

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

Howard W. Levinson

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

Vadim A. Markel

Departments of Radiology and Bioengineering and Graduate Group in Applied Mathematics and Computational Science, University of Pennsylvania, Philadelphia, Pennsylvania 19104

Abstract. 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 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. It is shown that DCTMC is a viable method for solving strongly nonlinear inverse problems with large data sets.

Submitted to: *Inverse Problems*

1. Introduction

This paper is Part II of the series on solving nonlinear inverse problems (ISPs) by data-compatible T-matrix completion (DCTMC). Part I of the series [1] contains theory and here we report initial numerical simulations for the three-dimensional diffraction tomography with scalar waves.

We note that DCTMC can be applied in a very similar manner to the problems of inverse electromagnetic scattering [2–6] and diffuse optical tomography [7, 8]. 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. Since, at this initial stage of research, our primary focus is on the effects of nonlinearity, no noise has been added to the data. We have used the streamlined iteration cycle described in Sec.5.3 of [1]. Physical constraints and checks for sparsity were applied between Steps 1 and 2 of this algorithm.

The remainder of this paper is organized as follows. In Sec. 2, we describe the mathematical 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. 3. Numerical results are adduced in Sec. 4. Finally, Sec. 5 contains a brief discussion of the obtained results.

2. 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 [9, 10]. A related method for the scalar diffusion equation was described by us in [11].

We start by re-writing (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

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

where

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

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,

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

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 ‡:

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

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

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

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

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

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

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

‡ 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.

We then obtain

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

where

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

The imaginary part of the right-hand side of (14) is the first non-vanishing 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 [12, 13]. 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 [14, 15]. 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.

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 equation (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 re-write (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 (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

§ Here we briefly summarize the statement of the nonlinear ISP for the discretized equation (17). A more general and detailed description is given in [1]. See Fig. 1 of that reference and the related discussion for the definition of the measurement surfaces. Also, we use the same notations as in [1]. In particular, we use Dirac notations for the elements of finite-dimensional vector spaces.

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.

3. Details of the numerical algorithms used

3.1. Choice of the fundamental 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 non-vanishing 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 or many 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 iteratively by DCTMC. In the process, we apply a physical constraint and take account of the target sparsity (although we do not make any *a priori* assumptions about the latter) as is discussed in detail below. Finally, when the DCTMC process has converged or the maximum number of iterations has been reached, we convert the reconstructed set of α_n 's to χ_n 's by using (25) and display the latter quantities in all figures.

3.2. 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 physical constraints usually stem from some *a priori* knowledge that is based on general physical principles. For example, if the medium is passive (that is, non-amplifying), then $\text{Im}\chi_n \geq 0$. If it is also known that the medium is transparent (non-absorbing), 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's 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 non-absorbing 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.

3.3. 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 non-interacting voxels are located. 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 3 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 non-interacting and are excluded from the computational domain. When this happens, we repeat the initial setup of Sec.5.3 of [1], 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 non-interacting voxel is reduced to the factor of 60, and after 600 iterations the relative threshold is reduced to the factor of 40.

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 majority of computational experiments we have performed, such incorrect assignment did not occur. However, in the very strong nonlinearity regime it can happen. We note however that such occurrences can be easily "diagnosed" by monitoring the error of the matrix equation (24) as is illustrated below. If an incorrect assignment does occur, one can potentially alleviate the problem by adjusting the constants used by the algorithm.

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 [8]. 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 zero with the sole purpose of being able to use the mathematically-simple function given in (6).

3.4. Model targets and arrangement of sources and detectors

We have used two kinds of targets in the simulations, which we refer to as "small" and "large". 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 small target is discretized on a $16 \times 16 \times 9$ grid ($N_v = 2,304$) and the large target is discretized on a $30 \times 30 \times 15$ grid ($N_v = 13,500$). The shapes of both 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 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

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

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 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'll 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. For the small target, we have used the following two arrangements:

Near-field zone source-detector arrangement for the small target. In this case, 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. Since $kh/2 = 0.1$, we say that, for this particular arrangement, the source/detector planes are located in the near-field zone 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/detector meshes does not improve this result any further.

