

Linear Reconstruction Techniques Applied to Scattering Media

Benjamin T. Cecchetto

cecchett@cs.queensu.ca

School of Computing, Queen's University
Kingston, Ontario, Canada

James Stewart

jstewart@cs.queensu.ca

School of Computing, Queen's University
Kingston, Ontario, Canada

ABSTRACT

The goal of reconstruction or tomographic techniques is to solve for material parameters from boundary information. Linear reconstruction techniques such as ART or SIRT are desirable because of their efficient performance. The derivation of these methods do not take into account scattering media, which is non-linear in nature. We present a summary of linear reconstruction techniques applied to scattering media. We also evaluate using photon distributions as a novel algebraic reconstruction technique matrix. We show the clear benefit of using the randomized reconstruction techniques with many passes over their non-randomized counterparts. We show a marginal improvement in all linear reconstruction techniques with a moderate amount of scattering. We also demonstrate the poor performance of the linear techniques with scattering media, even when using known photon distributions.

CCS CONCEPTS

• **Computing methodologies** → **Image-based rendering; Computational photography; Reflectance modeling; Volumetric models.**

KEYWORDS

inverse scattering, tomography, reconstruction, imaging, physics

1 INTRODUCTION

In this paper we examine reconstructing a heterogeneous scattering medium given a set of observations using linear reconstruction algorithms. Biological tissue consists of scattering media in the visible spectrum. Using visible light to reconstruct biological tissue is desirable to prevent usage of X-rays or other radiation that can damage the tissue.

We focus on extending the well known and commonly used algebraic reconstruction technique (ART), which tends to be very efficient in practice. This leads to a closer initial estimate that could be used in other solvers. ART reconstructs the attenuation coefficients from each observation one at a time [5]. These coefficients are iteratively computed considering a different ray, i , with each iteration using the equation

$$\sigma_t^{(k+1)} = \sigma_t^{(k)} + \mathbf{a}_i \frac{o_i - \mathbf{a}_i \sigma_t^{(k)}}{\mathbf{a}_i \mathbf{a}_i^T}, \quad (1)$$

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICBSP'20, September 27–29, 2020, Kitakyushu, Japan

© 2020 Copyright held by the owner/author(s).

ACM ISBN 1-58113-000-0/00/0010.

<https://doi.org/http://dx.doi.org/10.1145/12345.67890>

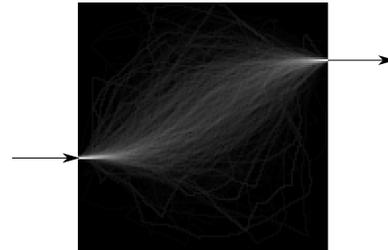


Figure 1: An example photon distribution for a scattering medium. The intensity of each pixel in the image is proportional to the fraction of photons passing through that pixel, which enter and exit in the shown poses (arrows).

where $\sigma_t^{(k+1)}$ is the reconstructed image after the $(k+1)$ th iteration, o_i is the captured attenuation for observation i as defined above, and \mathbf{a}_i is the i th row vector from the matrix A . The initial reconstruction image $\sigma_t^{(0)}$ is set to all zeros.

Given a set of observed intensities \mathbf{o} , the equation allows us to reconstruct a set of attenuation coefficients over the volume σ_t . The reconstructed volume σ_t is solved one path at a time, or in batches of all paths (with Simultaneous Algebraic Reconstruction Technique [1]). The drawback of ART is that each light path is assumed to be a straight line. However, this assumption does not physically hold true for scattering media, since the contributions for each observation do not necessarily come from a straight line path. A more realistic path, or *photon distribution*, is illustrated in Figure 1 and will be discussed later in Section 3.

Extending ART to work for non-linear physical models is not trivial because there is no analytical formulation of light transport through a non-homogeneous scattering medium. Another linear method is the SIRT, which utilizes the same matrix as in ART, but alternates between forward and backward projection at each iteration [11]. We examine the effectiveness of ART, randomized ART, SIRT, and our proposed distribution-based variants of those techniques in largely scattering media. The distribution-based variants assume that the light scatters in the same manner throughout the medium, *i.e.* $\sigma_s(\mathbf{x})$ and $g(\mathbf{x})$ are constant for all \mathbf{x} , but $\sigma_a(\mathbf{x})$ is not. The proposed approaches use the same iterative algorithms as the aforementioned techniques but modify the A matrix to store photon distributions in place of straight lines.

