

Reconstruction of Heterogeneous Scattering Media Using Stochastic Search

Benjamin T. Cecchetto and James Stewart

{cecchett,jstewart}@cs.queensu.ca

School of Computing, Queen's University

Kingston, Ontario, Canada

Abstract - A heterogeneous scattering medium has different material properties in different areas. For such a medium, we present an algorithm to reconstruct the interior materials from measurements of light at the boundary of the medium. The algorithm uses a novel hierarchical stochastic search over the space of materials to find an arrangement of inner materials that best matches the boundary conditions. The algorithm performs a combination of depth- and breadth-first search, choosing random permutations of materials at each step. Validation is performed with a variety of difficult material combinations. We have shown the proposed algorithm is capable of reconstructing a variety of heterogeneous media over the full domain of Henyey-Greenstein materials if the materials are known a priori. We have also demonstrated the algorithm's capability to reconstruct the same materials without knowing the materials a priori.

Index Terms - inverse scattering, rendering, reconstruction, imaging, physics

I. INTRODUCTION

Inverse rendering reconstructs the material parameters in a space using controlled incoming lighting and captured outgoing lighting at the boundary of the material space. Knowing material parameters is of importance for various applications such as medical imaging, atmospheric science, or rendering. The problem is difficult, since a single element in the space can affect the whole part of the boundary. Conversely, the information at a single point on the boundary can potentially come from any element on the interior. Other solutions employ a stochastic gradient based approach. This works quite well, however relies on a good initial estimate. We will show, even for simple scattering reconstruction problems, that the space is riddled with local minima. We provide an algorithm to generate a coarse solution that can better pose a more fine-tuning algorithm such as stochastic gradient descent.

Let the material be defined over a two-dimensional area Ω with boundary $\partial\Omega$. The behaviour of the light inside the material is governed by the Stationary Rendering Equation [6]:

$$(\omega^T \nabla)L(\mathbf{x}, \omega) = Q(\mathbf{x}, \omega) - \sigma_t(\mathbf{x})L(\mathbf{x}, \omega) + \sigma_s(\mathbf{x}) \int_{\psi \in \mathbb{S}^2} p_s(\mathbf{x}, \omega, \psi)L(\mathbf{x}, \psi) d\Phi(\psi), \quad (1)$$

where $L(\mathbf{x}, \omega)$ is *lightfield* (i.e. the radiance at position \mathbf{x} with direction ω), Q is the radiance being emitted, σ_t is the attenuation coefficient, σ_s is the scattering coefficient, $p_s(\mathbf{x}, \omega, \psi)$ is the phase

function at a position \mathbf{x} with entry light direction ω and exiting light direction ψ , and Φ is the spherical integration metric. Also define $\sigma_t = \sigma_a + \sigma_s$, where σ_a is the fraction of light absorbed.

We assume that the phase function is isotropic and thus cylindrically symmetric, so $p_s(\mathbf{x}, \omega, \psi) = p_s(\mathbf{x}, \theta)$ for θ the angle between ω and ψ . Finally, $L(\mathbf{x}, \omega)$ for $\mathbf{x} \in \partial\Omega$ is the boundary lightfield.

The problem is to recover the material parameters, $\sigma_a(\mathbf{x})$, $\sigma_s(\mathbf{x})$, and $p_s(\mathbf{x}, \theta)$ for $\mathbf{x} \in \Omega$, from the boundary lightfield.

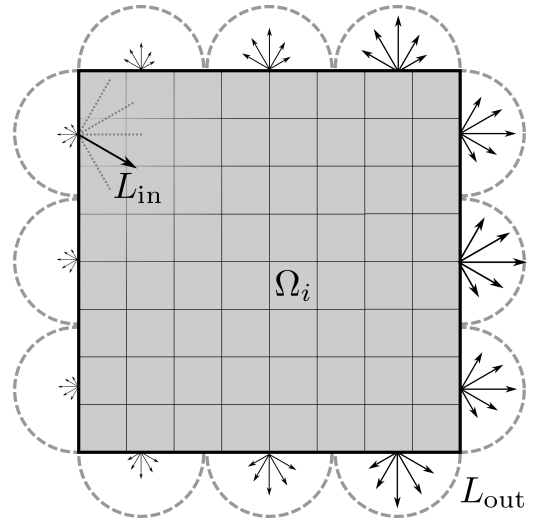


Figure 1: One observation consists of an ideal incoming light source L_{in} (a ray) and outgoing lightfield L_{out} . The reconstruction algorithm determines the material in each region, Ω_i , that generates L_{out} from L_{in} for multiple observations.

If the only unknown is the lightfield L over the whole interior region $\Omega \setminus \partial\Omega$, the equation has been shown to have a unique solution for L under certain practical boundary conditions [4]. Algorithms exist to approximate the solution by discretizing the Stationary Rendering Equation's directional derivative with a small step size [11]. This yields a solution to $L(\mathbf{x}, \omega)$, particularly for $\mathbf{x} \in \partial\Omega$, which can be used to simulate different capture setups with camera models.

To reconstruct the material parameters in Ω , the boundary $\partial\Omega$ is illuminated (i.e. with an incoming lightfield) and the outgoing lightfield is measured. This is done with many incoming lightfields. The ideal light source has only one Dirac position and direction to illuminate the boundary and generates a unique outgoing lightfield [10].