Intermediate-field zone source-detector arrangement for the small target. In the second arrangement, the meshes of sources and detectors are 40×40 and displaced from the sample surface by $5h$ (from each of the two 16×16 faces of the sample). In this case, there are twelve rows and columns of sources or detectors extending past the sample surface in each direction. Since $5(kh) = 1$, we say that, in this case, the sources and detectors are placed in the intermediate-field zone of the sample. The size of the data set for this arrangement is as follows: $N_s = N_d = 1,600$ and the total number of data points is $N_s N_d = 2,560,000$. Note that data sets of this size are usually considered as problematic for solving nonlinear ISPs numerically.

By performing many numerical experiments with the small target, we have found that in the weak and intermediate nonlinearity regimes there is a slight (but not a dramatic) loss of information when we move the source/detector grids from the near-field zone to the intermediate-field zone of the sample. However, in the very strong nonlinearity regime the nonlinear inversion can break down for the intermediate-zone arrangement (in this case, due to incorrect assignment of noninteracting voxels) while, for the near-field zone arrangement, the nonlinear reconstruction is still stable. As was mentioned above, this problem can be overcome by adjusting the prescription for assigning the noninteracting voxels. However, we have used only near-field arrangement for the large target at this stage of research.

Near-field zone source-detector arrangement for 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/detector meshes are displaced from each of the two 30×30 faces of the large target by $h/2$, as is also the case in the near-field zone arrangement for the small target. The data set size is defined by the following numbers: $N_s = N_d = 1,444$ and $N_s N_d = 2,085,136$.

