



HAL
open science

Solution of the nonlinear inverse scattering problem by T-matrix completion. I. Theory

Howard Levinson, Vadim Markel

► **To cite this version:**

Howard Levinson, Vadim Markel. Solution of the nonlinear inverse scattering problem by T-matrix completion. I. Theory. *Physical Review E*, 2016, 94 (4), pp.043317. 10.1103/PhysRevE.94.043317. hal-01282076v1

HAL Id: hal-01282076

<https://hal.science/hal-01282076v1>

Submitted on 3 Mar 2016 (v1), last revised 26 Oct 2016 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Solution of the nonlinear inverse scattering problem by T-matrix completion. I. Theory

Howard W. Levinson

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

E-mail: levh@sas.upenn.edu

Vadim A. Markel ‡

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

E-mail: vmarkel@fresnel.fr

Abstract. We propose a conceptually new method for solving nonlinear inverse scattering problems (ISPs). The method is inspired by the theory of nonlocality of physical interactions and utilizes the relevant mathematical formalism. We formulate the ISP as a problem whose goal is to determine an unknown and, generally, nonlocal interaction potential V from external scattering data. We then utilize the one-to-one correspondence between V and the T-matrix of the problem, T . An iterative algorithm is proposed in which we seek T that is (i) compatible with the data and (ii) corresponds to an interaction potential V that is as diagonally-dominated as possible. We refer to this algorithm as to the data-compatible T-matrix completion (DCTMC). This paper is Part I in a two-part series and contains theory only. Numerical examples are given in Part II.

Submitted to: *Inverse Problems*

1. Introduction

Inverse scattering problems (ISPs) are known to be ill-posed and nonlinear [1, 2]. Nevertheless, with the use of suitable regularization, many nonlinear ISPs can be solved with reasonable precision. Examples include optical diffusion tomography [3, 4], diffraction tomography [5, 6], electrical impedance tomography (Calderon problem) [7–9] and near-field [10–12] and far-field [13, 14] tomographic electromagnetic imaging, and in seismic tomography [15, 16]. Solving nonlinear ISPs is a difficult computational task, especially in three dimensions. This is even more true for problems involving large data

‡ On leave from the Department of Radiology, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA

sets that are available with the use of modern experimental techniques. Developing efficient algorithms for solving nonlinear ISPs remains an important challenge.

Nonlinear ISPs are amply reviewed in the literature [1, 2, 17–19]. The mainstream approach to solving these problems numerically is Newton’s method and its variants such as Levenberg-Marquardt method, iteratively regularized Gauss-Newton method, Newton-Kantorovich method and steepest descent (Landweber iteration). These methods (except for Newton-Kantorovich) are succinctly explained in [20]. Newton-Kantorovich iterations are closely related [21] to the method of inverse Born series [9, 22, 23]. A different class of non-deterministic inversion approaches that make use of some form of prior knowledge about the medium is based on Bayesian inference [24]. The common feature of all these approaches (except for the inverse Born series) is that a certain cost function is minimized and updated iteratively and that this cost function depends on all available measurements (data points). In the case of inverse Born series, the solution is obtained as an analytically-computable functional of the data.

The method proposed in this paper is conceptually different from the methods reviewed above and is based on a digression into a seemingly unrelated field of physics, namely, into the theory of nonlocality. This theory accounts for the fact that certain physical processes occurring at the point \mathbf{r} in space can be influenced by the field in some finite vicinity of that point. For example, in local electrodynamics, Ohm’s law is written as $\mathbf{J}(\mathbf{r}) = \sigma(\mathbf{r})\mathbf{E}(\mathbf{r})$. In a nonlocal theory, this linear relation is generalized by writing $\mathbf{J}(\mathbf{r}) = \int V(\mathbf{r}, \mathbf{r}')\mathbf{E}(\mathbf{r}')d^3r'$. Of course, we expect on physical grounds that $V(\mathbf{r}, \mathbf{r}') \rightarrow 0$ when $|\mathbf{r} - \mathbf{r}'| > \ell$, where ℓ is the characteristic scale of nonlocality (the radius of influence), which is usually much smaller than the overall size of the sample. If the electric field $\mathbf{E}(\mathbf{r})$ does not change noticeably on the scale of ℓ , we can define the local conductivity as $\sigma(\mathbf{r}) = \int V(\mathbf{r}, \mathbf{r}')d^3r'$ and use Ohm’s law in its local form. This is all well known in physics. However, implications of nonlocality for nonlinear ISPs have not been considered so far.

Let us assume that we want to find $\sigma(\mathbf{r})$ from the measurements of voltage drop for a direct current injected into the sample by two point-like electrodes attached to its surface at various points (Calderon problem). It turns out that it is much easier to find a nonlocal kernel $V(\mathbf{r}, \mathbf{r}')$ that is consistent with the measurements. Of course, $V(\mathbf{r}, \mathbf{r}')$ can not be determined uniquely from a typical data set because the number of unknown parameters (degrees of freedom) in $V(\mathbf{r}, \mathbf{r}')$ is usually much larger than the number of measurements. However, as explained above, we also expect that $V(\mathbf{r}, \mathbf{r}')$ should be approximately diagonal. We then proceed as follows:

- (1) First, we define a class of kernels $V(\mathbf{r}, \mathbf{r}')$ that are compatible with the data. This is the only instance when the data are used, and it turns out that the size of the data set is not a limiting factor for this step.
- (2) Then we iteratively reduce the off-diagonal norm of $V(\mathbf{r}, \mathbf{r}')$ while making sure that $V(\mathbf{r}, \mathbf{r}')$ remains within the class of “data-compatible” kernels.
- (3) Once the ratio of the off-diagonal and diagonal norms of $V(\mathbf{r}, \mathbf{r}')$ is deemed

sufficiently small, we compute $\sigma(\mathbf{r}) = \int V(\mathbf{r}, \mathbf{r}') d^3r'$. This gives an approximate numerical solution to the nonlinear ISP.

The above algorithm can be generalized to other ISPs. We refer to it as to the data-compatible T-matrix completion (DCTMC). The role of the T-matrix in solving nonlinear inverse problems has been recognized in several prior publications [15, 16, 25, 26]. The advantages of using the T-matrix in solving the inverse problem is two-fold. First, it has been recognized, in particular, in the references quoted herein, that computation of the T-matrix yields the forward solution for any number of arbitrarily placed sources and detectors. In other words, the T-matrix is source and detector-independent. This can be very advantageous when the number of sources and detectors is very large. For example, finite-difference and finite elements forward solves must be run anew for each source used. The T-matrix approach is free of this requirement. The price that is being paid for this simplification is that computing the T-matrix involves inversion of dense matrices while finite difference and finite element methods employ only sparse matrices. In this work, we utilize the fact that there exists a one-to-one correspondence between the T-matrix T and the interaction operator V . However, using this correspondence in practice requires that we treat V as a general matrix (not necessarily diagonal). We further show that the T-matrix based approach allows one to obtain a useful data reduction that is applicable to both linear and nonlinear regimes. Therefore, our results are useful for solving linear inverse problems as well.

