

Solution of the nonlinear inverse scattering problem by T -matrix completion. II. Simulations

Howard W. Levinson*

Department of Mathematics, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA

Vadim A. Markel†

Aix-Marseille Université, CNRS, Centrale Marseille, Institut Fresnel UMR 7249, 13013 Marseille, France

(Received 22 April 2016; published 25 October 2016)

This is Part II of the paper series on data-compatible T -matrix completion (DCTMC), which is a method for solving nonlinear inverse problems. Part I of the series [H. W. Levinson and V. A. Markel, *Phys. Rev. E* **94**, 043317 (2016)] contains theory and here we present simulations for inverse scattering of scalar waves. The underlying mathematical model is the scalar wave equation and the object function that is reconstructed is the medium susceptibility. The simulations are relevant to ultrasound tomographic imaging and seismic tomography. It is shown that DCTMC is a viable method for solving strongly nonlinear inverse problems with large data sets. It provides not only the overall shape of the object, but the quantitative contrast, which can correspond, for instance, to the variable speed of sound in the imaged medium.

DOI: [10.1103/PhysRevE.94.043318](https://doi.org/10.1103/PhysRevE.94.043318)**I. INTRODUCTION**

This paper is Part II of the series on solving nonlinear inverse scattering problems (ISPs) by data-compatible T -matrix completion (DCTMC). Part I [1] contains theory and here we report initial numerical simulations for the three-dimensional diffraction tomography with scalar waves. This problem arises, in particular, in tomographic ultrasound imaging [2,3] and in seismology [4–6]. We note that DCTMC can be applied in a very similar manner to the problems of inverse electromagnetic scattering [7–11] and diffuse optical tomography [12,13]. However, in this paper we probe the medium with scalar propagating waves. By doing so, we avoid, on one hand, the additional complexity related to the vectorial nature of electromagnetic fields and, on the other hand, the severe ill posedness of the ISP associated with the exponential decay of diffuse waves. Therefore, we can focus on the effects of nonlinearity of the ISP. For the same reason, no noise will be added to the simulated data.

We have studied convergence and computational performance of several modifications of the DCTMC algorithm. We start with the streamlined iteration cycle with both computational shortcuts as described in [1] used. This approach minimizes computational time per one iteration of DCTMC. However, it is not the most efficient algorithm in terms of the number of iterations required for convergence. We then implement several modifications and improvements of DCTMC. In particular, we utilize weighted summation to the diagonal for the “force-diagonalization operator” $\mathcal{D}[\dots]$, which is used to reduce the iteratively updated interaction matrix V_k to a diagonal matrix D_k . This makes the use of Computational Shortcut 2 impossible. However, weighted summation to the diagonal uses explicitly the off-diagonal elements of V_k and is, therefore, more in line with the main idea

of DCTMC. Correspondingly, the convergence of DCTMC is significantly improved. With all the improvements, the number of required iterations in a representative test case was reduced from 900 (for the standard algorithm) to 75. We also demonstrate how physical constraints and checks for sparsity can be embedded in the DCTMC algorithm. Application of these checks can be viewed as a method of regularizing the nonlinear ISP.

The remainder of this paper is organized as follows. In Sec. II, we describe the procedure for discretization of the scalar wave equation. Technical details of the numerical procedures, description of the targets and of the source-detector arrangements used are given in Sec. III. Numerical results for the DCTMC algorithm without any modifications are shown in Sec. IV. In Sec. V, we illustrate several methods to improve the rate of convergence of DCTMC. In Sec. VI, we show that DCTMC can provide good results when the Gauss-Newton method fails. Finally, Sec. VII contains a brief discussion of the obtained results.

II. DISCRETIZATION OF THE SCALAR WAVE EQUATION

Consider a scalar field $u(\mathbf{r})$ (e.g., the pressure wave in ultrasound imaging) and the wave equation

$$[\nabla^2 + k^2\epsilon(\mathbf{r})]u(\mathbf{r}) = -4\pi k^2 q(\mathbf{r}), \quad (1)$$

where $q(\mathbf{r})$ is the source function, $\epsilon(\mathbf{r}) = 1$ outside of the bounded domain Ω (the sample), and the factor $-4\pi k^2$ in the right-hand side has been introduced for convenience. We work in the frequency domain and assume that the wave number $k = \omega/c$ is fixed, where c is the velocity of waves in free space, that is, outside of Ω .

Since we wish to implement a numerical procedure for reconstructing $\epsilon(\mathbf{r})$ from external measurements, the problem must be suitably discretized and the unknown function must be represented by a finite number of degrees of freedom (voxels). To this end, we follow the general approach of the discrete dipole approximation that was originally developed for Maxwell’s equations [14,15]. A related method for the scalar diffusion equation was described by us in [16].

*Currently with the Department of Mathematics, University of Michigan; levh@umich.edu

†On leave from the Department of Radiology, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA; vmarkel@fresnel.fr, vmarkel@mail.med.upenn.edu

We start by rewriting (1) identically as

$$(\nabla^2 + k^2)u(\mathbf{r}) = -4\pi k^2[\chi(\mathbf{r})u(\mathbf{r}) + q(\mathbf{r})], \quad (2)$$

where

$$\chi(\mathbf{r}) \equiv \frac{\epsilon(\mathbf{r}) - 1}{4\pi}. \quad (3)$$

This quantity can be referred to as the susceptibility of the medium. Inverting the differential operator in the left-hand side of (2), we obtain Lippmann-Schwinger equation

$$u(\mathbf{r}) = u_{\text{inc}}(\mathbf{r}) + \int G_0(\mathbf{r}, \mathbf{r}')\chi(\mathbf{r}')u(\mathbf{r}')d^3r', \quad (4)$$

where

$$u_{\text{inc}}(\mathbf{r}) = \int G_0(\mathbf{r}, \mathbf{r}')q(\mathbf{r}')d^3r' \quad (5)$$

is the incident field and

$$G_0(\mathbf{r}, \mathbf{r}') = k^2 \frac{\exp(ik|\mathbf{r} - \mathbf{r}'|)}{|\mathbf{r} - \mathbf{r}'|} \quad (6)$$

is the free-space Green's function of the wave equation, which satisfies

$$(\nabla^2 + k^2)G_0(\mathbf{r}, \mathbf{r}') = -4\pi k^2\delta(\mathbf{r} - \mathbf{r}'). \quad (7)$$

Note that the second term in the right-hand side of (4) is the scattered field $u_{\text{scatt}}(\mathbf{r})$, that is,

$$u_{\text{scatt}}(\mathbf{r}) = \int G_0(\mathbf{r}, \mathbf{r}')\chi(\mathbf{r}')u(\mathbf{r}')d^3r'. \quad (8)$$

The total field is given by a sum of the incident and scattered components, $u = u_{\text{inc}} + u_{\text{scatt}}$.

We now proceed with discretization of the integral equation (4). Let the sample be rectangular and discretized into cubic voxels \mathbb{C}_n ($n = 1, \dots, N_v$) of volume h^3 each, where N_v is the total number of voxels. We then make the following approximation [17]:

$$\chi(\mathbf{r}) = \chi_n \quad \text{AND} \quad u(\mathbf{r}) = u_n \quad \text{IF} \quad \mathbf{r} \in \mathbb{C}_n. \quad (9)$$

Here, χ_n and u_n are constants. Obviously, with this approximation used, (4) can not hold for all values of \mathbf{r} . However, by restricting attention to the points $\mathbf{r} = \mathbf{r}_n$, where \mathbf{r}_n are the centers of the respective voxels \mathbb{C}_n , we obtain a set of algebraic equations

$$u_n = e_n + \sum_{m=1}^{N_v} \chi_m u_m \int_{\mathbf{r} \in \mathbb{C}_m} G_0(\mathbf{r}_n, \mathbf{r})d^3r, \quad (10)$$

where

$$e_n \equiv u_{\text{inc}}(\mathbf{r}_n). \quad (11)$$

Equation (10) can be reasonably expected to have a solution, which then gives a discrete approximation to the continuous field $u(\mathbf{r})$.

In principle, we can compute the integrals in the right-hand side of (10) analytically or numerically with any degree of precision. However, doing so is not practically useful because there is already an approximation involved in writing (10). Therefore, the conventional approach is to approximate the integrals by expressions that are easy to compute if G_0 is

known analytically. To obtain such expressions, we consider the two cases $n \neq m$ and $n = m$ separately.

For $n \neq m$, the standard approximation is

$$\int_{\mathbf{r} \in \mathbb{C}_m} G_0(\mathbf{r}_n, \mathbf{r})d^3r \approx h^3 G_0(\mathbf{r}_n, \mathbf{r}_m), \quad n \neq m. \quad (12)$$

This approximation is not very good for neighboring cells but generally believed to be adequate for cubic lattices (it would be precise for spherical region). However, it should not be applied unaltered for orthorhombic lattices with primitive lattices of nonequal length [18].

For $n = m$, the integrand contains a singularity and a more careful evaluation of the integral is required. It is standard to assume that $kh \ll 1$ (otherwise, the discretization is not a valid approximation to the underlying differential equation) so that the exponent in (6) can be approximated as

$$\exp(ik|\mathbf{r} - \mathbf{r}'|) \approx 1 + ik|\mathbf{r} - \mathbf{r}'|. \quad (13)$$

We then obtain

$$\int_{\mathbf{r} \in \mathbb{C}_n} G_0(\mathbf{r}_n, \mathbf{r})d^3r \approx \int_{-h/2}^{h/2} dx \int_{-h/2}^{h/2} dy \int_{-h/2}^{h/2} dz \times \left(\frac{k^2}{\sqrt{x^2 + y^2 + z^2}} + ik^3 \right) = (kh)^2(\xi + ikh), \quad (14)$$

where

$$\xi = \ln(26 + 15\sqrt{3}) - \pi/2 \approx 2.38. \quad (15)$$

The imaginary part of the right-hand side of (14) is the first nonvanishing radiative correction to the voxel polarizability. As is known for the electromagnetic case, accounting for this correction is important to ensure energy conservation of the scattering process [19,20]. Note that making higher-order approximations in (13) is known to produce higher-order *dynamic corrections*. For the electromagnetic case, the relevant results are given in [21,22]. Also, an alternative estimate of the parameter ξ can be obtained if we replace integration over a cubic voxel by integration over a ball of equal volume. This approach results in a much simpler integral and does not affect the radiative correction; however, it gives a slightly different value for ξ , namely, $\xi = (9\pi/2)^{1/3} \approx 2.42$. The two approaches to computing ξ can be applicable to different physical situations. For example, integration over a ball may be more appropriate if we wish to describe scattering by a collection of spherical particles rather than voxelization of a medium on a cubic grid. A review of different definitions of the parameter ξ for the electromagnetic case and a summary of several relevant expressions, including (15), are given in [23].

