

# An Evaluation of Computational Imaging Techniques for Heterogeneous Inverse Scattering

Ioannis Gkioulekas<sup>1</sup>, Anat Levin<sup>2</sup>, Todd Zickler<sup>1</sup>

<sup>1</sup>Harvard University    <sup>2</sup>Weizmann Institute of Science

**Abstract.** Inferring internal scattering parameters for general, heterogeneous materials, remains a challenging inverse problem. Its difficulty arises from the complex way in which scattering materials interact with light, as well as the very high dimensionality of the material space implied by heterogeneity. The recent emergence of diverse computational imaging techniques, together with the widespread availability of computing power, present a renewed opportunity for tackling this problem. We take first steps in this direction, by deriving theoretical results, developing an algorithmic framework, and performing quantitative evaluations for the problem of heterogeneous inverse scattering from simulated measurements of different computational imaging configurations.

**Keywords:** inverse scattering, computational imaging

## 1 Introduction

We consider the heterogeneous inverse scattering problem, where light can be controllably injected and measured at the boundary of a volume, and where we want to infer the scattering material parameters that vary internally. In the general form that we consider, this is an extreme multi-path problem. The volume can include parts of varying thickness, so that low-order scattering, mid-order scattering, and high-order scattering can all contribute substantially to the measurements; and the internal material varies spatially in terms of both absorption and angular scattering, so that the unknown variables number in the hundreds of thousands. Finding reliable solutions to this general problem would extend three-dimensional imaging to many types of turbid volumes (deep tissues, many gemstones, thick smoke and clouds) that cannot yet be accurately measured by any non-invasive, non-destructive means.

Despite decades of work on inverse scattering, the problem has yet to be considered at this level of generality and scale. But the growth of processing power and the accelerating development of computational imaging techniques, which allow unprecedented control and measurement of light at the boundary, are making this previously-intractable problem more interesting.

The most important questions to answer about the general heterogeneous inverse scattering problem are how to determine when a set of measurements is sufficient to reconstruct the internal volume; what internal ambiguities exist in the absence of sufficient measurements; and how to formulate and solve the

massive optimization problem. While we do not provide definite answers to these questions in this paper, we make progress in several directions.

First, we describe a mathematical model of the problem that encompasses the types of measurements that are obtainable using almost any form of computational imaging, including structured lighting, spatial probing, and transient imaging. Second, we use this model to derive theoretical results about aspects of internal material information that are un-recoverable, and measurement configurations that can reduce internal ambiguities among the material parameters. Third, we generalize recent optimization frameworks [14, 30], which are based on Monte Carlo rendering and stochastic gradient descent, to accommodate our more general material and measurement spaces. Fourth and finally, we use simulations to evaluate the utility of different computational imaging configurations for several heterogeneous inverse scattering problems. Our code and supplementary material are available at the project page [1].

## 2 Related Work

**Inverse scattering.** Inverse radiative transport is studied in graphics, physics, chemistry, and biomedical sciences. A review can be found in [4]. Existing algorithms for volumetric reconstruction of scattering materials can be roughly classified into three categories. Methods based on the *diffusion* approximation consider optically thick media where high-order scattering is dominant. This allows for simpler inference and has been used for the acquisition of both homogeneous [28, 9, 39] and heterogeneous materials [46]. However, it also introduces ambiguities between different scattering parameters [49, 51]. At the other extreme, methods based on the *single scattering* approximation assume that the unknown medium is so optically thin that all photons scatter only once. This allows directly measuring scattering parameters of media such as smoke and thin or dilutable liquids [20, 12, 17, 34]. A third class of methods seek to use all orders of scattering when solving appearance matching objectives to infer scattering parameters [42, 2, 14, 32, 30]. Our method falls in this category, by extending the algorithms of [14, 30] to apply to general heterogeneous media and different types of imaging techniques. Orthogonal to the above are techniques that, instead of volumetric reconstructions, recover surface-based descriptions of scattering materials (BSSRDF), which can be spatially-varying [8, 15, 41, 10].

**Computational imaging.** Different imaging techniques can be categorized as different ways to decompose photon contributions. *Pathlength decomposition* techniques, also referred to as *transient imaging*, separate photons in terms of the distance they travel from source to camera, and have been implemented using combinations of pulsed lasers with ultra-fast cameras [45, 48, 47], time-of-flight sensors [21, 29, 37, 22], and optical coherence tomography [24, 13]. *Spatial probing* techniques use camera-projector systems [38, 36] or interferometry [13] to decompose photons in terms of the endpoints of their paths. Techniques from the two categories have been combined into imaging systems that simultaneously

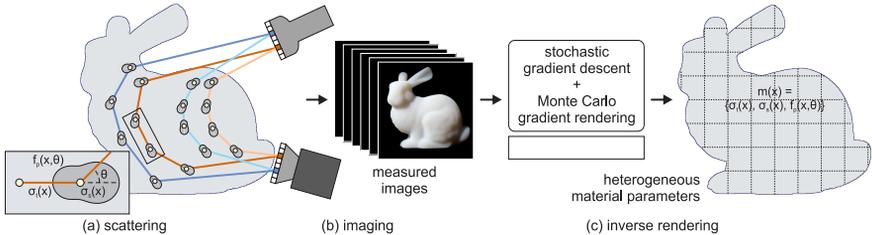


Fig. 1: (a) Photons traveling inside scattering materials perform random walks that depend on the material parameters. (b) Computational imaging techniques capture measurements of such materials by collecting different subsets of photons, depending on the paths they follow. (c) We present an inverse rendering algorithm that uses such measurements to recover the spatially-varying scattering material parameters.

decompose photons in terms of both pathlength and endpoints [37, 13]. Finally, structured light systems can be used to separate photons based on the number of times they bounce, into a direct (single bounce) and global (multiple bounces) component [35, 18, 43]. With respect to inverse scattering applications, pathlength decomposition has been used together with the diffusion approximation for measuring heterogeneous tissues [3, 5]; whereas direct-global separation has been combined with single-scattering analysis to simplify appearance matching problems for scattering materials [33]. In our work, we describe all these types of decompositions in a unified theoretical framework that allows us to utilize them for heterogeneous inverse scattering, while accounting for all orders of scattering.

### 3 Theoretical Background

We begin with background on the *radiative transfer* framework for describing light in scattering materials. We use lower-case bold letters for points  $\mathbf{x}$  in the Euclidean space  $\mathbb{R}^3$  and directions  $\boldsymbol{\omega}$  in the unit sphere  $\mathbb{S}^2$ . We use capital-case bold letters for position-direction pairs  $\mathbf{X} = (\mathbf{x}, \boldsymbol{\omega})$ , and the notation  $\mathbf{x}(\mathbf{X})$  and  $\boldsymbol{\omega}(\mathbf{X})$  to refer to the position or direction component of such pairs, respectively.

We assume that  $\mathcal{M}$ , a subset of  $\mathbb{R}^3$ , is occupied by a scattering medium with uniform index of refraction  $\eta$ , corresponding to speed of light  $c = c_o/\eta$  inside the medium. We use  $\partial\mathcal{M}$  for the boundary of  $\mathcal{M}$ , and at every boundary point  $\mathbf{x} \in \partial\mathcal{M}$  we use  $\hat{\mathbf{n}}(\mathbf{x})$  for the outward normal vector. We also define sets  $\Gamma_i = \{(\mathbf{x}, \boldsymbol{\omega}) \in \partial\mathcal{M} \times \mathbb{S}^2 : \boldsymbol{\omega} \cdot \hat{\mathbf{n}}(\mathbf{x}) < 0\}$  and  $\Gamma_o = \{(\mathbf{x}, \boldsymbol{\omega}) \in \partial\mathcal{M} \times \mathbb{S}^2 : \boldsymbol{\omega} \cdot \hat{\mathbf{n}}(\mathbf{x}) > 0\}$  of position-direction pairs on the boundary  $\partial\mathcal{M}$  pointing in or out, respectively.