We underscore that physical interactions are never truly local and some small degree of nonlocality exists in all physical systems. However, the radius of influence ℓ is typically so small (e.g., equal to the atomic scale) that the nonlocality can be safely ignored for most practical purposes. In our approach, we relax this condition and allow V to be off-diagonal on much larger scales. Of course, we will seek to find V that is as diagonal as possible. However, we do not expect to eliminate all off-diagonal terms that are separated by more than one atomic scale, not to mention that such fine discretization of the medium is practically impossible. Thus, the non-locality of V that is utilized in DCTMC is not an intrinsic physical property but rather a physically-inspired mathematical trick that is used to facilitate the solution of nonlinear ISPs. In other words, we simplify the solution process by relaxing the underlying physical model.

This paper is Part I of a two-part series, wherein we focus our attention on theory. Numerical examples for the nonlinear inverse diffraction problem are given in Part II [27]. The remainder of this paper is organized as follows. In Sec. 2 we state the general algebraic formulation of the nonlinear ISP that is applicable to many different physical scenarios. In Sec. 3 we introduce the data-compatible T-matrix, which is a central idea of the proposed method. In Sec. 4 we define the basic iterative algorithm of DCTMC. In Sec. 5, we introduce "computational shortcuts", which combine analytically several steps of the DCTMC algorithm into a single step with reduced computational complexity. Variations and enhancements of DCTMC are discussed in Sec. 6. DCTMC algorithm in the linear regime is discussed in Sec. 7. Sec. 8 contains a brief discussion. Auxiliary information is given in several appendices. Summary of linearizing approximations

(first Born, first Rytov and mean-field) is given in Appendix A. Appendix B contains a derivation that establish the correspondence between DCTMC and the conventional methods in the linear regime. Finally, definitions and properties of several functionals used in this paper are summarized in Appendix C.

2. General formulation of the ISP

Consider a linear operator \mathcal{L} and the equation

$$\mathcal{L}u(\mathbf{r}) = q(\mathbf{r}) , \quad (1)$$

where $u(\mathbf{r})$ is a physical field and $q(\mathbf{r})$ is the source term. Note that (1) does not contain time but can depend parametrically on frequency. It can be said that we work in the frequency domain. Moreover, we consider only a single fixed frequency. Using different working frequencies as additional degrees of freedom for solving an ISP can be very useful but is outside of the scope of this paper.

Let $\mathcal{L} = \mathcal{L}_0 - V$, where \mathcal{L}_0 is known and V is the unknown interaction operator that we seek to reconstruct. As discussed above, we assume at the outset that V is an integral operator with the kernel $V(\mathbf{r}, \mathbf{r}')$ but, eventually, the computed image will be obtained as a function of \mathbf{r} only. We also assume that $V(\mathbf{r}, \mathbf{r}') \neq 0$ only if $\mathbf{r}, \mathbf{r}' \in \Omega$, where Ω is a spatial region occupied by the sample. Our goal is to recover V from the measurements of u performed outside of the sample, assuming that it is illuminated by various external sources. We can not perform measurements or insert sources inside the sample, which would have greatly simplified the ISP solution if it was physically possible.

The inverse of \mathcal{L} is the *complete* Green's function of the system, denoted by $G = \mathcal{L}^{-1}$. The formal solution to (1) is then $u = Gq$. We know that G exists as long as the forward problem has a solution. This is usually the case if V is physically admissible. Likewise, the inverse of \mathcal{L}_0 is the *unperturbed* Green's function, denoted by $G_0 = \mathcal{L}_0^{-1}$. The field $u_{\text{inc}} = G_0q$ is the incident field, in other words, it is the field that would have existed everywhere in space in the case $V = 0$. Nonzero V gives rise to a scattered field u_{scatt} , and the total field is a sum of the incident and scattered components, $u = u_{\text{inc}} + u_{\text{scatt}}$. A straightforward algebraic manipulation yields the following result:

$$u_{\text{scatt}} = (G - G_0)q = G_0(I - VG_0)^{-1}VG_0q , \quad (2)$$

where I is the identity operator.

A single *data point* $\Phi(\mathbf{r}_d, \mathbf{r}_s)$ is obtained by illuminating the medium with a localized source of unit strength, $q(\mathbf{r}) = \delta(\mathbf{r} - \mathbf{r}_s)$, and measuring the scattered field by a detector at the location \mathbf{r}_d §. By scanning \mathbf{r}_d and \mathbf{r}_s on the measurement surfaces Σ_d and Σ_s outside of the sample, we measure a function of two variables $\Phi(\mathbf{r}_d, \mathbf{r}_d)$, which is coupled

§ If the source is not of unit strength, measurements should be divided by the source amplitude.

to $V(\mathbf{r}, \mathbf{r}')$ by the equation

$$G_0(I - VG_0)^{-1}VG_0 = \Phi . \quad (3)$$

All product and inversion operations in (3) should be understood in the operator sense. The ISP can now be formulated as follows: Given a measured function $\Phi(\mathbf{r}_d, \mathbf{r}_s)$, where $\mathbf{r}_d \in \Sigma_d$ and $\mathbf{r}_s \in \Sigma_s$, find an ‘‘approximately diagonal’’ kernel $V(\mathbf{r}, \mathbf{r}')$, where $\mathbf{r}, \mathbf{r}' \in \Omega$. We do not need to define ‘‘approximate diagonality’’ precisely at this point, but in the case of matrices that are inevitably used in all computations, this requirement implies a sufficiently small ratio of the off-diagonal and diagonal norms.

It is important to note that G_0 in (3) is the same operator in all instances where it appears, but for the purpose of computing the operator products and inverses, its kernel $G_0(\mathbf{r}, \mathbf{r}')$ is differently restricted. This is illustrated graphically in Fig. 1. Thus, for the first term G_0 in the left-hand side of (3), $\mathbf{r} = \mathbf{r}_d \in \Sigma_d$ and $\mathbf{r}' = \mathbf{r}'_1 \in \Omega$. For the second term (inside the brackets) $\mathbf{r} = \mathbf{r}_1 \in \Omega$ and $\mathbf{r}' = \mathbf{r}_2 \in \Omega$. For the last term, $\mathbf{r} = \mathbf{r}'_2 \in \Omega$ and $\mathbf{r}' = \mathbf{r}_s \in \Sigma_s$. We emphasize that the imaging geometry shown in Fig. 1 is representative but not very general. In particular, the measurement surfaces Σ_d and Σ_s can be larger or smaller than the face of the cube, or curve, or even be regions of space of finite volume rather than surfaces \parallel . The sample volume Ω does not have to be cubic and, in an extreme case, it can be a two-dimensional surface. All this has no bearing on the method of this paper. The only requirement that we impose, which is physical rather than mathematical, is that Σ_d and Σ_s do not overlap with Ω . However, Σ_d can overlap with Σ_s .