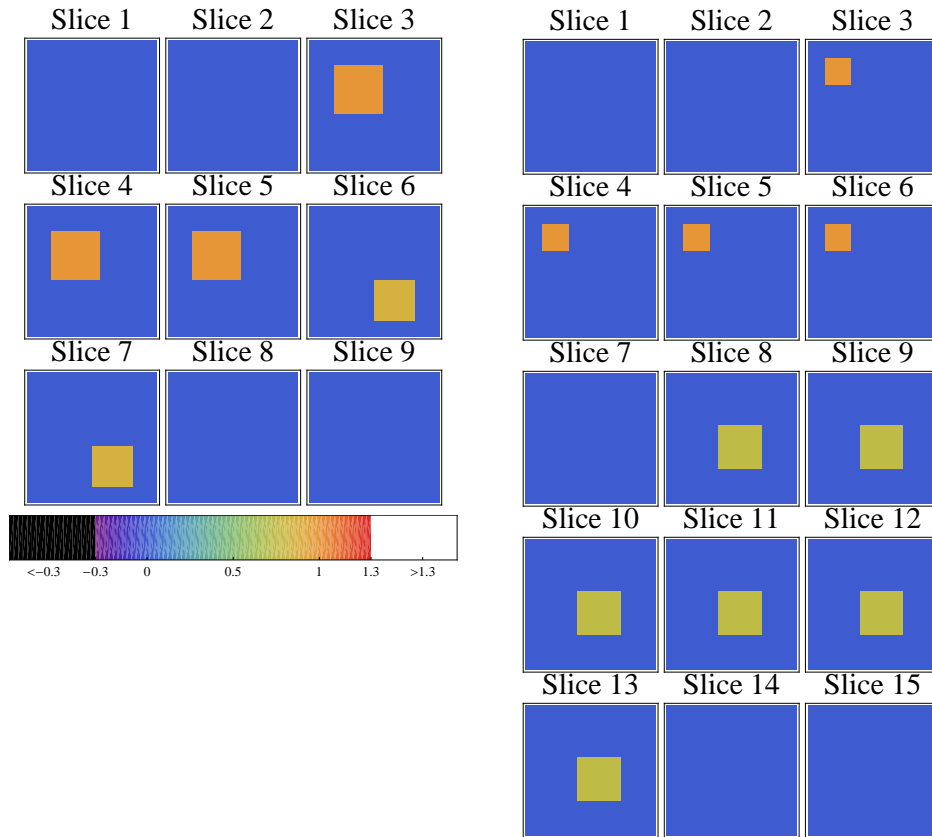


Figure 1. The shapes of the small (left) and the large (right) targets. The small target consists of $16 \times 16 \times 9$ voxels and is represented by 9 slices, each slice being of the dimension 16×16 . These slices are shown in the figure consecutively. Similarly, the large target consists of 15 slices of the size 30×30 each. Even though the shape function satisfies the condition $0 \leq \Theta \leq 1$, the color scale of this figure 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. Note that the color scale used in this and all other figures of this paper has two cutoffs. Values that are smaller than -0.3 are shown by the uniform black color and values that are larger than 1.3 are shown as white.

3.5. 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 and Eq. 6 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 112Gb of RAM (in single precision). One can follow the approach of [16] 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 inverting Eq. 34 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. Note that the derivation of Eq. 34 from Eq. 6 (both numbers refer to equations of [1]) is given in Appendix B of that reference.

All linearized reconstructions shown below were obtained by regularizing and solving Eq. 34 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, the computational formula (Eq. 36 of Ref. [1]) 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 the first Born, first Rytov and the mean-field approximations as described in Appendix A of Ref. [1]. The choice of the linearization method only affects the way in which the data is 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.

3.6. Error measures

To quantify the convergence of the method, we use normalized root mean square errors η_χ and η_Φ , which are defined by the following equations:

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

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

Here η_χ quantifies the error of the reconstruction while η_Φ quantifies the error of the equation (24). Note that η_χ is computed after Step 1 of the streamlined iteration cycle (Sec. 5.3 of [1]). The error η_Φ is computed after Step 5 of this algorithm. After Step 6, the T-matrix is fully "data-compatible" and the error in question is zero (up to the numerical precision of a 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*.

4. Numerical results

Reconstructed images of the small target after 900 DCTMC iterations are shown in Fig. 2 for the near-field zone arrangement of sources and detectors and 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 $\chi_0 = 0.1$. 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. For example, one can change the initial threshold condition for assigning the noninteracting voxels to be 1/200 of the maximum rather than 1/100. Note that it is possible to learn whether some noninteracting voxels have been assigned incorrectly by looking at the convergence curve for the error η_Φ (see Fig. 4 below). Doing so does not require any *a priori* knowledge of the target. A detailed discussion of the convergence curves is given below.

Reconstructions of the small target for the intermediate-field zone source/detector arrangement are shown in Fig. 3. In this case, some information contained in the evanescent waves is lost and the reconstructions look slightly worse even in the case of weak nonlinearity. This is true for both linearized reconstructions and DCTMC. For the strongest contrast used, $\chi_0 = 1.75$, DCTMC reconstruction is obviously not useful. This