#### 3.1 Light Transport in Scattering Media

In the radiative transfer framework, light propagation inside a scattering medium  $\mathcal{M}$  is described in terms of idealized light particles, often called “photons”, that perform random walks consisting of stochastic reflection, refraction, absorption

and scattering events. These interactions are determined by a set of material parameters, as shown in Figure 1(a): At every boundary point  $\mathbf{x} \in \partial\mathcal{M}$ , the *bidirectional scattering distribution function* (BSDF)  $f_s(\mathbf{x}, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i)$  controls refraction and reflection events. At every interior point  $\mathbf{x} \in \mathcal{M}$ , the medium is characterized by the scattering parameters  $m(\mathbf{x}) = \{\sigma_a(\mathbf{x}), \sigma_s(\mathbf{x}), f_p(\mathbf{x}, \cos\theta)\}$ . The *scattering coefficient*  $\sigma_s(\mathbf{x})$  and *absorption coefficient*  $\sigma_a(\mathbf{x})$  determine the amount of light that is scattered or absorbed, respectively, at every scattering event. The *extinction coefficient*  $\sigma_t(\mathbf{x}) = \sigma_s(\mathbf{x}) + \sigma_a(\mathbf{x})$  determines the spatial frequency of such events. The *phase function*  $f_p(\mathbf{x}, \cos\theta)$  determines the amount of light that scatters towards direction  $\boldsymbol{\omega}_o$  relative to the incident direction  $\boldsymbol{\omega}_i$ . As is usual, we make the assumption that the phase function is cylindrically-symmetric and invariant to rotations of  $\boldsymbol{\omega}_i$ , and is therefore a function of only  $\cos\theta = \boldsymbol{\omega}_i \cdot \boldsymbol{\omega}_o$ .

When a temporally-varying light source is applied at  $\partial\mathcal{M}$ , the photon random walks are described by the *time-dependent radiative transfer equation* (RTE),

$$\frac{1}{c} \frac{\partial L(\mathbf{x}, \boldsymbol{\omega}, t)}{\partial t} + \boldsymbol{\omega} \cdot \nabla L(\mathbf{x}, \boldsymbol{\omega}, t) = -\sigma_t(\mathbf{x}) L(\mathbf{x}, \boldsymbol{\omega}, t) + \sigma_s(\mathbf{x}) \int_{\mathbb{S}^2} f_p(\mathbf{x}, \boldsymbol{\omega} \cdot \boldsymbol{\psi}) L(\mathbf{x}, \boldsymbol{\psi}, t) d\boldsymbol{\psi}, \quad (1)$$

subject to BSDF-dependent boundary conditions on  $\Gamma_i$  and  $\Gamma_o$  [7, 25]. Note that more common in computer vision and graphics is the stationary form of the RTE, where radiance  $L$  is time-independent and there is no time derivative. We discuss the relationship between the two forms in the supplement [1].

The *time-dependent Green's function*  $\mathcal{T}_m(\mathbf{X}_o, \mathbf{X}_i, t)$  is the solution of the RTE at  $\mathbf{X}_o \in \Gamma_o$  and time  $t$ , for an input pulse of infinitesimal duration  $\delta(t)$  and unit radiance at  $\mathbf{X}_i \in \Gamma_i$ . We use the subscript  $m$  to denote explicitly that  $\mathcal{T}_m$  depends on the material parameters. The Green's function can also be defined inside the medium, but we restrict its point-direction arguments  $\mathbf{X}_o, \mathbf{X}_i$  to  $\Gamma_o, \Gamma_i$ , as we assume that we can only inject and measure light at the medium's boundary. For the same reason, we omit volumetric sources from Equation (1).

The change of variables  $\tau = ct$  converts time to (optical) pathlength. Then, the resulting *pathlength-resolved Green's function*  $\mathcal{T}_m(\mathbf{X}_o, \mathbf{X}_i, \tau)$  is equal to the radiance produced by accumulating contributions only from photons that travel paths starting at boundary point  $\mathbf{x}(\mathbf{X}_i)$  with direction  $\boldsymbol{\omega}(\mathbf{X}_i)$ , ending at boundary point  $\mathbf{x}(\mathbf{X}_o)$  with direction  $\boldsymbol{\omega}(\mathbf{X}_o)$ , and having total length  $\tau$ . The pathlength-resolved Green's function  $\mathcal{T}_m(\mathbf{X}_o, \mathbf{X}_i, \tau)$  is the continuous equivalent of the pathlength-resolved *light transport matrix* [37, 13]. In the sequel, we use the terms Green's function and light transport matrix interchangeably.

### 3.2 Imaging Scattering Media

The measurements produced using different imaging techniques can be described as different ways to sample the light transport matrix. Specifically, we can express them in terms of a *sampling function*  $W : \Gamma_o \times \Gamma_i \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ , as,

$$\mathcal{S}(\mathcal{T}_m, W) \triangleq \int_{\Gamma_o} \int_{\Gamma_i} \int_0^\infty W(\mathbf{X}_o, \mathbf{X}_i, \tau) \mathcal{T}_m(\mathbf{X}_o, \mathbf{X}_i, \tau) d\tau d\mathbf{X}_i d\mathbf{X}_o. \quad (2)$$

In this equation, we assume that  $\mathcal{T}_m$  and  $W$  are regular enough to allow changing the integration order. The sampling function  $W$  can typically be decomposed into three components: First, an *emittance function*  $W_i : \Gamma_i \rightarrow \mathbb{R}_{\geq 0}$ , which is non-zero on a subset of the inward boundary  $\Gamma_i$  and describes the incident illumination. Second, an *importance function*  $W_o : \Gamma_o \rightarrow \mathbb{R}_{\geq 0}$ , which is non-zero on a subset of the outward boundary  $\Gamma_o$  and corresponds to the rays accumulated by the sensor. Third, a *pathlength sampling function*  $W_\tau : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ , which is non-zero only for a subset of pathlength values. Then,

$$W(\mathbf{X}_o, \mathbf{X}_i, \tau) = W_o(\mathbf{X}_o) W_i(\mathbf{X}_i) W_\tau(\tau). \quad (3)$$

**Steady-state imaging.** Conventional imaging sensors measure all photons, regardless of the distance they have traveled inside the medium. This corresponds to using a pathlength sampling function  $W_\tau(\tau) = 1$ , for all values of  $\tau$ .

**Pathlength decomposition.** Pathlength decomposition discriminates between photons based on the pathlength they travel (Figure 1(b): low-saturation, short paths versus high-saturation, long paths). Ideally, they sample pathlength slices of the light transport matrix,  $W_\tau(\tau) = \delta(\tau - \tau_c)$ , for some  $\tau_c > 0$ . Real systems instead have finite pathlength resolution, with  $W_\tau(\tau)$  being, say, a Gaussian or a square function. Typically, pathlength decomposition techniques densely capture multiple such pathlength slices, each centered at a different  $\tau_c$ .