Further, it is convenient to introduce the ‘‘moments’’

$$d_n \equiv h^3 \chi_n u_n. \quad (16)$$

In terms of d_n , and with the account of the integration results (12) and (14), the system of equations (10) takes the form

$$d_n = \alpha_n \left(e_n + \sum_{m=1}^{N_v} \Gamma_{nm} d_m \right), \quad (17)$$

where

$$\alpha_n = \frac{h^3 \chi_n}{1 - (kh)^2(\xi + ikh)\chi_n} \quad (18)$$

is the polarizability of n th voxel and

$$\Gamma_{nm} = (1 - \delta_{nm})G_0(\mathbf{r}_n, \mathbf{r}_m) \quad (19)$$

is the interaction matrix. Note that Γ_{nm} has zeros on the diagonal.

If all polarizabilities α_n are known, then Eq. (17) is the statement of the forward scattering problem. Namely, given the incident field in all voxels, e_n , find the moments d_n by solving (17); then, use the formula

$$u_{\text{scatt}}(\mathbf{r}_d) = \sum_{n=1}^{N_v} G_0(\mathbf{r}_d, \mathbf{r}_n)d_n \quad (20)$$

[a discretized version of (8)] to find the scattered field at an arbitrary point of observation \mathbf{r}_d . Obviously, the forward problem is linear with respect to the unknowns d_n .

The inverse problem is stated differently and is, in general, nonlinear. Our goal is to use multiple measurements of the scattered field $u_{\text{scatt}}(\mathbf{r}_d)$ due to multiple external point sources of the form $q(\mathbf{r}) = \delta(\mathbf{r} - \mathbf{r}_s)$, where \mathbf{r}_d and \mathbf{r}_s are the positions of the detector and the source, to find all the voxel polarizabilities α_n . In the matrix form, the ISP is stated as follows. Let us define the interaction matrix V as the $N_v \times N_v$ matrix that contains α_n on the diagonal and zeros elsewhere. Then, we can rewrite (17) in matrix notations as

$$|d\rangle = V(|e\rangle + \Gamma|d\rangle), \quad (21)$$

where $|d\rangle$ and $|e\rangle$ are the vectors of the length N_v with the elements d_n and e_n , respectively. The formal solution to (21) is

$$|d\rangle = (I - V\Gamma)^{-1}V|e\rangle = T[V]|e\rangle, \quad (22)$$

where

$$T[V] \equiv (I - V\Gamma)^{-1}V \quad (23)$$

is the T matrix. Let us assume that we measure the scattered field at a set of points $\mathbf{r}_{dk} \in \Sigma_d$, $k = 1, \dots, N_d$, and let $f_k = u_{\text{scatt}}(\mathbf{r}_{dk})$. Then, $|f\rangle$ is a vector of the length N_d and we can use (20) to write $|f\rangle = A|d\rangle$, where $A_{kn} = G_0(\mathbf{r}_{dk}, \mathbf{r}_n)$. Similarly, we have $|e\rangle = B|q\rangle$, where $B_{nl} = G_0(\mathbf{r}_n, \mathbf{r}_{sl})$ and $\mathbf{r}_{sl} \in \Sigma_s$, $l = 1, \dots, N_s$ is the set of points from which we illuminate the medium. By turning on one source at a time and by measuring the scattered field at all points \mathbf{r}_{dk} for each source, we measure all elements of the data matrix Φ_{kl} . As follows from (22), the data matrix satisfies the equation

$$AT[V]B = \Phi. \quad (24)$$

This is equivalent to Eqs. (4) and (8) of [1]. As in [1], the matrices A and B are obtained directly by sampling the Green's function $G_0(\mathbf{r}, \mathbf{r}')$, for which we have given a specific expression (6) that corresponds to the physical problem considered. The only fine (and slightly nontrivial) point here is the way in which the matrix Γ was obtained. For $n \neq m$, we still have $\Gamma_{nm} = G_0(\mathbf{r}_n, \mathbf{r}_m)$, i.e., the off-diagonal matrix elements were obtained by sampling G_0 . However, the diagonal elements of Γ can not be obtained by straightforward sampling due to the singularity of G_0 . The discretization procedure described above results in $\Gamma_{nn} = 0$ provided that we use the renormalized polarizabilities α_n [Eq. (18)] to quantify the response of individual voxels to external excitation.

The above approach to obtaining the renormalized polarizabilities is rather standard in physics and goes back to the Lorentz's idea of local-field correction (in electromagnetics). However, it is useful to keep in mind that the local-field correction and the related renormalization are mathematically related to the singularity of the free-space Green's function. From the purely algebraic point of view, we can point out that the renormalization in question is a special case of the identical transformation

$$(I - V\Gamma)^{-1}V = (I - V'\Gamma')^{-1}V',$$

where $V' = PV$, $\Gamma' = \Gamma - V^{-1}(I - P^{-1})$, and P is any invertible matrix (although V^{-1} appears in the above expression, invertibility of V is not really required). In particular, if Γ has the diagonal part D , then we can take $P = (I - VD)^{-1}$ (assuming that this inverse exists) and the renormalized matrix Γ' will have a zero diagonal. It also should be pointed out that the renormalized polarizabilities derived here are adequate for cubic lattices but require modification in the case of orthorhombic lattices with primitive vectors of unequal length [18].

III. DETAILS OF THE NUMERICAL ALGORITHMS USED

A. Definition of the unknowns

The above inverse problem has been formulated in terms of the unknown voxel polarizabilities α_n . However, the physical quantity of interest is the voxel permittivity ϵ_n or the susceptibility χ_n . To relate the latter to the former, we can use (3) and (18) to obtain in a straightforward manner

$$\chi_n = \frac{\alpha_n/h^3}{1 + (kh)^2(\xi + ikh)(\alpha_n/h^3)}. \quad (25)$$

The permittivities ϵ_n can be trivially related to χ_n by using (3), and we will focus on the latter quantities in the remainder of this paper.

Equation (25) is an analog of the Maxwell-Garnett formula written for scalar waves (with the account of the first nonvanishing radiative correction) and its inverse is the Clausius-Mossotti relation. It can be seen that the relation between χ_n and α_n is itself nonlinear. In other words, we have removed *some* nonlinearity of the ISP analytically. The nonlinearity in question is solely due to the self-interaction of a given voxel and it was accounted for by the procedure of renormalization of the polarizability. In other words, we can view the numerator $h^3\chi_n$ in the right-hand side of (18) as the bare polarizability and the complete expression as the renormalized polarizability. However, the nonlinearity of the ISP that is due to the interaction of different voxels can not be removed so simply and we will have to deal with it numerically.

Still, the above discussion gives some validity (although not a rigorous proof) to the proposition that it is advantageous to formulate the ISP in terms of the polarizabilities α_n rather than in terms of the susceptibilities χ_n . Indeed, consider the case when there is only one voxel and we wish to determine its properties (either α or χ) from external measurements by means of some generic iterative scheme that is applicable to a more general problem (e.g., involving many interacting voxels). If we view α as the fundamental unknown, we

arrive at a well-posed linear equation of the form $A\alpha = b$ ($A \neq 0$ and b are numbers), which can be solved in just one iteration. However, if we view χ as the fundamental unknown, then we will be solving iteratively an equation of the form $A\chi/(1 - \beta\chi) = b$, which can take several iterations depending on the value of the coefficient β .

In the simulations reported below, we formulate the model and display the reconstructions in terms of the susceptibilities χ_n , which we assume to be real valued. However, to generate the data function, we compute the set of α_n 's according to (18). Then, we “pretend” that α_n 's are unknown and, viewing these quantities as the fundamental unknowns, solve the ISP. We then convert the reconstructed values of α_n 's to χ_n 's by using (25) and display the latter quantities in all figures.

B. Application of physical constraints

As discussed in [1], the iterative procedure of DCTMC can benefit from applying physical constraints to the unknowns as a form of regularization. The constraints are usually based on general physical principles but can also accommodate some *a priori* knowledge about the medium. In this work, only fundamental constraints are used. Specifically, if the medium is passive (that is, nonamplifying), then we know that $\text{Im}\chi_n \geq 0$. If it is also known that the medium is transparent (nonabsorbing), then $\text{Im}\chi_n = 0$. If we view the polarizabilities α_n as the fundamental unknowns in the iterative DCTMC process, we can apply this physical constraint in the following manner. We first notice that

$$\text{Im}\left(\frac{h^3}{\alpha_n}\right) = -\left[\frac{\text{Im}\chi_n}{|\chi_n|^2} + (kh)^3\right] \leq -(kh)^3. \quad (26)$$

Let us say, a numerical iteration of DCTMC has produced a set of α_n . To enforce the condition (26), we can apply the following transformation:

$$\alpha_n \longrightarrow \frac{1}{\text{Re}(1/\alpha_n) - i \max[-\text{Im}(1/\alpha_n), k^3]}. \quad (27)$$

If, in addition, we know that the sample is nonabsorbing so that $\text{Im}\chi_n = 0$ for all voxels and strict equality in the last relation in (26) holds, then we can apply an even sharper constraint by writing

$$\alpha_n \longrightarrow \frac{1}{\text{Re}(1/\alpha_n) - ik^3}. \quad (28)$$

In the simulations reported below, we have assumed that the sample is transparent and used the constraint (28) at each iteration of DCTMC. We note that in many cases considered, DCTMC produces very similar results even without applying the constraint, but then the convergence is somewhat slower. The only instances in which we found that the physical constraint is critical were the cases with very strong nonlinearity.