The two-dimensional area Ω is discretized into regions, $\Omega_1, \Omega_2, \dots, \Omega_{n_{reg}}$ such that $\Omega = \bigcup_i \Omega_i$ and different pieces intersect

only on their common boundary. In practice, Ω is a rectangular image and the Ω_i are pixels of the image, as shown in Fig. 1.

The problem is restricted to a small dictionary of materials, the optical properties of which are known *a priori*. Region Ω_i contains exactly one of these materials, $\mathbf{m}_i = (\sigma_a^i, \sigma_s^i, g^i)$, where g^i is the parameter for the Henyey-Greenstein function, which describes the phase function p_s . Let $M = \{\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_{n_{reg}}\}$.

The forward rendering function, R , maps incoming light fields and materials to outgoing lightfields:

$$L_{out} = R(L_{in}, M). \quad (2)$$

But our problem is to find M as a function of n_{obs} sets of L_{out} and L_{in} , corresponding to the captured observations ($\{L_{out_i}\}_i^{n_{obs}}$) for every controlled light source ($\{L_{in_i}\}_i^{n_{obs}}$). This paper describes an algorithm to solve this problem using a joint breadth- and depth-first search with a maximum likelihood estimator. The algorithm is evaluated in terms of reconstruction error and running time.

II. RELATED WORK

Solving a discretized version of Equation 1 for scattering volumes is computationally expensive, since the corresponding matrix formulation can be very large. Some previous research constrained the equation to different specific materials and scenarios, such as heterogeneous translucent materials with a diffusion approximation [26]. The approximation assumes the light follows a diffusion approximation for the physical model, as opposed to the full radiative transfer equation. This approximation allows the physical differential equation to be approximated using a grid, which accelerates computation, but may not work well for non-analytic phase functions.

Some approaches to reconstruction in two-dimensional slices of media use a light source behind the medium and a camera axis oriented perpendicular to the slice [5, 7, 17, 18, 20, 24]. These methods are able to separate both multiple and single scattering components, but only for homogeneous materials. However, these methods only consider a two-dimensional slice, ignoring out-of-slice (three-dimensional) scattering effects.

The behaviour of scattering in liquid media is a matter of particular interest. In order to obtain the necessary optical parameters, we may dilute the liquid with water and observe the change in output. This is because diluting with water reduces scattering. This method has been used to estimate one- and two-parameter phase functions [19].

Photo-acoustic imaging has been used with a diffusion approximation in the interesting case of heated media. As a medium is heated, a sound wave is generated which can be captured [8]. This method is known to be effective to a depth of about 30 mm, though it assumes that the scattering coefficients are known beforehand.

An ultra-fast laser pulse can yield important temporal information, since it provides wavefront visualization at femtosecond resolution, rather than the usual visualization of the time integral. These laser pulses can be used to separate coherent and incoherent components of forward scattering [1, 28]. This type of imaging, although promising to simplify the reconstruction of scattering media, can be expensive. A formulation to reconstruct scattering media using temporal information was proposed by Gkioulekas *et al.* [10].




Figure 2: Sum of Squared Errors over the material space by varying the Henyey-Greenstein coefficients g_1, g_2 , for a square scattering material inside another scattering material. The materials share the same $\sigma_s = 0.5$ and $\sigma_a = 0.05$. The local minima are represented by black circles. The ground truth parameters are $g_1 = 0.9$ and $g_2 = 0.4$ indicated by the black square in the upper right.

A survey of optical tomographic reconstruction techniques [3] discusses how the “general reconstruction from optical scattering problem” is severely ill-posed since the errors grow exponentially. But this survey does not consider constrained problems, such as when materials are known *a priori*, or when polarization is affected by the material. Despite the reconstruction problem being ill-posed according to the authors, there exists a unique solution to the inverse problem [4] given certain conditions. This solution means that there is a unique global minimum for optimization problems.

However, the search space is riddled with local minima. This is especially true with the addition of noise in the measurement and simulation data. Consider a medium with two materials, an outer square band with $(\sigma_s, \sigma_a, g_1) = (0.1, 0.01, 0.9)$ for the outer material and $(\sigma_s, \sigma_a, g_2) = (0.1, 0.01, 0.4)$ for the inner square. Suppose we don’t know the two Henyey-Greenstein coefficients and simply search over the whole space by varying the two parameters. This is depicted in Fig. 2 as a contour plot of the sum of squared errors for the simulated lightfield with chosen (g_1, g_2) parameters with the initial target with the ground truth. Even in this simple example, there are many local minima. For heterogeneous media and arbitrary phase functions this would be even more complex.

There exist local *similarity relations* such that changing the material at a given point inside the medium will not affect the local radiance at that point and thus the boundary lightfields would be identical [27, 29]. However, we make the assumption that we do not have these types of materials as the only materials in the medium. Addressing these materials may require extending our approach

to an optimization method that yields multiple solutions, such as a genetic algorithm [15].

