrates into the volume and reemerges at different points. The distribution of this scattering is crucial to believable appearance of materials such as human skin, wax, and marble.

Recreating and matching scattering properties by hand is a nearly impossible task, especially for heterogeneous appearance, and we wish to free the user from the burden of specifying the exact distribution of basis materials necessary to reproduce a particular degree of translucency. We therefore introduce a novel end-to-end process for characterizing the available materials, simulating the appearance of a given layering of these materials, and automatically determining how to multiplex them to give the user the ability to specify a target appearance.

We use a data-driven approach in which we measure scattering properties of the base materials, using a novel measurement procedure and setup that allows us to efficiently recover material transmission and reflection profiles. Next, we devise a simulation method that correctly predicts the behavior of a *multi-layer material* built by stacking materials with known reflection and transmission profiles, extending the previous work of Kubelka and Munk [1931] and Donner and Jensen [2005]. We validate this approach, demonstrating a match between simulation and real materials.

This simulation comprises the inner loop of a novel optimization-based (inverse) design strategy. In this process, the user specifies the desired material properties (e.g. a desired reflection profile). Our algorithm then efficiently searches the space of all stacked material combinations to produce one that matches the user's specifications. Since the search space grows exponentially with the number of layers, we devise efficient strategies to speed up the search process by pruning branches that yield poor solutions. Putting everything together, we use the system to explore the output space for a multi-material 3D printer and to print a variety of materials with complex scattering properties.

2 Related Work

Models for scattering. The general model for scattering behavior in materials with subsurface light transport is the *Bidirectional Scattering-Surface Reflectance Distribution Function (BSSRDF)*, as specified by Nicodemus et al. [1977]. It is a function of eight parameters (positions and angles of incidence and exitance), and as such has proven unwieldy to capture or handle in rendering. Accordingly, a number of approximations and simplifications have been found.

One of the earlier examples is the model by Kubelka and Munk [1931], which describes light transport in layered materials. While the model considers mutual interaction among different layers, it assumes that light is not transported laterally. In the context of computer graphics, it was used by Haase and Meyer [1992] for the simulation of pigmented paint. Hanrahan and Krueger [1993] considered a one-dimensional light transport theory able to handle translucent materials, emphasizing practical applicability. Stam [1995] modeled these materials with a diffusion process, while Pharr and Hanrahan [2000] investigated thin slices. Stam [2001] also considered the special case of skin, modeling it as a layered material with rough boundary surfaces.

Jensen et al. [2001] applied a dipole approximation for subsurface light transport and provided a measurement setup for homogeneously scattering materials. Donner et al. [2005] returned to the Kubelka-Munk model and extended it for lateral light transport, expressing the results in frequency space; this is also the model that we will employ in Section 5. More recently, Donner et al. [2009] derived an analytical function in six parameters from BSSRDF simulations that addresses directionality and forward scattering. In the same year, Song et al. [2009] found a representation permitting efficient editing of BSSRDFs through a decoupling of global effects.

Acquisition of scattering. Goesele et al. [2004] were the first to acquire a full BSSRDF of inhomogeneous materials, sampling with a laser projector. Peers et al. [2006] reduced the cost with a factored matrix representation and, by projecting a grid pattern, accelerated the measurement process. Wang et al. [2008] use a finite difference approach to solve the diffusion equation for rendering subsurface scattering, and introduce an inverse process to fit volume parameters to heterogeneous measurements; this is related to our inverse problem, but leads to continuous rather than discrete optimization.

Tong et al. [2005] address quasi-homogeneous materials that can be expressed as global variations of a basis material, which they captured. Hawkins et al. [2005] recorded the scattering phase function of smoke separately from its density distribution, facilitating the acquisition of a time-variant volume data set. Some researchers have proposed models optimized for certain material types. Human skin has received attention from Donner et al. [2008], who modeled it with several inhomogeneous layers. Human faces were modeled by Tariq et al. [2006], with scattering profiles recorded with stripe patterns; by Weyrich et al. [2006], using specialized devices for different aspects of scattering; and by Ghosh et al. [2008], exploiting polarization effects. For liquids, Narasimhan et al. [2006] employed dilution to simplify the recording of scattering behavior.

Fabrication. In concurrent work, Dong et al. [2010] investigate an alternative approach for fabricating surfaces with subsurface scattering. One major difference is that they also explore using a milling machine and a 2D printed color texture, as opposed to our system, which uses only a 3D printer. This enhances the achievable gamut, but is not applicable to arbitrary 3D geometry. There are also significant differences in representation: our approach consistently uses the framework of scattering profiles for all stages of acquisition, representation, and optimization, while Dong et al. convert among several representations (profiles, dipole and multipole models, diffusion equation).

Apart from this, few publications have addressed the fabrication of materials with desired reflectance and scattering properties. Matusik et al. [2009] proposed a method for materials without subsurface light transport, employing inks with non-trivial reflectance properties. Weyrich et al. [2009] used precision CNC milling on metallic surfaces. For transmissive light transport, Fuchs et al. [2008] proposed a method that includes even non-local light transport such as shadows, but is limited by the possible resolution and contrast.

3 Overview

Figure 2 shows an overview of our system for fabricating objects with desired subsurface scattering. Our process solves the following challenges:

- Measurement of the scattering of the base materials (e.g. the ones available in the 3D printer). In particular, we capture the reflection and transmission profile of a thin layer of each base material (Section 4).
- Prediction of composite appearance: Given the reflection and transmission profiles of base materials at a representative thickness, we compute profiles for a different thickness, and for composites made of multiple layers of different materials; we term this the *forward problem* (Section 5).
- Matching the appearance of a target material: Given the reflection profile of a material, we solve the *inverse problem* of finding an assignment of materials to layers and the layer thicknesses that result in a close match to the desired reflection. We achieve this through non-linear discrete optimization (Section 6).
- Extending the results from flat homogeneous slabs to heterogeneous materials and 3D objects by running the optimization on many surface points, spatially varying the resulting layer thicknesses, and wrapping the layering around the object (Section 7).
- Using a 3D printer to fabricate physical results (Section 8).