Further, in all practical implementations, the data are sampled rather than measured continuously and the medium is voxelized. An example of such discretization is given in [27]. At this point we proceed under the assumption that (3) can be suitably discretized and converted to a matrix equation. In this case, it is logical to use different notations for the matrices that are obtained by different restriction and sampling of the kernel $G_0(\mathbf{r}, \mathbf{r}')$. Indeed, the matrices obtained in this manner are different and can even be of different size. We will denote the matrices obtained by sampling the first, the second, and the last terms G_0 in (3) by A , Γ and B , respectively. These notations are also illustrated in Fig. 1. Then the discretized version of (3) takes the following form:

$$A(I - V\Gamma)^{-1}VB = \Phi . \quad (4)$$

In (4), A , B and Γ are known theoretically, Φ is measured, and we seek to find the unknown V .

Eq. (4) is the main nonlinear equation that is discussed in this paper. It is, in fact, very general and encompasses many different problems of imaging and tomography. The underlying physical model is encoded in the operator G_0 and in the matrices A , B and Γ that are obtained by sampling this operator. The following three important remarks about this equation can be made:

\parallel There is no real distinction between the two cases if the data are sampled.




Figure 1. (color online) Illustration of the imaging geometry. The symbols A , B , Γ and V in blue frames denote the matrices obtained by restricting and sampling the kernels $G_0(\mathbf{r}, \mathbf{r}')$ and $V(\mathbf{r}, \mathbf{r}')$. The scattering diagram corresponds to the second order term $G_0 V G_0 V G_0$ in the formal power-series expansion of the left-hand side in (3). Note that in the local limit $V(\mathbf{r}, \mathbf{r}') = D(\mathbf{r})\delta(\mathbf{r} - \mathbf{r}')$, the two green arrows contract to two vertexes at $\mathbf{r}_1 = \mathbf{r}'_1$ and $\mathbf{r}_2 = \mathbf{r}'_2$.

Remark 1: Noninvertibility of A and B . If matrices A and B were invertible in the ordinary sense, the nonlinear ISP would be solvable exactly by three operations of matrix inversion. Unfortunately, A and B are almost never invertible. To construct A and B of sufficiently high rank, one needs to perform measurements inside the medium. As was noted above, this is usually impossible. The typical sizes of all matrices involved will be discussed below (see Fig. 2 and its discussion).

Remark 2: Linearization. One may seek a linearization of the ISP by approximating the left-hand side in (4) with various expressions that allow an analytical linearizing transformation. The three main approaches to achieve this end are first Born, first Rytov and mean-field approximations, and they are briefly summarized in Appendix A. In the mathematical formulation of the ISP, the three approximations differ only by the transformation that is applied to the data matrix Φ , while the general form of the linearized equation is in all cases

$$AVB = \Psi[\Phi], \quad (5)$$

where $\Psi[\Phi]$ is the appropriate transformation of the data matrix; in the simplest case of first Born approximation, $\Psi[\Phi] = \Phi$.

Remark 3: Matrix unrolling for the linearized problem. The linearization approaches described in Appendix A do not require or enforce by design the diagonality of V . However, in the conventional treatments of the problem, it is typical to assume that V is strictly diagonal and to operate with the vector $|v\rangle$ § composed of the diagonal elements of V . Accordingly, the matrix Ψ is unrolled into a vector $|\psi\rangle$ by the matrix operation known as `vec`, that is, by stacking the columns of Ψ into one column-vector. The resultant equation has the form

$$K|v\rangle = |\psi\rangle , \quad (6)$$

where K is a matrix obtained by multiplying the elements of A and B according to the rule $K_{(mn),j} = A_{mj}B_{jn}$ and (mn) is a composite index. The important point here is that the conventional methods often treat K in (6) as a matrix of the most general form. In contrast, DCTMC algorithm takes account of the special algebraic structure of K and, therefore, can be used advantageously even in the linear regime. This is discussed in more detail in Appendix B.

3. T-matrix and its representations; “experimental” T-matrix

The basic definition of the T-matrix (which is, actually, an operator) is through the relation between the complete and the unperturbed Green’s functions: $G = G_0 + G_0 T G_0$. By direct comparison with (2) we find that

$$T = (I - V G_0)^{-1} V . \quad (7)$$

We will not use different notations for the operator T and its discretized version, which is truly a matrix. Consequently, Eq. (4) can be rewritten as

$$ATB = \Phi , \quad (8)$$

where

$$T = \mathcal{T}[V] \equiv (I - V \Gamma)^{-1} V = V (I - \Gamma V)^{-1} . \quad (9)$$

Here we have defined the nonlinear functional $\mathcal{T}[\cdot]$, which contains Γ as a parameter.

We can view (9) as a matrix formulation of the forward problem. If V is known, we can use (9) to compute T , and once this is accomplished, we can predict the result of a measurement by any detector due to any source by matrix multiplication according to (8). Therefore, computation of T yields the most general solution to the forward

§ Here and below, we use Dirac notations in algebraic expressions involving elements in finite-dimensional vector spaces. Briefly, $|f\rangle$ is a column-vector with the entries $f_n = \langle n|f\rangle$, $\langle f|$ is the Hermitian conjugate of $|f\rangle$ (a row-vector with the elements f_n^*), $\langle f|g\rangle = \langle g|f\rangle^*$ is the scalar product of two vectors of the same length, $|f\rangle\langle g|$ is a matrix (not necessarily square) and $\langle f|A|g\rangle = \sum_{nm} f_n^* A_{nm} g_m$.




Figure 2. (color online) Block diagram of Eq. (8) with sizes of all matrices indicated. here N_d and N_s are the numbers of detectors and sources and N_v is the number of voxels.

problem. The forward solution is usually known to exist if V is physically admissible. In the iterative process of DCTMC, we can ensure physical admissibility of V every time before the transformation $\mathcal{T}[V]$ is used (see Sec. 6.3). In this case, the matrix inversion involved in computing $\mathcal{T}[V]$ is always well-defined.