2 BACKGROUND

The most seminal work in reconstructing heterogeneous scattering media was completed by Gkioulekas *et al.* [4]. A stochastic gradient descent algorithm was used to determine the parameters over the

medium. This algorithm would benefit from a closer initial solution. Other recent work involves using time-dependent imaging setups [10], satellite data [8], and applying tomographic techniques to cloth [13]. These methods are dependent heavily on the experimental setup, which may be expensive, or make assumptions of the material being reconstructed. Instead of deriving a specific gradient for the scattering optimization, we will consider using photon distributions in place of the traditional ART matrix. This has the benefit of taking into account scattering, without doing a computationally intensive optimization.

A *pose* is a tuple of a position and direction. The entry pose is where the light enters the medium. The exit pose is where the light exits the medium. A *sensor* is a one dimensional array of bins that captures flux at each *bin*. A bin is a subregion of the sensor that captures light at a certain position and in a certain direction, *i.e.* the exit pose. The radiant flux that each bin collects and its pose is called an *observation*. The collection of bin observations is called the *sensor image*. The sensor is placed at the flat boundary on the opposite side of the emitter to measure light leaving the medium. An example for 10 entry poses and 10 exit poses is illustrated in Figure 2.

In order to solve Equation 1, we need to capture a sufficiently large number of observations. For a 256×256 resolution material image, we require at least that many observations. We accomplish this by varying the input pose position, as well as capturing light at different exit pose positions. We fix the directions of both the entry and exit poses perpendicular to the flat boundary, and hence parallel to each other. The number of input poses is the same as the number of output poses, in this case 256 to reach the required number of observations.

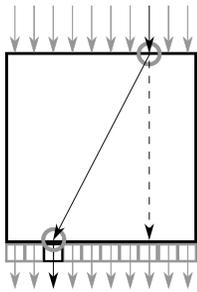
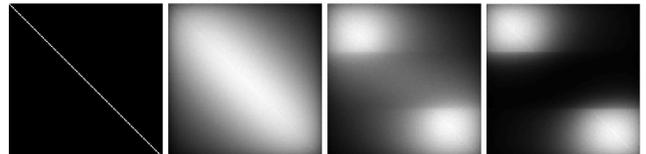


Figure 2: The experimental setup. The arrows at the top indicate 10 different entry poses. The arrows at the bottom indicate 10 different exit poses at each sensor bin illustrated as squares. The dashed line indicates the single path a photon travels in a transmissive medium. The black line represents the most probable path a photon travels through a uniform scattering medium.

As we collect sensor images from different entry poses, we construct a two-dimensional image of observations. This two-dimensional image is a *sinogram*, SI . A pixel at position (row i , column j) in the sinogram image, $SI(i, j)$, corresponds to a single observation with an entry pose i and exit pose j .

Examples of sinogram images for different media are illustrated in Figure 3. In a transmissive medium the sinogram becomes a

diagonal image since the light emitted from an entry pose i travels directly to the corresponding exit pose i (Figure 33(a)). In contrast, light in scattering media travels from a single entry pose to many exit poses. This results in a sinogram image with very few non-zero elements (Figure 33(b)). In a medium with a partially opaque object in the centre, the light is absorbed by the object and the sensor image becomes darker (Figure 33(c)). The reduction in flux is exaggerated with an opaque object in the same place (Figure 33(d)).



(a) Transmissive media (b) Scattering media (c) Scattering media with partially opaque object (d) Scattering media with opaque object

Figure 3: The input to the reconstruction techniques, sinograms for varying media.

Instead of storing the direct captured flux at each pixel of the sinogram image, we process the flux to match the observations of Equation 1. Each captured flux is stored as $-\log \frac{I_{\text{out}}}{I_{\text{in}}}$, where I_{out} is the captured bin flux and I_{in} is the emitter's input flux. This forms the left-hand side of the Beer-Lambert law:

$$-\log \frac{I_{\text{out}}}{I_{\text{in}}} = \sum_{p=1}^N \sigma_t(\mathbf{x}_p) d(\mathbf{x}_p, \mathbf{x}_{p+1}), \quad (2)$$

where the sum of the attenuation coefficients is equal to $-\log \frac{I_{\text{out}}}{I_{\text{in}}}$.