3.1 Background: BSSRDFs and scattering profiles

The BSSRDF $S(\vec{x}_i, \vec{\omega}_i, \vec{x}_o, \vec{\omega}_o)$ describes the appearance of a scattering material [Nicodemus et al. 1977; Jensen et al. 2001]. It is an eight-dimensional function that relates the outgoing radiance from a given point in a given direction to the incoming flux from all other positions and directions in the scene:

$$L_o(\vec{x}_o, \vec{\omega}_o) = \int_A \int_{\Omega^+} L_i(\vec{x}_i, \vec{\omega}_i) \cdot S(\vec{x}_i, \vec{\omega}_i, \vec{x}_o, \vec{\omega}_o) \cos\theta_i \, d\vec{\omega}_i \, d\vec{x}_i.$$

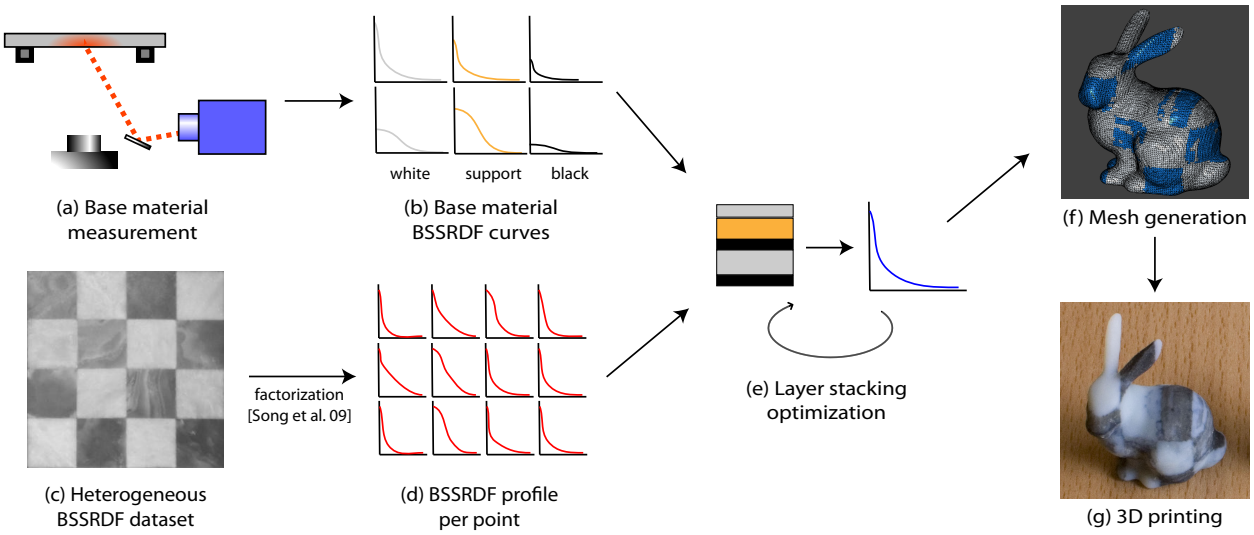

Figure 2: *Overview of our system. The scattering properties of the base materials available in the Object Connex 3D printer are measured (a), resulting in the reflection and transmission profiles (b). A 4-dimensional heterogeneous BSSRDF (c) is factorized into a set of "local" reflection profiles for every surface point (d). An optimization process (e) finds, for a given target reflection curve, a layer stacking of the base materials that most closely approximates it. Once the layer thicknesses at every surface point are known, multi-material meshes can be generated (f) and fabricated by 3D printing (g).*

By introducing the common assumptions of a homogeneous material dominated by multiple scattering, and mostly uniform illumination, the BSSRDF becomes a 1D (radial) function that only depends on the distance between the incoming and outgoing point:

$$S(r) = S(\|\vec{x}_o - \vec{x}_i\|). \tag{1}$$

Integrating over the exitant light directions, we can define the reflection profile $R(r)$ and the transmission profile $T(r)$ of a thin slab as

$$R(r) := \pi \cdot S^{+}(r), \qquad T(r) := \pi \cdot S^{-}(r), \tag{2}$$

where the + and - signs denote whether the incoming and outgoing point are on the same or opposite sides of the slab. This is a convenient representation, because the appearance of a composite layer can be computed from the profiles of the base layers using a series of convolutions [Donner and Jensen 2005].

4 Measurement

Our measurements of the base materials aim to establish both a reflective and a transmissive scattering profile. Even though the materials we use do not perfectly adhere to the assumptions of directionally uniform scattering behavior and spatial homogeneity, we can gain useful measurements by recording the average case. Brute-force averaging would require us to record many different combinations of directions and locations, which would be expensive. Therefore, we split the recording process into two distinct steps: first, we record the profiles for few different surface locations and incident directions, then we record an approximation of their integral in smooth illumination, and use this to re-scale averaged profiles to match.

Reflective albedo. Ideally, we would create directionally uniform illumination and average over the outgoing directions. This is difficult to achieve, though, as the recording sensor would conflict with the light source. As an approximation, we place our samples on a sheet of black cloth on the floor of a windowless room with white walls and extended ceiling illumination. We acquire a picture with a digital SLR camera while a gray card with known albedo is in view; using its average pixel value as reference, we can recover the respective albedos of the base materials. As we do not investigate wavelength dependence of scattering profiles, we use only the green channel of the camera.

Transmissive albedo. For the transmissive albedo, uniform conditions can be more easily created. As incident illumination, we use an LCD computer screen displaying a constant white image, switching off the remaining room illumination. We place the material sample directly in contact with the screen and a thin diffuser sheet directly in contact with the sample, in order to average the outgoing radiance directionally. Then, we record a picture with a digital camera. The recorded average emittance divided by the screen's emittance measured at the same spot yields the average transmittance.

Reflective profile. The reflective profiles of the base materials have high contrast and a steep decay over the very first millimeter; however, appearance also depends on scattering far from the incident position, at which the BSSRDF may be on the order of 10^{-5} of its value at the center. Therefore, we require a precisely defined illumination situation in an environment that interacts as little as possible with the sample through global illumination.

Our solution is depicted in Figure 3 (left). We place a digital SLR camera and the sample into a tent built of black, opaque cloth. The illumination comes from a DLP projector displaying a black image with a single pixel set to white. A mirror deflects its light to the plane in which the sample resides, where the pixel is in maximal focus. In order to reduce unwanted illumination from the dark projector pixels, the mirror is chosen so that only a small fraction of the projector's frame actually hits the sample. We suppress all room lighting outside the tent, eliminating unwanted stray light.

