

Inverse Volume Rendering with Material Dictionaries

Ioannis Gkioulekas
Harvard University

Shuang Zhao
Cornell University

Kavita Bala
Cornell University

Todd Zickler
Harvard University

Anat Levin
Weizmann Institute




Figure 1: Acquiring scattering parameters. Left: Samples of two materials (milk, blue curacao) in glass cells used for acquisition. Middle: Samples illuminated by a trichromatic laser beam. The observed scattering pattern is used as input for our optimization. Right: Rendering of materials in natural illumination using our acquired material parameter values.

Abstract

Translucent materials are ubiquitous, and simulating their appearance requires accurate physical parameters. However, physically-accurate parameters for scattering materials are difficult to acquire. We introduce an optimization framework for measuring bulk scattering properties of homogeneous materials (phase function, scattering coefficient, and absorption coefficient) that is more accurate, and more applicable to a broad range of materials. The optimization combines stochastic gradient descent with Monte Carlo rendering and a material dictionary to invert the radiative transfer equation. It offers several advantages: (1) it does not require isolating single-scattering events; (2) it allows measuring solids and liquids that are hard to dilute; (3) it returns parameters in physically-meaningful units; and (4) it does not restrict the shape of the phase function using Henyey-Greenstein or any other low-parameter model. We evaluate our approach by creating an acquisition setup that collects images of a material slab under narrow-beam RGB illumination. We validate results by measuring prescribed nano-dispersions and showing that recovered parameters match those predicted by Lorenz-Mie theory. We also provide a table of RGB scattering parameters for some common liquids and solids, which are validated by simulating color images in novel geometric configurations that match the corresponding photographs with less than 5% error.

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

Keywords: scattering, inverse rendering, material dictionaries

Links: [DL](#) [PDF](#) [WEB](#)

1 Introduction

Scattering plays a critical role in the appearance of most materials. Much effort has been devoted to modeling and simulating its visual effects, giving us precise and efficient scattering simulation algorithms. However, these algorithms produce images that are only as accurate as the material parameters given as input. This creates a need for acquisition systems that can faithfully measure the scattering parameters of real-world materials.

Collecting accurate and repeatable measurements of scattering is a significant challenge. For homogeneous materials—which is the primary topic of this paper—scattering at any particular wavelength is described by two scalar values and one angular function. The *scattering coefficient* σ_s and *absorption coefficient* σ_a represent the fractions of light that are scattered and absorbed, and the *phase function* $p(\theta)$ describes the angular distribution of scattering. Measurement is difficult because a sensor almost always observes the combined effects of many scattering and absorption events, and these three factors cannot be easily separated. Indeed, for deeply-scattering geometries, similarity theory [Wyman et al. 1989] proves that one can analytically derive distinct parameter-sets that nonetheless produce indistinguishable images.

Most existing acquisition systems address the measurement challenge using a combination of two strategies (e.g., [Hawkins et al. 2005; Narasimhan et al. 2006; Mukaigawa et al. 2010]). First, they manipulate lighting and/or materials to isolate single-scattering effects; and second, they “regularize” the recovered scattering parameters by relying on a low-parameter phase function model, such as the Henyey-Greenstein (HG) model. These approaches can provide accurate results, but both of the employed strategies have severe limitations. The single-parameter HG model limits applicability to materials that it represents well; and this excludes some common natural materials [Gkioulekas et al. 2013]. Meanwhile, isolating single scattering relies on either: (a) diluting the sample [Hawkins et al. 2005; Narasimhan et al. 2006], which cannot be easily applied to solids or to liquids that have unknown dispersing media; or, (b) using structured lighting patterns [Mukaigawa et al. 2010], which provide only approximate isolation [Holroyd and Lawrence 2011; Gupta et al. 2011] and therefore induce errors in measured scattering parameters that are difficult to characterize.

We introduce an optimization framework that allows measuring homogeneous scattering parameters without these limitations. Our optimization undoes the effects of low-order scattering by inverting the radiative transfer equation (i.e., by inverting a random walk) using a combination of Monte Carlo rendering and stochastic gradient descent. We evaluate our optimization framework by creating a volumetric scanner that uses a camera and narrow-beam RGB sources to collect a handful of images of a material sample that resides in a box-shaped transparent glass cell (Figure 1, left). Once calibrated, this scanner provides images of low-order scattering in which the geometry is precisely known. Using these images we optimize a dictionary-based set of scattering parameters so that they produce re-rendered images that match the acquired ones.

We validate our results in two ways: one, we measure prescribed dispersions of nano-scale particles whose scattering parameters can be computed using Lorenz-Mie theory and show that our recovered parameters are in close agreement; two, we re-render color images of these materials in novel lighting configurations and show that they numerically match, within 5%, the new captured images.

The benefits of our approach are:

(1) We can accurately measure scattering for a broader set of liquids *and* solids since there is no need for precisely-isolated single-scattering effects. The thickness of the glass cell can be selected appropriately for each material; but it does not need to be chosen with excessive care since our optimization succeeds for a relatively broad range of thicknesses (anywhere between 0.1 to 10 times the mean free path of the material being measured).

(2) We can support general phase function measurement, and are not restricted to HG or other low parameter models, because our optimization incorporates a large material dictionary that allows phase functions to be any convex combination of hundreds of dictionary elements. These general phase functions are particularly visually important for accurate visual appearance of objects with thin features.

Our contributions include:

- An optimization framework to invert volumetric scattering using MC rendering and stochastic gradient descent.
- An acquisition setup to acquire homogeneous material scattering properties with physically accurate parameters.
- A table of RGB scattering parameters for a variety of common materials, both liquid and solid, as well as a publicly-available collection of tabulated phase functions.

The availability of physically-accurate scattering parameters with general phase functions can improve simulations of translucent material appearance to better match that in the real world.

2 Related Work

Inverse radiative transport is studied in graphics, as well as in geophysical, biomedical, and chemistry domains; Bal [2009] provides a comprehensive review. Problems can be grouped into three categories, according to the ratio of the material’s mean free path—the average distance a photon travels before it is scattered—to the size of the scattering volume. We discuss these three categories in turn, and then we discuss phase function models and surface-based appearance models.

Deep scattering and approximation by diffusion. Inverse problems in this category consider media that are optically-thick, so that photons scatter many times before being measured. Radiative transport is then modeled using the diffusion approximation, where the angular variation of the internal radiance is limited and the radiative

transport equation reduces to a partial differential equation [Ishimaru 1978]. The advantage of this approach is that it simplifies the inference problem, allowing efficient acquisition and rendering systems [Jensen et al. 2001; Donner and Jensen 2005] and, as demonstrated by Wang et al. [2008], the estimation of spatially-varying structure within a medium. In physics, the diffusion approximation is employed by diffusing-wave spectroscopy [Pine et al. 1990], which is used for applications such as particle sizing or the measurement of molecular weight. The diffusion approximation applies when high-order scattering is dominant, causing the phase function to be confounded with other scattering parameters [Wyman et al. 1989] and therefore reasonable to ignore. It is not appropriate for our application, where we seek material-specific phase functions that can accurately predict appearance for shapes that have arbitrary thin and thick parts.

Single scattering and direct methods. At the other extreme are optically-thin situations, where photons scatter only once before being measured. Scattering parameters can often be measured directly in these cases, using techniques like static or dynamic light scattering [Johnson and Gabriel 1994]. For graphics, Hullin et al. [2008] use fluorescent dyes to make qualitative observations on the scattering parameters of optically thin media exhibiting mostly single scattering. Hawkins et al. [2005] use a laser to measure albedo and a tabulated phase function of a sparse homogeneous aerosol. For liquids, Narasimhan et al. [2006] successively dilute samples with water until they are sparse enough to infer from single-scattering an HG phase function. Single-scattering has also been exploited to capture time-varying and spatially-varying wisps of smoke and sparse mixing liquids [Hawkins et al. 2005; Fuchs et al. 2007; Gu et al. 2008]. All of these techniques rely on manipulating materials so that single-scattering dominates, and while dilution can be used for aerosols and some liquids, it cannot be easily applied to solids, or to liquids whose dispersing medium is unknown and significantly different from water. This limitation motivates methods for suppressing multiple scattering without dilution. Techniques for particle sizing or molecular weight, for example, exploit cross-correlation properties of multiple temporal measurements [Pusey 1999], but these are specific to those applications and do not easily extend to our problem. For graphics, Mukaigawa et al. [2010] use high frequency lighting patterns to isolate single-scattering effects [Nayar et al. 2006], allowing direct access to the mean free path and a good initialization for an indirect (multi-scattering) optimization over an HG phase function parameter. Such lighting-based isolations of single scattering are potentially quite useful, but as discussed in the context of 3D surface reconstruction [Holroyd and Lawrence 2011; Gupta et al. 2011], they provide only approximate isolation, and there is currently no analysis of how this affects the accuracy of inferred scattering parameters.