In scattering media, photons with the same initial position and direction can take many different paths. It is therefore useful to classify these types of paths. *Ballistic photons* travel in a straight path through the medium and retain all of their coherency. *Snake photons* travel in a mostly-straight path, but may scatter a small amount while retaining some degree of coherency. All other photons are called *diffuse photons*. If we can distinguish between snake, ballistic, and diffuse photons, then we can detect objects hidden in turbid media with a femtosecond laser pulse [9]. Snake and ballistic photons have previously been captured using a time-resolved single photon detector array to detect endomicroscopy fibres with centimetre resolution [25].

Previous work has attempted to deal with small amounts of multiple scattering, though this work is limited in terms of reconstruction [21]. A diffusion approximation (rather than the full radiative transport equation) can lead to large errors in regions where there is very little scattering present [22]. Techniques that have gaps in reconstructing spatially-varying phase functions use a regularized approach, which smooths out the reconstructed phase function in unknown regions [2].

Inverse rendering techniques, of which the method of this paper is one, generate a simulated output image by simulating photon travel through the medium [11, 12, 14]. The simulated output image is then compared to the corresponding physically measured image, and an optimization algorithm attempts to determine the material properties of the medium that produce the most accurate simulated output image. One method successfully determined these properties by reconstructing a piecewise linear phase function at each point in the medium, rather than using a one- or two-parameter model, which might not take into account multiple scattering directions in one function [11].

Other work describes a general form for reconstructing optical properties of a domain from measurements performed at the domain’s boundary in different scenarios [4]. This work mentions that polarization has been useful for inverse transport problems, such as measuring atmospheric effects and seismic waves. This work also provides a theoretical framework for reconstructing materials with scattering parameters, including uniqueness results under certain assumptions about the boundary region. More recently, researchers have proposed a method for reconstructing scattering media using a time-dependent radiative transfer equation, which decomposes each observation into different path lengths depending on the time the observation takes to arrive at the sensor [10]. This work also formulates a time-independent variant, but only demonstrated a result for a limited type of scattering media. Obtaining three-dimensional tomography for clouds from satellite data has also been investigated [16]. Also, a similar technique was used to estimate scattering parameters in surface volumetric data of cloth [30].

Our algorithm attempts to use all the possible incoming and outgoing light information in the scene. For a two-dimensional medium (having a one-dimensional boundary), light at the boundary is encoded with two parameters: one for position and one for direction. The algorithm proposed in this paper takes as input the outgoing light captured on the boundary of the medium that corresponds to each position, direction, and intensity of an individual incoming

light source. We call this set of output light from input light the *observation*. The scene is rendered with a virtual test material, and the rendered observation is compared to the target observation. The algorithm attempts to minimize the difference between the two observations by iteratively modifying the virtual test material. This feedback loop is *inverse rendering*.

Our reconstruction algorithm uses a known, discrete, finite dictionary of materials with optical parameters $(\sigma_s, \sigma_a, p_s(\theta))$. The reconstruction algorithm assumes that the volume contains only one of these specified materials in each discrete region; this assumption simplifies the problem space substantially. The algorithm additionally considers only two-dimensional slices of a volume and divides the material space into a two-dimensional grid of pixels.

Our work differs from previous work in that we do not use a gradient-based approach. Previous work derived and computed a gradient and used a multiscale stochastic gradient descent to come to a solution [10]. Our work differs from that work by not using a derived gradient for this particular physical model, yielding a more general approach. Instead, we modify the materials to search for a material whose simulated output lightfield matches the input lightfield. We compare the two lightfields using a novel similarity metric to determine how well the simulated observations for a particular M agree with the input lightfield observations.

We are motivated by probabilistic satisfiability solvers [23], which randomly flip bits of literals in a heuristic manner. We propose to perform local search by randomly examining nearby branches both in breadth and depth. This is similar to stochastic gradient descent. However, we permit much wider variances and do not use a computed gradient, which may yield a solution in one of many local minima. In addition, since we do not need a specialized gradient, our optimization may work with other similarly dense problems by directly substituting the simulation step. Our work compares a captured target to a simulated target by modifying the material space, which abstracts the simulation itself. Using a small number of known materials may have immediate applications in medical imaging, as the number of tissue materials may be known *a priori*. We show this approach to work with a variety of dictionary materials as well as demonstrating the feasibility for future work with randomized dictionaries.

We have noted that previous research for reconstructing heterogeneous scattering media only deals with the phase parameter g being in the range $[0, 0.7]$. Our approach works with materials with g in the range $[-1, 1]$. This works for back-scattering materials as well as the commonly tested forward-scattering materials. We demonstrate it for a coarse set of materials as it is intended to precondition for the more appropriate fine-tuned gradient descent.

III. APPROACH

The lightfield is discretized into a number of oriented bins, consisting of uniformly spaced and sized extents on the boundary of the medium and uniformly spaced and sized solid angles above each extent. An observation, $L(L_{in}, \mathbf{b}_{out})$, is the fraction of light leaving the medium in oriented bin \mathbf{b}_{out} that entered the medium due to ideal light L_{in} .

Many optical measurements are made. Let $L_{real}^i(L_{in}^i, \mathbf{b}_{out}^i)$ be the i^{th} real observation. Let $L_{sim}^i(L_{in}^i, \mathbf{b}_{out}^i, M)$ be the corresponding

simulated observation for the current set of materials, M , under consideration.

Our approach is hierarchical: We find the best material for a 1×1 subdivision of Ω , then for a 2×2 subdivision, then 4×4 , and so on. The material chosen for a given cell is used as an initial estimate for the materials in the children of that cell.