We repeat this process for the four boundary edges of the square medium to generate four sinogram images. The side used to generate each sinogram image defines a *view*, *i.e.* top, bottom, left, or right. These four sinogram images are the input to each reconstruction technique. The output of the reconstruction is an image of attenuation coefficients $\sigma_t(\mathbf{x})$ for all \mathbf{x} in the two-dimensional medium. Recall that the attenuation coefficient includes the scattering and absorption coefficient, *i.e.* $\sigma_t(\mathbf{x}) = \sigma_s(\mathbf{x}) + \sigma_a(\mathbf{x})$.

Reconstructions of similar heterogeneous media have been studied using the assumption that the light scatters only once in the medium. This is called a *one scattering approximation* as described by Mukaigawa *et al.*, who performed a survey of different similar methods [9]. In the one scattering approximation, light travels from an entry pose to a given position somewhere in the medium, changes direction due to a scattering event, then travels to an exit pose. The position of the scattering event is fixed by the intersection of the lines formed by the two poses. As the light travels, it is attenuated through a heterogeneous medium. Gu *et al.* have reconstructed this type of media using structured light patterns projected onto a volume [6]. In their implementation, they reconstruct a volume density field using an iterative, nonlinear inverse rendering approach which tends to be computationally expensive. In addition, the one scattering approximation is not appropriate for dense media since the probability that the light takes the above path is very low. Also, for our particular input, the one scattering approximation

does not account for flux from an entry pose arriving at a different numbered exit pose. This is because the lines formed by the two poses do not intersect. Thus the above model does not model our desired media.

Using a similar approximation, Atcheson *et al.* reconstruct a volume of refractive indices for a gas medium with great success [2]. They assume the light is emitted from a background into a medium, refracts once near the middle of the medium, then arrives at an exit pose. A large set of sparse linear equations is solved to perform the reconstruction. This approach works well, but does not consider scattering.

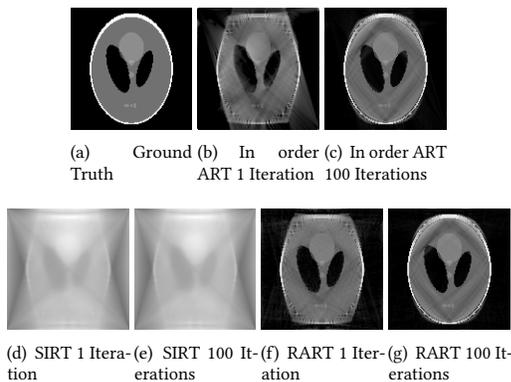


Figure 4: Comparison of reconstruction techniques with our proposed experimental setup and transmissive medium.

Performing ART in order of the observations creates a bias in the reconstructed image, as shown in Figure 4(b). The bias is shown as streaks along the initial reconstructed rays through the medium. To address this bias, instead of performing the reconstruction in order, we randomly order *all* observations from *all* views and perform the reconstruction in that order. Exhausting all possible observations yields one iteration of randomized ART (RART). Using a randomized order significantly improves the output as illustrated in Figure 4(f). Performing more iterations further improves the quality of the reconstructed medium as illustrated in Figures 4(c)4(g). We also consider the reconstruction from SIRT, which is depicted in Figures 4(d)4(e). SIRT introduces the issue of ghosting over the reconstructed medium, likely from being a least squares solution. Additionally, the convergence for these techniques are depicted in Figure 5 for a per-pixel intensity difference (the L2 norm) and a structural similarity metric [12]. The RART yields a marginal improvement over ART in this experimental setup; in contrast SIRT does not perform very well.

3 APPROACH

Recall that for ART, the matrix A is constructed such that each row \mathbf{a}_i is an encoding of the anti-aliased light path through the medium. For distribution-based reconstruction techniques, we propose to use the photon distribution instead of a straight line. An example of the photon distribution is shown in Figure 1. Each photon distribution can be measured as described by Cecchetto *et al.* [3], or can be computed with a Monte Carlo ray tracing algorithm that