Low-order scattering and indirect methods. Our approach is in this category, where the goal is use moderately-thick samples to infer scattering parameters by iteratively refining them until they predict measurements that are consistent with the acquired ones (e.g., [Singer et al. 1990; Mukaigawa et al. 2010]). Most of these approaches focus like we do on homogeneous materials, since this is already very challenging. (A notable exception is Antyufeev [2000], who used regularized estimation to infer spatially-varying phase functions.) A common strategy in biomedical and physics domains is to simplify calculations by using planar slabs and spatially-uniform lighting that reduces the relevant scene geometry from three dimensions to two (e.g., [Chen et al. 2006; Prah et al. 1993; McCormick and Sanchez 1981]). But this has the significant disadvantage of limiting access to angular scattering information, thereby increasing reliance on restrictive low-parameter phase function models. In contrast, our optimization applies to any geometrical configuration with any incident light field, as long as both

are precisely calibrated. This allows using narrow-beam illumination, which improves access to angular scattering information and allows considering a much richer space of phase function models.

Phase function models. Most existing approaches to inverse radiative transfer use the single parameter Henyey-Greenstein [1941] model or other low-parameter models [Reynolds and McCormick 1980], but these can only be accurate for materials they represent well. The shape of the phase function is important for appearance, especially for objects with thin parts, and as recently shown by Gkioulekas et al. [2013], there are common materials that are not well-represented by the HG model. We avoid the restrictions of low-parameter models through the use of a phase function dictionary with hundreds of elements. This is similar in spirit to dictionary-based BRDF representations used to analyze and edit opaque scenes without being restricted to any particular analytic BRDF model (e.g., [Lawrence et al. 2006; Ben-Artzi et al. 2008]).

Surface reflectance fields and BSSRDF. There are a number of acquisition systems devoted to recovering surface-based descriptions of light transport through translucent objects [Debevec et al. 2000; Goesele et al. 2004; Tong et al. 2005; Peers et al. 2006; Donner et al. 2008]. These provide mappings between the input and output light fields on a specific object’s surface, and they do so without explicitly estimating all of the internal scattering parameters. They have the advantage of being very general and providing accurate appearance models for heterogeneous objects with complex shapes. Our goal is very different. We seek scattering material parameters that are independent of geometry, so that we can easily edit these materials and accurately predict their appearance when sculpted into *any* geometric shape.

3 Volume light transport

Scattering occurs as light propagates through a medium and interacts with material structures. There are many *volume events* that cause absorption or change of propagation direction. This process has been modeled by the *radiative transfer equation* (RTE) [Chandrasekhar 1960; Ishimaru 1978]:

$$\begin{aligned} (\omega^T \nabla) L(\mathbf{x}, \omega) &= Q(\mathbf{x}, \omega) - \sigma_t L(\mathbf{x}, \omega) \\ &+ \sigma_s \int_{\mathbb{S}^2} p(\omega, \psi) L(\mathbf{x}, \psi) d\mu(\psi), \end{aligned} \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^3$ is a point in the interior or boundary of the scattering medium; $\omega, \psi \in \mathbb{S}^2$ are points in the sphere of directions and μ is the usual spherical measure; $Q(\mathbf{x}, \omega)$ accounts for emission from light sources; and $L(\mathbf{x}, \omega)$ is the resulting light field radiance at every spatial location and orientation. The material is characterized by the triplet of macroscopic bulk parameters $k = \{\sigma_t, \sigma_s, p(\theta)\}$. Specifically, the *extinction coefficient* σ_t controls the spatial frequency of scattering events, and the *scattering coefficient* $\sigma_s \leq \sigma_t$ the amount of light that is scattered. The difference $\sigma_a = \sigma_t - \sigma_s \geq 0$ is known as the *absorption coefficient* and is the amount of light that is absorbed. Finally, the *phase function* p is a function on $\mathbb{S}^2 \times \mathbb{S}^2$ determining the amount of light that gets scattered towards each direction ψ relative to the incident direction ω . The phase function is often assumed to be invariant to rotations of the incident direction and cylindrically symmetric; therefore, it is a function of only $\theta = \arccos(\omega \cdot \psi)$ satisfying the normalization constraint

$$2\pi \int_{\theta=0}^{\pi} p(\theta) \sin(\theta) d\theta = 1. \quad (2)$$

We also adopt this assumption in the remaining of the paper and consider only phase functions of this type (we discuss some of the technical details related to this assumption in the Supplementary

Appendix C). Complementary to the above quantities, the following two parameters are also used for describing scattering behavior: the *mean free path*, equal to $d = 1/\sigma_t$, and the *albedo*, equal to $a = \sigma_s/\sigma_t$. In the following, we will use the parameter triplet k as an interchangeable term for scattering material.

We consider homogeneous materials in which scattering parameters do not depend on spatial location. Scattering parameters also exhibit perceptually dominant spectral dependency [Fleming and Bülthoff 2005; Frisvad et al. 2007] but for notational clarity we omit wavelength dependency.

3.1 Operator-theoretic formulation

We present the operator-theoretic formulation of the RTE that we will use to setup a tractable optimization procedure for volume rendering inversion. The specific formulation we use is tailored toward our optimization algorithm, but approximations of similar nature have been considered for rendering applications [Rushmeier and Torrance 1987; Bhate and Tokuta 1992]. In the following, we consider only points \mathbf{x} in the interior of the scattering medium; we discuss boundary conditions and other details for this formulation in the Supplementary Appendix A.

We begin by considering a finite difference approximation for the directional derivative [LeVeque 2007]

$$(\omega^T \nabla) L(\mathbf{x}, \omega) \approx \frac{1}{h} (L(\mathbf{x} + h\omega, \omega) - L(\mathbf{x}, \omega)). \quad (3)$$

Defining $L_i(\mathbf{x}, \omega) = hQ(\mathbf{x} - h\omega, \omega)$, after simple algebraic manipulation we obtain from Equation (1)

$$\begin{aligned} L(\mathbf{x}, \omega) &= L_i(\mathbf{x}, \omega) + (1 - h\sigma_t) L(\mathbf{x} - h\omega, \omega) + \\ &h\sigma_s \int_{\mathbb{S}^2} p(\omega^T \psi) L(\mathbf{x} - h\omega, \psi) d\mu(\psi). \end{aligned} \quad (4)$$

We define the following linear operator in terms of the material parameters $k = \{\sigma_t, \sigma_s, p(\theta)\}$, that acts on functions on $\mathbb{R}^3 \times \mathbb{S}^2$

$$\begin{aligned} \mathcal{K}_k(L)(\mathbf{x}, \omega) &\triangleq (1 - h\sigma_t) L(\mathbf{x} - h\omega, \omega) \\ &+ h\sigma_s \int_{\mathbb{S}^2} p(\omega^T \psi) L(\mathbf{x} - h\omega, \psi) d\mu(\psi). \end{aligned} \quad (5)$$

Intuitively, the action of \mathcal{K}_k can be viewed as a single step in the temporal propagation of a photon inside the medium. After traveling a distance of length h , the photon will transit in one of the following ways: 1) keep the same direction unaffected by the medium (probability $1 - h\sigma_t$); 2) scatter towards a new direction determined by the phase function p (probability $h\sigma_s$); 3) absorbed (probability $h(\sigma_t - \sigma_s)$). Consecutive applications of \mathcal{K}_k describe the time-resolved random walk the photon performs as it travels through the medium.

Using the \mathcal{K}_k operator, we can rewrite the RTE (4) in the form

$$L = L_i + \mathcal{K}_k L. \quad (6)$$

Solving Equation (6) for L , we can express the light field L as the result of a radiative transfer process \mathcal{R}_k on the input light field L_i :

$$L = \mathcal{R}_k L_i, \quad (7)$$

with

$$\mathcal{R}_k \triangleq (\mathcal{I} - \mathcal{K}_k)^{-1} = \sum_{j=0}^{\infty} \mathcal{K}_k^j. \quad (8)$$

The second equality in Equation (8) follows by applying the Neumann series expansion, as it applies for the inverse of $\mathcal{I} - \mathcal{K}_k$.

(We discuss the invertibility of $\mathcal{I} - \mathcal{K}_k$ in the Supplementary Appendix C.) It implies that the light field L resulting from the radiative transfer process of Equation (3) can be expressed as the sum of all orders of consecutive applications \mathcal{K}_k to the input light field L_i . That is, L corresponds to the asymptotic density of photons, when accounting for all of the intermediate positions of each photon after an arbitrary number of random walk steps. We refer to \mathcal{K}_k as the *single-step propagation operator* and to \mathcal{R}_k as the *rendering operator* for material k .

The validity of the above formulation for volume light transport relies on the accuracy of the approximation in Equation (3). In the limit that h goes to zero, the operator \mathcal{R}_k and the light field L of Equation (7) converge to the usual *volume rendering operator* [Jensen 2001] and the true light field inside the volume. In the Supplementary Appendix A, we discuss in detail this relationship, as well as analogies between operator \mathcal{K}_k and Equation (8), and their counterparts derived from the volume rendering equation. As explained there, how small h needs to be depends on the extinction coefficient σ_t of the medium: $h \ll \varepsilon/\sigma_t$. In our experiments, we use $\varepsilon = 0.01$.