C. Account of sparsity

In many practical cases we can expect the target to be in some sense sparse. This means that many of the voxels have zero (or relatively small) susceptibilities χ_n , but we do not know *a priori* where these voxels are located or how many

such voxels exist in the computational domain. We will refer to such voxels as “noninteracting” as the moments d_n of these voxels are relatively small, do not interact effectively with the moments of other voxels, and produce a negligible input to the scattered field.

In DCTMC, it is possible to take the sparsity into account in an adaptive manner without actually knowing whether the target is sparse or not or where the noninteracting voxels are located. In the simulations of Sec. IV below, we have used the following rather *ad hoc* algorithm:

1. Run 50 iterations normally.
2. Then, every 20 iterations check whether some susceptibilities χ_n satisfy $|\chi_n| < \chi_{\max}/100$, where $\chi_{\max} = \max_n |\chi_n|$.
3. If a given voxel satisfies the above condition three checks in a row, the corresponding χ_n is set to zero.
4. The voxels with zero χ_n (as determined in the previous step) are declared to be noninteracting and are excluded from the computational domain. When this happens, we repeat the initial DCTMC setup, but now for a smaller number of interacting voxels N_v . This results in a smaller computational time per a subsequent iteration.
5. The process is repeated with the following modifications. After 200 iterations, checks are made every 10 iterations, and after 400 iterations, the relative threshold for determining a noninteracting voxel is reduced to the factor of 60, and after 600 iterations the relative threshold is reduced to the factor of 40.

Note that all integer constants used in the above algorithm are adjustable. In Sec. V, we have used somewhat different values of these constants (as noted in each particular case) and in some simulations we did not use sparsity checks at all.

The procedure of selecting the noninteracting voxels can be described as iterative “roughening” of the target. Any iterative numerical reconstruction is expected to produce small but nonzero values in the regions of the computational domain where the reconstructed function is really zero. Keeping these small values in subsequent iterations is nothing but a computational burden. The procedure described here removes this burden without affecting the final images in a significant way as long as the noninteracting voxels are not assigned incorrectly. In the simulations shown below, such incorrect assignment occurred in the strong nonlinearity regime. We note that such occurrences can be easily “diagnosed” by monitoring the error of the matrix equation (24) as is explained below. If an incorrect assignment does occur, one can potentially alleviate the problem by adjusting the constants used by the algorithm or by not using sparsity checks at all.

Finally, note that the homogeneous background against which the roughening is performed can be different from zero. Thus, in optical tomography, it is conventional to assume that the background properties of the medium are known yet different from those of free space [13]. In this case, a more complicated Green’s function G_0 must be used. We emphasize that G_0 can be computed analytically for many regular geometries of the background medium. In this paper, we have taken the homogeneous background to be free space (vacuum) with the sole purpose of being able to use the mathematically simple function given in (6).

D. Model targets and arrangement of sources and detectors

We have used three kinds of targets in the simulations, which we refer to as large, small, and tiny. A target of a particular kind has always the same “shape” but can have varying degrees of contrast. The contrast is defined in terms of the susceptibility $\chi(\mathbf{r})$. Mathematically, this means that $\chi(\mathbf{r})$ for a given target can be written as

$$\chi(\mathbf{r}) = \chi_0 \Theta(\mathbf{r}), \quad (29)$$

where $0 \leq \Theta(\mathbf{r}) \leq 1$ is the shape function (always the same for a target of a given kind) and $\chi_0 > 0$ is the variable amplitude. Obviously, the larger is χ_0 , the stronger is the nonlinearity of the ISP. Note that all reconstructions shown below display the ratio χ_n/χ_0 , which, ideally, should coincide with the shape function. However, we did not use any *a priori* knowledge about χ_0 to obtain the reconstructions. The normalization to χ_0 was performed *a posteriori*, which allowed us to use the same color scale in all figures.

The majority of simulations shown below are performed for the small target, which is discretized on a $16 \times 16 \times 9$ grid ($N_v = 2304$). DCTMC reconstructions were also demonstrated for the large target, which is discretized on a $30 \times 30 \times 15$ grid ($N_v = 13\,500$). The tiny target ($8 \times 8 \times 4$, $N_v = 256$) is used only in Sec. VI below for the purpose of comparison of DCTMC with the Gauss-Newton methods since the latter did not yield useful results for the small target and were too computationally inefficient to be applied for the large target.

The shapes of all targets are shown in Fig. 1. It can be seen that the small target contains two rectangular inclusions in a homogeneous zero background. One inclusion is of the size $6 \times 6 \times 3$ (in units of h) and has $\Theta = 1.0$ and the second inclusion is of the size $5 \times 5 \times 2$ and has $\Theta = 0.857$. Note that the smaller and the larger inclusions touch at one corner. The large target contains one rectangular inclusion of the size $6 \times 6 \times 4$ and with $\Theta = 1$ and another inclusion of the size $10 \times 10 \times 6$ and with $\Theta = 0.75$. These two inclusions do not quite touch. The tiny target contains an inhomogeneity in the form of two squares of the size $4 \times 4 \times 1$ in located in the two interstitial slices, as is shown in the figure. The two square regions overlap in the central 2×2 area when viewed from top. In terms of the shape function, we have $\Theta(\mathbf{r}) = 1$ inside the inhomogeneity and $\Theta(\mathbf{r}) = 0.5$ in the background. Thus, all voxels in the tiny target are different from free space.

It can be remarked that the tiny target does not represent an accurate discretization of the underlying differential equation. However, it poses a nonlinear ISP in its own right, and various inversion methods can be tested using this model.

The relation between the discretization step h and the free-space wave number k is $kh = 0.2$ in all cases. We can use this numerical value to estimate roughly the degree of nonlinearity of the ISP. We generally expect the ISP to be very nonlinear when the phase shift between two waves, one propagating through an inhomogeneity (say, in the z direction) and another propagating through the background medium, becomes of the order of $\pi/2$ or larger. This phase shift is given by

$$\Delta\varphi = (kh)n(\sqrt{1 + 4\pi\chi_0\Theta_i} - 1), \quad (30)$$

where n is the inhomogeneity depth in units of h , χ_0 is the contrast, and Θ_i is the value of the shape function inside the

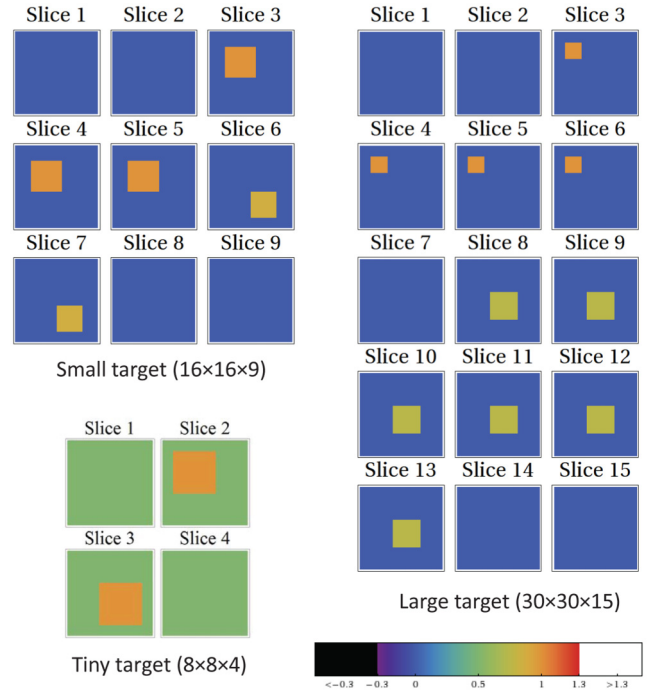


FIG. 1. The shapes of all three targets and the color scheme used in all figures below. Even though all model shape functions satisfy (by definition) the condition $0 \leq \Theta \leq 1$, the color scale can be used to represent any quantity in the range $[-0.3, 1.3]$. The allowances are made because the reconstructed values of χ_n/χ_0 can become negative or larger than unity. Reconstructed values that are smaller than -0.3 will be shown by the uniform black color and values that are larger than 1.3 will be shown as white.

inhomogeneity. For example, the largest contrast we used for the small target is $\chi_0 = 1.75$. The shape function inside the inhomogeneity that is $n = 3$ voxels deep is $\Theta_i = 0.857$. This corresponds to $\Delta\varphi \approx 0.66\pi$. The nonlinearity in this case is expected to be very strong. Of course, this analysis does not take into account multiple scattering between different inhomogeneities. In fact, we will see that the linearized reconstructions break down for a much smaller contrast, i.e., at $\chi_0 \approx 0.1$.

We now turn to the source-detector arrangements. In the case of the small target, the mesh of sources is a 22×22 rectangular grid with the same spacing h as was used to discretize the sample. The plane of sources is centered symmetrically near one of the 16×16 faces of the sample so that there are three rows and columns of sources extending past the sample surface in each direction. The plane of sources is displaced by $h/2$ from the physical surface of the sample. The mesh of detectors is identical in dimension and located on the other side of the sample. Regarding the size of the data set, we have $N_s = N_d = 484$ and the total number of data points is $N_s N_d = 234\,256$. Note that the additional rows and columns of sources and detectors extending past the sample surface are needed to achieve the best possible linear reconstruction in the weak nonlinearity regime. Adding even more rows and columns into the source or detector meshes does not improve this result any further.

In the case of the large target, the grids of sources and detectors are of the size 38×38 so that four rows and columns of sources or detectors are extending past the sample surface in each direction. The source or detector meshes are displaced from each of the two 30×30 faces of the large target by $h/2$, as is also the case for the small target. The data set size is defined by the following numbers: $N_s = N_d = 1444$ and $N_s N_d = 2\,085\,136$.