**Spatial probing.** Cameras typically have multiple sensor elements (pixels) that capture parallel sets of measurements  $\{\mathcal{S}(\mathcal{T}_m, W_o^p \cdot W_i \cdot W_\tau), p = 1, \dots, P\}$ , with  $P$  the number of pixels. Conventionally, the measurements in such a set use different importance functions  $W_o^p$  and a common illumination  $W_i$ ; but an alternative is to use spatial probing techniques that allow different pixels on the same sensor to capture measurements corresponding to different importance-emittance pairs,  $\{\mathcal{S}(\mathcal{T}_m, W_o^p \cdot W_i^p \cdot W_\tau), p = 1, \dots, P\}$ . With reference to Figure 1(b), this allows the orange and blue camera pixels to only measure photons that begin at the orange and blue source pixels, respectively. Note that equivalent measurements can be obtained by capturing multiple images sequentially, each time using a different source  $W_i^p$  and discarding unneeded pixels. Therefore, rather than providing a fundamentally different way to sample the light transport matrix, spatial probing allows reducing acquisition time through temporal multiplexing.

**Types of sources.** It is useful to define an *ideal source* that emits an ideally collimated and narrow beam, corresponding to emittance

$$W_i(\mathbf{X}_i) = \delta(\mathbf{x}(\mathbf{X}_i) - \mathbf{x}_l) \delta(\boldsymbol{\omega}(\mathbf{X}_i) - \boldsymbol{\omega}_l), \quad (4)$$

for some  $(\mathbf{x}_l, \boldsymbol{\omega}_l) \in \Gamma_i$ . Note that such a source cannot be realized physically: The wave nature of light implies that a source cannot be perfectly concentrated in both the spatial and angular domains simultaneously [16]; and any source ideally concentrated in either domain would have zero etendue and power [6].

The utility of ideal sources lies in their convenience for analysis [4], and for synthesizing other, more realistic sources. For instance, a perfectly collimated

area source can be created by combining ideal sources with different  $\mathbf{x}_l$  but sharing the same  $\boldsymbol{\omega}_l$ . A diffuse point source is created by combining ideal sources that share the same  $\mathbf{x}_l$ , for each  $\boldsymbol{\omega}_l \in \mathbb{S}^2$ . Finally, a physical source with non-zero spatio-angular extent  $E_i \subset \Gamma_i$  can be created as a weighted combination of ideal sources for each  $(\mathbf{x}_l, \boldsymbol{\omega}_l) \in E_i$ , with the weights depending on intrinsic (source power distribution) and extrinsic (geometry) factors.

**Types of cameras.** Similar to the ideal source, we can define an *ideal sensor element* with an importance function

$$W_o(\mathbf{X}_o) = \delta(\mathbf{x}(\mathbf{X}_o) - \mathbf{x}_s) \delta(\boldsymbol{\omega}(\mathbf{X}_o) - \boldsymbol{\omega}_s), \quad (5)$$

for some  $(\mathbf{x}_s, \boldsymbol{\omega}_s) \in \Gamma_o$ . Arrays of such sensors can be used to construct different types of cameras. An *orthographic camera* has multiple ideal sensor elements, each with a different  $\mathbf{x}_s$  but all sharing the same  $\boldsymbol{\omega}_s$ . As with sources, real sensors cannot be perfectly concentrated in either the spatial or angular domain, but can still be expressed as weighted combinations of ideal sensor elements. Simultaneously using an ideal sensor, source, and pathlength decomposition corresponds to sampling a single value  $\mathcal{T}_m((\mathbf{x}_s, \boldsymbol{\omega}_s), (\mathbf{x}_l, \boldsymbol{\omega}_l), \tau_c)$  of the light transport matrix.

## 4 Imaging Design

Given the range of imaging options available, we discuss theoretical results and empirical observations that can be used to guide the design of imaging systems for inverse scattering applications. All proofs are shown in the supplement.

**Local ambiguities.** In steady-state imaging, there exist *similarity relations* which, under certain conditions, allow changing the scattering parameters at a point  $\mathbf{x} \in \mathcal{M}$  without changing the radiance at that point [49, 51]. We prove that these similarity relations also hold for pathlength-decomposed measurements.

**Lemma 1.** *Let  $\{a_{n,l}(\mathbf{x}, \tau), n > 0, -n \leq l \leq n\}$  be the coefficients of the spherical harmonics expansion of the solution  $L(\mathbf{x}, \boldsymbol{\omega}, \tau)$  of Equation (1) at some point  $\mathbf{x} \in \mathcal{M}$ , and  $\{f_{p,n}(\mathbf{x}), n > 0\}$  the coefficients of the Legendre expansion of  $f_p$  at that point. If there exists  $N > 0$  such that  $a_{n,l}(\mathbf{x}, \tau) = 0$  for all  $n > N$ , then two materials  $m, m^*$  will produce equal values  $L(\mathbf{x}, \boldsymbol{\omega}, \tau)$  if, for  $1 \leq n \leq N$ ,*

$$\sigma_a(\mathbf{x}) = \sigma_a^*(\mathbf{x}), \quad (6)$$

$$\sigma_s(\mathbf{x})(1 - f_{p,n}(\mathbf{x})) = \sigma_s^*(\mathbf{x})(1 - f_{p,n}^*(\mathbf{x})). \quad (7)$$

These similarity relations are *local*, as they describe ambiguities at one point  $\mathbf{x} \in \mathcal{M}$ . To reduce these ambiguities, Lemma 1 suggests maximizing the angular frequency of  $L(\mathbf{x}, \boldsymbol{\omega}, \tau)$  throughout the medium. For a given shape of  $\mathcal{M}$ , the only way to control this is through the incident illumination  $W_i(\mathbf{X}_i)$ , which provides a lower bound to the angular frequency of  $L(\mathbf{x}, \boldsymbol{\omega}, \tau)$  by way of spherical convolution with the phase function in the RTE (1) acting as a low-pass filter. The best light sources will have high angular and spatial frequencies, such as

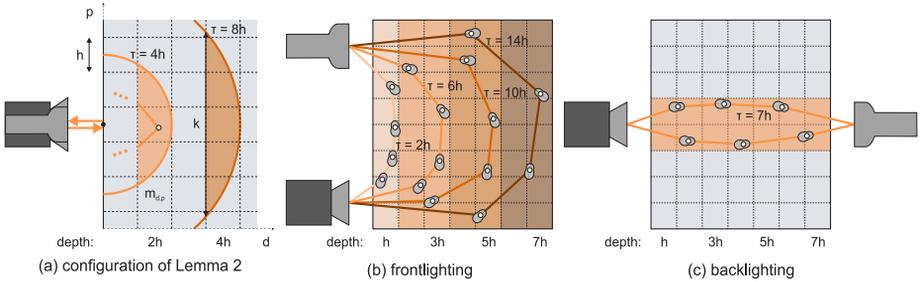


Fig. 2: (a) Configuration used for Lemma 2. (b) Frontlighting allows decomposing the inverse scattering problem into sub-problems for depth-wise layers, where at each layer the material parameters of lower-depth layers are known. (c) Backlighting allows measuring deep layers at shorter pathlengths  $\tau$ , and thus higher SNR, than frontlighting.

the ideal sources of Equation (4), or collimated area sources with high-frequency spatial profiles. These types of sources already feature prominently in previous theoretical and empirical work [33, 15, 14].