To extend this to three dimensions, we would use a 3D voxel grid with similar subdivisions. It would also be necessary to render the light in a three-dimensional medium. The complexity of this increases as there is another dimension of materials to sample over. However, for the three-dimensional case there would be four dimensions of observations on the boundary to determine three dimensions of materials. This is different from the two-dimensional case which has two dimensions of observations, so the three-dimensional variant may have a more constrained solution space.

A. Similarity Metric

A similarity metric quantifies the similarity between the n_{obs} real observations, $\{L_{real}^i\}$, and the corresponding simulated observations, $\{L_{sim}^i\}$.

Since light can potentially travel throughout the whole scattering medium, each observation should be considered to be equally important. The L_1 , L_2 , and L_∞ norms are poor similarity measures, as they can cause an optimizer to concentrate on brighter observations (since reducing the error in those observations has the greatest effect) instead of over all observations equally.

A similarity metric that measures the fraction of observations that are within an error threshold will consider each observation equally, since the error of an observation is either below the threshold, or is not. For threshold ϵ , let

$$\text{within}(a, b, \epsilon) = \begin{cases} 1 & \text{if } |a - b| < \epsilon \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Then the similarity measure is

$$S(L_{real}, L_{sim}, \epsilon) = \frac{1}{n_{obs}} \sum_{i=1}^{n_{obs}} \text{within}(L_{real}^i, L_{sim}^i, \epsilon) \quad (4)$$

In practice, a good value for ϵ can be determined by taking observations of a known homogeneous material and comparing those observations to the simulated observations of the same material. Previous work used an inverse intensity weighted L2 norm to avoid emphasis in brighter regions of the image [10]. However, we wanted a metric to deal with a known capture error as well, hence our metric.

We cannot just set ϵ to be the maximum of these error measurements since there may be noise. We collect the absolute error for all observations, rank them, then set ϵ to the error at the k^{th} percentile, for some k . We have done this through experimentation and set $k = 99.9$ as it is the best result as seen in Fig. 3.

B. Simulated Rendering

The ideal incoming light sources (with single position and direction) are uniformly spaced around the boundary of the medium and have directions uniformly spaced in $(-\frac{\pi}{2}, \frac{\pi}{2})$ relative to the inward pointing normal. In our experiments, we use 128 positions on the




Figure 3: Reconstructions of an opaque occluder with varying iterations and percentile of ground truth error. Each column represents 15, 30, 45, 60, and 75 iterations. Each row represents 90, 99, 99.9, 99.99, 99.999, and 100th percentile of errors used as the error threshold. Setting $k = 100$ is the same as setting ϵ to the maximum error. The best results occur with this setting as well as with more iterations. The bottom-right most figure is the accurately reconstructed ground truth material.




Figure 4: Smoothed running times (hours) for the results in Fig. 3 using an NVIDIA GeForce 780. As expected, running time increases with number of iterations and error threshold.

boundary and 9 orientations, for a total of 1152 incoming light sources.

Observations are also evenly distributed, as described above, and consist of $bins$ on the boundary of the medium. Each bin stores the fraction of the entering radiance that leaves through that bin. For outgoing light, we also use bins at 128 positions and 9 orientations, for a total of 1.3 million ($= 1152^2$) observations.

We modified C++ code from Jaques *et al.* [13] to solve Equation 1. The code uses the discretization of the medium into $\Omega_1, \Omega_2, \dots, \Omega_{n_{reg}}$, as described above. The code has been modified to use NVIDIA’s CUDA platform and to be highly parallelized on the GPU. Each GPU thread operates on 512 photons (this number is optimized for our hardware) from a particular incoming oriented light source bin. When a photon from the light source, after travelling through the medium, reaches the medium’s boundary, the photon’s energy is accumulated in the appropriate oriented bin.

A *simulation iteration* calls CUDA to fire 512 photons from every light source, storing each exiting photon in a bin associated with its corresponding light source.

Inverse rendering performs 100 simulation iterations and evaluates the similarity metric on the aggregate result. This causes 51,200 photons to be sent from each light source, for a total of 58.9 million photons, and takes 9.8 seconds on a GeForce 770.

C. Inverse Rendering to Determine Materials

The goal is to determine a material, \mathbf{m}_i , for each region, Ω_i , such that the similarity metric is maximized. Note that the similarity metric measures similarity of *observations*, not similarity of materials.

The algorithm proceeds hierarchically through a quad tree over the regions, with a 1×1 region equal to Ω at the root. A single material is found that maximizes the similarity metric at the root of the quad tree. This is done by performing inverse rendering with each material, in turn, covering all of Ω .

When proceeding to the next level, the material of each parent region is initially assigned to its 2×2 descendant regions. Then an optimization procedure is performed over the *entire* next level to maximize the similarity metric at that level. This is repeated until the finest level is reached.

Optimizing Materials in One Level Level ℓ of the quad tree contains 4^ℓ regions, which have their region materials initialized from their respective parents.

The algorithm considers a number (usually about ten) of random permutations of the level’s regions. For a particular permutation, $\{\pi_1, \pi_2, \dots, \pi_{4^\ell}\}$, region materials are optimized in a greedy manner in the permutation order: Material \mathbf{m}_{π_1} is optimized, then \mathbf{m}_{π_2} is optimized, then \mathbf{m}_{π_3} , and so on.