4 Inverting volume scattering

We are interested in using images of an unknown material, to recover its scattering parameters $k = \{\sigma_t, \sigma_s, p(\theta)\}$. Formally, we consider a known 3D shape filled with the unknown material, illuminated by calibrated light sources L_i^m , and imaged by calibrated cameras to produce images I^m . Using the volume scattering Equation (6), we can write

$$I^m = \mathcal{S}^m L = \mathcal{S}^m (\mathcal{I} - \mathcal{K}_k)^{-1} L_i^m, \quad (9)$$

where the sampling operator \mathcal{S}^m describes the combination of light field rays measured by the corresponding camera. The terms L_i^m and \mathcal{S}^m fold in complete information about the 3D shape of the material volume, the relative light source and camera positions, as well as light interactions at the interface between the material volume and its surroundings. Modeling these light interactions involves a knowledge of the materials' refraction indices, and an assumption there is no further scattering between the material and the camera.

Given a set of illuminations $\{L_i^m, m = 1, \dots, M\}$ and their corresponding measurements $\{I^m, m = 1, \dots, M\}$, we cast the problem of inferring the material properties in an *appearance matching* framework: find the material parameters k that best reproduce the measurements in the least-squares sense

$$\min_k \sum_{m=1}^M w_m (\mathcal{S}^m (\mathcal{I} - \mathcal{K}_k)^{-1} L_i^m - I^m)^2. \quad (10)$$

where k is any permissible material parameter triplet $\{\sigma_t, \sigma_s, p(\theta)\}$. We weight the error for each radiance measurement with $w_m = \max\{c, (\bar{I}^m)^\alpha\}^{-1}$, $c, \alpha > 0$, to prevent the solution from overfitting only the brightest measurements. We selected $c = 0.01$ and $\alpha = 3$ using the experiments on synthetic data described in Section 6.

In the rest of this section we derive an optimization algorithm for the appearance matching problem. We express the material parameters as a convex linear combination of the material dictionary. We then differentiate the appearance matching error with respect to the mixing weights and derive an efficient optimization scheme based on stochastic gradient descent and Monte-Carlo rendering. Despite the highly non-linear problem, we show through simulations that the error surface is smooth without local minima and allows accurate reconstruction of material parameters.

4.1 Dictionary representation of materials

To better parameterize the search space in Equation (10), we use a *dictionary representation* for the materials. Specifically, consider a dictionary set of materials $D = \{k_n, n = 1, \dots, N\}$, where $k_n = \{\sigma_{t,n}, \sigma_{s,n}, p_n(\theta)\}$, and their corresponding single-step propagation operators $\{\mathcal{K}_{k_n}\}$. Then, for any weight vector π in the N -dimensional simplex Δ^N ,

$$\pi = [\pi_n] \in \mathbb{R}^N, \pi_n \geq 0, n = 1, \dots, N, \sum_{n=1}^N \pi_n = 1, \quad (11)$$

we can represent a novel *mixture material* $k(\pi) = \{\sigma_{t,\pi}, \sigma_{s,\pi}, p_\pi(\theta)\}$, in terms of the dictionary atoms as:

$$\sigma_{t,\pi} \triangleq \sum_{n=1}^N \pi_n \sigma_{t,n}, \quad \sigma_{s,\pi} \triangleq \sum_{n=1}^N \pi_n \sigma_{s,n}, \quad (12)$$

$$p_\pi(\theta) \triangleq \frac{\sum_{n=1}^N \pi_n \sigma_{s,n} p_n(\theta)}{\sum_{n=1}^N \pi_n \sigma_{s,n}}. \quad (13)$$

It is easy to see that if each $p_n(\theta)$ satisfies the normalization condition of Equation (2), so does $p_\pi(\theta)$.

In the following, we denote by $\mathcal{K}(\pi)$ and $\mathcal{R}(\pi)$ the single-step propagation (Equation (5)) and rendering (Equation (8)) operators, respectively, for the material $k(\pi)$. We denote the appearance matching error of the mixing weights π as

$$E(\pi) = \sum_{m=1}^M w_m (\mathcal{S}^m (\mathcal{I} - \mathcal{K}(\pi))^{-1} L_i^m - I^m)^2. \quad (14)$$

Then, we search for a convex combination of the material atoms in the dictionary D which best reproduces the captured images, by minimizing $E(\pi)$ over $\pi \in \Delta^N$.

To justify our use of convex combinations π and the mixing Equations (12)-(13), we present the following lemma.

Lemma 1. *For any vector $\pi \in \Delta^N$, the single-step propagation operator $\mathcal{K}(\pi)$ for the mixed material $k(\pi)$ defined by Equations (12)-(13), is a convex combination of the single-step propagation operators of the individual atoms, with the exact same mixing weights,*

$$\mathcal{K}(\pi) = \sum_{n=1}^N \pi_n \mathcal{K}_{k_n}. \quad (15)$$

Proof. Denoting $f = L(\mathbf{x} - h\boldsymbol{\omega}, \boldsymbol{\omega})$, and using Equation (5),

$$\begin{aligned} \sum_{n=1}^N \pi_n \mathcal{K}_{k_n}(L) &\stackrel{(5)}{=} \underbrace{\left(\sum_{n=1}^N \pi_n - h \underbrace{\left(\sum_{n=1}^N \pi_n \sigma_{t,n}\right)}_{\triangleq \sigma_t}\right)}_{=1} f \\ &+ h \int_{\mathbb{S}^2} \left(\sum_{n=1}^N \pi_n \sigma_{s,n} p_n(\boldsymbol{\omega}^T \boldsymbol{\psi})\right) f(\mathbf{x}, \boldsymbol{\psi}) \, d\mu(\boldsymbol{\psi}). \end{aligned} \quad (16)$$

For the expression of Equation (16) to be an operator of the form of Equation (5), we express the integral as:

$$\underbrace{\left(\sum_{n=1}^N \pi_n \sigma_{s,n}\right)}_{\triangleq \sigma_{s,\pi}} \int_{\mathbb{S}^2} \underbrace{\left(\frac{\sum_{n=1}^N \pi_n \sigma_{s,n} p_n(\boldsymbol{\omega}^T \boldsymbol{\psi})}{\sum_{n=1}^N \pi_n \sigma_{s,n}}\right)}_{\triangleq p_\pi(\theta)} f(\mathbf{x}, \boldsymbol{\psi}) \, d\mu(\boldsymbol{\psi}). \quad (17)$$

From Equations (16) and (17), we get exactly the single-step propagation operator of the material $k(\pi)$ in Equation (15). \square




Figure 2: Gradient computation (Equation (19)) for the dictionary-based appearance matching optimization problem (Equation (14)). For each pair of lighting and viewing directions, the gradient with respect to the weight of the n -th dictionary material requires computing a cascade of three rendering operations: a full volume rendering using the mixture material; then a single-step propagation using the corresponding atom material; then another full volume rendering using the mixture material. The resulting light fields are then sampled to produce the images used algebraically in Equation (19).

Lemma 1 shows that a convex combination of materials k_n can be directly identified with a convex combination of single-step propagation operators \mathcal{K}_{k_n} . As a result, the objective function of the optimization problem of Equation (14) has a much simpler functional dependence on the parameters π we optimize over, allowing us to derive a tractable optimization strategy as discussed in the following subsection. This property is the key motivator for our use of the finite-difference approximation in the operator-theoretic formulation of Section 3.1. As shown in the Supplementary Appendix A deriving an analogous result from a volume rendering formulation, in contrast, requires that the extinction coefficient be known beforehand and fixed for all the atoms in dictionary D .

4.2 Optimization Algorithm

To minimize the appearance matching error of Equation (14), we differentiate it with respect to the mixing weights π using the following lemma:

Lemma 2. For the operator $\mathcal{K}(\pi)$ defined in Equation (15), the following differentiation rule holds

$$\frac{\partial}{\partial \pi_n} (\mathcal{I} - \mathcal{K}(\pi))^{-1} = (\mathcal{I} - \mathcal{K}(\pi)) \mathcal{K}_{k_n} (\mathcal{I} - \mathcal{K}(\pi)). \quad (18)$$

This is a well-known result in the case of finite-dimensional matrices. In the Supplementary Appendix C, we provide a precise statement and proof of the lemma for the case of infinite-dimensional linear operators.