**Non-local ambiguities.** On top of local ambiguities, a heterogeneous medium can exhibit *non-local* ambiguities involving material parameters at different points  $\mathbf{x}$ . In the following, we argue that pathlength decomposition combined with a specific input-output geometry can help reduce such non-local ambiguities. For this, we consider the scene in Figure 2(a): A cubic volume is discretized into  $h$ -sized voxels indexed by coordinates  $[d, p]$ . We call the coordinate  $d$  the *depth* of a voxel, and we call the set of voxels with the same  $d$  the *layer* at depth  $d$ . We assume that material parameters  $m[d, p]$  are constant within each voxel, and that  $h$  is also the spatial and pathlength resolution at which we can image.

We use an ideal source (Equation (4)) co-located with an ideal sensor (Equation (5)). Then, the following lemma states that, assuming all material parameters at layers  $[h, 2h, \dots, (n-1)h]$  are known from previous measurements, pathlength decomposed measurements at  $\tau = 2nh$  and  $\tau = (2n+1)h$  provide *linear* equations in the parameters of layer  $d = nh$ .

**Lemma 2.** *Using pathlength decomposition, the configuration of Figure 2(a) provides measurements of the form*

$$I_\tau = \begin{cases} Q_\tau + \sum_{p \in k(n)} \sigma_s [nh, p] \int_0^\pi f_p([nh, p], \theta) R_{\tau, p}(\theta) d\theta, & \tau = 2nh, \\ S_\tau + \sum_{p \in k(n)} \sigma_t [nh, p] T_{\tau, p}, & \tau = (2n+1)h, \end{cases} \quad (8)$$

where  $k(n)$  is the subset of voxels in layer  $d = nh$  that are intersected by a source-centered circle of radius  $nh$ ;  $Q_\tau$  and  $R_{\tau, p}(\theta)$  are functions of material parameters  $\{m[d, p], d < nh\}$ ; and  $S_\tau$  and  $T_{\tau, p}$  of  $\{m[d, p], d < nh; \sigma_s [nh, p]; f_p([nh, p], \theta)\}$ .

Using Figure 2, we can understand the first line of Equation (8) as follows. Measurements of the form  $I_{2nh}$  are sums of contributions from paths of length

$\tau = 2nh$ . All these paths are contained inside a circle of radius  $nh$  centered at the source/sensor location. We can split the paths into two categories: those that stay within layers  $[h, 2h, \dots, (n-1)h]$ , and those that reach layer  $d = nh$ . Most paths fall into the first category, and their contributions are included in the term  $Q$ . A fraction of the paths are in the second category, and when they reach layer  $d = nh$ , their interaction is limited to a single bounce somewhere within a particular subset of the voxels in that layer (denoted  $k(n)$  and shaded orange in Figure 2(a)). The parts of these paths that are contained in layers  $d < nh$  correspond to the term  $R$  in Equation (8), and their single-bounce interactions within layer  $d = nh$  lead to linear dependence on parameters  $\sigma_s f_p(\theta)$ . Thus, when the materials in layers  $d < nh$  are known, the terms  $Q$  and  $R$  can be computed; and we can build a system of linear equations in the parameters  $\sigma_s f_p(\theta)$  for all voxels in layer  $d = nh$  by changing the source/sensor location and, as discussed in the supplement, by shifting the sensor relative to the source.

Similarly, the second line of Equation (8) refers to measurements  $I_{(2n+1)h}$  of contributions from paths that have an odd number of steps,  $\tau = (2n+1)h$ . As explained in detail in the supplement, these paths may include up to two scattering events in layer  $d = nh$ , but only one that involves  $\sigma_t[d, p]$ . Therefore  $I_{(2n+1)h}$  will depend linearly on  $\sigma_t[d, p]$ , and because the other parameters  $\sigma_s[d, p]$  and  $f_p([d, p], \theta)$  can be separately estimated from the even-step measurements, we can create a linear system of equations in  $\sigma_t$  for all voxels in layer  $d = nh$  by combining odd-step measurements from different source and sensor locations.

Overall, Lemma 2 suggests a recursive, layer-wise procedure for inferring scattering parameters from pathlength-resolved measurements (Figure 2(b)): Assuming parameters at layers  $d < nh$  have been estimated from previous measurements, use measurements of the form of  $I_{2nh}$  and  $I_{(2n+1)h}$  to estimate parameters at layer  $d = nh$ , through a linear system in  $\sigma_s[d, p] f_p([d, p], \theta)$  and another linear system in  $\sigma_t[d, p]$ .

**Noise considerations.** So far, we have shown that pathlength decomposed measurements in a *frontlighting* configuration, where there source and sensor are at the same side of the medium, helps reduce non-local ambiguities. To complete the picture, we must also consider the signal-to-noise ratio (SNR) of the measurements in Equation (8) as pathlength  $\tau$  increases. One observation is that their magnitude, and therefore SNR, decreases exponentially with  $\tau$  due to volumetric attenuation (see Equation (15)). Another is that, from Figure 2(a), the ratio of the circle's area contained in layers  $d < nh$  versus that in layer  $d = nh$  increases with  $\tau$ . This implies that the second terms in Equation (8), which contain all information about parameters at depth  $d = nh$ , become smaller relative to the terms  $Q_\tau, S_\tau$ , which are independent of those parameters.

As a result of these two factors, the information available for inferring material parameters becomes progressively worse at greater depths. One way to ameliorate this noise is to use a supplementary *backlighting* configuration (Figure 2(c)), where sources and sensors are on opposite sides of the medium. These backlighting measurements cannot be used in the recursive procedure described above, but they can provide cleaner measurements of deeper layers.

While these noise considerations also apply to the case of steady-state measurements from frontlighting and backlighting configurations, the preceding analysis of non-local ambiguities cannot be directly extended to that case. In place of theoretical analysis, we evaluate the relative utility of steady-state versus pathlength-resolved measurements quantitatively: We first introduce an inverse rendering algorithm for inferring heterogeneous scattering parameters from both types of measurements (Section 5), then we use this algorithm to perform inverse scattering experiments with simulated volumes and measurements (Section 6).

## 5 Inverse Rendering Algorithm

Measurements of the light transport matrix  $\mathcal{T}_m$  provide information about internal scattering parameters. Given a measurement set  $\{\bar{I}^n, n = 1, \dots, N\}$  that is calibrated, meaning that sampling functions  $W^n$  are known, we can try to recover the parameters by solving an *appearance matching* optimization problem,

$$\min_{\boldsymbol{\pi}} \sum_{n=1}^N \frac{1}{2} (\bar{I}^n - \mathcal{S}(\mathcal{T}_m(\boldsymbol{\pi}), W^n))^2, \quad (9)$$

where  $\boldsymbol{\pi}$  is an appropriate  $K$ -dimensional parameterization of the material  $m$  (we will be omitting the dependence on  $\boldsymbol{\pi}$  for notational simplicity). In the following, we introduce a framework for efficiently solving this inference problem. We do this by extending the inverse rendering algorithms introduced in [14, 30], to apply to any possible set of measurements from the light transport matrix  $\mathcal{T}_m$ .