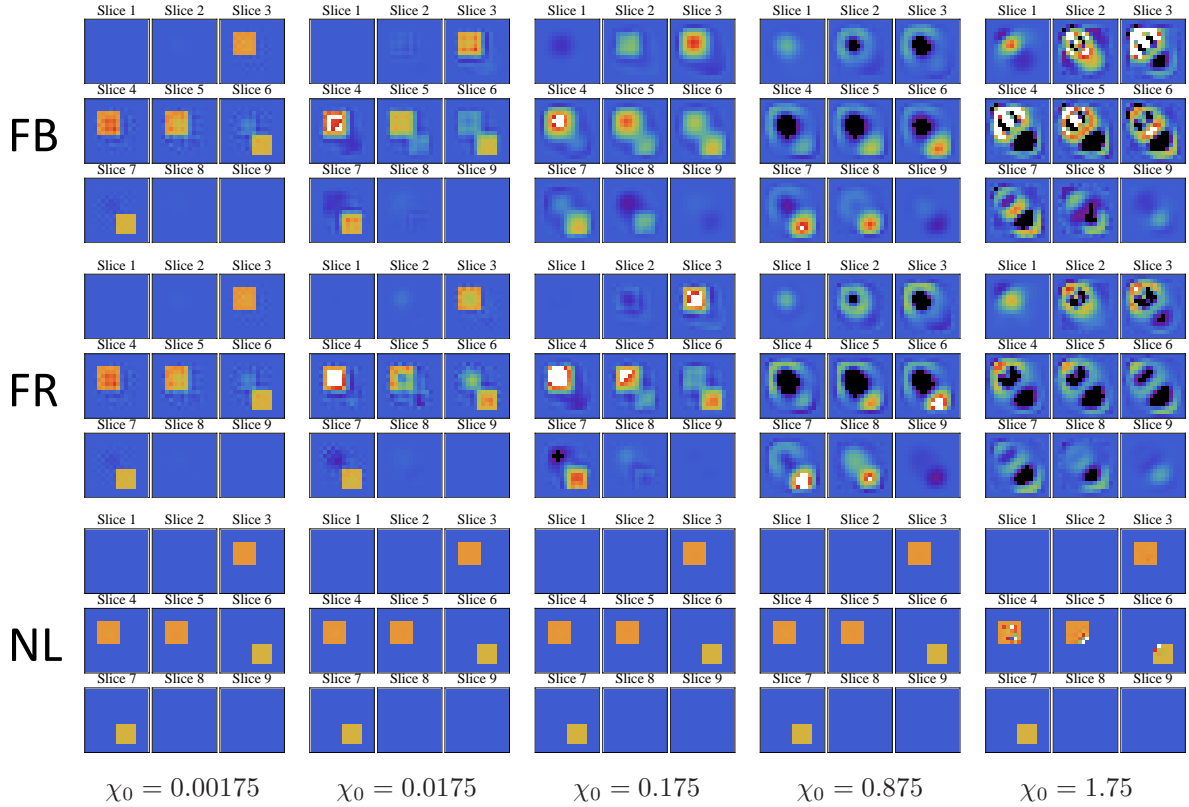


Figure 2. Linear (top and middle rows, marked FB and FR) and nonlinear (bottom row, marked NL) reconstructions of the small target for different levels of contrast χ_0 . The source/detector planes are in the near-field zone of the sample. The quantity shown by the color scale in each plot is χ_n/χ_0 , 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.

reason for the breakdown is that the algorithm has assigned many non-interacting voxels incorrectly. We emphasize again that the algorithm for assigning noninteracting voxels can be adjusted, but in this work we have intentionally applied the same prescription to all cases considered.

Errors η_χ and η_ϕ for the small target and for both the near-field and the intermediate-field zone arrangements of sources and detectors are shown in Fig. 4 as functions of the iteration number, i . We note that the curves $\eta_\chi(i)$ all look very similar (for a given source/detector arrangement) and are almost independent of χ_0 , except for $\chi_0 = 1.75$. The latter case will be discussed separately. We now consider the curves $\eta_\chi(i)$ for $\chi_0 < 1.75$. It can be seen that these curves describe three distinct convergence regimes. The classification is obvious for the near-field zone source/detector arrangement but can also be discerned in the intermediate-field zone arrangement.