Using Equations (15) and (18), we can write the gradient of $E(\pi)$ with respect to each coordinate of the mixture weight vector π as

$$\begin{aligned} \frac{\partial E}{\partial \pi_n}(\pi) &= \sum_{m=1}^M 2w_m \left(\mathcal{S}^m \underbrace{(\mathcal{I} - \mathcal{K}(\pi))^{-1} L_i^m}_{\triangleq L_1^m} - \bar{I}^m \right) \\ &\quad \cdot \mathcal{S}^m \left(\underbrace{(\mathcal{I} - \mathcal{K}(\pi))^{-1} \mathcal{K}_{k_n}}_{\triangleq L_2^m} \underbrace{(\mathcal{I} - \mathcal{K}(\pi))^{-1} L_i^m}_{\triangleq L_1^m} \right). \end{aligned} \quad (19)$$

Equation (19) is crucial for our optimization. It shows that the gradient computation simplifies to rendering and sampling operations:

1. Render a light field L_1^m starting from input radiance L_i^m , and using the rendering operator $\mathcal{R}(\pi)$.
2. Apply the single-step propagation operator \mathcal{K}_{k_n} , corresponding to the n -th material k_n in the dictionary D , on L_1^m , producing a light field L_2^m .
3. Render a light field L_3^m , by applying the full rendering operator $\mathcal{R}(\pi)$ on L_2^m .
4. Apply the sampling operator \mathcal{S}^m on L_1^m and L_3^m , and evaluate their product (Equation (19)).

These gradient evaluation steps are summarized in Algorithm 1, and visualized in Figure 2.

Rendering. The fact that the appearance error gradient can be expressed as a sequence of rendering steps has an important practical implication: it allows us to evaluate it efficiently using Monte-Carlo rendering techniques [Dutré et al. 2006]. For the first stage, in our implementation we use the traditional volume rendering operator (described in the Supplementary Appendix A), as for small enough h it produces equivalent results to the rendering operator $\mathcal{R}(\pi)$ of our finite-difference formulation. We estimate separately the direct illumination term (which depends only on the assumed known radiance sources in the scene and is easy to compute) and the indirect component. For the latter, we use a Monte-Carlo particle tracing process to estimate I_1 , while simultaneously caching all intermediate particle positions in a set C_1 to form an approximation of L_1 . This process is described in Algorithm 3. Then the application of \mathcal{K}_k on L_1 is stochastically approximated as described in Algorithm 4: particles are uniformly sampled from C_1 , propagated by h , and then either scattered or absorbed. The results are cached in a set C_2 as an approximation of L_2 . Finally, samples from C_2 are used as sources for another full particle tracing process that directly estimates I_3 , without further caching. This is performed similar to Algorithm 3, but with the initialization of \mathbf{x} and ω in Step 2 replaced by an initialization from a particle drawn uniformly from C_2 and with the caching Step 10 omitted. We discuss further details about the algorithm we use to render $\frac{\partial E}{\partial \pi_n}$, including handling of Fresnel reflection and refraction, in the Supplementary Appendix B.

Stochastic gradient descent. The availability of stochastic estimates of the gradient makes *stochastic gradient descent* (SGD) algorithms attractive for minimizing Equation (14). Similar to standard gradient descent, SGD algorithms perform iterations of steps

Algorithm 1 ComputeGradient.

```

Input:  $\pi \in \Delta^N$ .
1: for  $n = 1$  to  $N$  do
2:    $g_n \leftarrow 0$ .
3:   for  $m = 1$  to  $M$  do
4:     Render  $L_1^m \leftarrow \mathcal{R}(\pi) L_i^m$ .
5:     Apply single-step propagation  $L_2^m \leftarrow \mathcal{K}_{k_n} L_1^m$ .
6:     Render  $L_3^m \leftarrow \mathcal{R}(\pi) L_2^m$ .
7:     Sample  $I_1^m \leftarrow \mathcal{S}^m L_1^m$ .
8:     Sample  $I_3^m \leftarrow \mathcal{S}^m L_3^m$ .
9:      $g_n \leftarrow g_n + 2w_m (I_1^m - \bar{I}^m) I_3^m$ .
10:  end for
11: end for
12: return  $g$ .

```

proportional to the negative of the stochastic estimates of the gradient. These algorithms only require that the estimates be unbiased; even if they are otherwise noisy, there exist convergence guarantees analogous to those of standard gradient descent, with the noise variance only affecting convergence speed. This behavior has an important practical implication: we can reduce the number of particles in Monte-Carlo evaluations of the gradient, and speed computation at the cost of introducing noise to the gradient estimate. As long as the rendering algorithm is unbiased such noisy gradient estimates are valid inputs to SGD. The noise due to the reduced number of particles somewhat increases the number of iterations. However, it is still possible to reduce the number of particles quite drastically, and achieve a significant speedup in terms of overall computation time [Bottou and Bousquet 2008].

As the vector π is constrained to lie on the simplex, we use *projected stochastic gradient descent* (PSGD). We denote by $g \in \mathbb{R}^N$ consecutive noisy estimates of the gradient of $E(\pi)$ obtained through Monte-Carlo rendering such that

$$\mathbb{E}[g_n] = \frac{\partial E(\pi)}{\partial \pi_n}. \quad (20)$$

We use them to iterate

$$\pi^{(t+1)} = \mathcal{P}_{\Delta^N} \left(\pi^{(t)} - \eta^{(t)} g^{(t)} \right), \quad (21)$$

where \mathcal{P}_{Δ^N} denotes the Euclidean projection operator to the simplex Δ^N [Duchi et al. 2008]. The step size is often chosen equal to $\eta^{(t)} = \frac{c}{\sqrt{t}}$. Though the speed of convergence depends on the proportionality constant c , in practice SGD is known to be robust to this selection. To cancel noise in individual steps, SGD returns as its final output the average of all T iterations, $\pi^{\text{opt}} = \frac{1}{T} \sum_{t=0}^T \pi^{(t)}$. This procedure is summarized in Algorithm 2.

The optimization problem we solve is highly non-linear and essentially involves inversion of the photon random walk process. Despite the non-linear formulation, our simulations in Section 6.2 show it allows an accurate reconstruction of material parameters. While an exact proof of this property is a subject for further research, all our simulations indicate that the error surface is very smooth and does not suffer from local minima, explaining the good convergence we are able to achieve.

4.3 Dictionary

The dictionary-based formulation of the appearance matching problem in Equation (14) can be used with any dictionary choice¹. Our own simple dictionary is described below. We start with the phase

¹We use the term “dictionary” because the phase function sets we use can be under- or over-complete and not strictly “bases” in the technical sense.

Algorithm 2 Solve appearance matching optimization problem.

```

1: Initialize  $\pi_n^{(0)} \leftarrow 1/N$ .
2: while not converged do
3:    $g^{(t)} \leftarrow \text{ComputeGradient}(\pi^{(t)})$ .
4:    $\pi^{(t+1)} \leftarrow \mathcal{P}_{\Delta^N} \left( \pi^{(t)} - \eta^{(t)} g^{(t)} \right)$ .
5: end while
6: return  $\pi^{\text{opt}} = \frac{1}{T} \sum_{t=0}^T \pi^{(t)}$ .

```

Algorithm 3 Adjoint particle tracing for computing L_1 and I_1 .

```

1: Let  $x_0$  be the location where the laser hits  $\partial S$ .
2:  $x \leftarrow x_0, \omega \leftarrow \omega_L, t \leftarrow 1, C_1 \leftarrow \emptyset$ .
3: while true do
4:    $t \leftarrow t \cdot a$ .
5:   Sample  $s$  from pdf  $p(s) = \sigma_t \exp(-\sigma_t s)$ .
6:    $x' \leftarrow x + s \cdot \omega$ .
7:   if  $x' \notin S$  then
8:     break
9:   end if
10:  Cache the particle location  $C_1 \leftarrow C_1 \cup \{(x, \omega)\}$ .
11:  Let  $\psi_1$  be the direction connecting  $x'$  and the camera.
12:  Let  $y_1$  be the intersection of ray  $(x', \psi_1)$  and the image sensor.
13:   $v_1 \leftarrow t \cdot p(\psi_1^T \omega) \cdot \exp(-\sigma_t \|x' - y_1\|) \cdot P_0 c/A$ ,
    where  $c$  is the total number of pixels on the sensor,  $A$  is the sensor’s surface area, and  $P_0$  is the source power.
14:  Add  $v_1$  to the corresponding pixel on the image sensor.
15:  Sample a direction  $\psi_2$  according to the phase function  $p$ .
16:   $x \leftarrow x', \omega \leftarrow \psi_2$ .
17: end while

```

Algorithm 4 Importance sampling for computing L_2 .

```

1:  $C_2 \leftarrow \emptyset$ 
2: Uniformly sample a pair  $(x_0, \omega_0) \in C_1$ .
3:  $x \leftarrow x_0, \omega \leftarrow \omega_0$ .
4:  $x \leftarrow x + h\omega$ .
5: Sample  $u$  uniformly in  $(0, 1)$ .
6: if  $u < h(\sigma_{t,k} - \sigma_{s,k})$  then
7:   terminate
8: end if
9: if  $h(\sigma_{t,k} - \sigma_{s,k}) < u < h\sigma_{t,k}$  then
10:  Sample a direction  $\psi$  according to the phase function  $p_k$ .
11:   $\omega \leftarrow \psi$ .
12: end if
13: Cache the particle location  $C_2 \leftarrow C_2 \cup \{(x, \omega)\}$ .