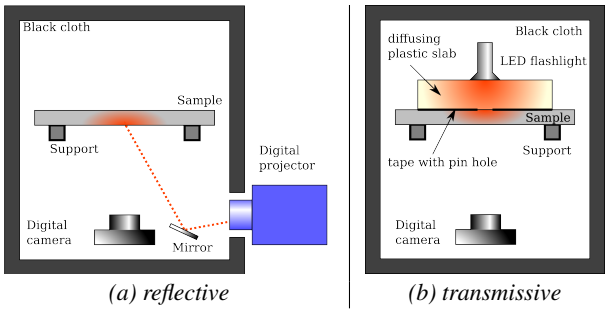

Figure 3: *Construction of the measurement setup.*

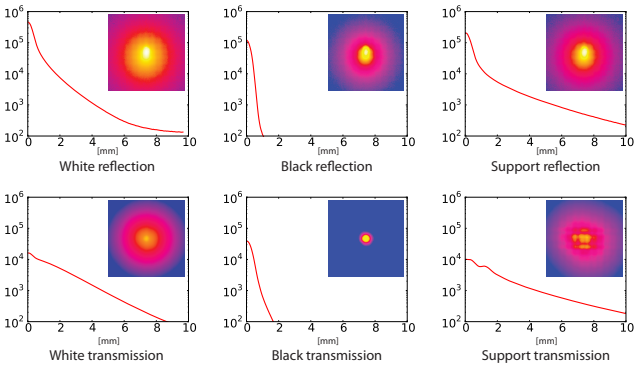

Figure 4: *Reflection and transmission profiles of the three materials used in our results, measured at 3mm, 0.75mm and 5.6mm respectively. A visualization of one HDR input images used in each measurement is included.*

We record multiple exposures ranging from $1/15$ to 30 seconds, and reconstruct a high dynamic range image. We also capture an additional series of exposures with all projector pixels set to black, and subtract their contribution from the image to obtain a precise reconstruction of the surface appearance when illuminated by the single projector pixel.

We detect the pixel location of highest radial symmetry in the input image, and by radial averaging we obtain the reflective profile with a sample spacing of 44 μm. Figure 4 shows logarithmic plots of the extracted profiles. We rescale the curves such that the 2D integral of their rotational extrusions match the measured albedo.

Transmissive profile. The recording of transmissive profiles faces challenges similar to the reflective case with respect to contrast and precision. However, due to forward scattering in the materials, it is even more important to create an illumination condition that is spatially concentrated and directionally uniform.

We use a thick, multiply scattering plastic slab to directionally disperse the light of a white LED flashlight, and black, opaque tape with a pin hole on the slab's bottom to restrict the light to a small aperture. Thus, we can achieve a better estimate of the overall profile than previous, directional light approaches [Tong et al. 2005; Goesele et al. 2004] can supply. Since the diffusing slab is in direct contact with the sample, the slab provides the illumination with the desired small spatial extent. Figure 3 (right) illustrates the construction. The processing of the acquired data is analogous to the reflective profile case.

5 Forward Problem

In this section, we describe how to compute the appearance of a composite material made from layers of homogeneous materials, given a description of the homogeneous materials and their thicknesses. While many of the ideas are adapted from previous work, our formulation allows any number of layers with arbitrarily variable thicknesses while staying within the framework of scattering profiles (without conversion to scattering/absorption coefficients and back). We use the Hankel transform [Weisstein 2010], instead of the inefficient 2D Fourier transform which requires N^2 instead of N coefficients to represent an N-sample scattering profile.

5.1 Representation of homogeneous materials

We represent materials by radially symmetric transmission and reflection profiles, $T_h(r)$ and $R_h(r)$, for any given material thickness h. Similarly to Donner and Jensen [2005], we assume the profiles do not depend on the angle of incidence and exitance, and additionally we ignore Fresnel effects. We can think of the profiles $R_h(r)$ and $T_h(r)$ as turning irradiance (uniform incoming light) to radiosity (uniform outgoing light). This is appropriate even for materials exhibiting forward scattering or high absorption, such as the black material used in our results.

Volume representation. We could also consider representing the input materials by specifying volume parameters such as σ_a (absorption), σ_s (scattering), $p(\omega)$ (phase function), and η (index of refraction). This has good accuracy even for optically thin materials, but the main problem of this approach is that the resulting appearance is more difficult to compute efficiently. Furthermore, every conversion between scattering profiles and volume parameters can only be approximate.

5.2 Composition of layers

As noted in [Donner and Jensen 2005], combining two slabs with scattering profiles given by T_1, R_1 and T_2, R_2 will result in a new slab with profiles that can be computed via series of convolutions. For example, the top-down transmission of the composite will be:

$$T_{12} = T_1 * T_2 + T_1 * R_2 * R_1 * T_2 + T_1 * R_2 * R_1 * R_2 * R_1 * T_2 + \cdots$$

where $*$ denotes 2D convolution. This is because for the light to get through the two slabs, it can either transmit straight through both, or inter-reflect between the layers on the way any number of times. This infinite series of convolutions can be conveniently computed by applying the Hankel transform, mapping them into a space where convolutions of radially symmetric functions become element-wise multiplications of their 1D profiles:

$$\mathcal{T}_{12} = \mathcal{T}_1\mathcal{T}_2 + \mathcal{T}_1\mathcal{R}_2\mathcal{R}_1\mathcal{T}_2 + \cdots = \frac{\mathcal{T}_1\mathcal{T}_2}{1 - \mathcal{R}_1\mathcal{R}_2}.$$

and a similar idea can be applied to reflection. The Hankel transform is given by

$$\hat{f}(k) = 2\pi \int_0^\infty f(r)\, J_0(2\pi r k)\, r\, dr,$$

where $J_0(x)$ is the 0-th order Bessel function of the first kind. We use the discrete Hankel transform [DHT 2010] to compute this integral, which is equivalent to a 2D Fourier transform (up to a tiny error caused by resampling to the roots of the Bessel function), but has the advantage that only N rather than N^2 coefficients are required in the frequency domain. Thus a homogeneous slab in our

framework is given by two N-dimensional vectors of Hankel frequency coefficients, one each for reflection and transmission.

When stacking two general layers, each of which can be a composite of more than one homogeneous layer, the equations become slightly more complicated [Donner et al. 2008]:

$$\mathcal{R}_{12}^{\cup} = \mathcal{R}_1^{\cup} + \frac{\mathcal{T}_1^2 \mathcal{R}_2^{\cup}}{1 - \mathcal{R}_1^{\cap} \mathcal{R}_2^{\cup}} \qquad \mathcal{R}_{12}^{\cap} = \mathcal{R}_2^{\cap} + \frac{\mathcal{T}_2^2 \mathcal{R}_1^{\cap}}{1 - \mathcal{R}_1^{\cap} \mathcal{R}_2^{\cup}} \tag{3}$$

$$\mathcal{T}_{12} = \frac{\mathcal{T}_1 \mathcal{T}_2}{1 - \mathcal{R}_1^{\cap} \mathcal{R}_2^{\cup}} \tag{4}$$

Here we denote top reflection as $\mathcal{R}^{\cup}$ and bottom reflection as $\mathcal{R}^{\cap}$. Thus the appearance of a stacking of multiple homogeneous layers is given by a triple of N-dimensional vectors of Hankel coefficients: $\mathcal{R}^{\cup}, \mathcal{R}^{\cap}$ and $\mathcal{T}$.