We can also formally invert \mathcal{T} and write

$$V = \mathcal{T}^{-1}[T] \equiv (I + TT)^{-1}T = T(I + \Gamma T)^{-1} . \quad (10)$$

Much less is known about the existence of the inverse in (10). In other words, we do not know the conditions of physical admissibility of T apart from the general but not very useful symmetry property $T_{ij} = T_{ji}$. Certainly, $\mathcal{T}^{-1}[T]$ does not exist for all arguments T . In DCTMC, one of the possible approaches is to update V iteratively by using (10). In this case, existence of the inverse is required. While we do not possess a general proof, numerical simulations for the inverse diffraction problem have encountered no singularities in (10). More importantly, the problem of invertibility of \mathcal{T} does not arise at all if Computational Shortcut 2 is used (see Sec. 5.2 below).

A block diagram of Eq. (8) with all matrix sizes indicated is shown in Fig. 2. Here N_d and N_s are the numbers of detectors and sources used (not necessarily equal) and N_v is the number of volume voxels. For a practical estimate of these numbers, refer to Fig. 1. Let the measurement surfaces Σ_d and Σ_s be identical squares located on the opposite sides of a cubic sample. Let the detectors and sources be scanned on an $L \times L$ square grids and let the sample be discretized on a $L \times L \times L$ cubic grid with the same pitch. Then $N_d = N_s = L^2$, $N_v = L^3$. These estimates are typical but, admittedly, not very general. Still, in many practical cases we can expect that

$$N_d, N_s \ll N_v \ll N_d N_s . \quad (11)$$

The first inequality in (11) illustrates Remark 1 of Sec. 2 because the matrices A and B are in this case clearly not invertible. The second inequality is important if we wish to compare DCTMC to some of the traditional approaches. For example, the conventional formulation of the linearized ISP starts from equation (6). As is explained

in Remark 3, it is commonly assumed that K in (6) is a general matrix of the size $N_d N_s \times N_v$ ($L^4 \times L^3$). However, the sizes of A and B are $N_d \times N_v$ ($L^2 \times L^3$) and $N_v \times N_s$ ($L^3 \times L^2$), respectively. Computing numerically the pseudoinverse of K (if we do not account for its special algebraic structure as is described in Appendix B) is a much more computationally-intensive task than computing the pseudoinverses of A and B . Therefore, the relaxation of the strict requirement of diagonality of V allows one to work with two much smaller “weight matrices” A and B instead of one large “weight matrix” K .

We now turn to the central idea of DCTMC, namely, to the concept of data-compatibility of the T-matrix. To formulate the mathematical constraints that equation (8) places on T in a computationally-tractable form, consider the singular value decompositions of A and B :

$$A = \sum_{\mu=1}^{N_d} \sigma_{\mu}^A |f_{\mu}^A\rangle \langle g_{\mu}^A|, \quad B = \sum_{\mu=1}^{N_s} \sigma_{\mu}^B |f_{\mu}^B\rangle \langle g_{\mu}^B|. \quad (12)$$

Here σ_{μ}^A , $|f_{\mu}^A\rangle$ and $|g_{\mu}^A\rangle$ are the singular values and right and left singular vectors of A , and similarly for B . Note that $|f_{\mu}^A\rangle$ and $|g_{\mu}^B\rangle$ are vectors of length N_d and N_s , respectively, while $|g_{\mu}^A\rangle$ and $|f_{\mu}^B\rangle$ are both of length N_v , and we have assumed in (12) that $N_d, N_s \leq N_v$. Using the orthogonality of singular vectors, we obtain from (8) and (12)

$$\sigma_{\mu}^A \sigma_{\nu}^B \tilde{T}_{\mu\nu} = \tilde{\Phi}_{\mu\nu}, \quad 1 \leq \mu \leq N_d, \quad 1 \leq \nu \leq N_s, \quad (13)$$

where

$$\tilde{T}_{\mu\nu} \equiv \langle g_{\mu}^A | T | f_{\nu}^B \rangle, \quad 1 \leq \mu, \nu \leq N_v; \quad (14a)$$

$$\tilde{\Phi}_{\mu\nu} \equiv \langle f_{\mu}^A | \Phi | g_{\nu}^B \rangle, \quad 1 \leq \mu \leq N_d, \quad 1 \leq \nu \leq N_s. \quad (14b)$$

By \tilde{T} we denote the T-matrix in *singular-vector representation* while T that was used previously is the T-matrix in *real-space representation*. The two representations are related to each other by the transformation

$$\tilde{T} = R_A^* T R_B \equiv \mathcal{R}[T], \quad T = R_A \tilde{T} R_B^* \equiv \mathcal{R}^{-1}[\tilde{T}], \quad (15)$$

where R_A is the unitary matrix whose columns are the singular vectors $|g_{\mu}^A\rangle$ while R_B is the unitary matrix whose columns are the singular vectors $|f_{\mu}^B\rangle$. Equation (15) defines the functional $\mathcal{R}[\cdot]$. We note that $\mathcal{R}[\cdot]$ is linear and invertible even though it is not equivalent to a conventional rotation because $R_A \neq R_B$. It is useful to keep in mind that $\tilde{\Phi} \neq \mathcal{R}[\Phi]$. As can be seen from the definition (14), $\tilde{\Phi}$ is related to Φ by a similar transformation but with different unitary matrices.

We now can write the solution to (13) as follows:

$$\tilde{T}_{\mu\nu} = \begin{cases} \frac{1}{\sigma_{\mu}^A \sigma_{\nu}^B} \tilde{\Phi}_{\mu\nu} & , \text{ if } \sigma_{\mu}^A \sigma_{\nu}^B > \epsilon^2; \\ \text{unknown} & , \text{ otherwise.} \end{cases} \quad (16)$$

Here ϵ is a small positive constant. If computations could be performed with infinite precision, we could have set $\epsilon = 0$. In practice, we should take ϵ to be small but at least larger than the smallest positive floating-point constant for which a particular implementation of numerical arithmetic adheres to the IEEE standard. We note that under the assumptions stated above, the condition $\sigma_\mu^A \sigma_\nu^B > \epsilon^2$ can be satisfied only for $1 \leq \mu \leq N_d$ and $1 \leq \nu \leq N_s$. Singular values σ_μ^A and σ_ν^B with indexes outside of these ranges are identically zero.