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

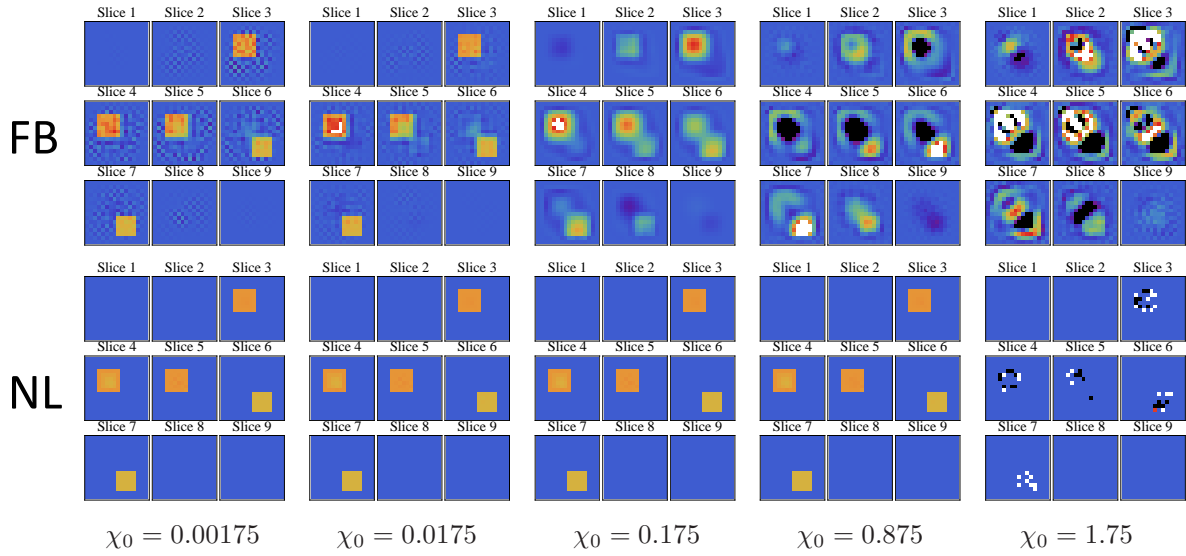


Figure 3. Same as in Fig. 2 but the source/detector planes are located in the intermediate-field zone of the sample.

that the initial iterations simply solve the linearized inverse problem by the iterative algorithm described in Sec. 7 of [1] (Richardson first-order iteration). Naturally, convergence of this iterative process is expected to be slow because we did not use Tikhonov regularization as in modified Step 5, Sec.7 of [1]. However, in principle, this initial slow convergence regime can be bypassed and 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. 3.5, this can be done 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 . Note that we have already implemented this approach and the preliminary results indicate that it works as expected; however, a detailed investigation of this and other improvements of the method will be reported separately.

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 with the actual exponent depending on the source/detector arrangement. This exponential convergence regime sets in faster and is more pronounced in the case of near-field zone source/detector arrangement. We note that, in this convergence regime, all noninteracting voxels have been determined correctly and the algorithm improves its estimate of the amplitudes of the interacting voxels. Apparently, for the near-field zone source/detector arrangement,

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. Consider first the near-field zone source/detector arrangement, in which case only three noninteracting voxels are assigned incorrectly. It can be seen that, 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 . In the case of intermediate-field zone source/detector arrangement, too many noninteracting voxels have been assigned incorrectly early on so that the method is completely broken; the error η_χ increases and fluctuates around some very large values.

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 (in each source/detector arrangement). 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 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 .

Now it was mentioned 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 non-monotonous dependence of η_Φ on χ_0 is a manifestation of the rather complicated nonlinearity of the inverse problem, and is also 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 non-monotonous dependence was observed for the large target as well (see below).

The curves $\eta_\phi(i)$ are rather interesting in the strong nonlinearity case $\chi_0 = 1.75$. For the near-field zone source/detector arrangement, the error χ_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 $i \approx 600$. In Figs. 2 and 3, we compare the reconstructions after exactly the same number of iterations, $i = 900$. However, below we provide a comparison of reconstructions obtained at different iteration numbers, one of which is deemed to be optimal from the analysis of the $\eta_\phi(i)$, for the large target (in a similar convergence regime).

For the intermediate-field zone source/detector arrangement, when many noninteracting voxels are incorrectly determined at the early stages of the iterations, the curve $\eta_\phi(i)$ appears to be chaotic and does not decay significantly. It can be clearly distinguished from the other curves. This behavior of $\eta_\phi(i)$ gives one an unambiguous indication that the obtained DCTMC reconstruction is not useful.