For the tiny target, the grids of sources and detectors were of the size 10×10 ($N_d = N_s = 100$) and centered about the sample on all sides.

E. Methods for linearized reconstruction

Applying a nonlinear solver to an ISP is meaningful only if the linearized solution to the same problem has failed. Therefore, we compare the results of nonlinear reconstruction to those obtained by linearized inversion. We have used two different approaches to this problem.

First, we have used the traditional approach of computing the pseudoinverse of the matrix K (see Remarks 2 and 3 in [1]). This approach leads to a linear problem whose size grows very rapidly with the size of the data set and the number of voxels. While still feasible in the case of the small target, application of this method for the large target is already problematic. Indeed, the matrix K for the large target has the total number of elements $N_s N_d N_v = 28\,149\,336\,000$. Storing a matrix of this size in computer memory requires at least 112 Gb of RAM (in single precision). One can follow the approach of [24] where we have computed the product K^*K iteratively by storing only sufficiently small blocks of K in computer memory. Then, the product K^*K was Tikhonov regularized by applying the operation $K^*K \rightarrow K^*K + \lambda^2 I$ and the resultant system of equations was solved by the conjugate gradient descent. The bottleneck of this approach is the computation of K^*K , which can still take very considerable time.

On the other hand, solving the same linearized problem by the method described in Appendix B of [1] is a much simpler task, and this approach yields, essentially, the same result. Indeed, in all cases we have considered, the two results were visually indistinguishable. This fact illustrates the proposition that considering the matrices A and B separately rather than combining them into one large matrix K is computationally advantageous even in the linear regime. One can understand this improvement as a result of *data reduction*, that is, defining a transform of the data that is smaller in size but still contains all essential information.

All linearized reconstructions shown below were obtained by regularizing and solving the equation $W|v\rangle = |v_{\text{expt}}\rangle$ (Sec. VII of [1]). The positive-definite matrix W of this equation was Tikhonov regularized by the operation $W \rightarrow W + \lambda^2 I$. The regularization parameter λ was adjusted manually to obtain the best linear reconstruction in each case considered. The equation was then solved by the conjugate-gradient descent method. We note that the computational complexity of computing W (which is exactly of the same size as K^*K , that is, $N_v \times N_v$) is much smaller than the numerical complexity of computing K^*K . Indeed, computing W requires only the singular vectors of the much smaller matrices A and B .

Therefore, the main computational bottleneck of the linearized inversion is removed in this approach.

We have applied different linearization methods, including first Born, first Rytov, and mean-field approximations as described in Appendix A of [1]. The choice of a linearization method only affects the way in which the data are computed and not the matrices K or W . Extensive simulations have revealed that reconstructions based on first Rytov or mean-field approximations do not provide any noticeable advantages when compared to first Born. First Born and first Rytov reconstructions are compared in Fig. 2 below for illustrative purposes, but in all other cases only first Born based reconstructions are shown.

F. Error measures

To quantify the convergence of the method, we use normalized root mean square errors η_χ (error of the solution) and η_Φ (error of the equation). These quantities are defined as

$$\eta_\chi^2 = \frac{1}{N_v \chi_0^2} \sum_{n=1}^{N_v} [\chi_n^{(\text{Reconstructed})} - \chi_n^{(\text{True})}]^2, \quad (31a)$$

$$\eta_\Phi^2 = \frac{1}{N_d N_s \chi_0^2} \sum_{i=1}^{N_d} \sum_{j=1}^{N_s} [\Phi_{ij} - (ATB)_{ij}]^2. \quad (31b)$$

Note that η_χ is computed after Step 1 of the streamlined iteration cycle with the use of Computational Shortcut 2 or after Step 2 for the algorithm without the use of Computational Shortcut 2 (as defined in [1]). The error η_Φ is computed after Step 4 of this algorithm. After Step 5, the T matrix is fully data compatible and the error in question is zero up to the numerical precision of the computer. It should be kept in mind that η_χ can be computed only if the target is known *a priori*, which in practical applications is almost never the case. However, η_Φ can be computed even if the target is not known *a priori*.

IV. NUMERICAL RESULTS

A. Small target

Reconstructed images of the small target after 900 DCTMC iterations are shown in Fig. 2 for various degrees of contrast. Images obtained by linearized inversion are also shown for comparison. It can be seen that the linearized image reconstruction methods break down between $\chi_0 = 0.01$ and 0.1. First Rytov approximation does not provide a visible advantage over first Born, and the same is true for the mean-field approximation (data not shown). However, DCTMC yields reconstructions that are close to the correct result for all values of the contrast used. At $\chi_0 = 1.75$, the DCTMC reconstruction starts to visibly break down. The reason for this is incorrect assignment of three noninteracting voxels, as can be seen from the figure with sufficient magnification. The vicinity of these three voxels is also reconstructed incorrectly. So, the breakdown in this case is due to the *ad hoc* algorithm for assigning the noninteracting voxels. This problem is not fatal for DCTMC as the algorithm can be adjusted. Note that it is possible to learn whether some noninteracting voxels have been assigned incorrectly by looking at the convergence curve

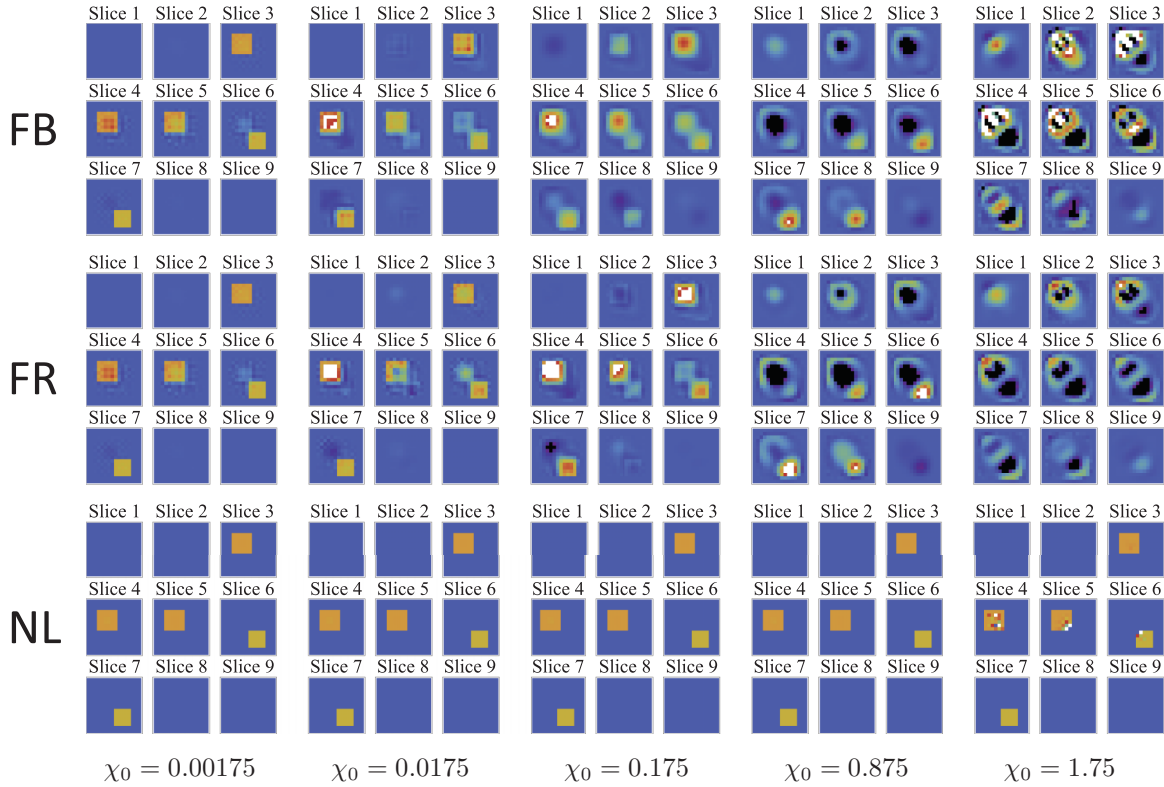


FIG. 2. Linear (top and middle rows, marked FB and FR) and nonlinear (bottom row, marked NL) reconstructions of the small target for different levels of contrast χ_0 . The quantity shown by the color scale in each plot is χ_n/χ_0 , where χ_n is the reconstructed susceptibility of the n th voxel (real under the assumptions used) and χ_0 is the amplitude of the shape function. For the linearized reconstructions, FB denotes first Born approximation and FR denotes first Rytov approximation.

for the error η_ϕ (see Fig. 3). Doing so does not require any *a priori* knowledge of the target.




FIG. 3. Convergence data for the small target. Errors η_χ and η_ϕ are plotted vs the iteration number i . Different line types correspond to different contrast χ_0 as follows: $\chi_0 = 0.00175$ (1), $\chi_0 = 0.0175$ (2), $\chi_0 = 0.175$ (3), $\chi_0 = 0.875$ (4), and $\chi_0 = 1.75$ (5).

The errors η_χ and η_ϕ for the small target are shown in Fig. 3 as functions of the iteration number i . We start the discussion of convergence with the error of reconstruction η_χ . The curves $\eta_\chi(i)$ all look very similar and are almost independent of χ_0 , except for $\chi_0 = 1.75$. The latter case will be discussed separately. For $\chi_0 < 1.75$, the function $\eta_\chi(i)$ has three distinct convergence regimes, as described below.

First, there is a region of slow convergence. While DCTMC is in this regime, the reconstructed values χ_n (or α_n) are all very small. Therefore, it can be concluded that the initial iterations simply solve the linearized inverse problem by the iterative algorithm described in [1] (Richardson first-order iteration). Naturally, convergence of this iterative process is expected to be slow. In principle, this initial slow convergence regime can be completely avoided by solving the linearized problem directly or by a fast iterative method such as conjugate-gradient descent, and then using the result as the initial guess for DCTMC. As is explained in Sec. III E, linearized inversion can be obtained relatively fast with the use of data reduction that is inherent in treating the matrices A and B separately rather than combining them into one large matrix K . The computational advantage obtained by this approach is illustrated in Sec. V A.