Eq. (16) summarizes our knowledge about the system that is contained in the data. There are few matrix elements of \tilde{T} that are known with certainty. These matrix elements can be computed by the first expression in (16). The other matrix elements can not be determined from equation (8). We can vary these unknown elements arbitrarily and the error of (8) will not noticeably change. The number of known elements of \tilde{T} can not exceed $N_d N_s$ but can, in principle, be smaller, e.g., if the rank of A is less than N_d or the rank of B is less than N_s , although this situation is not typical even for severely ill-posed ISPs. In any event, $N_d N_s$ is usually much smaller than the total number of the matrix elements of \tilde{T} , which is equal to N_v^2 . Using the previously introduced estimates, $N_d N_s / N_v^2 \sim 1/L^2$. Therefore, only a small fraction of the elements of \tilde{T} are known. In what follows, we assume that the singular values of A and B are arranged in the descending order and that the known elements of \tilde{T} can be collected into the upper-left rectangular block of the size $M_A \times M_B$ (see Fig. 3), where $M_A \leq N_d$ and $M_B \leq N_s$. We emphasize again that, in many practically-important cases, equalities will hold in the above expressions. However, it is possible to arrange the sources in such a way that the rank of B is less than N_s (and similarly for detectors and A), at least up to the numerical precision of the computer ¶. Moreover, the region of known matrix elements can be of a more general shape than a rectangle, as is shown in Fig. 4. It is not conceptually difficult to account for this fact. However, we will proceed under the assumption that the region is rectangular in order to shorten the discussion. Besides, in the numerical simulations of [27], this region was, in fact, rectangular.

Even though we can not gain any knowledge about the matrix elements of \tilde{T} outside of the shaded area shown in Fig. 3 by using equation (8) alone, we can make an initial guess for \tilde{T} , which we denote by \tilde{T}_{exp} (the "experimental" T-matrix). We define T_{exp} (in real-space representation) as the matrix that satisfies (8) in the minimum norm sense and has the smallest entry-wise norm $\|T\|_2$. This matrix is uniquely defined by the equation

$$T_{\text{exp}} = A^+ \Phi B^+ , \quad (17)$$

where "+" denotes Moore-Penrose pseudoinverse. If A and B are rank-deficient or invertible, (8) is satisfied by T_{exp} exactly so that $\|AT_{\text{exp}}B - \Phi\|_2 = 0$. The experimental T-matrix in singular-vector representation is obtained from (17) by the transformation (15). In fact, the experimental T-matrix is more easily characterizable in singular-vector

¶ This possibility is closely related to the existence of non-radiating sources. Indeed, to force the rank of B to be less than N_s , we need to arrange the sources in space and to assign them relative amplitudes and phases so that they produce almost no field in Ω .

$$\tilde{T} = \begin{array}{|c|c|} \hline \begin{array}{c} \tilde{\Phi}_{\mu\nu} \\ \sigma_\mu^A \sigma_\nu^B \end{array} & M_A \\ \hline M_B & \text{UNKNOWN} \\ \hline \text{UNKNOWN} & \text{UNKNOWN} \\ \hline \text{UNKNOWN} & \text{UNKNOWN} \\ \hline \end{array} \quad \tilde{T}_{\text{exp}} = \begin{array}{|c|c|} \hline \begin{array}{c} \tilde{\Phi}_{\mu\nu} \\ \sigma_\mu^A \sigma_\nu^B \end{array} & M_A \quad \mathbf{0} \\ \hline M_B & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{0} \\ \hline \end{array}$$

Figure 3. (color online) Left panel: Elements of \tilde{T} inside the shaded block can be computed from the data by using (16). Elements outside of the shaded block are not known and can not be in any way inferred from the data. Right panel: The initial guess for the T-matrix, T_{exp} . In this initial guess, we set the unknown elements of \tilde{T} (in singular-vector representation) to zero.

$\sigma_1^A \sigma_1^B$	$\sigma_1^A \sigma_2^B$	$\sigma_1^A \sigma_3^B$	$\sigma_1^A \sigma_4^B$	$\sigma_1^A \sigma_5^B$	$\sigma_1^A \sigma_6^B$	$\sigma_1^A \sigma_7^B$
$\sigma_2^A \sigma_1^B$	$\sigma_2^A \sigma_2^B$	$\sigma_2^A \sigma_3^B$	$\sigma_2^A \sigma_4^B$	$\sigma_2^A \sigma_5^B$	$\sigma_2^A \sigma_6^B$	$\sigma_2^A \sigma_7^B$
$\sigma_3^A \sigma_1^B$	$\sigma_3^A \sigma_2^B$	$\sigma_3^A \sigma_3^B$	$\sigma_3^A \sigma_4^B$	$\sigma_3^A \sigma_5^B$	$\sigma_3^A \sigma_6^B$	$\sigma_3^A \sigma_7^B$
$\sigma_4^A \sigma_1^B$	$\sigma_4^A \sigma_2^B$	$\sigma_4^A \sigma_3^B$	$\sigma_4^A \sigma_4^B$	$\sigma_4^A \sigma_5^B$	$\sigma_4^A \sigma_6^B$	$\sigma_4^A \sigma_7^B$
$\sigma_5^A \sigma_1^B$	$\sigma_5^A \sigma_2^B$	$\sigma_5^A \sigma_3^B$	$\sigma_5^A \sigma_4^B$	$\sigma_5^A \sigma_5^B$	$\sigma_5^A \sigma_6^B$	$\sigma_5^A \sigma_7^B$

Figure 4. (color online) Assuming that the singular values σ_μ^A and σ_μ^B are arranged in the descending order, this sketch shows an example of a more general shape (compared to Fig. 3) of the region in which the elements of \tilde{T} are known. The numbers above the thick red line satisfy $\sigma_\mu^A \sigma_\nu^B > \epsilon^2$. In the general case, the boundary line can only go from left to right and from bottom to top if followed from the left-most boundary of the matrix.

representation. Indeed, all the elements of \tilde{T}_{exp} in the unshaded area of the diagram in Fig. 3 (right panel) are equal to zero. This is expressed mathematically by writing

$$\left(\tilde{T}_{\text{exp}}\right)_{\mu\nu} = \begin{cases} \frac{1}{\sigma_\mu^A \sigma_\nu^B} \tilde{\Phi}_{\mu\nu} & , \text{ if } \sigma_\mu^A \sigma_\nu^B > \epsilon^2 ; \\ 0 & , \text{ otherwise } . \end{cases} \quad (18)$$

This expression is equivalent to (17).

We conclude this section with two important observations about the experimental T-matrix:

Remark 4: Lack of sparsity of T_{exp} . The matrix \tilde{T}_{exp} is sparse but the same is not true for T_{exp} .

Remark 5: Lack of symmetry of \tilde{T}_{exp} . It is known theoretically that the correct T-matrix is symmetric in real-space representation. However, this is not generally true for T_{exp} . Indeed, $T_{\text{exp}} = \mathcal{R}^{-1}[\tilde{T}_{\text{exp}}]$, and in \tilde{T}_{exp} a large fraction of the elements are replaced by zeros. The resultant T_{exp} is not likely to be symmetric.

4. Basic iteration cycle

In this section we describe a computational algorithm in which the matrices T and V are continuously updated so that T is kept data-compatible and V becomes increasingly diagonally-dominated. Our goal is to fill the unknown elements of \tilde{T} (the white areas in the left panel of Fig. 3) in such a way that the corresponding interaction matrix V , computed according to (10), is approximately diagonal. This is a general formulation of the problem of matrix completion, although the constraint that we apply to \tilde{T} is not the same as in the conventional statement of the problem.