```




Figure 3: Phase functions in a tent dictionary with $N = 10$ atoms (each atom is colored uniquely for better visualization). To be valid probability distributions on the sphere, the atoms are normalized to satisfy Equation 2, and thus have varying magnitudes. The atoms centered at 0° and π° are shown cropped.

function component of the materials and then address the extinction and scattering components.

Phase functions. To allow the dictionary to be as general as possible, we aim to be able to express any phase function, that is, any cylindrically symmetric probability distribution on the sphere. Hence, we simply model the phase function as a piecewise linear

function of θ . Then, we can use a set of tent (triangular) functions, spaced equally over angular domain $\theta \in [0, \pi]$, to approximate it. Denoting the bins number by N and the bin spacing by $\theta_s = \pi / (N - 1)$, we use tent functions of width $2\theta_s$ and centered at points $\theta_n = (0, \theta_s, 2\theta_s, \dots, \pi)$. Each of the tent functions is normalized to satisfy the constraint of Equation (2). In Figure 3, we show the phase functions in a tent dictionary with $N = 10$. In our experiments, we use $N = 200$, which corresponds to a discretization step of 0.9° .

To avoid high frequency artifacts in the phase function solution we include in Equation (14) a quadratic regularization on its derivatives $\sum_n (\pi_n - \pi_{n+1})^2$.

Extinction and scattering coefficients. The definition of the atoms' extinction and scattering coefficients should ensure that the dictionary can represent materials with a wide range of σ_t, σ_s values. We select an upper bound $\sigma_{t,\max}$ on the desired extinction coefficients. Note that each scattering function of the form $\{\sigma_t, \sigma_s, p(\theta)\}$ with $0 \leq \sigma_t \leq \sigma_{t,\max}$ and $0 \leq \sigma_s \leq \sigma_t$, lies on the simplex spanned by $\{\sigma_{t,\max}, \sigma_{t,\max}, p(\theta)\}$, the purely absorptive atom $\{\sigma_{t,\max}, 0, \emptyset\}$, and an atom of the form $\{0, 0, \emptyset\}$ describing scattering-free propagation in vacuum. We use the symbol \emptyset to indicate that the last two atoms are independent of phase function (the phase function is undefined for these two media).

Therefore, we create a scattering dictionary including 200 atoms of the form $\{\sigma_{t,\max}, \sigma_{t,\max}, p_n(\theta)\}$ with the $p_n(\theta)$ defined above, plus the two purely absorptive atoms $\{\sigma_{t,\max}, 0, \emptyset\}$ and $\{0, 0, \emptyset\}$. In our experiments we set $\sigma_{t,\max} = 200 \text{ mm}^{-1}$, based on the results from Section 6.2.

Other parameterizations: Our specific choice of dictionary is aimed to represent any general phase function shape. Depending on the application, other dictionaries may be more appropriate, and some examples include: zonal spherical harmonics for low-frequency phase functions, phase functions derived from Mie theory [Bohren and Huffman 1983; Frisvad et al. 2007] when measuring dispersions, and compact dictionaries such as a set of a few Henyey-Greenstein and von Mises-Fisher functions [Gkioulekas et al. 2013] when a simple phase function model is sufficient. A small adaptation can also allow differentiating directly with respect to the single parameter of a Henyey-Greenstein function (the average cosine). Our optimization framework is quite attractive even for retrieving such simpler phase functions, since it alleviates the need for input measurements which isolate single scattering events.

5 Acquisition Setup Design

The optimization strategy described above is general enough to be applied to captured data with any geometry, as long as we can calibrate the 3D shape of the material, the relative position of the camera and light source, and the indices of refraction of the scattering material and its surroundings. Below we describe the physical acquisition setup we built, which is motivated by the simplicity of this calibration process and by some considerations related to the stability of the optimization problem. Inspiration is also drawn from analogous designs in [Jensen et al. 2001; Goesele et al. 2004; Wang et al. 2008] and physics [Johnson and Gabriel 1994]. A schematic and a photograph of our acquisition setup are shown in Figure 4. Further implementation details are provided in the Supplementary Appendix D.

Geometry. We cast the material we are interested in measuring into glass cells of variable thickness w . This allows us to create box-shaped material samples whose exact shape is known with very high accuracy. Furthermore, using micron-accurate smooth glass surfaces means that transition and refraction at the various material interfaces (material and glass, glass and air) can be easily simulated




Figure 4: Setup for scattering parameter acquisition. A sample cell is illuminated by collimated beams and imaged by a camera. Rotation stages to achieve arbitrary combinations of lighting and viewing directions. Top: schematic; bottom: implementation.

using Fresnel refraction and reflection laws. In our experiments, we use cells of widths $w = 1, 2.5, 5, \text{ and } 10 \text{ mm}$.

Imaging and lighting. We use an approximately orthographic camera with a high magnification macro lens (4.3° subtended angle and $1 : 3$ reproduction ratio) to sample the light field produced by the material volume. We use narrow (1 mm diameter) collimated beams to illuminate the sample. We use a configuration that allows illuminating either the sample surface imaged by the camera (frontlighting), or its opposite (backlighting). Through a combination of two motorized rotation stages, we can achieve different combinations of front and back lighting directions θ_f, θ_b and viewing directions θ_o . In our experiments, we use all possible combinations of $\theta_f, \theta_b, \theta_o \in \{5^\circ, 15^\circ, 25^\circ\}$, resulting in a set of 18 measurements per sample, 3 viewing directions times 3 frontlighting plus 3 backlighting directions.

The above combination of sample shape, camera, and illumination lends itself to accurate calibration. In the Supplementary Appendix C we justify the use of collimated beams mathematically. Similar to the BRDF arguments of Ramamoorthi and Hanrahan [2001], we argue that to maximize angular information the configuration should have broadband angular frequency content, and hence be as close as possible to a delta function. The use of both frontlighting and backlighting is motivated by the understanding that a backlighting beam produces measurements dominated by high-order scattering; such measurements are intuitively useful for determining the optical thickness of the material. Conversely, frontlighting results in measurements where low-order scattering is significant, and therefore is informative for the recovery of the material phase function.

Multi-chromatic measurements. Scattering parameters vary as functions of wavelength, and this spectral dependency can create perceptually important effects in appearance [Fleming and Bülthoff 2005; Frisvad et al. 2007]. To capture spectral variations we use monochromatic laser light at three RGB wavelengths,

488, 533, 635 nm, and solve the optimization problem of Equation (14) independently for each wavelength.

Index of refraction. To calibrate for the unknown material’s index of refraction, we use a set of additional measurements with backlighting such that $\theta_0 = \theta_b$ (corresponding to direct observation of the source in the absence of a medium). By measuring the shift in the location of the point-spread-function peak caused by refraction in these images, we can easily solve for the material index of refraction at each of the three wavelengths we use. We discuss this process in more detail in the Supplementary Appendix D. We have found our measurement procedure to be adequately accurate for our purposes, but if necessary more accurate measurements of the material index of refraction can be obtained using a refractometer. Additionally, in experiments on synthetic data, we found that small perturbations of the index of refraction (± 0.1) did not affect recovered scattering parameters considerably.

6 Experiments

We now demonstrate and validate our approach for acquiring scattering parameters. We begin with evaluations on synthetic data aimed at understanding the characteristics of our optimization problem. We then show results on two sets of measured materials. The first is a “validation set” of carefully-constructed nano-dispersions whose scattering parameters can be computed using Lorenz-Mie theory; this set provide a means for quantitative validation. The second set consists of everyday materials that are evaluated by their ability to produce accurate rendered images for novel geometries.

6.1 Capture and computation time

We first provide some quantitative information for the acquisition and inversion stages of our measurement pipeline. At the acquisition stage, as described in Section 5, for a single material we take measurements at three wavelengths and a set of 18 different scene configurations, for a total of 54 measurements. Each of these measurements is a high-dynamic range (HDR) image, composited from low-dynamic range images captured at 19 different exposures. In addition, for every material we measure, we capture a set of low-dynamic range calibration images. This process results in a total capture time of approximately 75 minutes per material. We provide more details about the calibration and high-dynamic range imaging procedures in the Supplementary Appendix D.

At the inversion stage, we solve the optimization problem of Section 4 on Amazon EC2 clusters of 100 nodes, with 32 computational cores and at least 20 GB of memory per node (required for the caching of intermediate light fields, as described in Section 4.2). We use the nodes to distribute the outer loop of Algorithm 1, that is, the gradient computation for each dictionary atom (for a dictionary of $N = 200$ atoms, each node is responsible for two atoms). The results are accumulated at a single master node, which then performs the gradient step of Algorithm 2, and the process is repeated for the number of iterations required until convergence is achieved. We found that processing one set of measurements requires approximately 200 iterations of the SGD algorithm. Overall, fitting one wavelength for a single material requires three to six hours, depending on the density of the material. We use our own C/C++ implementation, which we have optimized through experiments on synthetic data. However, computation could be reduced by further fine-tuning the various parameters involved, such as dictionary, camera spatial resolution, number of iterations, number of samples per rendering, and so on.