Second, there is a region of fast convergence. Assignment of noninteracting voxels occurs in this range of iteration indexes. If the curves are viewed with sufficient magnification, it can be seen that they are not smooth but contain “kinks”, that is, points where the slope changes abruptly. This change of slope occurs when one or more noninteracting voxels are determined correctly.

Third, there is the last region in which the rate of convergence is lower than in the second region but still much larger than in the first, slow convergence region. Clearly, the error $\eta_\chi(i)$ decreases in the third region according to an exponential law. In this convergence regime, all noninteracting voxels have been determined correctly and the algorithm improves its estimate of the amplitudes of the remaining interacting voxels. The error η_χ continues to decrease to some very small values. This means that the reconstructions can be made very precise. However, the exponential convergence can not continue indefinitely because the error η_χ can not decrease to an arbitrarily small value; it is bounded from below either by the ill posedness of the inverse problem or by round-off errors.

The above discussion is valid for the contrasts $\chi_0 < 1.75$. For $\chi_0 = 1.75$, the pattern is quite different. The reason is that three noninteracting voxels have been assigned incorrectly in this case. As a result, in the first, slow convergence region, the function $\eta_\chi(i)$ actually increases. These jumps take place when noninteracting voxels are assigned incorrectly (apparently, too early). Then, there still exists the fast convergence region, wherein many noninteracting voxels are assigned correctly. Finally, the third, exponential convergence region does not exist for $\chi_0 = 1.75$. This is so because the incorrectly assigned noninteracting voxels set a relatively large lower bound for the error η_χ .

We next discuss the error η_Φ . This error can be computed even if the target is not known *a priori*. First, we note that the dependence $\eta_\Phi(i)$ can also be classified into three different regimes (slow, fast, and exponential), just as it was done for $\eta_\chi(i)$. However, unlike in the case of $\eta_\chi(i)$, the curves $\eta_\Phi(i)$ depend noticeably on the contrast, χ_0 . The dependence is, of course, weak for small χ_0 . Thus, the two curves for $\chi_0 = 0.00175$ and $\chi_0 = 0.0175$ are barely distinguishable. However, the curves for $\chi_0 = 0.0175$, $\chi_0 = 0.175$, and $\chi_0 = 0.875$ can be easily distinguished. The difference is most pronounced in the exponential convergence region. The exponent appears to be the same but the overall factor depends on χ_0 . This dependence of η_Φ on χ_0 is a manifestation of the nonlinearity of the inverse problem. Indeed, it can be easily shown that, in the linear regime, $\chi_0 \rightarrow 0$, η_Φ is independent of χ_0 . We note that the dependence on χ_0 can also be visible in the slow convergence region of the iteration indexes. This does not contradict the previously made observation that, in the slow convergence regime, DCTMC solves the linearized problem iteratively. The reason is that the definition of η_Φ involves the data matrix Φ , which, in general, is not proportional to χ_0 .

It was mentioned above that the overall factor of the function $\eta_\Phi(i)$ in the exponential convergence region depends on χ_0 . Initially, this factor increases with χ_0 . More generally, the curve $\eta_\Phi(i)$ for $\chi_0 = 0.875$ goes higher than the curve for $\chi_0 = 0.175$ and the curve for $\chi_0 = 0.175$ goes higher than the curve for $\chi_0 = 0.0175$, etc. However, at larger values of χ_0 , this tendency is reversed. Thus, the curve for $\chi_0 = 1.75$ starts lower than the curve for $\chi_0 = 0.00175$ even before any noninteracting voxels have been assigned. This nonmonotonous dependence of η_Φ on χ_0 is a manifestation of the rather complicated nonlinearity of the inverse problem and is suggestive of a resonance phenomenon. That is, when we increase the value of χ_0 , it passes close to a pole in the complex

plane where the T matrix has a singularity. We note that this nonmonotonous dependence was observed for the large target as well (see below).

The curves $\eta_\Phi(i)$ are rather interesting in the strong nonlinearity case $\chi_0 = 1.75$. The error η_Φ initially increases due to the incorrect assignment of the three noninteracting voxels. Then, the error drops rapidly when many noninteracting voxels are assigned correctly. Then, the error undergoes exponential decay with the same exponent as for the smaller contrasts. Interestingly, the error η_χ in this range of iteration indexes is nearly constant while the error η_Φ steadily decreases. This is so because η_χ is dominated at this point by the three incorrectly assigned noninteracting voxels while the algorithm still improves the accuracy of χ_n in the remaining majority of voxels. Finally, the exponential decay of $\eta_\Phi(i)$ crosses over to exponential growth. Why this happens is not entirely clear; one could expect $\eta_\Phi(i)$ to plateau. We can conjecture that this crossover to exponential growth occurs due to a complex interplay of the physical constraint (that is still being applied at each iteration) and the use of an incomplete computational domain due to the incorrect exclusion of some of the voxels. We note in passing that the best reconstruction result for $\chi_0 = 1.75$ would have been obtained if we stopped the iterations at $i \approx 600$.

We conclude that the characteristic behavior of $\eta_\Phi(i)$ gives one an unambiguous indication that noninteracting voxels have been assigned incorrectly or (in more severe cases) that the obtained DCTMC reconstruction is not useful.

B. Large target

We now turn to the large target. In Fig. 4, we show the reconstructions for five different values of the contrast, from $\chi_0 = 0.002$ to 2. Of course, the computational domain and the size of the inhomogeneities are in this case larger than in the small target and we can expect the onset of nonlinearity to occur at smaller values of the contrast. Also, the linearized inverse problem is more ill posed because there are voxels in the interior of the large target that are quite far from any source or detector and not effectively probed by incident evanescent waves. Indeed, the linearized reconstructions of the large target are not as good as the similar reconstructions of the small target. We emphasize that these are the best linearized reconstructions we were able to obtain by tuning the Tikhonov regularization parameter. Also, and as was the case for the small target, neither first Rytov nor the mean-field approximation provide a noticeable improvement of the linearized reconstructions (data not shown).

We note, however, that the DCTMC reconstructions of the large target in the weak nonlinearity regime (e.g., $\chi_0 = 0.002$) are considerably better than the linearized reconstructions. Moreover, it is not possible to improve the image quality of the linearized reconstruction by setting to zero the amplitudes of all voxels that are less in magnitude than, say, $\frac{1}{40}$ of the maximum (recall that 40 is the smallest thresholding factor used by us in this paper to determine noninteracting voxels in DCTMC). In fact, a voxel with the reconstructed susceptibility $\chi_n = \chi_0/40$ is not visually distinguishable from zero in the color scheme used in this paper. Therefore, the difference in quality between DCTMC and linearized methods

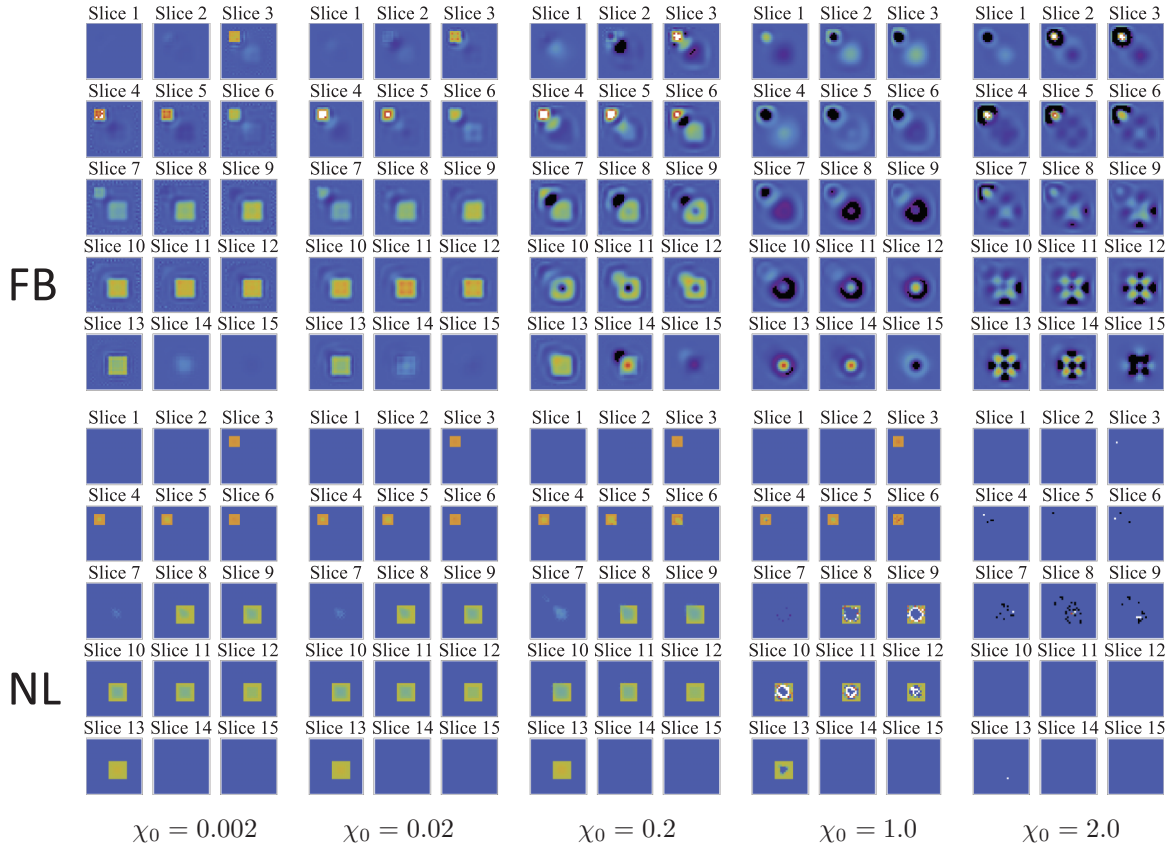


FIG. 4. Same as in Fig. 2, but for the large target and a somewhat different set of contrasts χ_0 . Utilization of first Rytov approximation for linearized reconstruction does not provide any improvements over first Born approximation, and the corresponding results are not shown in this figure.

is not a trivial consequence of image “roughening”. The result may appear counterintuitive. Indeed, in the limit $\chi_0 \rightarrow 0$, the linearized inversion methods and DCTMC are solving exactly the same problem. However, the two approaches involve different regularization methods. In linearized inversions, Tikhonov regularization is used. In DCTMC, regularization is afforded by applying the physical constraint at each iteration. Apparently, the latter approach is much better at reproducing sharp edges.