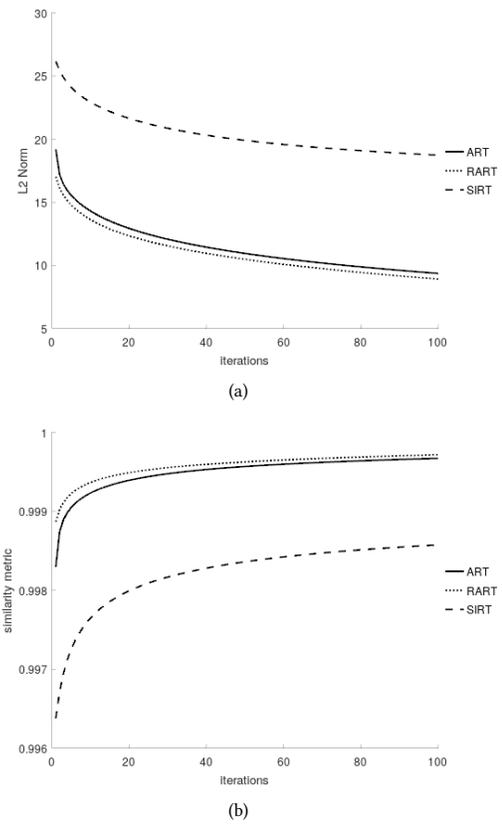


Figure 5: Convergence rates for different reconstruction techniques using the Shepp-Logan phantom

incorporates the scattering parameters of the medium, as will be discussed later in this section.

The photon distribution at a region in the medium is the flux passing through that region normalized by the output magnitude for a given (entry pose, exit pose) pair. We propose to construct A such that each component of \mathbf{a}_i is set to the fraction of the total flux that passes through the corresponding pixel. In classical ART, this weight is usually the length of the intersection of the photon path with the pixel. Some implementations may simplify this to a value that is either 0 or 1, depending upon whether the photon path intersects the pixel. These weights are summed at each pixel for *all* photons, yielding the photon distribution.

In order to generate and use photon distributions, as well as validate results with a realistic model, we need to use an accurate simulation. To accomplish this, Monte Carlo sampling is used to approximate the Radiative Transfer Equation [7]. Photons are emitted from an input pose, travel through the medium in a piecewise linear path generated by sampling the equation, and are aggregated at the sensor.

For a given (entry pose i , exit pose j) pair, we store a photon distribution D_{ij} as illustrated in Figure 6. The array of distributions D has size $n_{\text{pixels}}^2 \times n_{\text{poses}}^2$ where n_{pixels} is the resolution of the square reconstruction image and n_{poses} is the number of the entry

poses, which is equal to the number of exit poses. Each photon distribution D_{ij} is stored in the corresponding reconstruction row, \mathbf{a}_i , in Equation 1 similar to Section 2. The distribution array D is computed *a priori* with a known background material.

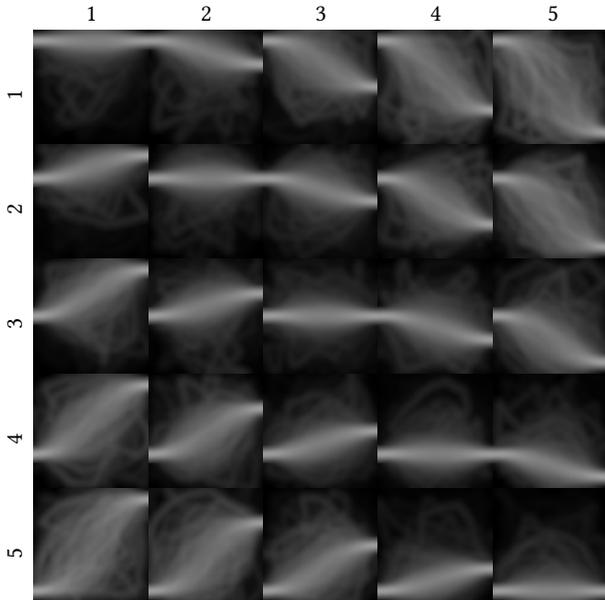


Figure 6: An array of simulated photon distributions D used in distribution-based reconstruction techniques. Each (row entry pose i , column exit pose j) pair yields a two-dimensional photon distribution, D_{ij} .