Before proceeding, we need to introduce several additional operators. First, define the masking operators $\mathcal{M}[\cdot]$ and $\mathcal{N}[\cdot]$:

$$\left(\mathcal{M}[\tilde{T}]\right)_{\mu\nu} \equiv \begin{cases} 0, & \sigma_{\mu}^A \sigma_{\nu}^B > \epsilon^2 ; \\ \tilde{T}_{\mu\nu}, & \text{otherwise} . \end{cases} \quad \left(\mathcal{N}[\tilde{T}]\right)_{\mu\nu} \equiv \begin{cases} \tilde{T}_{\mu\nu}, & \sigma_{\mu}^A \sigma_{\nu}^B > \epsilon^2 ; \\ 0, & \text{otherwise} . \end{cases} \quad (19)$$

We note that $\mathcal{M}[\tilde{T}] + \mathcal{N}[\tilde{T}] = \tilde{T}$. Then the operator of enforcing data-compatibility of \tilde{T} (in singular-vector representation) $\mathcal{O}[\cdot]$ can be defined as follows:

$$\mathcal{O}[\tilde{T}] \equiv \mathcal{M}[\tilde{T}] + \tilde{T}_{\text{exp}} = \tilde{T} - \mathcal{N}[\tilde{T}] + \tilde{T}_{\text{exp}} . \quad (20)$$

It can be seen that the action of $\mathcal{O}[\tilde{T}]$ is to overwrite (hence the notation \mathcal{O}) the elements of \tilde{T} in the shaded area of Fig. 3 with the elements of \tilde{T}_{exp} and to leave all other elements unchanged. The operator $\mathcal{O}[\cdot]$ is defined for any $N_v \times N_v$ matrix but in the iterations discussed below we always apply this operator to the T-matrix in singular-vector representation.

Next, we will need to define a diagonal approximation to V , which we denote by D . The corresponding operator is $\mathcal{D}[\cdot]$, where

$$D = (\mathcal{D}[V])_{ij} \equiv \delta_{ij} V_{ii} . \quad (21)$$

Unlike the operator $\mathcal{O}[\cdot]$, $\mathcal{D}[\cdot]$ will always be applied in real-space representation.

We are now ready to describe the iterative process of DCTMC. The iteration steps will be defined in terms of the operators $\mathcal{R}[\cdot]$ and $\mathcal{T}[\cdot]$, $\mathcal{D}[\cdot]$ and $\mathcal{O}[\cdot]$. A summary of all operators used in this paper is given in Appendix C. We assume that the SVD decompositions of matrices A and B and the experimental T-matrix \tilde{T}_{exp} (18) have been precomputed. Consider the case when the iterations start from an initial guess for the T-matrix. We then set $k = 1$, $\tilde{T}_1 = \tilde{T}_{\text{exp}}$ and run the following iterations:

- 1: $T_k = \mathcal{R}^{-1}[\tilde{T}_k]$

This transforms the T-matrix from singular-vector to real-space representation.

Both \tilde{T}_k and T_k are data-compatible.




Figure 5. (color online) Basic flowchart of the DCTMC iteration process for the case when the iterations start with an initial guess for \tilde{T} . Other options are discussed in Sec. 6.1. Numbered iteration steps are defined in Sec. 4. Computational shortcuts are described in Sec. 5. The iteration step labeled "The Bottleneck" is the step in comparison to which the complexity of all other steps can be made insignificant. Eliminating the bottleneck by accounting for sparsity of V is discussed in Sec. 6.2. Matrix representations are abbreviated by SV (singular-vector) and RS (real-space).

2: $V_k = \mathcal{T}^{-1}[T_k]$

This gives k -th approximation to the interaction matrix V . V_k is data-compatible but not diagonal. Compute the off-diagonal and diagonal norms of V_k . If the ratio of the two is smaller than a predetermined threshold, exit; otherwise, continue to the next step.

3: $D_k = \mathcal{D}[V_k]$

Compute the diagonal approximation to V_k , denoted here by D_k . D_k is diagonal but not data-compatible.

4: $T'_k = \mathcal{T}[D_k]$

Compute the T-matrix that corresponds to the diagonal matrix D_k . Unlike T_k , T'_k is no longer data-compatible.

5: $\tilde{T}'_k = \mathcal{R}[T'_k]$

Transform T'_k to singular-vector representation. Here \tilde{T}'_k is still not data-compatible.

$$6: \tilde{T}_{k+1} = \mathcal{O}[\tilde{T}'_k]$$

Advance the iteration index by one and overwrite the elements of \tilde{T}'_k that are known from data with the corresponding elements of \tilde{T}_{exp} . This will restore data-compatibility of \tilde{T}_{k+1} . Then go to Step 1.

The computational complexity of Steps 1,2,4,5 is $O(N_v^3)$. However, the complexity can be dramatically reduced with the use of the computational shortcuts that are described in the next section, with the only exception of Step 4. Therefore, Step 4 is the true computational bottleneck of the method. It's complexity can be reduced by accounting for sparsity of V as is described in Sec. 6.2. However, if no *a priori* knowledge about sparsity of V is available, then the computational complexity of Step 4 is the limiting factor of DCTMC, at least to the best of our current understanding.

5. Computational shortcuts

5.1. Shortcut 1: Fast rotations

Consider iteration Steps 5,6,1 written sequentially:

$$\begin{array}{lll} 5 : & \tilde{T}'_k = \mathcal{R}[T'_k] & O(N_v^3) = O(L^9) , \\ 6 : & \tilde{T}_{k+1} = \mathcal{O}[\tilde{T}'_k] & \leq O(N_d N_s) = O(L^4) , \\ 1 : & T_{k+1} = \mathcal{R}^{-1}[\tilde{T}_{k+1}] & O(N_v^3) = O(L^9) . \end{array}$$

To the right, we have indicated the computational complexity of each step and used the previously introduced estimates for N_d , N_s and N_v in terms of the grid size L . The complexity of Step 6 is equal to, at most, $N_d N_s$. Therefore, the complexity of Steps 5 and 1 is dominating. Now, let us combine the steps by writing

$$T_{k+1} = \mathcal{R}^{-1}[\mathcal{O}[\mathcal{R}[T'_k]]] = \mathcal{R}^{-1}[\mathcal{R}[T'_k] - \mathcal{N}[\mathcal{R}[T'_k]] + \tilde{T}_{\text{exp}}] , \quad (22)$$

where in the second equality we have used the definition of $\mathcal{O}[\cdot]$ (20). We now use the linearity and invertibility of \mathcal{R} to rewrite (22) as

$$T_{k+1} = T'_k + T_{\text{exp}} - \mathcal{R}^{-1}[\mathcal{N}[\mathcal{R}[T'_k]]] . \quad (23)$$