As we move to the strong nonlinearity regime, the linearized reconstructions break down. At $\chi_0 = 0.2$, the linearized reconstruction is not useful while DCTMC still provides a quantitatively accurate result. The contrast level $\chi_0 = 1$ is borderline for DCTMC. It can be seen that the smaller, higher contrast inhomogeneity is still reconstructed correctly. However, the interior region of the larger inhomogeneity is not properly reconstructed (for the most part, underestimated). The boundaries of the larger inhomogeneity are nevertheless clearly visible. Note that we have encountered a similar situation in the case of nonlinear inversion by inverse Born series [25]. As in the case of the small target, the reason for the incorrect reconstruction of the larger inhomogeneity is the incorrect assignment of noninteracting voxels. However, the reconstruction is not yet entirely broken for the contrast $\chi_0 = 1$. At $\chi_0 = 2$, too many noninteracting voxels have been assigned incorrectly and the resulting reconstruction is not useful.

Convergence data for the large target are shown in Fig. 5. The behavior of the errors is qualitatively similar to what was observed in the case of the small target, although the overall rate of convergence is obviously lower. We note that the results for $\chi_0 = 1$ and 2 are qualitatively different from those for smaller values of the contrast. This is explained by the incorrect assignment of noninteracting voxels at the higher levels of contrast.

We note that the curve $\eta_\Phi(i)$ for the borderline case $\chi_0 = 1$ has a well pronounced minimum at $i = 590$. We have observed a similar nonmonotonous convergence for the smaller target as well. We have compared the reconstructions of the large target with $\chi_0 = 1$ obtained at the “optimal” iteration index $i = 590$ and at $i = 900$ (data not shown). As expected, the result is visibly but not dramatically better at $i = 590$. This result confirms our conjecture that monitoring the error $\eta_\Phi(i)$ is a useful approach to deciding when the iterations should stop.

V. METHODS TO IMPROVE CONVERGENCE

This section investigates several improvements to the “standard” DCTMC algorithm. It will be shown that, with a few simple modifications, we can substantially reduce the number of iterations required to obtain accurate reconstructions. All simulations shown below concern the small target.




FIG. 5. Convergence data for the large target. The different line types correspond to different contrast χ_0 as follows: $\chi_0 = 0.002$ (1), $\chi_0 = 0.02$ (2), $\chi_0 = 0.2$ (3), $\chi_0 = 1$ (4), and $\chi_0 = 2$ (5).

A. Starting from the linear reconstruction

As was shown in [1], DCTMC in the linear regime is equivalent to Richardson first-order iteration, which can have slow convergence. The regions of initial slow convergence are shown in Figs. 3 and 5. Two approaches can be used to cope with this problem. The first approach is to regularize the ISP as is explained in [1]. This, however, can result in image degradation beyond what is warranted by the intrinsic ill posedness of the problem. The second approach is to start iterations with the linearized solution to the ISP, V_{lin} , as the initial guess. This is computationally efficient because we have a fast linearized solver at our disposal as is explained in Sec. III E.

In Fig. 6, we compare two initial guesses. The first initial guess is obtained by $V_1 = \mathcal{D}[(I + T_{expt}\Gamma)^{-1}T_{expt}]$. The second initial guess is $V_1 = V_{lin}$. Both initial guesses were obtained for the small target and $\chi_0 = 0.175$. It can be seen that the first initial guess is very close to zero. It takes many Richardson iterations to transform it to the image resembling the one shown in the right panel. After this is accomplished, DCTMC starts to actually solve the nonlinear ISP. But, the initial computational effort to arrive from the image in the left panel to the one in the right panel is unnecessary and can be avoided.

To show that the number of necessary DCTMC iterations can be significantly reduced by using the initial guess $V_1 = V_{lin}$, we have used the following algorithm:

1. Obtain the linearized reconstruction.
2. Using the result of Step 1 as the initial guess, run five iterations of DCTMC normally.
3. Then, every five iterations check whether some of the values of χ_n satisfy $|\chi_n| < \chi_{max}/500$, where $\chi_{max} = \max_n |\chi_n|$.

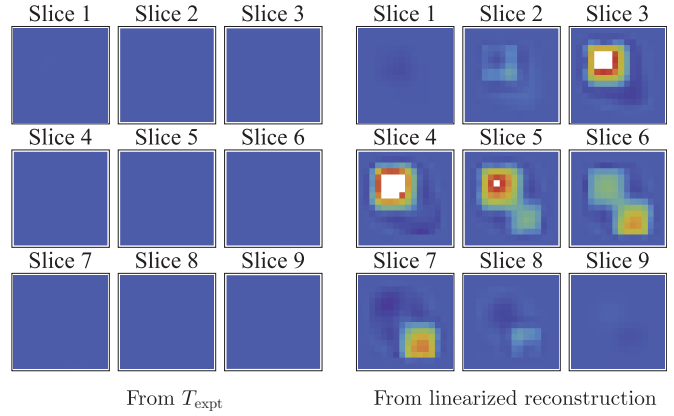


FIG. 6. Two different initial guesses (DCTMC iteration starting points) for the small target with $\chi_0 = 0.175$. The left panel is the initial guess for the T matrix, $T_1 = T_{expt}$; the image is then computed by $V_1 = \mathcal{D}[(I + T_1\Gamma)^{-1}T_1]$. The right panel is $V_1 = V_{lin}$, where the linearized reconstruction V_{lin} was obtained by using the first Born approximation.

4. If a given voxel satisfies the above condition three checks in a row, the corresponding χ_n is set to zero, and the computational domain is reduced.

5. The process is repeated with the following modifications. After 15 iterations, checks are made every 20 iterations and the relative threshold for determining a noninteracting voxel is reduced to the factor of 100. After 200 iterations, checks are made every 10 iterations, and after 400 iterations, the relative threshold is reduced to the factor of 60, and after 600 iterations the relative threshold is further reduced to the factor of 40.

Compared to the original algorithm described above, the interval between the sparsity checks is reduced at the beginning to utilize the noninteracting voxels found by the linear reconstruction, while simultaneously raising the relative threshold so as to not trust the linear reconstruction too much. Empirical evidence suggests that linearized reconstructions tend to capture the boundaries of an object more or less correctly but the interior and the regions close to the boundaries can be reconstructed with a large error [25]. Therefore, it can be expected that a linearized reconstruction would predict correctly the existing noninteracting voxels outside of the object and sufficiently far from its boundaries. However, noninteracting voxels in the interior or close to the boundaries can be predicted incorrectly. Therefore, we have modified the first 15 iterations of the algorithm to avoid such incorrect assignments. After the initial 15 iterations, the parameters of the algorithm are reset to match the original method.

The error plots illustrating the convergence of DCTMC when started from the two different initial guesses are shown in Fig. 7. A few immediate conclusions can be made from these plots. First, as we are starting from a much more reasonable guess, our initial error is much smaller. More importantly, as we can use sparsity checks accurately and relatively early in the iteration process, we have fast convergence almost from the start. Overall, for the contrast considered, we arrive at the final result in about 150 iterations less. This is significant but not yet dramatic.




FIG. 7. Convergence rate for the small target with $\chi_0 = 0.175$. Different curves correspond to using the initial guess shown in Fig. 6, as labeled.

Also, the above improvement is not obtained for the strongest contrast we have considered, $\chi_0 = 1.75$. Here, when we start from the linearized reconstruction, we do not outperform the original method (data not shown). This can be understood qualitatively by noting that the linearized reconstruction at this contrast is very far from the true target. In fact, it is further away in the L_2 norm from the true target than the zero initial guess. Moreover, the linearized reconstruction is very strongly affected by the choice of the Tikhonov regularization parameter; an attempt to tune the latter to receive the best or most reasonable image quality can be counterproductive. More research is needed to understand how the best initial guess should be formed in the very strong nonlinearity regime. However, we will see that combining *all the improvements* discussed in this section in one algorithm makes DCTMC convergence fast and reliable even at this high value of the contrast.

B. Using reciprocity of sources and detectors

For a vast majority of physical equations of interest, all measurements are reciprocal. In particular, this is true for the scalar wave equation considered here. This means that interchanging the source and detector yields the same measurement. For this reason, making physical measurements with interchanged sources and detectors does not provide any additional information about the target. Mathematically, reciprocity of measurements is equivalent to symmetry of the T matrix.

In spite of what has been said above, there exists a subtle point in DCTMC that makes explicitly accounting for the reciprocity useful. Namely, the experimental T matrix, T_{expt} is generally nonsymmetric. Therefore, it does not account for




FIG. 8. Convergence data for the small target with $\chi_0 = 0.0175$ comparing the original DCTMC algorithm and the one using reciprocity of sources and detectors.

the reciprocity of measurements. Consequently, including the data points with interchanged sources and detectors in the data matrix Φ will change T_{expt} and make it closer to being symmetric and closer to correspondence to a diagonal V . This account of reciprocity does not require any additional physical measurements. All that is needed is to include in Φ the data points that correspond to each source-detector pair being interchanged. In the traditional approaches to ISPs, such additional data points can be viewed as completely redundant and useless. However, in the context of DCTMC, inclusion of these data points makes T_{expt} more informative. In a sense, by including these redundant data points, we apply the additional condition that T_{expt} should be as symmetric as possible in the real-space representation.