Here we discuss our graphics processing unit (GPU) implementation for generating photon distributions. Initially all values in D are set to zero and the array is shared by all GPU cores. For a given entry pose i , we simulate a photon bundle using Monte Carlo sampling. The bundle arrives at an exit pose j with an output flux I_{out} . The line path generated by the bundle has an intensity I_{out} that is rasterized and added to the corresponding photon distribution D_{ij} . We simulate each photon bundle from different entry poses on different GPU cores, and the addition of the line path is atomically added (*i.e.* a mutex-locked addition) to the shared array D . We repeat these simulations on more cores for a fixed number of repetitions per pose. This adds a level of parallelism. After the simulation is complete, each element of D is normalized by the total number of photons, giving the fraction of photons passing through each pixel from each entry to exit pose. We obtain very efficient performance since all simulations for the entry poses and repetitions are computed simultaneously. Note that the chance of collision (adding to the same element) in a four-dimensional array is $\frac{1}{n_{\text{pixels}}^2 n_{\text{poses}}^2}$. This becomes quite low with higher resolution distributions.

From the distribution array D , we construct the matrix A which is then used in ART, RART, and SIRT. We denote the variants of these techniques by ART-D, RART-D, and SIRT-D, respectively. We validate the six reconstruction techniques by simulating light using the Radiative Transfer Equation for different test media composed

of many materials. We use three primary test media, and consider linearly interpolated media between them. Each medium has three varying optical parameters, $\sigma_s(\mathbf{x})$, $\sigma_a(\mathbf{x})$, and $g(\mathbf{x})$, as described earlier. We use the one-channel Shepp-Logan phantom image (illustrated in Figure 4(a)) to construct our primary media. This phantom has values in $[0, 1]$ and is commonly used as a one-channel image for $\sigma_a(\mathbf{x})$; however, we will use it to construct three-channel media. We denote the phantom image $\text{SLP}(\mathbf{x})$.

The first medium is mostly transmissive. We denote this medium M_t , where $\sigma_s(\mathbf{x}) = 0.1$, $\sigma_a(\mathbf{x}) = \text{SLP}(\mathbf{x})$, and $g(\mathbf{x}) = 0.99$. The second medium M_e is a uniform scattering phantom in a transmissive medium, thus $\sigma_s(\mathbf{x}) = 2 \cdot \text{SLP}(\mathbf{x})$, $\sigma_a(\mathbf{x}) = 0$, and $g(\mathbf{x}) = 0$. The third medium M_i is an inverted phantom to ensure the exterior of the phantom shape has scattering components. We set $\sigma_s(\mathbf{x}) = 2 \cdot (1 - \text{SLP}(\mathbf{x}))$, $\sigma_a(\mathbf{x}) = 0$, $g(\mathbf{x}) = 0.5$.

We blend between the media using linear interpolation with a parameter α that ranges from 0 to 1. We blend between M_t when $\alpha = 0$ and M_e when $\alpha = 1$. Similarly, we blend between M_t and M_i . The choice of medium for simulating the photon distributions used in the distribution-based techniques is the interpolated parameters without the shape of $\text{SLP}(\mathbf{x})$, *i.e.* $\text{SLP}(\mathbf{x}) = 1$ for M_t and M_e , and $(1 - \text{SLP}(\mathbf{x})) = 1$ for M_i . We make this substitution to hide the shape of the ground truth phantom from the distribution-based techniques.

We simulate the four views described earlier for each interpolated medium in a $10\text{mm} \times 10\text{mm}$ medium. We compare the reconstructed attenuation coefficients (σ_t) with the ground truth (from the scattering parameters $\sigma_t = \sigma_a + \sigma_s$) by using the L2 norm (for per pixel differences) and a structural similarity metric (for the overall medium structure) [12]. The similarity metric combines luminance, contrast, and structural differences into a single metric. Given two images, the similarity metric returns a value in $[-1, 1]$ where -1 means a totally dissimilar structure and 1 means the exact structure.

4 RESULTS AND DISCUSSION

Figure 7 depicts the results from blending between M_t and M_i in the first row, and M_t and M_e in the second row. The graphs in Figure 7 compare the six different techniques ART, RART, SIRT, ART-D, RART-D, and SIRT-D. The graphs 7(a)7(b) in each figure show the effect of the blending parameter α . In graphs 7(c)7(d), α is fixed to 0.33 (taken from a low error region) and shows the convergence for the different techniques. Finally, the ground truth and six final reconstruction images for the same fixed α are depicted in Figure 8.