5.3 Continuous change of layer thickness

Given profiles $\mathcal{T}_h$ and $\mathcal{R}_h$ of a homogeneous material, we describe how to compute the profiles for a different thickness of h', by extending the theory of Kubelka and Munk [1931]. The theory studies how the scattering/absorption properties of a layer of paint and the underlying substrate determine the albedo (reflectance) of the paint, ignoring lateral transport and only considering the variation of light across the depth of the layer. However, as noted by [Donner and Jensen 2005], the resulting equations are conceptually identical to the profile stacking equations, except the functions $\mathcal{T}_h$ and $\mathcal{R}_h$ are scalar-valued in Kubelka-Munk theory and vector-valued in our case. In other words, the theory applies separately to all N Hankel coefficients of our profiles. Therefore, $\mathcal{T}_h$ and $\mathcal{R}_h$ as functions of h will be given by expressions predicted by Kubelka-Munk theory [Cortat 2004], remembering that $\mathcal{T}_h, \mathcal{R}_h, \alpha, \beta$ and γ are vectors of size N:

$$\mathcal{R}_h = \frac{\sinh(\gamma h)}{\alpha \sinh(\gamma h) + \beta \cosh(\gamma h)}$$

$$\mathcal{T}_h = \frac{\beta}{\alpha \sinh(\gamma h) + \beta \cosh(\gamma h)},$$

where

$$\alpha = 1 + \frac{\mathcal{K}}{\mathcal{S}} \qquad \beta = \sqrt{\alpha^2 - 1} \qquad \gamma = \sqrt{\mathcal{K}\left(\mathcal{K} + 2\mathcal{S}\right)}.$$

The constants α and γ in turn depend on the "scattering coefficient" $\mathcal{S}$ and "absorption coefficient" $\mathcal{K}$ of the material:

$$\mathcal{S} = \lim_{h \to 0} \frac{\mathcal{R}_h}{h} \qquad \mathcal{K} = \lim_{h \to 0} \frac{1 - \mathcal{R}_h - \mathcal{T}_h}{h}$$

Note that $\mathcal{S}$ and $\mathcal{K}$ are not the standard scattering and absorption coefficients of the material, σ_a and σ_s; instead, they can be seen as their vector-valued extensions. The limits can be efficiently computed by iteratively applying the following equations that halve the thickness of a slab:

$$\mathcal{R}_h = \frac{\mathcal{R}_{2h}}{\mathcal{T}_{2h} + 1} \qquad \mathcal{T}_h = \sqrt{\mathcal{T}_{2h}(1 - \mathcal{R}_h^2)}$$

Note that the above equations require the Hankel coefficients to be between 0 and 1; the upper bound is implied by energy conservation, while we enforce the lower bound by truncating negative coefficients (rare in measured profiles).

6 Inverse Problem

We define a *stacking* to be a sequence of materials with associated thicknesses. Given m homogeneous materials $M_1, \ldots, M_m$, and a desired target reflection profile, we would like to find a stacking S that produces a reflection profile close to a desired target profile $\mathcal{R}_t$ (as a vector of Hankel coefficients). That is, we want to solve the optimization problem

$$\arg\min_S \|\mathcal{R}(S) - \mathcal{R}_t\|^2, \tag{5}$$

where $\mathcal{R}(\mathcal{S})$ is the reflection of the stacked composite. We will assume the L_2-norm in this section, but we also support a simple extension where the first Hankel coefficient (the "DC component") can be weighted by a user-specified amount to better match albedo. Finding the ideal error metric for matching BSSRDF profiles in frequency space requires future work.

We constrain the total number of layers to be less than n and the total thickness of the stacking to be less than $h_{\max}$. (We can assume that the stacking can always be padded by black material on the bottom if more physical thickness is required.) This is a discrete-continuous optimization problem; the assignment of materials to layers is combinatorial while their thicknesses can vary continuously. Moreover, it is a highly non-linear problem: the function $R(S)$ computes the reflection profile through the non-linear formulas of composition and thickness change described above.

6.1 Discrete problem formulation

We simplify the fully general discrete-continuous formulation to a purely discrete problem: find an assignment of materials to a stacking of at most n layers, where each layer has a constant thickness of $h_0 = h_{\max}/n$. Clearly, this approximation will approach the original problem as $n \to \infty$. The discrete result is converted to a continuous one simply by joining adjacent layers of the same material.

An exhaustive search of the space of all material assignments is itself impractical. However, we will introduce several pruning strategies that allow for finding solutions in several seconds. We will build the composite layer top-down (by adding new layers to the bottom), where by top we mean closer to the observer.

6.2 The optimization algorithm

We define a *block* $B_{i,k}$ to be a layer made from material M_i of thickness kh_0. We also assume that n is a power of 2. This yields the following set of $m(\log_2(n) + 1)$ blocks:

$$\mathcal{B} = \{B_{1,1}, \ldots, B_{m,1}, B_{1,2}, \ldots, B_{m,2}, \ldots, B_{1,n}, \ldots, B_{m,n}\}.$$

The reflection and transmission profiles for these blocks can easily be precomputed from measurements at any convenient thickness.

Our optimization algorithm (shown in Algorithm 1) is an iterative deepening depth-first search with *branch and bound*, which is a non-linear optimization approach that prunes parts of the search space based on the following observation: if the upper bound on the objective function over one subset is smaller than the lower bound on another subset, we can prune out the latter subset, since there is no way a more optimal solution can be found there.