We now investigate the doubling of the data set that is obtained by interchanging sources and detectors. In fact, the consequences of this operation can be quite dramatic. To illustrate the point, we show in Fig. 8 the convergence plots for the small target at the contrast level $\chi_0 = 0.0175$. The initial guess $T_1 = T_{\text{expt}}$ was used as the starting point. Two different data sets were used in the reconstructions. In the first case, sources are on one side of the sample and detectors are on the other side. In the second case, we supplement this data set by the “reciprocal” data points obtained by interchanging the source and the detector in each source-detector pair. The size of the data set is doubled by this procedure. We emphasize again that the addition of the “reciprocal” data points is not truly redundant in DCTMC. Indeed, the fraction of known (nonzero) entries in the experimental T matrix \tilde{T}_{expt} (in the singular function representation) is 4% and 17% in the two cases, respectively. That is, by using reciprocity, the dimension of the nonzero minor in \tilde{T}_{expt} that is used in overwriting the iteratively updated T matrix is roughly doubled. In other

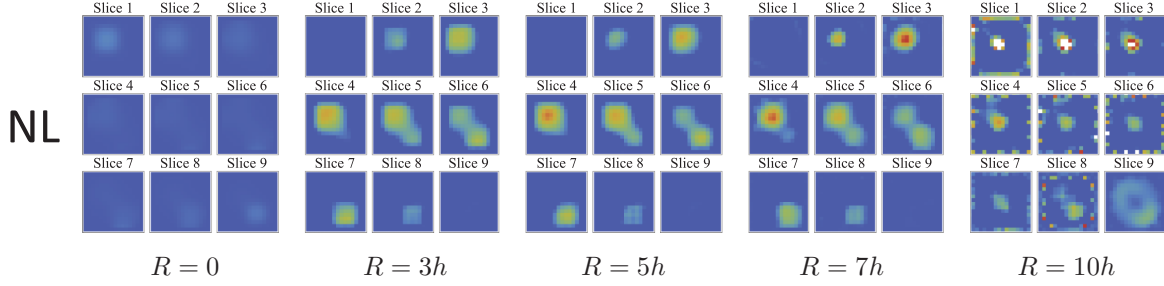


FIG. 9. DCTMC reconstruction of the small target after 150 iterations with varying degrees of R , where the row summing is done over voxels separated by no more than R .

words, the experimental T matrix is more informative in the second case.

It can be seen that the convergence speed is much higher with the use of reciprocity. Moreover, the converged results are much more precise with both errors η_χ and η_Φ approaching the limit determined by the numerical precision of the computer. This is a dramatic improvement for a simple change that does not require any additional information or processing power. The case represented in the figure is for $\chi_0 = 0.0175$, but all other levels of contrast that we have considered exhibit a similar behavior.

C. Using weighted summation to the diagonal for the operator $\mathcal{D}[\dots]$

Next, we abandon the use of the Computational Shortcut 2 (defined in Ref. [1]) and use instead the force-diagonalization operator $\mathcal{D}[\dots]$, viz.,

$$(\mathcal{D}[V])_{ij} \equiv \delta_{ij} \sum_k V_{ik} h(\ell_{ki}). \quad (32)$$

Correspondingly, we will use the iteration scheme ‘‘Main iteration without the use of Computational Shortcut 2’’ [1]. This algorithm is more in line with the nonlocal framework of DCTMC. Although the computational time per one iteration will be increased by approximately the factor of 2, we will see that the improvements obtained by this approach are the most promising in terms of significantly improving the convergence rate of DCTMC. It is worth emphasizing that the results in this section are obtained without the improvements related to starting from a better initial guesses and from doubling the data set by interchanging sources and detectors. We will also not apply the sparsity checks. Therefore, the results shown below can be attributed to DCTMC in its most basic form.

As an example, we use a simple step function for $\rho(\ell)$. That is, $\rho(\ell) = 1$ if $\ell \leq R$ and $\rho(\ell) = 0$ otherwise. For the contrast $\chi_0 = 0.0175$, 150 DCTMC iterations were performed without any sparsity checks. The reconstructions obtained after these 150 iterations for varying values of R are shown in Fig. 9. The immediate conclusion is that the case $R = 5$ is much better than the original case $R = 0$. It is also true that $R = 10h$ is too large, coupling voxels that are too far apart, which creates a characteristic aliasing.

The convergence data (for η_χ only and up to the iteration index $i = 150$) are shown in Fig. 10. It can be seen that a dramatic improvement of the convergence speed is achieved by using weighted summation to the diagonal with the appropriate

radius of influence. The case $R = 5$ is optimal, with $R = 2.5$ and 7.5 being close in performance. Furthermore, it is interesting that for the case $R = 10$, the first couple of iterations are quite effective, before something goes wrong and the error curve turns sharply upwards. This example is not suppressing as the nonlocal interaction of the far away voxels causes an unwanted behavior. Note that the case $R = 0$ corresponds to the ‘‘standard’’ algorithm and the Computational Shortcut 2 was used for this case. But, despite the fact that the $R = 0$ case completed in about half the time, looking at the error plot in Fig. 10 shows that there is no confusion on the preference between 75 iterations for $R = 5$ and 150 iterations of $R = 0$.

D. Putting it all together

We now combine all three improvements described above in one algorithm. We will demonstrate that for all values of the contrast used for the small sample, we can achieve acceptable either perfect or at least acceptable results in 75 iterations, compared to the 900 iterations previously performed. The algorithm used in this section proceeds as follows:

1. Run the linear reconstruction.
2. Using the result of step 1 as the initial guess, run five iterations without sparsity checks.
3. Then, every five iterations check whether some susceptibilities χ_n satisfy $|\chi_n| < \chi_{\max}/500$, where $\chi_{\max} = \max_n |\chi_n|$.
4. If a given voxel satisfies the above condition three checks in a row, the corresponding χ_n is set to zero, and the computational domain is reduced.




FIG. 10. Convergence data η_χ for the reconstruction shown in Fig. 9.

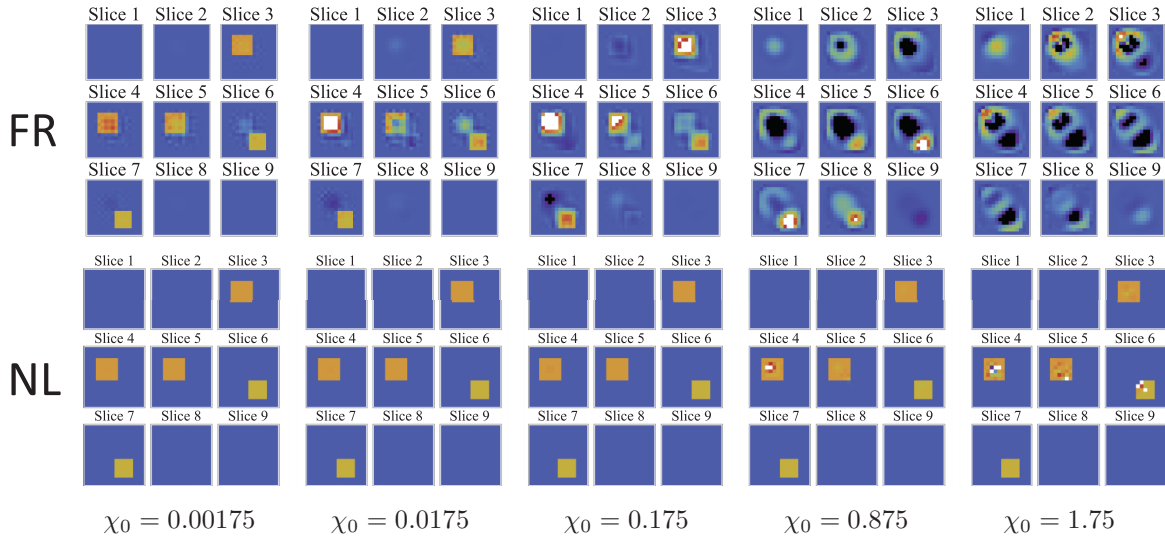


FIG. 11. Reconstructions of the small target by the algorithm combining all three improvements discussed in this section after 75 iterations for the five levels of contrast previously considered.

5. The process is repeated with the following modifications. After 20 iterations, the relative threshold for determining a noninteracting voxel is reduced to the factor of 100. After 40 iterations, the relative threshold is further reduced to the factor of 60.

Throughout the algorithm, reciprocity of sources and detectors was used, as well as weighted summation to the diagonal for the operator $\mathcal{D}[\cdot \cdot \cdot]$. For the weight function, we used the inverse distanced, namely, $\rho(\ell_{ik}) = \delta_{ik} + (1 - \delta_{ik})\ell_{ik}^{-1}$. Computational Shortcut 2 was not used. The reconstruction after 75 iterations is shown in Fig. 11 for all levels of contrast used previously. All reconstructions are nearly perfect at least for the largest value of the contrast, $\chi_0 = 1.75$. In the latter case, the reconstruction is not perfect but acceptable. The imprecision is again caused by three voxels that were incorrectly assigned as noninteracting. This happened even though we have increased the minimum relative threshold for determining noninteracting voxels from 40 to 60.

The convergence data for the algorithm employed in this subsection are shown in Fig. 12. It can be seen that the most dramatic improvement was obtained for the three lowest levels of contrast, as starting from the linearized inversion is of great assistance in these cases. In the intermediate case $\chi_0 = 0.875$, the reconstruction after 75 iterations is not as precise as the one produced by the original method after 900 iterations. This is clearly so because one voxel was erroneously assigned as noninteracting by the improved method. One incorrectly assigned noninteracting voxel does not degrade the reconstructed image seriously but bounds from below the error η_χ that can be reached. Clearly, more research on the application of sparsity checks in the strong nonlinearity regime is needed, especially when linearized inversion is used as the initial guess.