**Path formulation of light transport.** Section 3 describes entries of the light transport matrix  $\mathcal{T}_m$  and measurements  $\mathcal{S}(\mathcal{T}_m, W)$  as different accumulations of photon contributions based on their paths. This intuition has been formalized in computer graphics and is the foundation of path-based rendering algorithms. We can also use it to derive our inverse rendering algorithm. For notation, we define a path  $\bar{\boldsymbol{x}}$  as an ordered sequence of points in the medium  $\mathcal{M}$ ,

$$\bar{\boldsymbol{x}} = \boldsymbol{x}_0 \rightarrow \boldsymbol{x}_1 \rightarrow \dots \rightarrow \boldsymbol{x}_B, \quad (10)$$

for any finite integer  $B > 1$ . We denote the space of all such paths as  $\mathbb{P}$ . For each path segment  $\boldsymbol{x}_b \rightarrow \boldsymbol{x}_{b+1}$ , we denote by  $\boldsymbol{\omega}(\boldsymbol{x}_b \rightarrow \boldsymbol{x}_{b+1})$  its direction. For each path  $\bar{\boldsymbol{x}}$ , we denote by  $o(\bar{\boldsymbol{x}}) = \boldsymbol{x}_0$  and  $e(\bar{\boldsymbol{x}}) = \boldsymbol{x}_B$  its origin and end, by  $\boldsymbol{\omega}_o(\bar{\boldsymbol{x}}) = \boldsymbol{\omega}(\boldsymbol{x}_0 \rightarrow \boldsymbol{x}_1)$  and  $\boldsymbol{\omega}_e(\bar{\boldsymbol{x}}) = -\boldsymbol{\omega}(\boldsymbol{x}_{B-1} \rightarrow \boldsymbol{x}_B)$  its starting and ending directions, and by  $\tau(\bar{\boldsymbol{x}}) = \sum_{b=1}^B \|\boldsymbol{x}_b - \boldsymbol{x}_{b-1}\|_2$  its length.

Then, based on the path formulation of light transport [44, 40], every measurement of the light transport matrix can be written as,

$$\mathcal{S}(\mathcal{T}_m, W) = \int_{\mathbb{P}} W(\bar{\boldsymbol{x}}) \bar{f}_m(\bar{\boldsymbol{x}}) d\bar{\boldsymbol{x}}, \quad (11)$$

where we overload notation to make the sampling function of Equation (2) apply to paths, based on their endpoints and length:

$$W(\bar{\boldsymbol{x}}) \triangleq W((e(\bar{\boldsymbol{x}}), \boldsymbol{\omega}_e(\bar{\boldsymbol{x}})), (o(\bar{\boldsymbol{x}}), \boldsymbol{\omega}_o(\bar{\boldsymbol{x}})), \tau(\bar{\boldsymbol{x}})). \quad (12)$$

The *throughput function*  $\bar{f}_m$  determines the path's radiance contribution,

$$\bar{f}_m(\bar{\mathbf{x}}) = \prod_{b=1}^{B-1} f_m(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}), \quad (13)$$

$$f_m(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}) = a(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b) \sigma(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}), \quad (14)$$

$$a(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b) = \exp\left(-\int_{\mathbf{x}_{b-1}}^{\mathbf{x}_b} \sigma_t(\mathbf{x}) d\mathbf{x}\right), \quad (15)$$

$$\sigma(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}) = \begin{cases} f_s(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}), & \mathbf{x}_b \in \partial\mathcal{M}, \\ \left\langle f_s(\mathbf{x}_b) f_p(\mathbf{x}_b, \boldsymbol{\omega}(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b) \cdot \boldsymbol{\omega}(\mathbf{x}_b \rightarrow \mathbf{x}_{b+1})) \right\rangle, & \text{otherwise.} \end{cases} \quad (16)$$

Equation (15) is the volumetric attenuation along each path segment. Equation (16) corresponds to radiance transfer as direction changes at the end of a path segment. When  $\mathbf{x}_b \in \partial\mathcal{M}$ , the direction change is due to internal reflection at the medium boundary, and the amount of radiance transferred is determined by the material's BSDF. Otherwise, the direction change is due to scattering, and the radiance transferred is determined using the local phase function. If  $\mathcal{M}$  is not convex, a path may exit and re-enter the medium before reaching the sensor, in which case the attenuation  $a$  for its corresponding segments equals 1.

Following [30], from Equation (11), we can also formulate a path-based expression for the derivative of measurements of the light transport matrix with respect to any material parameter  $\pi_k$ . From the product form of  $\bar{f}_m$  in Equation (13), by applying the chain rule and re-arranging terms, we have that

$$\frac{\partial \mathcal{S}(\mathcal{T}_m, W)}{\partial \pi_k} \Big|_{\boldsymbol{\pi}=\boldsymbol{\pi}_o} = \int_{\mathbb{P}} W(\bar{\mathbf{x}}) \frac{\partial \bar{f}_m(\bar{\mathbf{x}})}{\partial \pi_k} \Big|_{\boldsymbol{\pi}=\boldsymbol{\pi}_o} d\bar{\mathbf{x}} = \int_{\mathbb{P}} W(\bar{\mathbf{x}}) \bar{f}_m(\bar{\mathbf{x}}) \bar{S}_{m,k}(\bar{\mathbf{x}}) d\bar{\mathbf{x}}, \quad (17)$$

$$\bar{S}_{m,k}(\bar{\mathbf{x}}) \triangleq \sum_{b=1}^{B-1} S_{m,k}(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}), \quad (18)$$

$$S_{m,k}(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}) \triangleq \frac{(\partial f_m(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1}) / \partial \pi_k) |_{\boldsymbol{\pi}=\boldsymbol{\pi}_o}}{f_m(\mathbf{x}_{b-1} \rightarrow \mathbf{x}_b \rightarrow \mathbf{x}_{b+1})}. \quad (19)$$

In statistics,  $S_{m,k}$  is known as the *score function* of  $f_m$  with respect to  $\pi_k$ .

**Monte Carlo integration.** The integrals of Equation (11) and Equation (17) can be estimated using Monte Carlo integration (Figure 3): We first use any probability distribution  $\mu$  on  $\mathbb{P}$  to sample a set of paths  $\{\bar{\mathbf{x}}_j, j = 1, \dots, J\}$ ; then we form the respective unbiased estimates,

$$I = \sum_{j=1}^J \frac{W(\bar{\mathbf{x}}_j) \bar{f}_m(\bar{\mathbf{x}}_j)}{\mu(\bar{\mathbf{x}}_j)}, \quad G_k = \sum_{j=1}^J \frac{W(\bar{\mathbf{x}}_j) \bar{f}_m(\bar{\mathbf{x}}_j) \bar{S}_{m,k}(\bar{\mathbf{x}}_j)}{\mu(\bar{\mathbf{x}}_j)}. \quad (20)$$

The paths  $\{\bar{\mathbf{x}}_j, j = 1, \dots, J\}$  can be chosen using the sampling strategies developed for physically accurate rendering, such as (volumetric) path tracing, bidirectional path tracing, or Metropolis Light Transport [44, 40]. Because

of the need to render individual elements of the light transport matrix (corresponding to the infinitesimal emittance and importance functions of Equations (4) and (5)), we use bidirectional path tracing. When rendering pathlength-decomposed measurements, we also use the local sampling modifications of [27]. Finally, most rendering algorithms use the product form of Equation (13) to efficiently compute the term  $\bar{f}_m$  in Equation (20) recursively while tracing a path. The same can be done for the term  $\bar{S}_{m,k}$ , using the sum form of Equation (18).