6.2 Experiments with Synthetic Data

The optimization problem of Section 4 involves the inversion of a random walk process that includes multiple scattering events and is




Figure 5: The error surface for 2D optimization problems. We consider a dictionary of three phase functions whose mixing weights lie on the simplex. The simplex is parameterized by its first two coordinates. Columns show results for three reference σ_t , and the rows show two different points as the correct reference in this space.

highly non-linear. Despite this, we almost always see in our experiments convergence to a solution that explains the measured data very well. This suggests that the error surface is fairly smooth. To provide more insight, we conduct a series of simulation experiments in which input image-sets are generated using small, artificial three-element dictionaries. Since the three mixing weights are constrained to a simplex, the set of phase functions spanned by three elements is a 2D space, allowing the entire cost surface to be visualized. For these experiments, we parameterize the 2D phase function space by the weights on the first two atoms (π_1, π_2), and in each experiment we choose a “ground truth” phase function (π_1^*, π_2^*) and compute for each (π_1, π_2) the \mathcal{L}^2 -difference between input images rendered with that phase function $\{I^m(\pi_1, \pi_2)\}_{m=1\dots M}$ and those rendered with the true one $\{I^m(\pi_1^*, \pi_2^*)\}_{m=1\dots M}$.

Results from six representative experiments are shown in Figure 5. Each row shows three separate experiments in which the true phase function is the same while the optical density σ_t differs. We find that the error surface has a clear minimum at the true value in all of these 2D experiments, and while an exact proof remains a subject for future research, the cost function appears to be very smooth and without spurious local minima, at least for these 2D problems.

In the next experiment with synthetic data, we compare accuracy on absorbing materials versus scattering materials, and on materials with varying optical densities. We consider a large set of artificial materials that are combinations of: (i) σ_t values sampled logarithmically in the interval $[0.01, \dots, 200 \text{ mm}^{-1}]$, for a total of 21 values; (ii) σ_s values corresponding, for each σ_t , to 21 albedo values, linearly sampled between $a = 0$ (purely absorptive) and $a = 1$ (purely scattering); and (iii) a set of eight different phase functions spanning a wide range of shapes. For each artificial material we render synthetic images using geometry that matches our setup (Section 5) with a sample width of $w = 1 \text{ mm}$. Sensor noise is an important consideration for this analysis, so we simulate image noise using photon (Poisson) noise with the parameters reported for two different commercial DSLR cameras [Hasinoff et al. 2010] (which is very large relative to the Monte Carlo rendering error). The noisy images are input to our optimization algorithm, and we measure error between the recovered parameters and the true ones. We use a tent dictionary with $N = 200$ atoms.

Figure 6 provides a summary of these experiments, by visualizing separately the relative error between the estimated and true values of (left to right): albedo $a = \sigma_s/\sigma_t$, extinction coefficient σ_t , and phase function $p(\theta)$. Each point in these tables corresponds to the percent error—averaged over the eight true phase function shapes—for distinct values of true albedo (horizontal axis) and extinction coefficient (vertical axis). These tables reveal which types of materials we can expect to measure accurately with our setup. Traveling from




Figure 6: Accuracy in the recovery of material parameters. The plots show recovery errors for albedo, extinction coefficient, and mean error for phase function, as a function of different values of albedo $a = \sigma_s/\sigma_t$ and mean free path $d = 1/\sigma_t$.

left to right in these tables makes a gradual transition from purely absorbing materials to purely scattering ones. Traveling from bottom to top moves through materials of increasing optical density, with the top being materials whose mean free path is two hundred times smaller than the sample width $d = 1/\sigma_t = w/200$, and the bottom being materials whose mean free path is one hundred times larger than the sample width $d = 1/\sigma_t = 100w$.

The first observation—based on the large, low-error regions in the center of the tables—is that estimation is accurate for a wide range of optical densities. This is a useful fact because it means the width of the glass cell need not be chosen with excessive care. We expect very accurate results as long as the sample width is within an order of magnitude of the material’s mean free path, and we expect graceful degradation when the width extends beyond this in either direction. For extremely optically-thin materials (lower rows in table), scattering events become very rare, and images are dominated by noise. For extremely optically-thick materials (top rows), the diffusion approximation [Jensen et al. 2001] becomes applicable, and recovering both the phase function and the scattering coefficient becomes ill-posed. In practice, we simply choose the width for each material sample from a discrete set of available glass cells (1, 2.5, 5, 10 mm) so that they look reasonably translucent under natural light; see examples in Figure 7.

Errors induced by extreme optical thinness and thickness at the top and bottom of these tables should be interpreted differently. If a material is excessively thin at sample width w , it is relatively easy to instead use a glass cell that is larger. This is less true for materials that are excessively dense, however, since it is physically challenging to cast materials into glass cells that are too small ($w < 1$ mm). Thus, in cases of extreme optical thickness, our setup will not provide material parameters that can accurately predict appearance on arbitrary geometries, but only for novel geometries at least as wide as the measured sample.

As expected, we also observe large errors in the estimated phase function when materials are extremely absorptive (left column of third table in Figure 6). These errors are somewhat of a computational artifact and have a limited impact on visual appearance. They occur because the appearance of these materials is dominated by attenuation due to absorption, so very little scattering is observed and there is little discernible information about the shape of the phase function. These errors do not impact our ability to predict material appearance, however, because the phase function makes little difference. Indeed, for purely absorptive materials (left-most column) there is no scattering at all, and the phase function can be defined arbitrarily without having any effect on appearance.



Figure 7: Measured materials in glass cells of width $w = 1, 10,$ and 2.5 mm, from left to right. It is not necessary for all of the cell to be filled, as long as there exists a homogeneous region of size comparable to the beam diameter (e.g., hand cream).

6.3 Validation materials

It is common in graphics to evaluate measured scattering parameters by demonstrating renderings of visually plausible results. This is an important benchmark, but it does not directly assess the accuracy of the recovered physical parameters. Since our goal is to produce parameters that are faithful to the true mean free path lengths and phase functions in an absolute sense, being able to directly validate the scattering parameters is crucial.

To achieve comparison to “ground truth” parameters, we capture liquid materials whose exact physical structure are known, similar to materials that are used to calibrate instruments used in a variety of domains for particle sizing or estimating molecular weight [Johnson and Gabriel 1994; Pine et al. 1990]. They are created by dispersing nano-scale spherical particles of known chemical composition into a homogeneous embedding medium of a different refractive index, using procedures that allow for very precise control of particle concentration, particle size distribution, and homogeneity. Given these parameters, Lorenz-Mie theory [Bohren and Huffman 1983; Frisvad et al. 2007] provides analytic expressions of the bulk material scattering parameters $\{\sigma_t, \sigma_s, p(\theta)\}$ at any wavelength.

We measure nanodispersions of two types. First, we measure dispersions of polystyrene spheres in water that are almost monodisperse (single particle size) and are precise enough to be traceable to NIST Standard Reference Materials. We measure three such dispersions², each having a 1% (w/v) concentration of particles at a different particle radius: 200, 500, or 800 nm. Second, we measure a spherical polydispersion of aluminum oxide particles (Al_2O_3) in water³, with an approximately known particle size distribution in

²Nanobead NIST Traceable Particle Size Standards, Polysciences, Inc.

³NanoArc Aluminum Oxide, Nanophase Technologies Corporation.

the range 20 – 300 nm and mean radius of 30 nm. We use glass cells of width $w = 1$ mm for all of these measurements, and instead of estimating the indices of refraction from image data, we use those predicted by Lorenz-Mie theory.

The results of our measurements are shown in Table 1. In all cases, the error in the recovered parameters is less than 5%. (Note that these materials are purely scattering, so $\sigma_t = \sigma_s$.) The largest error occurs for the aluminum oxide material, for which the particle size distribution is known much less precisely. Figure 8-left compares the green-channel phase functions recovered by our optimization (purple curves) to the ground truth phase functions predicted by Lorenz-Mie theory (dotted orange curves). We see that the matches are extremely close. As a reference, we compare both to Henyey-Greenstein phase functions; as the single parameter g of an HG phase function is equal to its average cosine, we plot (green curves) the HG phase function that have g values that are equal to the average cosine of the ground truth phase function. We note that their shapes deviate significantly from the ground-truth. This deviation is important for appearance, particularly for objects that have thin geometry with low-order scattering, where the phase function plays an important role visually. The middle columns demonstrate this by showing captured and fit (pseudo-colored) images of the materials under frontal laser illumination at a new angle (which was not used in optimization). The rightmost column shows cross-sections of the image intensities. The deviations of the HG fits from ground-truth lead to discernible differences between the images.

These experiments highlight the fact that simple, single-parameter phase function models can be insufficient for modeling the appearance of scattering materials, and it justifies our choice to fit higher-dimensional phase function models.