VI. COMPARISON WITH GAUSS-NEWTON METHOD

For comparison purposes, we have also implemented the regularized Gauss-Newton method. We used the standard

implementation of Levenberg-Marquardt with multiplicative damping. This method was too slow to be applied to the large target but we could apply it to the small target, even though one iteration of the Gauss-Newton method took much longer than one iteration of the Gauss-DCTMC. More importantly, the Gauss-Newton method did not converge to a reasonable result for the small target even at the lowest level of the contrast considered above. It is worth stressing that numerous modifications of the regularization and damping techniques were tried without producing any noticeably better results. Therefore, the Gauss-Newton reconstruction shown below is representative.

A typical Gauss-Newton reconstruction of the small target with $\alpha_0 = 0.00175$ is shown in Fig. 13. The corresponding error plots are shown in Fig. 14. It is obvious from the data of Fig. 14 that the Gauss-Newton iterations are converging




FIG. 12. Convergence data of η_χ for the algorithm combining all improvements discussed in this section. The horizontal lines show the respective error attained after 900 iterations of DCTMC without any of the improvements. Different line types correspond to different contrast χ_0 as follows: $\chi_0 = 0.00175$ (1), $\chi_0 = 0.0175$ (2), $\chi_0 = 0.175$ (3), $\chi_0 = 0.875$ (4), and $\chi_0 = 1.75$ (5).




FIG. 13. A typical Gauss-Newton (with Levenburg-Marquardt damping technique) reconstruction of the small target with $\alpha_0 = 0.00175$.

to a false solution (a local minimum of the cost function). Specifically, the error of the reconstruction η_χ is increasing with each iteration while the error of the equation η_ϕ is decreasing and reaching some very small values that are consistent with the limit imposed by the finite numerical precision of the computer.

To confirm that the Gauss-Newton method was programmed correctly, we have reduced the size of the target even further to a point where the method works well, at least, for some levels of contrast. The shape of this even smaller (tiny) target is shown in Fig. 1 and reconstructions are shown in Fig. 15. For the tiny target, the Gauss-Newton reconstructions work well except at the high level of contrast ($\chi_0 = 3$), when the method fails to localize the inhomogeneity correctly.




FIG. 14. Convergence data for the Gauss-Newton (with Levenberg-Marquardt damping technique) reconstruction shown in Fig. 13. Note that for $i > 5$, the error $\eta_\phi(i)$ continues to decrease, as is guaranteed by the algorithm, although this decrease is not visible due to the logarithmic vertical scale used in the figure.




FIG. 15. Gauss-Newton (with Levenburg-Marquardt damping technique) reconstructions of the tiny target for different levels of contrast as labeled.

Thus, we have produced numerical evidence that the DCTMC algorithm is able to solve nonlinear ISPs where other mainstream nonlinear inversion techniques fail. The results reported here further suggest that the size of the target is a critical factor. Specifically, the Gauss-Newton method works reasonably well for the tiny target but fails dramatically for the small target. A possible mathematical interpretation of this empirical observation is that the cost function used by the Gauss-Newton method develops a spurious local minimum for the small target. This minimum is connected by some path in the multidimensional space of solutions along which the cost function increases to the linearized solution. It follows therefore that the linearized solution is not a true minimum of the cost function. This statement may seem counterintuitive since we know for sure that the linearized solution minimizes the linearized cost function. However, the subtle but important point is that we do not really know whether the linearized solution minimizes the nonlinear cost function. It is also a realistic possibility that the linearized solution is in fact a minimum of the nonlinear cost function but this minimum is so shallow that, numerically, the iterations always overshoot and eventually descend towards the false solution. Clearly, more research in this direction is needed to fully understand the advantages and disadvantages of DCTMC compared to the standard methods.

VII. DISCUSSION

We have provided an initial numerical investigation of the data-compatible T -matrix completion (DCTMC) method for solving nonlinear inverse scattering problems (ISPs). DCTMC was implemented for the scalar wave equation of the form (2) which is relevant, in particular, to ultrasound imaging and seismic tomography. In many of the cases we have considered, DCTMC provided quantitatively accurate reconstructions, while any linearized inversion (including those based on first Born, first Rytov, or mean-field approximations) or the standard optimization methods (Gauss-Newton) have failed. This provides us with a proof-of-principle demonstration of the method's utility. Indeed, DCTMC does not utilize a cost function and does not suffer from the local minima and the associated false solution. Perhaps, this is the most important conceptual difference that distinguishes DCTMC from other nonlinear inversion methods. The method is especially useful for solving ISPs with large data sets. Even though we did

not use supercomputers or any massively parallel computational platforms, we have succeeded in solving a strongly nonlinear inverse problem with more than 2×10^6 data points. Still, many questions regarding DCTMC improvement, optimization, and applicability remain to be investigated either theoretically or numerically. Some of these topics for future research are listed below.

First, the obvious task is to investigate DCTMC with adjustable sparsity checks. These checks are very useful for improving the speed and convergence of DCTMC. However, incorrect assignment of noninteracting voxels has resulted in the method breakup in the strong nonlinearity regime. We emphasize that all cases of DCTMC not providing a useful reconstruction were due to these incorrect noninteracting voxel assignments. If the sparsity checks are suppressed, the method can be slowly converging but we have not encountered any divergences or convergence to false solutions.

Second, we did not investigate in much detail various approaches to regularization. It is obvious that, as the target size increases, the ISP becomes more ill posed. We can attempt a combination of Tikhonov regularization of the type that was discussed in [1], physical constraints and, potentially, other methods. However, complicated interaction between different types of regularization require a careful and systematic investigation.

Third, there are questions of computational efficiency and inclusion of larger targets. A combination of methods described in Sec. V is expected to increase the rate of

convergence or reduce the number of necessary iterations. Improving the computational bottlenecks is another important goal for future research. We note that some approximate approaches to speeding up this operations have been proposed by Jakobsen and Ursin [6] and these approaches can be utilized by DCTMC.

Fourth, weight functions $\rho(\ell)$ used in the definition of the force-diagonalization operator $\mathcal{D}[\dots]$ require a separate investigation

We finally note that the major challenge in understanding and optimizing the DCTMC is the complicated interplay of the various tunable parameters and algorithms that are inherent to the method. Nevertheless, the authors are optimistic that further progress can be made by a combination of analytical and numerical investigations.

ACKNOWLEDGMENTS

This work has been carried out thanks to the support of the A*MIDEX project (No. ANR-11-IDEX-0001-02) funded by the “Investissements d’Avenir” French Government program, managed by the French National Research Agency (ANR) and was also supported by the U. S. National Science Foundation under Grant No. DMS 1115616 and by the U. S. National Institutes of health under Grant No. P41 RR002305. The authors are grateful to J. C. Schotland, A. Yodh, and A. Sentenac for very useful discussions.

-
- [1] H. W. Levinson and V. A. Markel, *Phys. Rev. E* **94**, 043317 (2016).
 - [2] M. M. Bronstein, A. M. Bronstein, M. Zibulevski, and H. Azhari, *IEEE Trans. Med. Imag.* **21**, 1395 (2002).
 - [3] G. R. DeVore and B. Polanco, *J. Ultrasound* **24**, 1685 (2005).
 - [4] Q. Liu and Y. J. Gu, *Tectonophysics* **566-567**, 31 (2012).
 - [5] M. Jakobsen, *Stud. Geophys. Geod.* **56**, 1 (2012).
 - [6] M. Jakobsen and B. Ursin, *J. Geophys. Eng.* **12**, 400 (2015).
 - [7] P. S. Carney, R. A. Frazin, S. I. Bozhevolya, V. S. Volkov, A. Boltasheva, and J. C. Schotland, *Phys. Rev. Lett.* **92**, 163903 (2004).
 - [8] K. Belkebir, P. C. Chaumet, and A. Sentenac, *J. Opt. Soc. Am. A* **22**, 1889 (2005).
 - [9] K. Belkebir, P. C. Chaumet, and A. Sentenac, *J. Opt. Soc. Am. A* **23**, 586 (2006).
 - [10] G. Bao and P. Li, *Opt. Lett.* **32**, 1465 (2007).
 - [11] E. Mudry, P. C. Chaumet, K. Belkebir, and A. Sentenac, *Inverse Probl.* **28**, 065007 (2012).
 - [12] D. A. Boas, D. H. Brooks, E. L. Miller, C. A. DiMarzio, M. Kilmer, R. J. Gaudette, and Q. Zhang, *IEEE Signal Proc. Mag.* **18**, 57 (2001).
 - [13] S. R. Arridge and J. C. Schotland, *Inverse Probl.* **25**, 123010 (2009).
 - [14] E. M. Purcell and C. R. Pennypacker, *Astrophys. J.* **186**, 705 (1973).
 - [15] B. Draine and P. Flatau, *J. Opt. Soc. Am. A* **11**, 1491 (1994).
 - [16] V. A. Markel and J. C. Schotland, *Inverse Probl.* **23**, 1445 (2007).
 - [17] One can envisage a situation in which the first equality in (9) is not an approximation but holds exactly. However, this can not be known *a priori* when solving an ISP. On the other hand, the second equality in (9) is never exact.
 - [18] D. A. Smunov, P. C. Chaumet, and M. A. Yurkin, *J. Quant. Spectrosc. Radiat. Transfer* **156**, 67 (2015).
 - [19] B. T. Draine, *Astrophys. J.* **333**, 848 (1988).
 - [20] V. A. Markel, *J. Mod. Opt.* **39**, 853 (1992).
 - [21] A. Lakhtakia, *Optik* **91**, 134 (1992).
 - [22] B. T. Draine and J. Goodman, *Astrophys. J.* **405**, 685 (1993).
 - [23] M. A. Yurkin, *Handbook of Molecular Plasmonics* (Pan Stanford, Singapore, 2013), pp. 83–135.
 - [24] H. Y. Ban, D. R. Busch, S. Pathak, F. A. Moscatelli, M. Machida, J. C. Schotland, V. A. Markel, and A. G. Yodh, *J. Biomed. Opt.* **18**, 026016 (2013).
 - [25] V. A. Markel, J. A. O’Sullivan, and J. C. Schotland, *J. Opt. Soc. Am. A* **20**, 903 (2003).