**Stochastic optimization.** We now consider the appearance matching problem of Equation (9). Denoting by  $E(\boldsymbol{\pi})$  its loss function, we have for its gradient,

$$\left. \frac{\partial E}{\partial \pi_k} \right|_{\boldsymbol{\pi}} = \sum_{n=1}^N (\bar{I}^n - \mathcal{S}(\mathcal{T}_m, W^n)) \left. \frac{\partial \mathcal{S}(\mathcal{T}_m, W^n)}{\partial \pi_k} \right|_{\boldsymbol{\pi}}. \quad (21)$$

We can estimate the gradient using the estimates of Equation (20), as

$$g_k(\boldsymbol{\pi}) = \sum_{n=1}^N (\bar{I}^n - I^n) G_k^n. \quad (22)$$

This estimate  $g_k$  is unbiased if  $I^n$  and  $G_k^n$  are statistically independent, which can be achieved by rendering them using independently selected sets of paths. Following [14], we can combine these gradient estimates with stochastic gradient descent algorithms to solve the appearance matching problem of Equation (9).

Standard SGD uses iterations  $\boldsymbol{\pi}^{(t+1)} = \boldsymbol{\pi}^{(t)} - \beta^{(t)} \mathbf{g}^{(t)}$ , with common step size  $\beta^{(t)}$  for all unknown parameters. This converges slowly when gradient vectors are very sparse [11]. This applies to our problem because the gradient estimate  $g_k$  for some material voxel will be zero if, during the rendering operations of Equation (20), no sampled paths travel through the voxel. We have experimented with a number of SGD variants that use separate, per-parameter step sizes, each decrementing adaptively based on the magnitudes of the per-parameter gradients in previous iterations [31, 50, 11]. We chose to use ADADELTA (Algorithm 1), which we found empirically to have the best performance for solving (9).

**Initialization.** We initialize Algorithm 1 using a multi-resolution procedure, that progressively increases the spatial resolution of material parameters. For pathlength-resolved measurements, we also use a layer-wise recursive procedure analogous to Figure 2(b). We discuss both procedures in the supplement.

## 6 Experiments

**Implementation.** We implemented the inverse rendering framework of Section 5 on top of the Mitsuba physically based renderer [26]. We extended the bidirectional path tracing algorithm to support spatial probing and pathlength decomposition rendering, for both radiance and gradient estimation. The stochastic optimization layer distributes rendering tasks involved in gradient computation

---

**Algorithm 1** ADADELTA.
 

---

**Require:** decay rate  $\gamma < 1$ , constant vector  $\varepsilon$ .

- 1:  $\boldsymbol{\pi}^{(0)} \leftarrow \text{Initialize}()$ ,  $\boldsymbol{\sigma}_g^{(0)} \leftarrow \mathbf{0}$ ,  $\boldsymbol{\sigma}_\delta^{(0)} \leftarrow \varepsilon$ .
  - 2: **while** not converged **do**
  - 3:  $\mathbf{g}^{(t)} \leftarrow \text{ComputeGradient}(\boldsymbol{\pi}^{(t)})$ .
  - 4:  $\boldsymbol{\sigma}_{g,k}^{(t)} \leftarrow \gamma \boldsymbol{\sigma}_{g,k}^{(t-1)} + (1 - \gamma) \left( \mathbf{g}_k^{(t)} \right)^2$ .
  - 5:  $\boldsymbol{\delta}_k^{(t)} \leftarrow \frac{\boldsymbol{\sigma}_{\delta,k}^{(t-1)}}{\boldsymbol{\sigma}_{g,k}^{(t)}} \mathbf{g}_k^{(t)}$
  - 6:  $\boldsymbol{\pi}^{(t+1)} \leftarrow \boldsymbol{\pi}^{(t)} - \boldsymbol{\delta}^{(t)}$ .
  - 7:  $\boldsymbol{\sigma}_{\delta,k}^{(t)} \leftarrow \gamma \boldsymbol{\sigma}_{\delta,k}^{(t-1)} + (1 - \gamma) \left( \boldsymbol{\delta}_k^{(t)} \right)^2$ .
  - 8: **end while**
  - 9: **return**  $\boldsymbol{\pi}^{\text{opt}} = \frac{1}{T} \sum_{t=0}^T \boldsymbol{\pi}^{(t)}$ .
- 

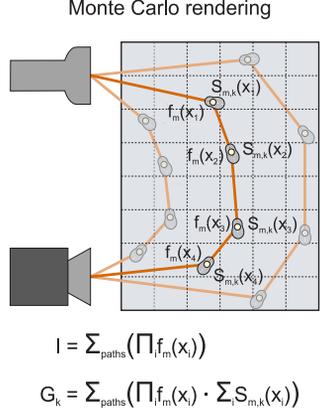


Fig. 3: Inverse rendering algorithm. Left: We use the ADADELTA variant of stochastic gradient descent to minimize the appearance matching objective of Equation (9). Right: We use a modified Monte Carlo rendering algorithm to compute stochastic gradient estimates. When shading a path, we compute for each segment both the usual throughput terms  $f_m$  and the score terms  $S_{m,k}$ . These are accumulated to determine the path’s contribution to the image  $I$  and its derivative  $G_k$  (Equations (13) and (18)). We repeat this process over many paths that are sampled to satisfy the sampling function  $W$  corresponding to the input image measurement (Equation (12)).

(multiple parameters and measurements) across a multi-CPU cluster. We ran all our experiments on 20-node Amazon EC2 clusters, with 36 cores per node.

**Comparison of imaging configurations.** We perform inverse scattering experiments on synthetic volumes, to evaluate the performance of different imaging configurations. Following Figure 2(a), we use a cubic medium of size  $10 \times 10 \times 10 \text{ mm}^3$  discretized at resolution 0.4 mm, resulting in 15625 voxels. To reduce the space of possible imaging configurations, we assume that the material parameters are characterized by a cross-section of the cube along the  $\{d, p\}$  plane, remaining constant across the third dimension, corresponding to  $3 \times 625 = 1875$  unknowns. We use an imaging resolution four times that of the material grid, corresponding to pixel size  $0.1 \times 0.1 \text{ mm}^2$  and pathlength resolution 0.1 mm.

We fix the medium refractive index to  $\eta = 1.3$  and assume a smooth dielectric BSDF at its boundary. We adopt the single-parameter Henyey-Greenstein model for the phase function [23]; therefore, each material voxel is associated with corresponding unknown values for  $\sigma_a$ ,  $\sigma_s$ , and the phase function parameter  $g$  (equal to the phase function’s first moment). We constrain  $\sigma_a, \sigma_s \in [1, 10] \text{ mm}^{-1}$ , and  $g \in [0, 0.6]$ . We generate volumes by modeling each parameter inside the medium as a mixture of two Gaussians of random mean and variance.