We are therefore left with the task of numerically evaluating an expression of the type $\mathcal{R}^{-1}[\mathcal{N}[\mathcal{R}[T]]]$, where we have dropped all indexes for simplicity. But this can be accomplished in much less than $O(N_v^3)$ operations due to the sparsity of $\mathcal{N}[\cdot]$. Indeed, consider first the computation of $\mathcal{N}[\mathcal{R}[T]]$. This operation is illustrated in Fig. 6. It can be seen that $\mathcal{N}[\mathcal{R}[T]] = P_A^* T P_B$, where the matrix P_A is obtained from R_A by overwriting all columns of R_A , except for the first M_A columns, with zeros, and P_B is defined analogously. The complexity of computing $P_A^* T P_B$ is $O(\min(M_A, M_B) N_v^2)$, which is less than N_v^3 by the factor of at least $O(L)$. Quite analogously, we can show that

$$\mathcal{R}^{-1}[\mathcal{N}[\mathcal{R}[T'_k]]] = P_A (P_A^* T P_B) P_B^* . \quad (24)$$




Figure 6. (color online) Schematics of computing $\mathcal{N}[\mathcal{R}[T]]$. Matrices P_A and P_B are obtained from R_A and R_B by setting all columns to zero except for the first M_A and M_B columns, respectively.

It should be kept in mind that $P_A P_A^* \neq I$ and $P_B P_B^* \neq I$ and that premultiplying these matrices is not a computationally efficient approach. Doing so will result in an expression of the type $Q_A^* T Q_B$, where $Q_A = P_A P_A^*$ and $Q_B = P_B P_B^*$ are not sparse. Instead, one should evaluate the right-hand side of (24) using the operator precedence implied by the parentheses. We conclude that the Steps 5,6,1 of the iterative procedure described above can be combined in the following single computational step:

$$T_{k+1} = T'_k + T_{\text{exp}} - P_A (P_A^* T'_k P_B) P_B^* . \quad (25)$$

In this formula, T_{exp} , P_A and P_B are precomputed and stored in memory. We finally note that the sparsity of \mathcal{N} can be used efficiently even if the region of "known" elements of \tilde{T} is not rectangular. Although matrices P_A and P_B can not be easily defined in this case, the use of appropriate masks in computing expressions of the type $\mathcal{N}[R_A^* T R_B]$ will achieve a similar reduction of computational complexity.

5.2. Shortcut 2: Fast $T \rightarrow D$ operation.

Next consider Steps 2 and 3:

$$\begin{array}{lll}
 2 : & V_k = \mathcal{T}^{-1}[T_k] & O(N_v^3) = O(L^9) , \\
 3 : & D_k = \mathcal{D}[V_k] & O(N_v) = O(L^3) ,
 \end{array}$$

To the right, we have indicated the computational complexity of each step. The goal of Steps 2 and 3 is to find a diagonal matrix D that in some sense "corresponds" to the previously-computed T -matrix. More specifically, we compute V that corresponds to T

exactly but is not diagonal in Step 2 and then seek a diagonal matrix D that minimizes the entry-wise norm $\|V - D\|_2$ in Step 3. Obviously, the latter task is achieved by simply nullifying all off-diagonal elements of V . An alternative approach is to seek a diagonal matrix D that satisfies the equation

$$T = D + D\Gamma T \quad (26)$$

in the minimum L_2 -norm sense (of course, (26) can not be satisfied exactly by any diagonal matrix D). The above is a classical minimization problem, which has the following analytical solution:

$$D_{ij} = \delta_{ij} \frac{T_{ii} + [(\Gamma T)^* T]_{ii}}{1 + [(\Gamma T)^* + (\Gamma T) + (\Gamma T)^* (\Gamma T)]_{ii}} . \quad (27)$$

It may seem that evaluation of (27) still requires $O(N_v^3)$ operations because it contains the matrix-matrix product ΓT . However, this is not so. The matrix $A \equiv \Gamma T$ can be updated iteratively during Computational Shortcut 1 by using (25) multiplied from the left by Γ , viz,

$$A_{k+1} = A'_k + A_{\text{exp}} - (\Gamma P_A) (P_A^* T'_k P_B) P_B^* . \quad (28)$$

Here $A_{k+1} = \Gamma T_{k+1}$, $A'_k = \Gamma T'_k$ and $A_{\text{exp}} = \Gamma T_{\text{exp}}$. The matrix ΓP_A can be precomputed and has exactly the same sparsity structure as P_A itself, that is, all of its columns except for the first M_A columns are zero. Therefore, computing the last term in (28) is of the same complexity as Computational Shortcut 1, that is, $O(\min(M_A, M_B)N_v^2)$. There remains the question of how A'_k is computed and whether this computation requires an extra matrix-matrix multiplication. The answer is, it can be precomputed at Step 4 of the k -th iteration without any additional matrix-matrix multiplications. Indeed, let us utilize the second formula in the definition of $\mathcal{T}[\cdot]$ (9) and write Step 4 of k -th iteration as follows:

$$T'_k = D_k(I - \Gamma D_k)^{-1} . \quad (29)$$

We then multiply from the left both sides of (29) by Γ and obtain

$$A'_k = \Gamma D_k(I - \Gamma D_k)^{-1} = (I - \Gamma D_k)^{-1} - I . \quad (30)$$

Now we can compute T'_k and A'_k without any additional complexity by using the following sub-steps:

- 1: Compute the product $\Delta_k \equiv \Gamma D_k$, which is fast because D_k is diagonal.
- 2: Compute the inverse $S_k \equiv (I - \Delta_k)^{-1}$, which has the complexity of N_v^3 .
- 3: Compute $A'_k = S_k - I$ [as follows from (30)].
- 4: Compute $T'_k = D_k S_k$ [as follows from (29)], which is again fast because D_k is diagonal.

Thus, in all the computations outlined above only inversion of $I - \Delta_k$ has the computational complexity of N_v^3 . This is, therefore, the true computational bottleneck of the algorithm.

5.3. Streamlined iteration cycle

The computational shortcuts can be integrated into a single streamlined iteration algorithm. Doing so requires careful consideration of the flowchart shown in Fig. 5 and of the associated data dependencies. However, the resulting algorithm is relatively simple. For ease of programming we have broken this algorithm into elementary computational steps as described below.