For a particular permutation, let

$$M^i = \{\mathbf{m}_1^i, \mathbf{m}_2^i, \dots, \mathbf{m}_{4^\ell}^i\}$$

be the optimized region materials *after* the i^{th} region material in the permutation (that is, material \mathbf{m}_{π_i}) has been optimized. M^0 is the set of materials inherited from the region parents.

Given materials M^{i-1} , materials M^i are determined as follows:

- Step 1. For region Ω_{π_i} , M^{i-1} is modified such that $\mathbf{m}_{\pi_i}^{i-1}$ is replaced by each material from the material dictionary, in turn. For each such material, inverse rendering is performed and the corresponding similarity metric is computed. Let \mathbf{m}_{opt_1} be the material resulting in the maximum similarity metric.
- Step 2. Separately, the *original* M^0 is modified such that $\mathbf{m}_{\pi_i}^0$ is replaced by each material from the material dictionary, in turn. For each such material, inverse rendering is performed and the corresponding similarity metric is computed. Let \mathbf{m}_{opt_2}

be the material resulting in the maximum similarity metric. (Step 2 is not performed for $i = 1$.)

If the maximum similarity of Step 1 is larger than that of Step 2, M^i is set to M^{i-1} with material $\mathbf{m}_{\pi_i}^{i-1}$ replaced by \mathbf{m}_{opt_1} . Otherwise, M^i is set to M^0 with material $\mathbf{m}_{\pi_i}^{i-1}$ replaced by \mathbf{m}_{opt_2} .

IV. RESULTS AND DISCUSSION

In each experiment reported in this section, a 16×16 square region was filled with objects of differing materials. “Real observations” were calculated through simulation for the 1152 ideal light sources and 1152 output bins. Then the algorithm was run using an error threshold of 99.9% and 10 permutations per level. Running times were very long, ranging from 3 to 35 hours.

In general (and discussed in detail below), the reconstruction performs very well when the materials are sparse in space with a different background material. The reconstruction also performs very well when there are not many absorbing materials. But it does not perform well when the materials are purely transmissive (as seen in Figures 7g and 7h).

A. Materials

Experiments using a wide variety of materials were performed. The materials are shown in Table 1. The reconstructions are shown in Fig. 7. In addition, three tests on a coarse Shepp-Logan phantom were performed, shown in Fig. 5. All tests used 100 iterations in the inverse rendering step, except for the “tunnel” experiments (Figures 7j, 7k, and 7l) which used 300 iterations.







Material	Color	σ_s (mm^{-1})	σ_a (mm^{-1})	g
Transmissive		0	0	0
Opaque		0	∞	0
Moderate Scattering		0.60	0.01	0.75
Mild Scattering		0.50	0.03	0.60
Uniform Scattering		0.80	0.01	0
Back Scattering		0.60	0.10	-0.50

Table 1: Known materials used for the simulations and reconstruction material dictionaries.

The algorithm performs well when all of the medium is purely scattering. If there is an opaque occluder (white in Fig. 7) the equation becomes ambiguous as the likelihood of a photon exiting near one of these regions is substantially reduced, as seen in Figures 7b, and 7g.

The algorithm seems not to perform well with significant amounts of transmissive media. This is seen in Figures 7h, 7e, and 7j. The “tunnel” geometry of Fig. 7j, in particular, is very challenging, as it cannot be reconstructed with purely transmissive media. This is because the floating opaque material is in a position that blocks all straight lines through the leftmost two materials. An ART/SART algorithm would be able to find the vertical gap. However, it would not be possible to reconstruct the horizontal one. It seems promising that, with enough iterations and simulation, the algorithm of this paper would be able to reconstruct this scene.




Figure 5: Averages of 5 sampled reconstruction images for a two-material image with moderate scattering as the background and mild scattering as the object. The ground truth was at a full 128x128 resolution. The algorithm performed three random permutations per level.




Figure 6: Average sum of squared error between material images as a function of number 20, 40, and 60 iterations for 5 reconstructions with a maximum of 3 iterations per level optimization. The error bars are one standard deviation.

B. Effect of Iterations and Error Threshold

Fig. 3 shows the results from an experiment varying the error threshold and the number of iterations.

As expected, more iterations usually give more accurate results. However, in some cases, further iterations cause the result to diverge (such as at the 99.99th percentile and 60 iterations, and at the 99.999th percentile and 75 iterations). In these cases, the similarity metric continued to decrease despite the material reconstruction being less accurate. (Recall that the similarity metric measures similarity of observations, not similarity of materials.) This could be caused when the algorithm leaves a local minimum caused by undersampling at a coarser level.

There also seems to be more variation with lower error thresholds. As the threshold is lowered, more observations are deemed to be correct at each iteration, giving the algorithm less guidance on what to optimize.

Fig. 4 shows running times for the scenes in Fig. 3. Running times generally increase with higher error thresholds (and with the number of iterations, of course).

We reconstructed the phantom of Fig. 5. Fig. 6 shows how the sum of squared error (SSE) converges as the number of iterations

increases for this reconstruction. This is the sum of squared differences between the ground truth and reconstructed material properties. Unlike other experiments, only three random permutations were performed at each level.