Overall, all of the techniques achieve a lower amount of error when there is some degree of scattering. This is shown in Figure 7 first row with $\alpha \in [0.2, 0.4]$ and in the second row with $\alpha \in [0.1, 0.6]$. This is likely since the sinogram has more non-zero elements compared to a transmissive medium, which provides more data to the techniques. With excessive scattering, the data becomes blurry which could account for the increase in error.

The techniques perform better with the uniform scattering phantom in an empty medium over the inverted scattering phantom. For higher values of α (and hence higher amounts of scattering) the performance of the distribution-based techniques appears to be better in some cases.

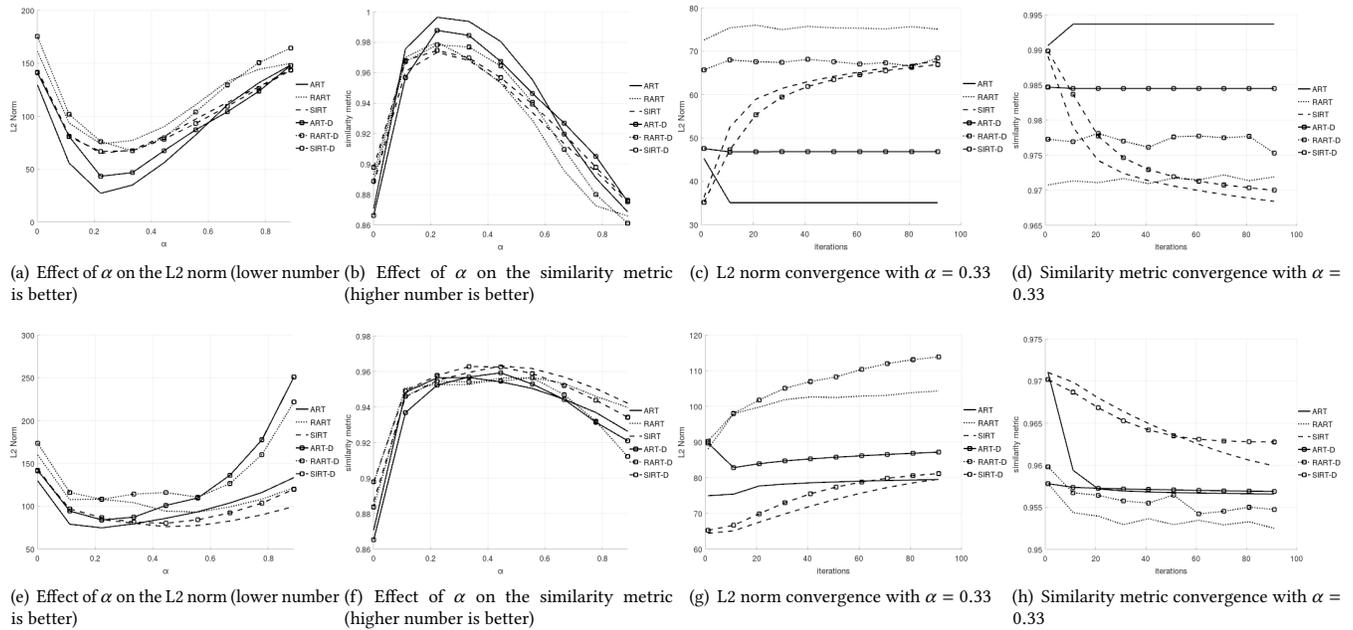


Figure 7: (a)(b) Reconstruction errors for blending between the transmissive phantom (M_t) to the inverted scattering phantom (M_i). 7(c)7(d) Convergence rates for a fixed $\alpha = 0.33$. (e)(f) Reconstruction errors for blending between the transmissive phantom (M_t) to the uniform scattering phantom (M_e). 7(g)7(h) Convergence rates for a fixed $\alpha = 0.33$.