In our case, we require bounds on the optimal error we can achieve by starting with a given partial stacking. An upper bound is trivially given by the best solution found so far; lower bounds can be computed using the "too bright" and "too dark" heuristics described

Algorithm 1 The discrete optimization algorithm.

```
optimalNode ← rootNode // initialize to empty stacking
for depth = 1 to maxDepth do
   search(rootNode, depth) // iterative deepening

function search(node, depth):
if objective(node) ¡ objective(optimalNode) then
   optimalNode ← node // found new optimum
if depth = 0 or bound(node) > objective(optimalNode) then
   return // prune this node
for block in B do
   if block.material() = node.bottomMaterial() then
      if block.thickness() ≥ node.bottomThickness() then
         continue // ordering constraint failed
   if block.thickness() + node.totalThickness() > h_max then
      continue // thickness constraint failed
   childNode ← extendStacking(node, block)
   search(childNode, depth-1) // branch and recurse

function bound(node):
tooBrightBound ← ‖ max(R^∪ − R_t, 0)‖^2
tooDarkBound ← ‖ max(R_t − R^∪ − T^2/(1 − R^∩), 0)‖^2
return max(tooBrightBound, tooDarkBound)
```

below. The branching, in our case, is equivalent to choosing the next material to add to the *bottom* of the current stacking.

The root of the search tree is an empty stacking. The best solution so far, and its objective function value, are stored in global variables. Branching in the tree consists of trying every block from $\mathcal{B}$ on the bottom of the current stacking, where the blocks are tried in order of decreasing thickness. If a newly created node creates a stacking thicker than $h_{\max}$), or the lower bound estimate is greater than the current best objective function value, the node is discarded.

6.3 Error bounds and pruning strategies

Top-down stacking is advantageous because it allows for efficient lower bounds. In contrast, there can never be a lower bound for bottom-up stacking, because by adding more layers on top we can arbitrarily change the resulting profile.

Too-bright pruning strategy. To compute the lower bound on the error achievable by extending a given partial stacking, we use an approach based on the following intuition: if a given stacking already reflects more light than the optimization target, then this reflection cannot be decreased by inserting another reflective slab under it (or "behind it", from the observer's point of view). Therefore, the difference between the current and desired reflection will never decrease, so it gives a lower bound on the error of all stackings constructed by extending the current one. More precisely, if the current stacking is given by the triple $(\mathcal{R}^{\cup}, \mathcal{R}^{\cap}, \mathcal{T})$ and the newly added slab has a reflection profile $\mathcal{R}'$, then the new reflection will be:

$$\mathcal{R}^{\cup}_{\text{new}} = \mathcal{R}^{\cup} + \frac{\mathcal{T}^2\mathcal{R}'}{1-\mathcal{R}^{\cap}\mathcal{R}'}. \tag{6}$$

The term $\mathcal{T}^2\mathcal{R}'/(1-\mathcal{R}^{\cap}\mathcal{R}')$ has all non-negative coefficients. Therefore, coefficients can only increase, and any extension of this stacking will lead to an error of at least $\|\max(\mathcal{R}^{\cup}-\mathcal{R}_t, 0)\|^2$.

Too-dark pruning strategy. Similarly, if the current stacking absorbs a significant amount of light, there is a limit to how much brighter the reflection can become by inserting new slabs under the stacking. This limit is achieved by adding a perfectly white Lambertian material to the bottom, giving us the maximum achievable values of the Hankel coefficients. By inserting a vector of ones for $\mathcal{R}'$ in equation (6), we find that the error will be at least $\|\max(\mathcal{R}_t - \mathcal{R}^{\cup} - \mathcal{T}^2/(1-\mathcal{R}^{\cap}), 0)\|^2$.

Pruning based on ordering. Since there are multiple blocks of each material in $\mathcal{B}$, there are multiple ways to achieve the same material assignment to layers. For example, thickness $3h_0$ can be achieved as $2h_0 + h_0$ or $h_0 + 2h_0$. It only makes sense to add a block of material M if there is no block of material M of the same or smaller thickness at the bottom of the stacking already, otherwise an equivalent stacking must have been considered earlier (because the blocks are tried in an order of decreasing thickness).

We found these pruning strategies to be quite efficient; by using them, we can match a single BSSRDF curve in about 3-4 seconds with a single-threaded prototype implementation, using a depth-6 search and $n = 64$; an exhaustive search would be impractical.

7 Extension to Heterogeneity and Meshes

All previous discussion assumed layering of flat, homogeneous, laterally infinite slabs of constant thickness. However, the range of appearances of such slabs is not too interesting. The distinctive appearance of many real-world scattering materials (marble, skin, etc.) comes from their heterogeneity. Furthermore, in practical applications, one would like to produce arbitrary geometric objects. In this section, we propose extensions that address these limitations of the flat slab model.

7.1 Heterogeneous BSSRDFs

A heterogeneous BSSRDF can be represented (ignoring directionality and Fresnel effects) by a function $R(\vec{x}_i, \vec{x}_o)$. This is a 4D function (points $\vec{x}_i$ and $\vec{x}_o$ are on 2D surfaces), and can be acquired using a projector/camera setup [Peers et al. 2006].

Our goal is to design and print a layer stacking with spatially varying thicknesses that approximates the appearance expressed by the 4D function R. However, it is computationally difficult to predict R given a varying stacking; it would likely make the inverse problem intractable, since it would have to run the slow forward simulation many times.

Factorization. Instead, we would like to answer the following question: what is the "local" BSSRDF profile for a given point $\vec{x}$? This is an ill-posed problem, but an elegant solution was proposed by the factorization technique of Song et al. [2009] in the context of editing BSSRDF datasets. The idea is to find, for each (discretized) surface point $\vec{x}$, a local profile $P(\vec{x}, r)$ such that

$$R(x_i, x_o) \approx \sqrt{P(\vec{x}_i, r)P(\vec{x}_o, r)}, \quad \text{where } r = \|\vec{x}_i - \vec{x}_o\|. \tag{7}$$

This reduces the dimensionality of the data from 4 to 3. The functions P can be found, in discretized form, by taking the logarithm of both sides and solving the resulting large sparse least-squares problem, with a regularization term to enforce the smoothness of the resulting curves.