C. The Case of Unknown Materials

A more difficult version of the problem is one in which no material dictionary is provided. In the absence of a dictionary of known materials, the material properties must be guessed.

Fig. 8 shows what happens if we consider n random materials at each level, where a “random material” is generated by choosing its three parameters uniformly randomly within known bounds. At each level, n new materials are generated. The error for these reconstructions is illustrated in Fig. 9.

This random sampling of materials yields poor, but plausible results. Note that the absorption coefficient is usually very small relative to the other coefficients, making it appear to be more random in nature. The effect of this parameter given the known bounds is negligible on the output lightfield in comparison to the other coefficients. In contrast, the phase parameter, which appears to have the most influence, converges faster than the other two parameters. The sum of squared error for these images as a function of the number of randomly sampled materials is depicted in Fig. 9. The sum of squared error is calculated as before using the difference between the ground truth and reconstructed material properties.

V. CONCLUSION

An algorithm has been presented to reconstruct a set of Henyey-Greenstein materials from the full output lightfield from a set of ideal incoming light sources. The algorithm performs well on small domains, except with purely transmissive media. Although presented in two dimensions, the extension to three dimensions is straightforward.

The algorithm could be extended to use arbitrary material dictionaries with more complex phase functions as well since we do not explicitly use the Henyey-Greenstein parameter in our optimization.

The major problem lies in the running time. As with any inverse rendering algorithm, the simulation at each step is very costly. The algorithm attempts to minimize the number of inverse renderings by performing a small number of greedy searches in the large search space. The results suggest that this strategy works.

To improve the accuracy and running time, we would have many more photons for the Monte Carlo simulation, and many more compute servers. The algorithm is highly parallelizable, making it suitable for render farms or multi-GPU computing. Our reconstructions ran on a consumer-grade NVIDIA GeForce 770, which has 133.9 GFLOPS for double-precision floating-point arithmetic, whereas the current top-of-the-line consumer card, the NVIDIA Titan X, has up to 380 GFLOPS and can be parallelized in SLI for additional performance.

Potential extensions to this algorithm include: limiting the number of views and positions; using known templates of materials to constrain the search space; and storing the computationally expensive simulations and performing analysis on them *a priori* using a machine learning approach.

The algorithm could eventually be used in medical imaging, computer vision, and computer graphics applications. In particular, it could be used to obtain two more parameters about materials, which could lead to better material identification, rendering, and diagnosis.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

AUTHOR CONTRIBUTIONS

Benjamin T. Cecchetto conducted the research and wrote the paper. James Stewart supervised the research and edited the paper. All authors had approved the final version.

REFERENCES

- [1] ALFANO, R., DEMOS, S., GALLAND, P., GAYEN, S., GUO, Y., HO, P., LIANG, X., LIU, F., WANG, L., WANG, Q., ET AL. Time-resolved and nonlinear optical imaging for medical applications. *Annals of the New York Academy of Sciences* 838, 1 (1998), 14–28.
- [2] ANTYUFEEV, V. *Monte Carlo method for solving inverse problems of radiation transfer*, vol. 20. Walter de Gruyter, 2000.
- [3] ARRIDGE, S. Optical tomography in medical imaging. *Inverse Problems* 15, 2 (1999), R41.
- [4] BAL, G. Inverse transport theory and applications. *Inverse Problems* 25, 5 (2009), 053001.
- [5] CECCHETTO, B. T., AND STEWART, J. Reconstruction of photon distributions in liquid scattering media. In *WSCG'2020 - 28. International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision'2020* (2020).
- [6] CHANDRASEKHAR, S. *Radiative transfer*. Courier Dover Publications, 1960.
- [7] CHEN, C., LU, J., DING, H., JACOBS, K., DU, Y., AND HU, X.-H. A primary method for determination of optical parameters of turbid samples and application to intralipid between 550 and 1630nm. *Optics Express* 14, 16 (2006), 7420–7435.
- [8] COX, B. T., ARRIDGE, S. R., KÖSTLI, K. P., AND BEARD, P. C. Two-dimensional quantitative photoacoustic image reconstruction of absorption distributions in scattering media by use of a simple iterative method. *Applied Optics* 45, 8 (2006), 1866–1875.
- [9] FARSIU, S., CHRISTOFFERSON, J., ERIKSSON, B., MILANFAR, P., FRIEDLANDER, B., SHAKOURI, A., AND NOWAK, R. Statistical detection and imaging of objects hidden in turbid media using ballistic photons. *Applied Optics* 46, 23 (2007), 5805–5822.
- [10] KGIOULEKAS, I., LEVIN, A., AND ZICKLER, T. An evaluation of computational imaging techniques for heterogeneous inverse scattering. In *European Conference on Computer Vision* (2016), Springer, pp. 685–701.
- [11] KGIOULEKAS, I., ZHAO, S., BAL, K., ZICKLER, T., AND LEVIN, A. Inverse volume rendering with material dictionaries. *ACM Transactions on Graphics* 32, 6 (2013), 162.
- [12] HAYAKAWA, C., SPANIER, J., BEVILACQUA, F., DUNN, A., YOU, J., TROMBERG, B., AND VENUGOPALAN, V. Perturbation Monte Carlo methods to solve inverse photon migration problems in heterogeneous tissues. *Optics Letters* 26, 17 (2001), 1335–1337.
- [13] JAQCUES, S. L., AND LI, T. mcxyz.c, a 3D Monte Carlo simulation of heterogeneous tissues. <http://omlc.org/software/mc/mcxyz/index.html>. Accessed: 2017-11-11.
- [14] KIM, A., HAYAKAWA, C., AND VENUGOPALAN, V. Estimating optical properties in layered tissues by use of the born approximation of the radiative transport equation. *Optics Letters* 31, 8 (2006), 1088–1090.
- [15] KOZA, J. R. Genetic programming.
- [16] LEVIS, A., SCHECHNER, Y. Y., AIDES, A., AND DAVIS, A. B. Airborne three-dimensional cloud tomography. In *Proceedings of the IEEE International Conference on Computer Vision* (2015), pp. 3379–3387.
- [17] MCCORMICK, N., AND SANCHEZ, R. Inverse problem transport calculations for anisotropic scattering coefficients. *Journal of Mathematical Physics* 22, 1 (1981), 199–208.
- [18] MUKAIGAWA, Y., YAGI, Y., AND RASKAR, R. Analysis of light transport in scattering media. In *IEEE Computer Vision and Pattern Recognition* (2010), pp. 153–160.
- [19] NARASIMHAN, S., GUPTA, M., DONNER, C., RAMAMOORTHY, R., NAYAR, S., AND JENSEN, H. Acquiring scattering properties of participating media by dilution. *ACM Transactions on Graphics* 25, 3 (2006), 1003–1012.
- [20] PRAHL, S., VAN GEMERT, M., AND WELCH, A. Determining the optical properties of turbid media by using the adding–doubling method. *Applied Optics* 32, 4 (1993), 559–568.
- [21] PUSEY, P. Suppression of multiple scattering by photon cross-correlation techniques. *Current Opinion in Colloid & Interface Science* 4, 3 (1999), 177–185.