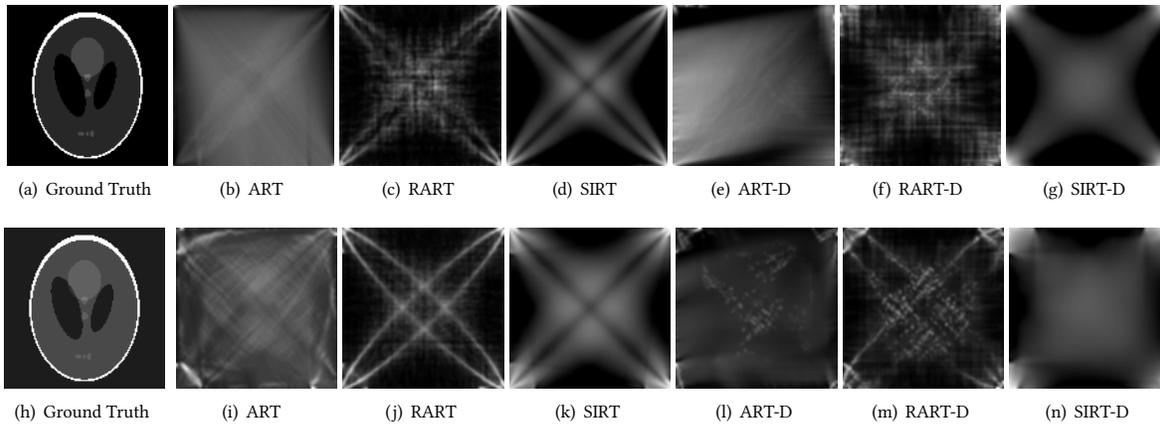


Figure 8: (a)-(g) Final reconstruction images for blending between the transmissive phantom (M_t) to the inverted scattering phantom (M_i). (h)-(n) Final reconstruction images for blending between the transmissive phantom (M_t) to the uniform scattering phantom (M_e). Blending parameter for both reconstruction sets is $\alpha = 0.33$.

All the ART-based techniques seem to converge with more iterations, with the exception of RART in Figure 7(g). Surprisingly, the SIRT-based techniques tend to increase in error with more iterations. The reconstructions from any technique are not guaranteed to be meaningful since the errors do not approach zero.

5 CONCLUSION

We have presented an extension to linear reconstruction techniques to reconstruct a field of attenuation coefficients for heterogeneous scattering media. Our proposed techniques yield a marginal improvement over traditional techniques in some scattering media. A limitation of linear reconstruction techniques is that we can only reconstruct the attenuation coefficient, σ_t . The other two scattering

parameters, σ_s and σ_a are not considered. There is also no guarantee to obtain a meaningful result since the reconstruction error does not approach zero.

A major issue with the distribution-based reconstruction techniques is the assumption that the medium does not affect the average photon distribution. Therefore, we cannot use the correct distribution for arbitrary media without knowing the medium we wish to reconstruct in advance.

A beneficial result of the above techniques is that they are a fast way to obtain a coarse reconstruction. This could be used as input to precondition more intensive algorithms that use accurate physical models. These results could narrow the search space rapidly, whereas more complex methods tends to take longer to achieve the same narrowing of the search space. The narrowing of the search space reduces the overall time of the accurate solution.

REFERENCES

- [1] ANDERSEN, A., AND KAK, A. Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm. *Ultrasonic Imaging* 6, 1 (1984), 81–94.
- [2] ATCHESON, B., IHRKE, I., HEIDRICH, W., TEVS, A., BRADLEY, D., MAGNOR, M., AND SEIDEL, H.-P. Time-resolved 3D capture of non-stationary gas flows. *ACM Transactions on Graphics* 27, 5 (2008), 132.
- [3] 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).
- [4] GKIOULEKAS, 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.
- [5] GORDON, R., BENDER, R., AND HERMAN, G. Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography. *Journal of theoretical Biology* 29, 3 (1970), 471–481.
- [6] GU, J., NAYAR, S., GRINSPUN, E., BELHUMEUR, P., AND RAMAMOORTHY, R. Compressive structured light for recovering inhomogeneous participating media. *IEEE Pattern Analysis and Machine Intelligence* 35, 3 (2013), 1–1.
- [7] HENYEV, L., AND GREENSTEIN, J. Diffuse radiation in the galaxy. *The Astrophysical Journal* 93 (1941), 70–83.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] TRAMPERT, J., AND LEVEQUE, J.-J. Simultaneous iterative reconstruction technique: physical interpretation based on the generalized least squares solution. *Journal of Geophysical Research: Solid Earth* 95, B8 (1990), 12553–12559.
- [12] WANG, Z., BOVIK, A. C., SHEIKH, H. R., AND SIMONCELLI, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612.
- [13] 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.