We use these volumes to compare three imaging configurations: 1) Pathlength decomposition where, for each ideal source, we measure the radiance exiting the

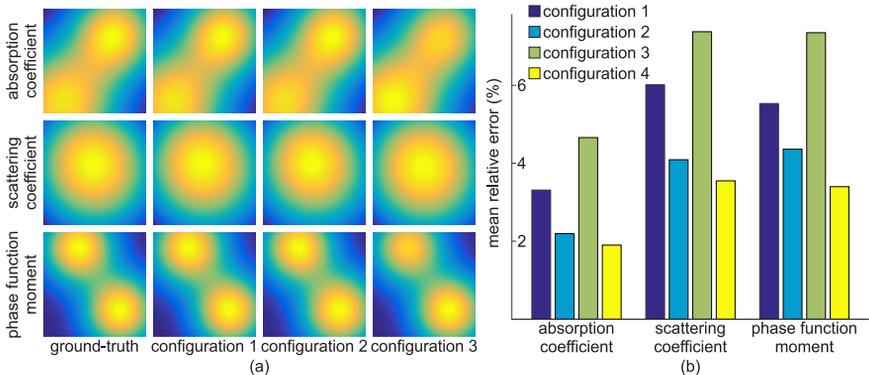


Fig. 4: Comparison of different imaging configurations. (a) Visualization of ground-truth and reconstructed material parameters as a function of location. (b) Mean relative reconstruction error for material parameters, averaged over multiple synthetic volumes.

volume at the opposite direction and from the same position, as well as its two spatial neighbors. We take measurements with the ideal source placed at every pixel on the medium boundary, including in sidelighting and backlighting positions. 2) Steady-state imaging where, for every ideal source, we use an orthographic camera to measure radiance exiting from all pixels in one surface of the cube. As before, we take measurements with the ideal source placed at every pixel on the medium boundary, and at every position at three different orientations. 3) Similar to (1), but instead of spatial shifts, we take measurements at multiple source orientations. Each of these configurations produces 120000 measurements, or 64 measurements per unknown. When rendering simulated measurements, we add sensor noise using [19].

In Figure 4(a), we visualize the reconstructed parameters for one of the synthetic volumes. We observe that all three configurations are generally able to reconstruct all three spatially varying parameters,  $\sigma_a$ ,  $\sigma_s$ , and  $g$ , within a mean relative error 7% and maximum relative error 15%, concentrated around areas of high absorption. In Figure 4(b), we compare the RMS error in the estimation of each parameter by each configuration, averaged across five synthetic volumes. We see that, the configuration using only steady-state measurements has a lower RMS error. We expect that this is due to the very low SNR of pathlength-decomposition measurements corresponding to large pathlength values. We show in Figure 4(b) the mean relative error obtained by a fourth configuration, created by replacing measurements in configuration (1) of magnitude comparable to the additive sensor noise, with steady-state measurements from configuration (2) corresponding to large camera-sensor distances. We see that this combination produces the lowest RMS error among all imaging configurations.

**3D reconstruction.** In Figure 5, we show volumetric reconstructions of a dense heterogeneous smoke volume, with smoothly spatially-varying scattering parameters, and assuming index of refraction equal to 1. We generate our own param-

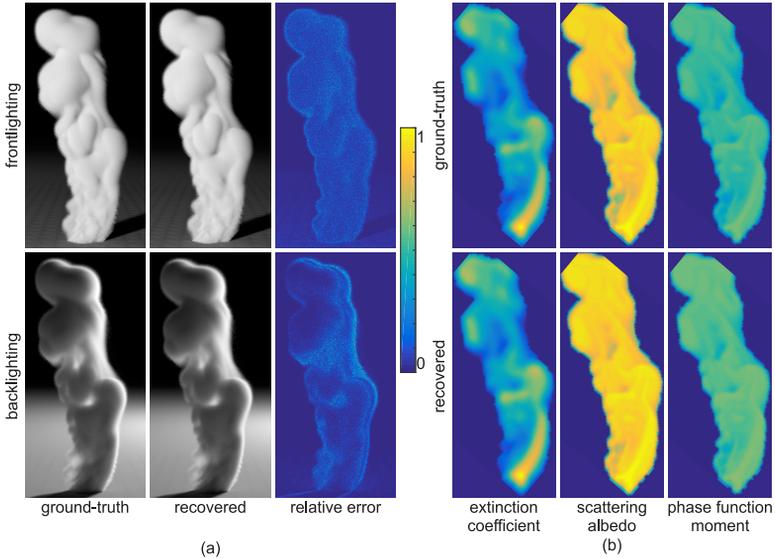


Fig. 5: (a) Renderings of a smoke volume using ground-truth and recovered material parameters under novel viewpoint and illumination conditions. (b) Ground-truth and recovered parameters for a vertical cross-section through the smoke volume.

eters for the volume mesh provided by [26]. We use pathlength-resolved measurements in frontlighting, sidelighting, and backlighting configurations.

Figure 5(a) compares renderings of the smoke volume using the ground-truth and recovered material parameters under novel imaging configurations (not used as input to the inverse rendering algorithm). In Figure 5(b), we compare ground-truth and recovered material parameters across a cross-section of the volume. Our algorithm accurately recovers all scattering parameters, with mean relative error 9.31% and maximum relative error 19.73%, and the recovered parameters can reproduce the appearance of the volume under new imaging conditions.

## 7 Conclusions

We have presented a theoretical and quantitative evaluation of various computational imaging techniques for the heterogeneous inverse scattering problem. Our theoretical results provide formal justification for the use of pathlength decomposition in applications requiring volumetric reconstruction of complex materials. Additionally, our experimental results suggest there are many different imaging configurations, including both steady-state and pathlength decomposition measurements, that can enable accurate recovery of heterogeneous scattering parameters. Our theoretical results and our optimization framework can be used to guide the design of new acquisition systems, such as when selecting from among various possible configurations, and when weighing practical considerations, such as hardware availability, exposure time, and geometry constraints.

**Acknowledgments:** This material is based upon work supported by: the US National Science Foundation under award IIS-1161564; the DARPA REVEAL program under Contract No. HR0011-16-C-0028; the European Research Council; the Israel Science Foundation; and research grants from Amazon Web Services.

## References

1. Project page, [http://vision.seas.harvard.edu/inverse\\_transient/](http://vision.seas.harvard.edu/inverse_transient/)
2. Antyufeev, V.: Monte Carlo method for solving inverse problems of radiation transfer, vol. 20. Inverse and Ill-Posed Problems Series, V.S.P. International Science (2000)
3. Arridge, S.R.: Optical tomography in medical imaging. Inverse problems (1999)
4. Bal, G.: Inverse transport theory and applications. Inverse Problems 25(5) (2009)
5. Boas, D.A., Brooks, D.H., Miller, E.L., DiMarzio, C.A., Kilmer, M., Gaudette, R.J., Zhang, Q.: Imaging the body with diffuse optical tomography. IEEE Signal Processing Magazine (2001)
6. Boyd, R.W.: Radiometry and the detection of optical radiation. John Wiley and Sons, New York, NY, USA (1983)
7. Case, K.M., Zweifel, P.F.: Linear Transport Theory. Addison-Wesley Pub. Co. (1967)
8. Debevec, P., Hawkins, T., Tchou, C., Duiker, H., Sarokin, W., Sagar, M.: Acquiring the reflectance field of a human face. In: Proceedings of SIGGRAPH 2000, Annual Conference Series (2000)
9. Donner, C., Jensen, H.: Light diffusion in multi-layered translucent materials. ACM Trans. Graph. 24(3) (2005)
10. Donner, C., Weyrich, T., d'Eon, E., Ramamoorthi, R., Rusinkiewicz, S.: A layered, heterogeneous reflectance model for acquiring and rendering human skin. ACM Trans. Graph. 27(5) (2008)
11. Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. The Journal of Machine Learning Research 12, 2121–2159 (2011)
12. Fuchs, C., Chen, T., Goesele, M., Theisel, H., Seidel, H.: Density estimation for dynamic volumes. Computers & Graphics 31(2) (2007)
13. Gkioulekas, I., Levin, A., Durand, F., Zickler, T.: Micron-scale light transport decomposition using interferometry. ACM Transactions on Graphics (2015)
14. Gkioulekas, I., Zhao, S., Bala, K., Zickler, T., Levin, A.: Inverse volume rendering with material dictionaries. ACM Transactions on Graphics (2013)
15. Goesele, M., Lensch, H., Lang, J., Fuchs, C., Seidel, H.: Disco: acquisition of translucent objects. ACM Trans. Graph. 23(3) (2004)
16. Goodman, J.W.: Introduction to Fourier Optics. McGraw-Hill Book Company (1968)
17. Gu, J., Nayar, S., Grinspun, E., Belhumeur, P., Ramamoorthi, R.: Compressive structured light for recovering inhomogeneous participating media. ECCV (2008)
18. Gupta, M. and Agrawal, A. and Veeraraghavan, A. and Narasimhan, S.G.: Structured light 3d scanning in the presence of global illumination. In: Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on. pp. 713–720 (June 2011)