We now turn to the large target. In this case, we have also used 900 iterations for each nonlinearity regime but considered only the near-field zone arrangement of sources and detectors. In Fig. 5 we show the reconstructions for five different values of the contrast, from $\chi_0 = 0.002$ to $\chi_0 = 2$. 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 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




Figure 4. Convergence data for the small target. Errors η_χ (a,b) and η_ϕ (c,d) are plotted vs the iteration number i for the near-field zone (NFZ) source-detector arrangement (a,c) and the intermediate-field zone (IFZ) source-detector arrangement (c,d). The different curves 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).

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, $1/40$ 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 is not a trivial consequence of image "roughening". The result may appear counter-intuitive. 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 and sparsity checks 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

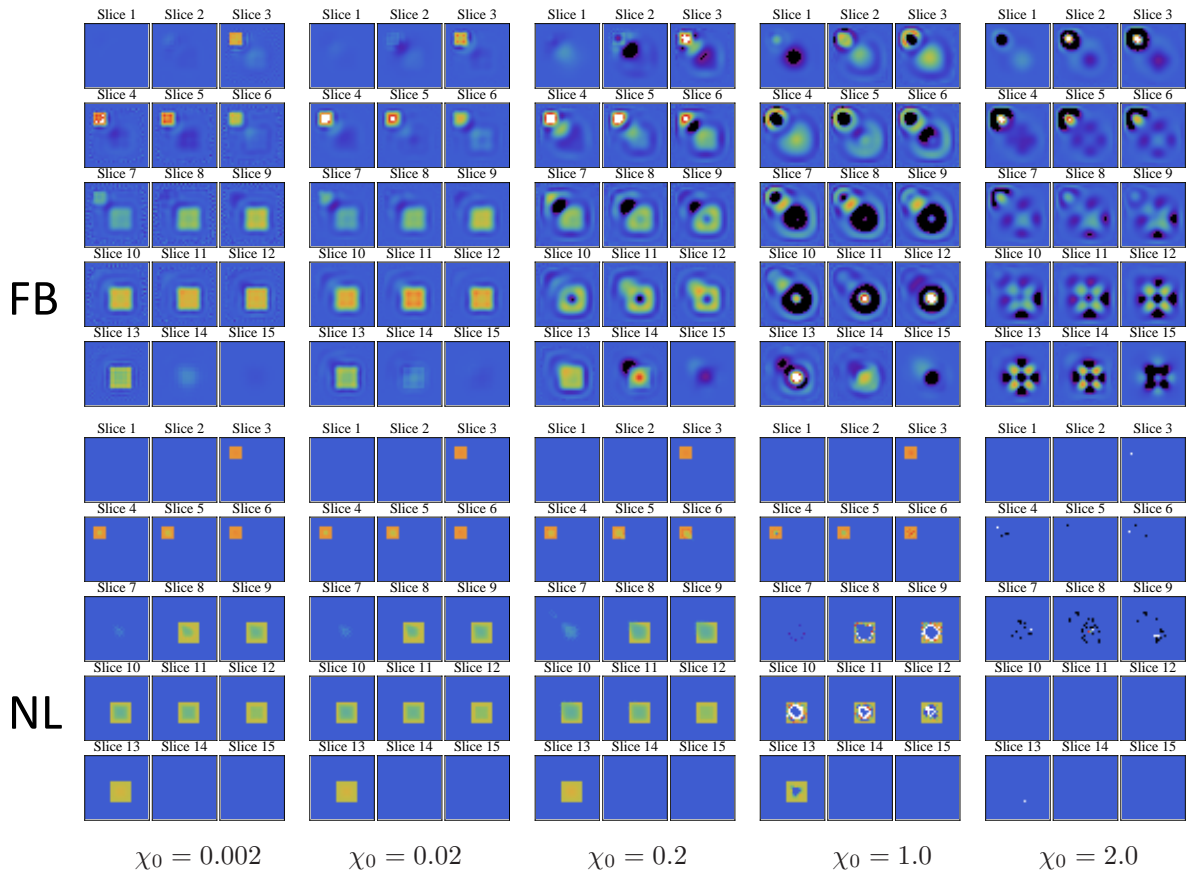


Figure 5. Same as in Fig. 2, but for the large target, near-field zone source/detector arrangement (see Sec. 3.4), and a somewhat different set of contrasts χ_0 . Utilization of first Rytov approximation for linearized reconstruction does not provide any improvements over first Born approximation, and the corresponding results are not shown in this figure.

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 [17]. 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. 6. 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 $\chi_0 = 2$ 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




Figure 6. Convergence data for the large target. The different curves 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).

minimum at $i = 590$. We have observed a similar non-monotonous convergence for the smaller target as well. Now we compare directly the reconstructions of the large target with $\chi_0 = 1$ obtained at $i = 590$ (the "optimal" iteration index) and $i = 900$ in Fig. 7. As expected, the result is visibly better at $i = 590$. The difference is not dramatic. However, this result clearly indicates that monitoring the error $\eta_\phi(i)$ is a useful approach to deciding when the iterations should stop.