6.4 Other materials

Next, we use our acquisition setup and optimization algorithm to measure several common materials. They can be grouped roughly into three categories:

- Highly scattering liquids of varying viscosities; including mustard, shampoo, hand cream, liquid designer clay, and different types of milk.
- Highly absorbing liquids with limited scattering; including coffee, robittussin, olive oil, blue curacao liquor, and red wine.
- Solids that can be molded into the glass cells; such as different types of soap.

By “eyeballing” each sample under natural light, we choose glass cell widths so that each sample looks reasonably translucent under ambient lighting. The results we report were captured using width $w = 1$ mm for materials in the first and third categories, except for glycerine soap; and $w = 10$ mm for the second category and glycerine soap. Photographs of samples in 1 mm, 2.5 mm, and 10 mm cells are shown in Figure 7. For each sample, we estimate the index of refraction as described in Section 5, and these range from values of 1.33 (for milk, reduced milk, milk soap, and the water soluble liquids) to 1.47 (for olive oil and glycerine).

The measured parameters are shown in Table 2. We quantitatively evaluate the quality of the recovered scattering parameters in two ways. First, we report the *fitting error*, which is the average \mathcal{L}_2 image difference between input images and the corresponding images rendered with the recovered parameters, normalized by the \mathcal{L}_2 -norms of the input images. Second, we compute a measure of *generalization error* by: i) using the recovered parameters to render laser-illumination images with different sample widths and lighting directions; and ii) comparing these simulated images to captured

photographs in these same novel configurations. For the novel configurations, we use lighting angles $\theta_f, \theta_b \in \{10^\circ, 20^\circ, 30^\circ\}$ and glass cells with widths $w \in \{2.5, 5 \text{ mm}\}$. The generalization error for each material is reported as the average relative \mathcal{L}_2 image difference over the set of all novel configurations. As shown in the right two columns of Table 2, fitting errors are less than 4% and generalization errors below 5%.

Figure 9 shows the measured phase functions, each superimposed with an HG phase functions whose g -value is equal to the average cosine of the phase function we measure. Some of these phase functions are well approximated by the HG model but others, including hand cream, liquid clay, and mixed soap, are not. This set of tabulated phase functions is available at the project website.

As qualitative evaluation, Figure 1-right shows an image rendered with our recovered material parameters under natural lighting. From left to right, are milk soap and glycerine soap (top and bottom, respectively), olive oil, blue curacao, and reduced milk. The soap geometry corresponds to scanned molded cubes made of the corresponding materials. We see that the recovered material parameters successfully reproduce the color variations that are critical to the translucent appearance of these materials. This is most notable in the glycerine soap, where blue wavelengths scatter first and cause a reddish glow in the middle of the object, but it is also visible on the left edge of the milk soap and the top-right corner of the milk. A high-resolution version of Figure 1 and a visualization that highlights the color variations are shown in the Supplementary Appendix E. The scene file used for this figure is available at the project website.

7 Conclusions

We present an optimization framework for inverting the effect of multiple-scattering to recover scattering properties of homogeneous volumes from a handful of images. The approach does not require precise isolation of single scattering, and this enables the measurement of a broader set of materials, including both solids and liquids. The optimization also incorporates a large material dictionary and thereby avoids the restrictions of low-parameter phase function models. Our analysis and experiments show that we can recover accurate physical scattering parameters for a variety of materials.

Our current setup and optimization framework do not account for polarization or fluorescence phenomena. Polarization can be important for the appearance of materials with strongly polarization-dependent scattering properties or index of refraction (birefringence), such as crystalline materials. Experimentally, our optimization has been unable to find scattering parameters that match our setup’s images of microcrystalline wax, and this may be due to some combination of our simulator’s ignorance of polarization and our use of partially-polarized (laser plus fiber) light. Regarding fluorescence, we have verified that it has negligible impact on our measurements of the materials listed in Section 6. Our setup can be easily modified to measure the strongly fluorescent behavior of other materials, by including in addition to the hyperspectral camera a mechanism to control wavelength at the source side.

While we proposed one possible scanning configuration, our optimization could be used to infer scattering parameters from images captured from a variety of scene geometries and incident light fields. The only requirement is that both lighting and geometry be precisely calibrated. Our setup combines the benefits of high-frequency angular lighting (for stable optimization) and precise, stable calibration (for repeatability), but it limits measurements to three wavelengths and to solids that can be cast into glass cells of thickness within an order of magnitude of the mean free path. In principle, our optimization could be applied to images of more general solid objects, but this would require enhancing our setup to also

dispersion	σ_s predicted			σ_s measured			σ_s error (%)			phase function error (%)		
	R	G	B	R	G	B	R	G	B	R	G	B
polystyrene, 200 nm	17.220	28.363	36.517	17.078	28.650	36.823	0.825	1.012	0.838	3.031	1.143	3.672
polystyrene, 500 nm	59.082	79.557	88.626	58.431	79.023	88.062	1.102	0.671	0.636	3.181	2.676	1.359
polystyrene, 800 nm	65.757	70.438	68.589	66.976	71.544	69.146	1.853	1.570	0.812	2.623	2.117	1.251
Al ₂ O ₃ , 30 nm	47.341	93.389	129.870	48.536	96.004	132.695	2.524	2.800	2.175	3.712	4.298	3.108

Table 1: Measurements of validation materials (controlled nano-dispersions). Values for σ_s are reported in (mm^{-1}). Phase function error is given as \mathcal{L}_2 difference normalized by the \mathcal{L}_2 -norm of the reference phase function. All four validation materials have negligible absorption, resulting on both the predicted and measured values for σ_t to agree with those we report for σ_s to the third decimal.




Figure 8: Measurements of validation materials (controlled nano-dispersions). Left: For each material, we show for the green wavelength the theoretically predicted (dashed orange) and recovered (purple) phase functions, as well as the best Henyey-Greenstein (green) phase function fit. The recovered phase functions are in close agreement with the correct ones and the purple and orange curves tightly overlap. As another visualization, we show the images for a novel configuration: under frontal collimated laser illumination ($\theta_f = 25^\circ, \theta_o = 0^\circ$). We compare our phase function and the best Henyey-Greenstein fit (images are color-mapped for better visualization). The rightmost column shows a crosssection through the captured and re-rendered images for this configuration.

recover the object shape and its surface microstructure (BSDF).

In addition, combinations of our optimization framework with more sophisticated imaging configurations could improve the optimization’s stability and convergence rate. In particular, it may be fruitful to apply our optimization to images captured with high-frequency illumination [Mukaigawa et al. 2010], basis illumination [Ghosh et al. 2007], adaptive illumination [O’Toole and Kutulakos 2010], or transient imaging [Wu et al. 2012]. It is also possible that optimization schemes like ours will allow exploiting such imaging modalities to solve more challenging inverse problems, such as measuring heterogeneous scattering media.

Acknowledgments

We thank Henry Sarkas at Nanophase for donating material samples and calibration data. Funding by the National Science Foundation (IIS 1161564, 1012454, 1212928, and 1011919), the European Research Council, the Binational Science Foundation, Intel ICRI-CI, and Amazon Web Services in Education grant awards. Much work was performed while T. Zickler was a Feinberg Foundation Visiting Faculty Program Fellow at the Weizmann Institute.

References

ANTYUFEEV, V. 2000. *Monte Carlo method for solving inverse*

material	σ_s			σ_a			first moment			fitting error (%)	generalization error (%)
	R	G	B	R	G	B	R	G	B		
whole milk	100.920	105.345	102.768	0.013	0.013	0.041	0.954	0.963	0.946	2.0460	3.6344
reduced milk	57.291	62.460	63.757	0.007	0.007	0.024	0.954	0.957	0.942	1.346	2.039
mustard	16.447	18.536	6.457	0.057	0.061	0.451	0.155	0.173	0.351	3.377	4.201
shampoo	8.111	9.919	10.575	0.178	0.328	0.439	0.907	0.882	0.874	3.962	4.752
hand cream	20.820	32.353	41.798	0.011	0.011	0.012	0.188	0.247	0.265	2.652	3.221
liquid clay	37.544	48.250	67.949	0.004	0.004	0.005	0.312	0.442	0.512	3.431	4.532
milk soap	7.625	8.004	8.557	0.003	0.004	0.015	0.164	0.167	0.155	1.895	2.956
mixed soap	3.923	4.018	4.351	0.003	0.005	0.013	0.330	0.322	0.316	1.474	3.316
glycerine soap	0.201	0.202	0.221	0.001	0.001	0.002	0.955	0.949	0.943	3.840	3.920
robitussin	0.009	0.001	0.001	0.012	0.197	0.234	0.906	0.977	0.980	1.379	3.998
coffee	0.054	0.051	0.049	0.275	0.309	0.406	0.911	0.899	0.906	1.957	2.199
olive oil	0.041	0.039	0.012	0.062	0.047	0.353	0.946	0.954	0.975	2.287	3.846
blue curacao	0.010	0.012	0.021	0.083	0.048	0.011	0.955	0.973	0.979	2.704	4.857
red wine	0.015	0.013	0.011	0.122	0.351	0.402	0.947	0.975	0.977	3.034	3.192

Table 2: Scattering parameters of materials measured using our proposed acquisition setup and inversion algorithm. Values for σ_s and σ_a are reported in (mm^{-1}) . The average cosine of the measured phase functions is reported, while the entire phase functions are shown in Figure 9. Fitting and generalization errors are given as % \mathcal{L}_2 difference normalized by the \mathcal{L}_2 -norm of the reference image, averaged across the fitting and novel captured images respectively.