- [22] REN, K., BAL, G., AND HIELSCHER, A. Transport-and diffusion-based optical tomography in small domains: a comparative study. *Applied Optics* 46, 27 (2007), 6669–6679.
- [23] SCHONING, T. A probabilistic algorithm for k-SAT and constraint satisfaction problems. In *Foundations of Computer Science, 1999. 40th Annual Symposium on* (1999), IEEE, pp. 410–414.
- [24] SINGER, J., GRÜNBAUM, F., KOHN, P., AND ZUBELLI, J. Image reconstruction of the interior of bodies that diffuse radiation. *Science* 248, 4958 (1990), 990–993.
- [25] TANNER, M., CHOUDHARY, T., CRAVEN, T., MILLS, B., BRADLEY, M., HENDERSON, R., DHALIWAL, K., AND THOMSON, R. Ballistic and snake photon imaging for locating optical endomicroscopy fibres. *Biomedical Optics Express* 8, 9 (2017), 4077–4095.
- [26] WANG, J., ZHAO, S., TONG, X., LIN, S., LIN, Z., DONG, Y., GUO, B., AND SHUM, H.-Y. Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Transactions on Graphics* 27, 1 (2008), 9.
- [27] WYMAN, D., PATTERSON, M., AND WILSON, B. Similarity relations for the interaction parameters in radiation transport. *Applied Optics* 28, 24 (1989), 5243–5249.
- [28] YOO, K., AND ALFANO, R. Time-resolved coherent and incoherent components of forward light scattering in random media. *Optics Letters* 15, 6 (1990), 320–322.
- [29] ZHAO, S., RAMAMOORTHY, R., AND BAL, K. High-order similarity relations in radiative transfer. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 104.
- [30] ZHAO, S., WU, L., DURAND, F., AND RAMAMOORTHY, R. Downsampling scattering parameters for rendering anisotropic media. *ACM Transactions on Graphics (TOG)* 35, 6 (2016), 166.

Copyright © 2020 by the authors. This is an open access article distributed under the Creative Commons Attribution License ([CC BY-NC-ND 4.0](https://creativecommons.org/licenses/by-nc-nd/4.0/)), which permits use, distribution and reproduction in any medium, provided that the article is properly cited, the use is non-commercial and no modifications or adaptations are made.



Benjamin T. Cecchetto was born in Toronto, ON, Canada in 1985. He attended the University of Toronto (Toronto, ON, Canada) for his Honours Bachelor of Science in computer science, completed in 2007. He obtained his Master’s of Science from the University of British Columbia (Vancouver, BC, Canada) in 2010 and completed his Doctorate of Philosophy from Queen’s University (Kingston, ON, Canada) in 2019. His research interests include reconstruction from sparse data, visual computing, computational photography, optimization, and physics.



James Stewart is a professor in the School of Computing at Queen’s University. He completed his Ph.D. in computing at Cornell University in 1992. His research program spans several areas within computer graphics and computer assisted surgery, including terrain rendering, visibility in terrains, surgical interfaces, automatic planning of surgery, computation and visualization of uncertainty in surgery. Dr. Stewart has supervised or co-supervised more than thirty graduate students. He has received the Canadian Human Computer Communications Society Service Award in 2008 and the Howard Stavelay Award for Teaching Excellence at Queen’s University in 2005.