Per-profile optimization. Given the factored curves, we can solve the inverse problem on each surface point separately. In effect, we are approximating the heterogeneous forward problem by assuming that local layer depths determine the final 4D transport

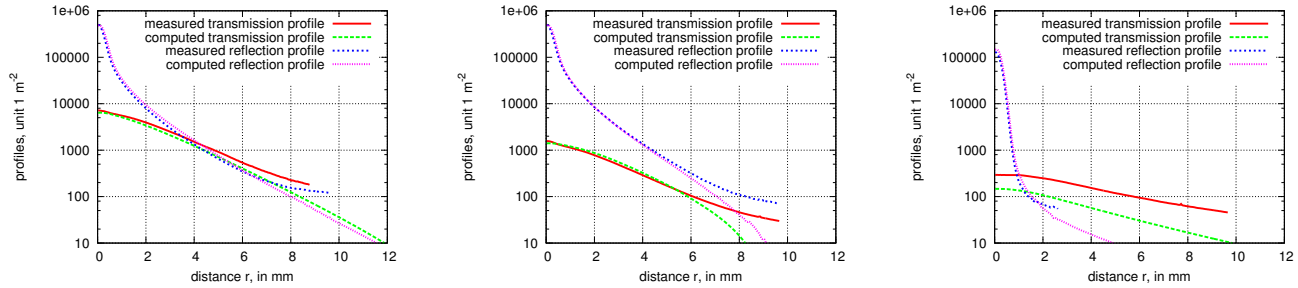

(a) Vero White, simulated two slabs of 2 mm each vs. a solid 4 mm slab

(b) Vero White, 4 mm in front of Vero Black, 0.25 mm vs. a manual stacking

(c) Vero Black, 0.50 mm in front of a 6 mm slab of support material in front of 2 mm Vero White

Figure 5: *Comparison of stacked material for recorded and simulated curves. There is a good match of curve shapes and scales for combinations of vero white and vero black materials. The match for the transmission of a stacking containing all three materials is off by a constant, which might be due to it being too dark to reliably measure, or due to strong forward scattering in the support material.*

through a combination of the stacking equations (3) and (4) and the factorization (7).

Running the optimization process for every surface point is time-consuming, but since many profiles P are very similar to each other, we run k-means clustering on the result of the factorization. We used $k = 100$ clusters; therefore, the optimization has to be run only 100 times per dataset (instead of over 76,000 times for our largest dataset).

Thickness interpolation. At this point, the layering for the output material is defined at discrete surface points, which we need to interpolate into continuous height fields to be passed to the printer. This requires the establishment of corresponding material layers in adjacent surface locations. We achieve this by inserting virtual layers of zero thickness until all stacks contain a periodic sequence of "white-support-black."

7.2 Layered meshes

Producing uniform layers on an arbitrary mesh is relatively straightforward. Note that the simplest approach of moving vertices along the normal direction often fails. We instead define a uniform grid around the object, compute the distance to the surface at each grid point using the 3D Euclidean distance transform [Saito and Toriwaki 1994] and construct the surface at all desired depths (including depth zero) by adding a constant to the grid values and repolygonizing using the marching cubes algorithm [Lorensen and Cline 1987]. This approach does not handle sharp creases on the object and will slightly blur them; we believe that using the L_p-averaged distance metric of Peng et al. [2004] would provide an effective solution for this problem.

Layers with varying thicknesses are somewhat more involved, but the above algorithm can be adapted as follows. The mesh is uniformly covered with points, and a vector of layer thicknesses is defined at each of these points. Then the pair of surfaces enclosing each layer is generated as follows: each distance grid point finds the closest surface point, from which it takes a pair of offsets corresponding to this layer. In the marching cubes algorithm, two polygons are generated for each cube, except in the case where the offsets are identical for all 8 vertices of the cube (handling the case of zero-depth layers).

8 Results

Hardware. We use the Objet Connex 500 3D printer for our results. This printer has the ability to use two different materials at a time. The materials used in our experiments are Vero White and Vero Black, both rigid materials of plastic appearance.

Support material. The printer uses a special *support material* for overhangs and to fill holes; some amount of the white material is mixed into the support material in a grid pattern, which gives it a more translucent appearance than white, but not completely clear. We use the support material as the third material in our experiments, simply by leaving holes in the meshes that are then automatically filled. We constrain the optimization process to never put the support material closer than 0.2 mm to the top of any stacking.

8.1 Stacking validation

In our first example, we investigate stackings of the base material samples, measure the resulting reflection and transmission curves, and compare them to the ones predicted by the forward process. The results are presented in Figure 5. In most cases, we have a very good match of both profile shapes and scaling; in some cases, the scale is off, but we believe this is due to the transmission of the composite being too low to reliably measure; the prediction is likely more accurate. Note that this level of accuracy was not easy to achieve, and some techniques used in our measurement setup (separate albedo measurements, diffusing plastic slab for transmission) were crucial in countering the effects of forward scattering.

8.2 Homogeneous slabs with varying scattering

To estimate and illustrate the gamut of possible BSSRDFs achievable with our three materials, we design the following experiment: We fabricate a translucent sample made from mostly support material (covered with a small amount of white material to ensure physical stability), and compute its reflection profile. Then we rescale the profile of the white material to match the albedo of the translucent sample, and use linear interpolation to create three in-between reflection curves. Then we run the optimization on these five curves, and fabricate the resulting stackings. As expected, the stackings vary in the amount of the more translucent support material. The top row of Figure 6 shows these five samples under room illumination, displaying a close match in albedos but no apparent difference in scattering. The second row shows the samples lit with a sharp shadow edge, clearly showing the differences in translucency.

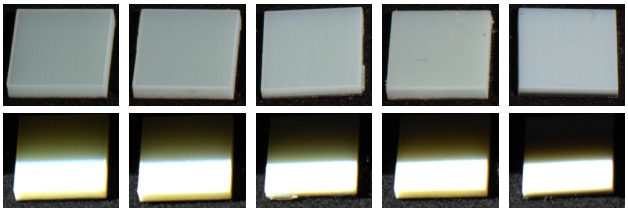

Figure 6: *Homogeneous slabs designed using our pipeline. The target curves have identical albedo but varying scattering width. The top row shows the samples under room illumination, illustrating a close match in albedo, while in the bottom row the samples are lit by a sharp shadow edge, displaying a decrease in translucency from left to right (see Section 8.2.)*

8.3 Slab fabrication from heterogeneous datasets

We use three heterogeneous BSSRDF datasets from previous work: chess and marble from [Peers et al. 2006], and artificial stone from [Song et al. 2009]. We convert these RGB datasets to luminance-only, since the 3D printer does not support color. We then apply the factorization of [Song et al. 2009] to find profiles $P(\vec{x}, r)$ for each point $\vec{x}$; our implementation takes about an hour per dataset.

We group the curves using k-means into 100 clusters, and run the optimizer on each. We constrain the stacking to a maximum of 5 mm, discretize into $n = 64$ layers, and use a maximum search depth of 6. We rescale the curves $P(\vec{x}, r)$ to maximum albedo of 0.6, which is close to the maximum achievable with 5 mm of material on a black background. We use the standard L_2-norm for chess and marble; for artificial stone, which is far outside of the printer gamut, we increase the DC coefficient weight to 10 to get a better albedo match. Using 100 clusters, the optimization process for a single dataset takes from 4.5 minutes (chess) to 6 minutes (stone), meaning that matching a single target curve takes about 3 seconds in our implementation. The multiple steps of the heterogeneous optimization process are illustrated in Figure 7.