5. Discussion

We have provided an initial numerical investigation of data-compatible T -matrix completion (DCTMC), which is an iterative numerical method for solving nonlinear inverse problems with large data sets. The method was implemented for the scalar wave equation of the form (2). We have used physical constraints and sparsity checks for regularization and improving the convergence rate of DCTMC. In many of the cases we have considered, DCTMC provided quantitatively accurate reconstructions whilst any linearized inversion (including those based on first Born, first Rytov or mean-field approximations) have failed. This provides us with a proof-of-principle demonstration of the method's utility. However, when the target contrast was increased beyond a certain point, DCTMC reconstructions were broken as well. The reason for this breakdown is incorrect and too early assignment of noninteracting voxels. This is not a flaw of

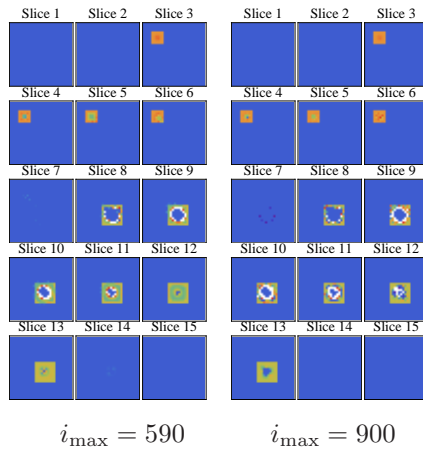


Figure 7. Comparison of the reconstructions of the large target with $\chi_0 = 1$ (Case 4 in Fig. 6) for different numbers of iterations i_{\max} , as labeled. Reconstruction with $i_{\max} = 490$ is marginally better than the reconstruction with $i_{\max} = 900$.

the DCTMC method itself but an artifact of the *ad hoc* prescription that was used to identify the noninteracting voxels. In principle, we can remove the sparsity checks from the algorithm altogether. This will definitely fix the problem associated with the incorrect assignment of noninteracting voxels, but in this case DCTMC convergence in the severe nonlocality regime can be slow. Therefore, appropriate sparsity checks are needed to balance the rate of convergence and the possibility of incorrect restriction of the computational domain. A reasonable approach to obtaining this balance appears to involve monitoring the error η_ϕ (31b) and adjusting the parameters in the prescription for identifying noninteracting voxels.

The main conclusion of this paper is that DCTMC is a viable method for solving strongly nonlinear inverse problems with extremely 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 $2 \cdot 10^6$ 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.