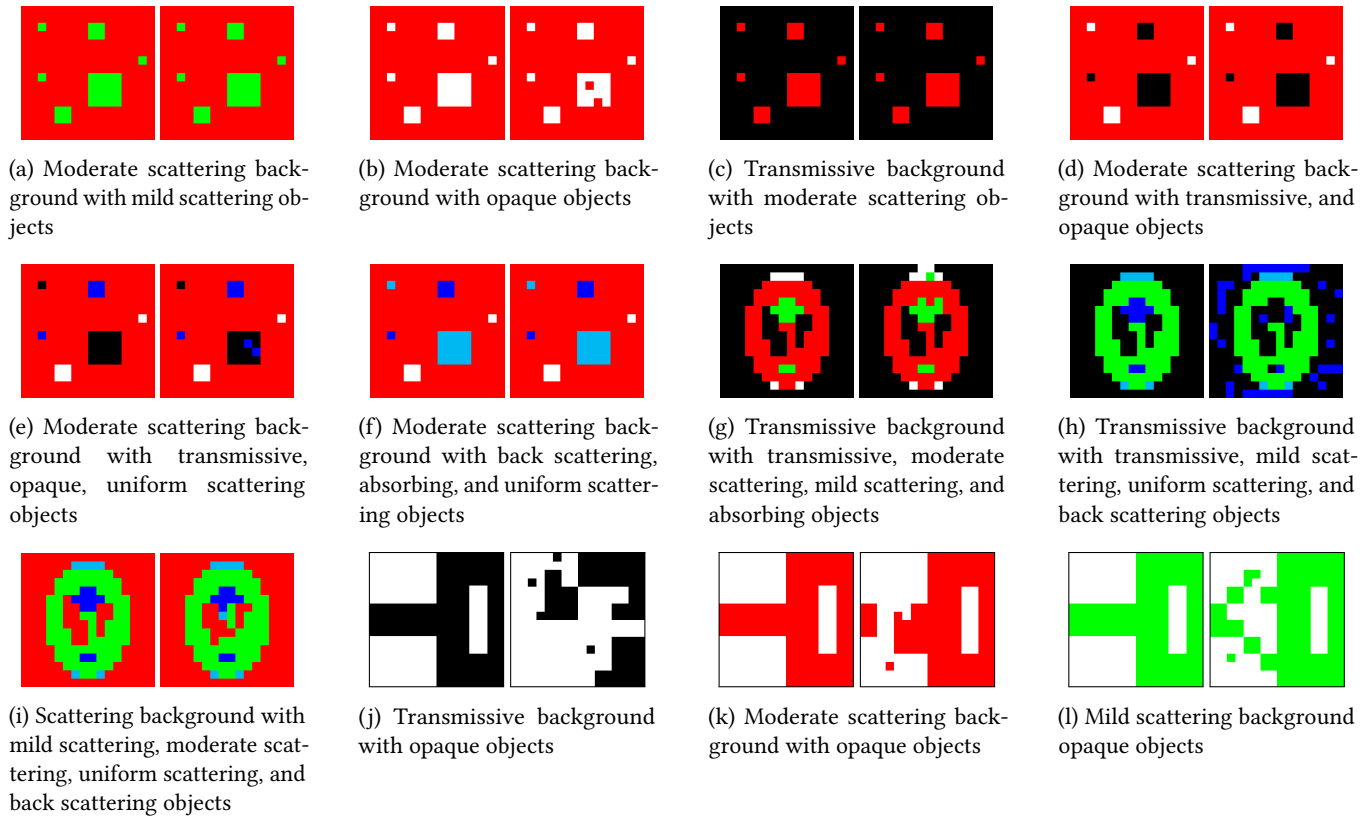


Figure 7: Various reconstructions with the ground truth on the left and the reconstruction on the right. Every dictionary for reconstruction is the same as the amount of materials shown in the ground truth image. The exception is subFig. 7i, which includes all 6 materials from Table 1 in the known dictionary.

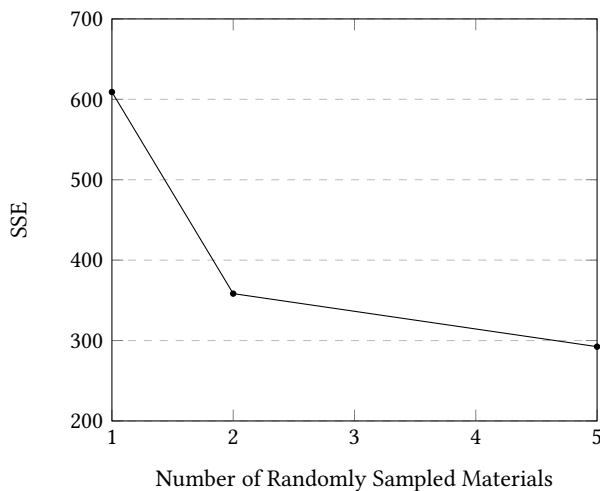


Figure 9: The sum of squared error for the materials reconstructed in Fig. 8.




Figure 8: Reconstructions with unknown materials, for different number of random materials, n , at each level. Each level is repeated 3 times, thus the $n = 1$ level has three random materials attempted, then $n = 2$ has three, and so on. B is the coefficient for the background material, P is the coefficient for the phantom material, and R is the range for the random materials and the images.