Figure 8 shows the results of printing the layered heightfields constructed from the optimized stackings. These samples are printed in the original size as measured by the authors of the 4D datasets. Even though the gamut achievable with our base materials is not large, our results already illustrate a plausible reproduction of the original materials.

8.4 3D meshes

Figure 1 demonstrates our complete pipeline for output of physical objects with heterogeneous scattering materials. The bunny meshes were wrapped with heterogeneous BSSRDFs, using the polygonization process described in Section 7.2. Though the results are compelling, they highlight a number of additional challenges beyond accurate BSSRDF reproduction. For example, sharp features on the meshes pose an additional difficulty, since there might not be enough thickness to construct the full layering; we believe future work will resolve this issue. Furthermore, meshes with large amounts of support material very close to the surface can be unstable, which can be addressed by additional optimization constraints.

9 Limitations and Future Work

The materials we used, not being specifically engineered for this application, have a limited gamut with respect to the possible achievable translucency. We anticipate that future output technologies will enable better consistency and the use of a richer variety of

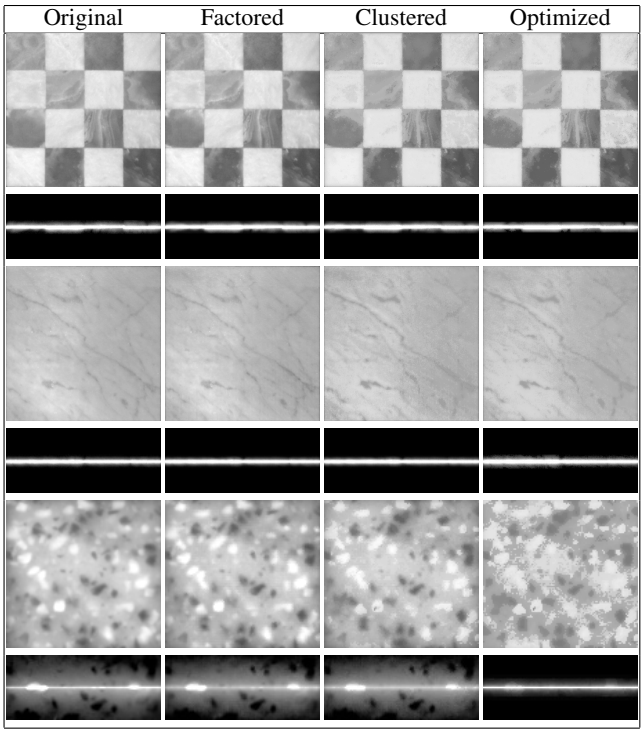

Figure 7: *An illustration of the steps in our heterogeneous pipeline, for the chess, marble and artificial stone datasets. The approximations introduced are usually acceptable, except in the case of artificial stone – here the translucency of the original sample is significantly outside of the gamut allowed by the current printer materials.*

basis materials, thus expanding the gamut of printable scattering behaviors, including the simultaneous control over translucency and color. More accurate forward simulation handling directional transmission and single scattering present in thin slabs will lead to more faithful prediction of appearance, while the incorporation of a directionally-sensitive model into our pipeline will accommodate a wider range of materials.

The printer we used already offers a choice between a more rough and a more shiny surface appearance of the created object; future devices might allow for finer control over the surface finish. Incorporating processing on a micro-geometric scale [Weyrich et al. 2009] might even further expand the creative possibilities in the design of object appearance.

10 Conclusions

Despite its limited material gamut, our prototype satisfies the goals of a flexible, easily-controlled system for output of materials with desired subsurface scattering. By incorporating measured basis material profiles, a goal-driven optimization, and effective strategies for pruning the search space, it efficiently maps user-specified spatially-varying translucency into a set of layers printable on a commercially-available 3D printer.

By freeing the user from the burden of exhaustively specifying material placement, our pipeline serves as an example of an emerging class of systems for bringing novel output devices into the computational realm. We are confident that the blueprint of our pipeline will shape these future systems: they will retain similar measurement,

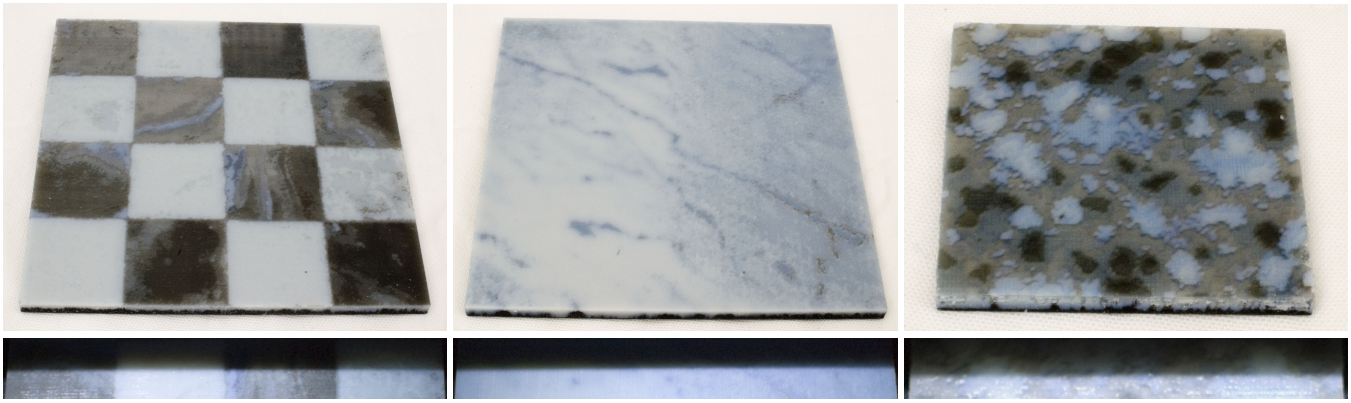

Figure 8: *Slabs fabricated by our process, approximating heterogeneous BSSRDF datasets: marble, artificial stone, and chess. The top row shows photographs under soft area illumination, clearly displaying albedo variation. In the bottom row, the samples are illuminated by a sharp shadow, demonstrating the translucency: the stone material is clearly more translucent than the marble, the chess data shows the variation in a single data set.*

simulation, and optimization stages while increasing the sophistication of each.