Figure 9: Phase functions of materials measured using our acquisition setup (purple), contrasted with the closest HG phase function (dashed green). See Table 2 for numerical values.

problems of radiation transfer, vol. 20. Inverse and Ill-Posed Problems Series, V.S.P. International Science.

BAL, G. 2009. Inverse transport theory and applications. *Inverse Problems* 25, 5.

BEN-ARTZI, A., EGAN, K., DURAND, F., AND RAMAMOORTHY, R. 2008. A precomputed polynomial representation for interactive BRDF editing with global illumination. *ACM Trans. Graph.* 27, 2.

BHATE, N., AND TOKUTA, A. 1992. Photorealistic volume rendering of media with directional scattering. In *Third Eurographics Workshop on Rendering*.

BOHREN, C., AND HUFFMAN, D. 1983. *Absorption and scattering of light by small particles*. Wiley-Vch.

BOTTOU, L., AND BOUSQUET, O. 2008. The Tradeoffs of Large Scale Learning. *NIPS*.

CHANDRASEKHAR, S. 1960. *Radiative transfer*. Dover.

CHEN, C., LU, J., DING, H., JACOBS, K., DU, Y., HU, X., ET AL. 2006. A primary method for determination of optical parameters of turbid samples and application to intralipid between 550 and 1630 nm. *Optics Express* 14, 16.

DEBEVEC, P., HAWKINS, T., TCHOU, C., DUIKER, H., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of SIGGRAPH 2000, Annual Conference Series*.

DONNER, C., AND JENSEN, H. 2005. Light diffusion in multi-layered translucent materials. *ACM Trans. Graph.* 24, 3.

DONNER, C., WEYRICH, T., D'ÉON, E., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2008. A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Trans. Graph.* 27, 5.

DUCHI, J., SHALEV-SHWARTZ, S., SINGER, Y., AND CHANDRA, T. 2008. Efficient projections onto the l_1 -ball for learning in high dimensions. *ICML*.

- DUTRÉ, P., BALA, K., AND BEKAERT, P. 2006. *Advanced global illumination*. AK Peters, Ltd.
- FLEMING, R., AND BÜLTHOFF, H. 2005. Low-level image cues in the perception of translucent materials. *ACM Transactions on Applied Perception (TAP)* 2, 3.
- FRISVAD, J., CHRISTENSEN, N., AND JENSEN, H. 2007. Computing the scattering properties of participating media using lorenz-mie theory. *ACM Trans. Graph.* 26, 3.
- FUCHS, C., CHEN, T., GOESELE, M., THEISEL, H., AND SEIDEL, H. 2007. Density estimation for dynamic volumes. *Computers & Graphics* 31, 2.
- GHOSH, A., ACHUTHA, S., HEIDRICH, W., AND O'TOOLE, M. 2007. BRDF acquisition with basis illumination. *IEEE CVPR*.
- GKIOULEKAS, I., XIAO, B., ZHAO, S., ADELSON, E., ZICKLER, T., AND BALA, K. 2013. Understanding the role of phase function in translucent appearance. *To appear in ACM Trans. Graph.* 32, 5.
- GOESELE, M., LENSCH, H., LANG, J., FUCHS, C., AND SEIDEL, H. 2004. Disco: acquisition of translucent objects. *ACM Trans. Graph.* 23, 3.
- GU, J., NAYAR, S., GRINSPUN, E., BELHUMEUR, P., AND RAMAMOORTHY, R. 2008. Compressive structured light for recovering inhomogeneous participating media. *ECCV*.
- GUPTA, M., AGRAWAL, A., VEERARAGHAVAN, A., AND NARASIMHAN, S. 2011. Structured light 3D scanning in the presence of global illumination. *IEEE CVPR*.
- HASINOFF, S., DURAND, F., AND FREEMAN, W. 2010. Noise-optimal capture for high dynamic range photography. *IEEE CVPR*.
- HAWKINS, T., EINARSSON, P., AND DEBEVEC, P. 2005. Acquisition of time-varying participating media. *ACM Trans. Graph.* 24, 3.
- HENYEV, L., AND GREENSTEIN, J. 1941. Diffuse radiation in the galaxy. *The Astrophysical Journal* 93.
- HOLROYD, M., AND LAWRENCE, J. 2011. An analysis of using high-frequency sinusoidal illumination to measure the 3d shape of translucent objects. *IEEE CVPR*.
- HULLIN, B., FUCHS, M., AJDIN, B., IHRKE, I., SEIDEL, H., AND LENSCH, H. 2008. Direct visualization of real-world light transport. *Vision, Modeling, and Visualization 2008*.
- ISHIMARU, A. 1978. *Wave propagation and scattering in random media*. Wiley-IEEE.
- JENSEN, H., MARSCHNER, S., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 2001, Annual Conference Series*.
- JENSEN, H. 2001. *Realistic image synthesis using photon mapping*. AK Peters, Ltd.
- JOHNSON, C., AND GABRIEL, D. 1994. *Laser light scattering*. Dover.
- LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.* 25, 3.
- LEVEQUE, R. 2007. *Finite Difference Methods for Ordinary and Partial Differential Equations, Steady-State and Time-Dependent Problems*. SIAM.
- MCCORMICK, N., AND SANCHEZ, R. 1981. Inverse problem transport calculations for anisotropic scattering coefficients. *Journal of Mathematical Physics* 22, 199.
- MUKAIGAWA, Y., YAGI, Y., AND RASKAR, R. 2010. Analysis of light transport in scattering media. *IEEE CVPR*.
- NARASIMHAN, S., GUPTA, M., DONNER, C., RAMAMOORTHY, R., NAYAR, S., AND JENSEN, H. 2006. Acquiring scattering properties of participating media by dilution. *ACM Trans. Graph.* 25, 3.
- NAYAR, S., KRISHNAN, G., GROSSBERG, M., AND RASKAR, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph.* 25, 3.
- O'TOOLE, M., AND KUTULAKOS, K. N. 2010. Optical computing for fast light transport analysis. *ACM Trans. Graph.* 29, 6.
- PEERS, P., VOM BERGE, K., MATUSIK, W., RAMAMOORTHY, R., LAWRENCE, J., RUSINKIEWICZ, S., AND DUTRÉ, P. 2006. A compact factored representation of heterogeneous subsurface scattering. *ACM Trans. Graph.* 25, 3.
- PINE, D., WEITZ, D., ZHU, J., AND HERBOLZHEIMER, E. 1990. Diffusing-wave spectroscopy: dynamic light scattering in the multiple scattering limit. *Journal de Physique* 51, 18.
- PRAHL, S., VAN GEMERT, M., AND WELCH, A. 1993. Determining the optical properties of turbid media by using the adding-doubling method. *Applied optics* 32, 4.
- PUSEY, P. 1999. Suppression of multiple scattering by photon cross-correlation techniques. *Current opinion in colloid & interface science* 4, 3.
- RAMAMOORTHY, R., AND HANRAHAN, P. 2001. A signal-processing framework for inverse rendering. In *Proceedings of SIGGRAPH 2001, Annual Conference Series*.
- REYNOLDS, L., AND MCCORMICK, N. 1980. Approximate two-parameter phase function for light scattering. *JOSA* 70, 10.
- RUSHMEIER, H., AND TORRANCE, K. 1987. The zonal method for calculating light intensities in the presence of a participating medium. In *Computer Graphics*, vol. 21.
- SINGER, J., GRÜNBAUM, F., KOHN, P., ZUBELLI, J., ET AL. 1990. Image reconstruction of the interior of bodies that diffuse radiation. *Science* 248, 4958.
- TONG, X., WANG, J., LIN, S., GUO, B., AND SHUM, H. 2005. Modeling and rendering of quasi-homogeneous materials. *ACM Trans. Graph.* 24, 3.
- WANG, J., ZHAO, S., TONG, X., LIN, S., LIN, Z., DONG, Y., GUO, B., AND SHUM, H. 2008. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Trans. Graph.* 27, 1.
- WU, D., O'TOOLE, M., VELTEN, A., AGRAWAL, A., AND RASKAR, R. 2012. Decomposing global light transport using time of flight imaging. *IEEE CVPR*.
- WYMAN, D., PATTERSON, M., AND WILSON, B. 1989. Similarity relations for the interaction parameters in radiation transport. *Applied optics* 28, 24.