19. Hasinoff, S., Durand, F., Freeman, W.: Noise-optimal capture for high dynamic range photography. *IEEE CVPR* (2010)
20. Hawkins, T., Einarsson, P., Debevec, P.: Acquisition of time-varying participating media. *ACM Trans. Graph.* 24(3) (2005)
21. Heide, F., Hullin, M.B., Gregson, J., Heidrich, W.: Low-budget Transient Imaging Using Photonic Mixer Devices. *ACM Trans. Graph.* 32(4), 45:1–45:10 (Jul 2013)
22. Heide, F., Xiao, L., Kolb, A., Hullin, M.B., Heidrich, W.: Imaging in scattering media using correlation image sensors and sparse convolutional coding. *Optics Express* 22(21), 26338–26350 (Oct 2014)
23. Henyey, L., Greenstein, J.: Diffuse radiation in the galaxy. *The Astrophysical Journal* 93 (1941)
24. Huang, D., Swanson, E., Lin, C., Schuman, J., Stinson, W., Chang, W., Hee, M., Flotte, T., Gregory, K., Puliafito, C., Fujimoto, G.: Optical coherence tomography. *Science* 254(5035), 1178–1181 (1991)
25. Ishimaru, A.: Wave propagation and scattering in random media. Wiley-IEEE (1978)
26. Jakob, W.: Mitsuba renderer (2010), <http://www.mitsuba-renderer.org>
27. Jarabo, A., Marco, J., Muñoz, A., Buisan, R., Jarosz, W., Gutierrez, D.: A Framework for Transient Rendering. *ACM Trans. Graph.* 33(6), 177:1–177:10 (Nov 2014)
28. Jensen, H., Marschner, S., Levoy, M., Hanrahan, P.: A practical model for sub-surface light transport. In: *Proceedings of SIGGRAPH 2001, Annual Conference Series* (2001)
29. Kadambi, A., Whyte, R., Bhandari, A., Streeter, L., Barsi, C., Dorrington, A., Raskar, R.: Coded Time of Flight Cameras: Sparse Deconvolution to Address Multipath Interference and Recover Time Profiles. *ACM Trans. Graph.* 32(6), 167:1–167:10 (Nov 2013)
30. Khungurn, P., Schroeder, D., Zhao, S., Bala, K., Marschner, S.: Matching real fabrics with micro-appearance models. *ACM Trans. Graph.* 35(1), 1:1–1:26 (2015)
31. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. *ICLR* (2015)
32. Levis, A., Schechner, Y., Aides, A., Davis, A.: Airborne three-dimensional cloud tomography. *IEEE International Conference on Computer Vision* (2015)
33. Mukaigawa, Y., Yagi, Y., Raskar, R.: Analysis of light transport in scattering media. *IEEE CVPR* (2010)
34. Narasimhan, S., Gupta, M., Donner, C., Ramamoorthi, R., Nayar, S., Jensen, H.: Acquiring scattering properties of participating media by dilution. *ACM Trans. Graph.* 25(3) (2006)
35. Nayar, S.K., Krishnan, G., Grossberg, M.D., Raskar, R.: Fast Separation of Direct and Global Components of a Scene Using High Frequency Illumination. *ACM Trans. Graph.* 25(3), 935–944 (Jul 2006)
36. O’Toole, M., Mather, J., Kutulakos, K.: 3D Shape and Indirect Appearance by Structured Light Transport. In: *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*. pp. 3246–3253 (June 2014)
37. O’Toole, M., Heide, F., Xiao, L., Hullin, M.B., Heidrich, W., Kutulakos, K.N.: Temporal Frequency Probing for 5D Transient Analysis of Global Light Transport. *ACM Trans. Graph.* 33(4), 87:1–87:11 (Jul 2014)
38. O’Toole, M., Raskar, R., Kutulakos, K.N.: Primal-dual Coding to Probe Light Transport. *ACM Trans. Graph.* 31(4), 39:1–39:11 (Jul 2012)
39. Papas, M., Regg, C., Jarosz, W., Bickel, B., Jackson, P., Matusik, W., Marschner, S., Gross, M.: Fabricating translucent materials using continuous pigment mixtures. *ACM Trans. Graph.* 32(4), 146:1–146:12 (Jul 2013)

40. Pauly, M., Kollig, T., Keller, A.: Metropolis light transport for participating media (2000)
41. Peers, P., vom Berge, K., Matusik, W., Ramamoorthi, R., Lawrence, J., Rusinkiewicz, S., Dutré, P.: A compact factored representation of heterogeneous subsurface scattering. *ACM Trans. Graph.* 25(3) (2006)
42. Prahl, S., van Gemert, M., Welch, A.: Determining the optical properties of turbid media by using the adding–doubling method. *Applied optics* 32(4) (1993)
43. Reddy, D., Ramamoorthi, R., Curless, B.: Frequency-space Decomposition and Acquisition of Light Transport Under Spatially Varying Illumination. In: *Proceedings of the 12th European Conference on Computer Vision*. pp. 596–610. *ECCV’12* (2012)
44. Veach, E.: Robust Monte Carlo methods for light transport simulation. Ph.D. thesis, PhD thesis, Stanford University (1997)
45. Velten, A., Wu, D., Jarabo, A., Masia, B., Barsi, C., Joshi, C., Lawson, E., Bawendi, M., Gutierrez, D., Raskar, R.: Femto-photography: Capturing and Visualizing the Propagation of Light. *ACM Trans. Graph.* 32(4), 44:1–44:8 (Jul 2013)
46. Wang, J., Zhao, S., Tong, X., Lin, S., Lin, Z., Dong, Y., Guo, B., Shum, H.: Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Trans. Graph.* 27(1) (2008)
47. Wu, D., Velten, A., OToole, M., Masia, B., Agrawal, A., Dai, Q., Raskar, R.: Decomposing Global Light Transport Using Time of Flight Imaging. *International Journal of Computer Vision* 107(2), 123–138 (2014)
48. Wu, D., Wetzstein, G., Barsi, C., Willwacher, T., Dai, Q., Raskar, R.: Ultra-fast Lensless Computational Imaging through 5D Frequency Analysis of Time-resolved Light Transport. *International Journal of Computer Vision* 110(2), 128–140 (2014)
49. Wyman, D., Patterson, M., Wilson, B.: Similarity relations for the interaction parameters in radiation transport. *Applied optics* 28(24) (1989)
50. Zeiler, M.D.: Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701* (2012)
51. Zhao, S., Ramamoorthi, R., Bala, K.: High-order similarity relations in radiative transfer. *ACM Transactions on Graphics* (2014)