Acknowledgments

We would like to thank the anonymous reviewers, Ying Song et al. and Pieter Peers et al. for kindly providing the heterogeneous BSSRDF datasets, and Rob Wood and the Harvard Microrobotics Laboratory / Wyss Institute for kindly letting us use their 3D printer. The first author was supported by the Federal Share of program income earned by Massachusetts General Hospital on C06 CA059267, Proton Therapy Research and Treatment Center. Support was also provided by the National Science Foundation, grant CCF-0702580.

References

CORTAT, F. 2004. The Kubelka-Munk theory, applications and modifications. *Presentation for the graduate course on Optical properties of Paper, Linkoping University*.

2010. Discrete Hankel Transforms. `http://www.gnu.org/software/gsl/manual/html_node/Discrete-Hankel-Transforms.html`.

DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating Spatially-Varying Subsurface Scattering. *ACM Transactions on Graphics*.

DONNER, C., AND JENSEN, H. W. 2005. Light Diffusion in Multi-Layered Translucent Materials. *ACM Transactions on Graphics, Vol. 24, No. 3*, 1032–1039.

DONNER, C., WEYRICH, T., D'EON, E., RAMAMOORTHI, R., AND RUSINKIEWICZ, S. 2008. A Layered, Heterogeneous Reflectance Model for Acquiring and Rendering Human Skin. *ACM Transactions on Graphics, Vol. 27, No. 5* (Dec.).

DONNER, C., LAWRENCE, J., RAMAMOORTHI, R., HACHISUKA, T., JENSEN, H. W., AND NAYAR, S. 2009. An Empirical BSSRDF Model. *ACM Transactions on Graphics*.

FUCHS, M., RASKAR, R., SEIDEL, H.-P., AND LENSCH, H. P. A. 2008. Towards passive 6D reflectance field displays. In *ACM Transactions on Graphics*, ACM, New York, NY, USA, 1–8.

GHOSH, A., HAWKINS, T., PEERS, P., FREDERIKSEN, S., AND DEBEVEC, P. 2008. Practical Modeling and Acquisition of Layered Facial Reflectance. *ACM Transactions on Graphics, Vol. 27*.

GOESELE, M., LENSCH, H. P. A., LANG, J., FUCHS, C., AND PETER SIEDEL, H. 2004. DISCO: Acquisition of translucent objects. *ACM Transactions on Graphics, Vol. 23, No. 3*, 835–844.

HAASE, C. S., AND MEYER, G. W. 1992. Modeling pigmented materials for realistic image synthesis. *ACM Transactions on Graphics, Vol. 11, No. 4*, 305–335.

HANRAHAN, P., AND KRUEGER, W. 1993. Reflection from Layered Surfaces due to Subsurface Scattering. In *Computer Graphics (Proceedings of SIGGRAPH 93)*, 164–174.

HAWKINS, T., EINARSSON, P., AND DEBEVEC, P. 2005. Acquisition of time-varying participating media. *ACM Transactions on Graphics, Vol. 24, No. 3*, 812–815.

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*, 511–518.

KUBELKA, P., AND MUNK, F. 1931. Ein Beitrag zur Optik der Farbanstriche. *Zeitschrift für technische Physik, Vol. 12*, 593–601. English translation by Steve Westin.

LORENSEN, W. E., AND CLINE, H. E. 1987. Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics (Proceedings of SIGGRAPH 87), Vol. 21, No. 4*, 163–169.

MATUSIK, W., AJDIN, B., GU, J., LAWRENCE, J., LENSCH, H. P., PELLACINI, F., AND RUSINKIEWICZ, S. 2009. Printing Spatially-Varying Reflectance. *ACM Transactions on Graphics, Vol. 28, No. 5*.

NARASIMHAN, S. G., GUPTA, M., DONNER, C., RAMAMOORTHI, R., NAYAR, S. K., AND JENSEN, H. W. 2006. Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics, Vol. 25, No. 3*, 1003–1012.

NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., GINSBERG, I. W., AND LIMPERIS, T. 1977. *Geometrical Considerations*

and Nomenclature for Reflectance. National Bureau of Standards.

PEERS, P., VOM BERGE, K., MATUSIK, W., RAMAMOORTHI, R., LAWRENCE, J., RUSINKIEWICZ, S., AND DUTRÉ, P. 2006. A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics, Vol. 25, No. 3*, 746–753.

PENG, J., KRISTJANSSON, D., AND ZORIN, D. 2004. Interactive modeling of topologically complex geometric detail. *ACM Transactions on Graphics, Vol. 23, No. 3*.

PHARR, M., AND HANRAHAN, P. 2000. Monte Carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proceedings of ACM SIGGRAPH 2000*, 75–84.

SAITO, T., AND TORIWAKI, J. I. 1994. New algorithms for Euclidean distance transformations of an n-dimensional digitized picture with applications. *Pattern Recognition, Vol. 27*, 1551–1565.

SONG, Y., TONG, X., PELLACINI, F., AND PEERS, P. 2009. SubEdit: A Representation for Editing Measured Heterogeneous Subsurface Scattering. *ACM Transactions on Graphics, Vol. 28, No. 3*.

STAM, J. 1995. Multiple scattering as a diffusion process. In *Rendering Techniques*, 41–50.

STAM, J. 2001. An Illumination Model for a Skin Layer Bounded by Rough Surfaces. In *Rendering Techniques*, 39–52.

TARIQ, S., GARDNER, A., LLAMAS, I., JONES, A., DEBEVEC, P., AND TURK, G. 2006. Efficient Estimation of Spatially Varying Subsurface Scattering Parameters. In *Vision, Modeling, and Visualization*.

TONG, X., WANG, J., LIN, S., GUO, B., AND YEUNG SHUM, H. 2005. Modeling and Rendering of Quasi-Homogeneous Materials. *ACM Transactions on Graphics, Vol. 24, No. 3*, 1054–1061.

WANG, J., ZHAO, S., TONG, X., LIN, S., LIN, Z., DONG, Y., GUO, B., AND SHUM, H.-Y. 2008. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Transactions on Graphics, Vol. 27, No. 1*, 1–18.

WEISSTEIN, E. W., 2010. Hankel Transform. From MathWorld – A Wolfram Web Resource. `http://mathworld.wolfram.com/HankelTransform.html`.

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, Vol. 25*, 1013–1024.

WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating Microgeometry for Custom Surface Reflectance. *ACM Transactions on Graphics, Vol. 28, No. 3*.