- It makes sense to use the linearized reconstruction as the initial guess for DCTMC. We have already implemented this approach in a number of cases and it appears to work as expected (by removing the slow convergence region completely), but a more systematic investigation is needed.
- It is obvious from the general reciprocity principle that interchanging the source and detector does not change the measurement. However, a somewhat counter-intuitive fact is that inclusion of an equation in which the source and detector are interchanged in (24) is not a redundant operation, unless we explicitly enforce symmetry of the T -matrix (in the real-space representation). But as was discussed in [1], this is not done in DCTMC. The "experimental" T -matrix T_{exp} is not

symmetric in the real-space representation and is not mathematically expected to be such. Under the circumstances, inclusion of the data points trivially obtained by source-detector interchange (this does not require any additional physical measurements) does provide an additional useful information about the T-matrix. The role of this information on the iterative image reconstruction needs to be investigated. We did obtain a preliminary indication that including this more "complete" data set results in a better starting point for the iterations (a more useful "experimental" T-matrix) and can cut the computational time significantly.

- So far, we did not investigate in much detail various approaches to regularization. It is obvious that, as the target size increases and as the source and detector surfaces are moved to the far-field zone of the sample, the inverse problem becomes more ill-posed. We can attempt a combination of Tikhonov regularization of the type that was discussed in Sec.7 of [1], physical constraints and potentially other methods. However, complicated interaction between different types of regularization can exist and a much more careful study of this question is required.
- There are questions of computational efficiency and inclusion of larger targets. A combination of methods described in the points above is expected to increase the rate of convergence or reduce the number of necessary iterations. Improving the computational bottleneck (currently, the matrix inversion in the operation of transforming V to T) is another important goal for future research.
- Finally, it would be interesting to apply the method to other physical models and correspondingly to other free-space Green's functions G_0 . Of particular interest is the Green's function of the diffusion equation, which can be obtained from (6) by the analytical continuation $k \rightarrow ik$.

Acknowledgments

This research was supported in part by the NSF under Grant DMS 1115616 and by the NIH under grant P41 RR002305. The authors are grateful to J.C.Schotland and A.Yodh for very useful discussions.

References

- [1] Levinson H W and Markel V A 2014 *Inverse Problems* Submitted
- [2] Carney P S, Frazin R A, Bozhevolnyi S I, Volkov V S, Boltasseva A and Schotland J C 2004 *Phys. Rev. Lett.* **92** 163903
- [3] Belkebir K, Chaumet P C and Sentenac A 2005 *J. Opt. Soc. Am. A* **22** 1889–1897
- [4] Belkebir K, Chaumet P C and Sentenac A 2006 *J. Opt. Soc. Am. A* **23** 586–595
- [5] Bao G and Li P 2007 *Opt. Lett.* **32** 1465–1467
- [6] Mudry E, Chaumet P C, Belkebir K and Sentenac A 2012 *Inverse Problems* **28** 065007
- [7] Boas D A, Brooks D H, Miller E L, DiMarzio C A, Kilmer M, Gaudette R J and Zhang Q 2001 *IEEE Signal Proc. Mag.* **18** 57–75
- [8] Arridge S R and Schotland J C 2009 *Inverse Problems* **25** 123010
- [9] Purcell E M and Pennypacker C R 1973 *Astrophys. J.* **186** 705–714

- [10] Draine B and Flatau P 1994 *J. Opt. Soc. Am. A* **11** 1491–1499
- [11] Markel V A 2007 *J. Quant. Spectrosc. Radiat. Transfer* **103** 428–429
- [12] Draine B T 1988 *Astrophys. J.* **333** 848–872
- [13] Markel V A 1992 *J. Mod. Opt.* **39** 853–861
- [14] Lakhtakia A 1992 *Optik* **91** 134–137
- [15] Draine B T and Goodman J 1993 *Astrophys. J.* **405** 685–697
- [16] Ban H Y, Busch D R, Pathak S, Moscatelli F A, Machida M, C S J, Markel V A and Yodh A G 2013 *J. Biomed. Opt.* **18** 026016
- [17] Markel V A, O’Sullivan J A and Schotland J C 2003 *J. Opt. Soc. Am. A* **20** 903–912