Initial setup:

- a: Permanently store in memory the analytically-known matrix Γ .
- b: Compute the SVD decomposition (12) of A and B . This will yield a set of singular values $\sigma_\mu^A, \sigma_\mu^B$ (some of which are identically zero) and singular vectors $|f_\mu^A\rangle, |f_\mu^B\rangle, |g_\mu^A\rangle, |g_\mu^B\rangle$.
- c: Use the previous result to construct and permanently store in memory the dense matrices R_A and R_B , sparse matrices P_A and P_B and $Q_A = \Gamma P_A$. Note: no additional memory allocation for P_A and P_B is required.
- d: Compute $\tilde{\Phi}_{\mu\nu}$ according to (14) and \tilde{T}_{exp} according to (18). Discard the real-space data function, the singular values and singular vectors, and deallocate the associated memory.
- e: Compute and store permanently in memory $T_{\text{exp}} = R_A \tilde{T}_{\text{exp}} R_B^* = P_A \tilde{T}_{\text{exp}} P_B^*$ and $\Lambda_{\text{exp}} = \Gamma T_{\text{exp}}$.
- f: Initialize iterations by setting $T_1 = T_{\text{exp}}$ and $\Lambda_1 = \Lambda_{\text{exp}}$.

Main iteration: For $k = 1, 2, \dots$, perform the following computations:

- 1: $(D_k)_{ij} = \delta_{ij} \frac{(T_k)_{ii} + (\Lambda_k^* T_k)_{ii}}{1 + (\Lambda_k^* + \Lambda_k + \Lambda_k^* \Lambda_k)_{ii}}$;
- 2: $\Delta_k = \Gamma D_k$;
- 3: $S_k = (I - \Delta_k)^{-1}$;
- 4: $T'_k = D_k S_k$, $\Lambda'_k = S_k - I$;
- 5: $T_{k+1} = T'_k + T_{\text{exp}} - P_A (P_A^* T'_k P_B) P_B^*$, $\Lambda_{k+1} = \Lambda'_k + \Lambda_{\text{exp}} - Q_A (P_A^* T'_k P_B) P_B^*$.

Operations whose order of execution is insignificant and which can run independently in parallel threads are shown on the same line. Note that computation of the terms $(\Lambda_k T_k)_{ii}$ and $(\Lambda_k^* \Lambda_k)_{ii}$ for $i = 1, \dots, N_v$ has the computational complexity of only $O(N_v^2)$. Therefore, the true bottleneck of each iteration is the operation of matrix inversion $S_k = (I - \Delta_k)^{-1}$ whose computational complexity is $O(N_v^3)$. Iterations can be stopped when the norm $\|T_k - D_k - D_k \Gamma T_k\|_2$ is smaller than a predetermined threshold.

6. Variations and improvements

6.1. Starting iterations from an initial guess for V

Starting from an initial guess for the T-matrix is not a requirement of DCTMC. We can as well start from an initial guess for V . For example, we can use the generalized solution to the linearized equation (6) to build the initial guess for V , D_L . Here D_L has the elements of $|v\rangle$ on the diagonal and zeros elsewhere. We then set $D_1 = D_L$ and skip Step 1 in the $k = 1$ iteration of the algorithm described in Sec. 5.3. A method to obtain the initial guess D_L that avoids the high computational complexity of finding the pseudo-inverse of K is described in Appendix B.

6.2. Accounting for sparse targets

The computational bottleneck of DCTMC is computing the matrix inverse $(I - \Delta_k)^{-1}$. The computational load of this step can be significantly reduced if some of the elements of D_k are zero. Indeed, since $\Delta_k = \Gamma D_k$, this occurrence would imply that some of the columns of Δ_k are zero and, consequently, that some of the columns of $(I - \Delta_k)$ contain unity on the diagonal and zeros elsewhere. Complexity of inverting such matrices scales as cube of the number of nonzero elements in D_k . We note that, if $(D_k)_{ii} = 0$, then i -th column and i -th row of T'_k are zero. The size of the T-matrix is effectively reduced in this case.

In practice, we can often expect that many diagonal elements of the unknown interaction matrix V are zero but the location of these zero elements is unknown. If we accidentally run into a matrix D_k that has many zeros on the diagonal, we would be able to reduce the computational complexity of inverting $(I - \Delta_k)$. However, it is highly unlikely that the iterative computation of $(D_k)_{ii}$ according to Step 1 of the algorithm of Sec. 5.3 would produce exact zeros for any i . To overcome this difficulty, we can apply an additional operation to D_k , which we refer to as "roughening". This operation entails setting to zero the diagonal elements of D_k that are smaller than a certain predetermined threshold, either relative or absolute. Such elements can be deemed as insignificant computational noise. Alternatively we can send a predetermined number of the smallest diagonal elements of D_k to zero. The approach was implemented in [27] and shown to significantly improve the efficiency and convergence rate of DCTMC.

6.3. Accounting for physical constraints

DCTMC allows one to apply additional nonlinear constraints to D_k . This can be accomplished immediately after D_k is updated in each iteration of the cycle. For example, a physical constraint may require that all elements of D_k be non-negative or have a non-negative imaginary part. This procedure can be viewed as a form of regularization of the ISP.

7. DCTMC in the linear regime and the questions of convergence and regularization

Consider the iteration cycle of Sec. 5.3 in the limit $\Gamma \rightarrow 0$. Omitting intermediate steps, we find that each iteration is reduced to the following two operations:

- 1: $D_k = \mathcal{D}[T_k]$;
- 2: $T_{k+1} = D_k + T_{\text{exp}} - (P_A P_A^*) D_k (P_B P_B^*)$,

where $\mathcal{D}[\cdot]$ is defined in (21). These two steps can be combined in the following simple iteration:

$$D_{k+1} = D_k + D_{\text{exp}} - \mathcal{D}[(P_A P_A^*) D_k (P_B P_B^*)] , \quad (31)$$

where $D_{\text{exp}} = \mathcal{D}[T_{\text{exp}}]$. Iteration (31) can be obtained simply by applying the operator $\mathcal{D}[\cdot]$ to the equation in Step 2 above. Let us now convert (31) to an equation with respect to the vector $|v_k\rangle$ that contains the diagonal elements of D_k . From the linearity of (31) we immediately find

$$|v_{k+1}\rangle = |v_{\text{exp}}\rangle + (I - W)|v_k\rangle , \quad (32)$$

where $|v_{\text{exp}}\rangle$ is the vector of diagonal elements of D_{exp} and matrix W has the elements

$$W_{ij} = (P_A P_A^*)_{ij} (P_B P_B^*)_{ji} . \quad (33)$$

It is easy to see that (32) is Richardson first-order iteration with the fixed point $|v_\infty\rangle = W^{-1}|v_{\text{exp}}\rangle$. Therefore, DCTMC in the linear regime simply provides an iterative way of solving the equation

$$W|v\rangle = |v_{\text{exp}}\rangle . \quad (34)$$

This equation can be derived independently from DCTMC and in a more straightforward manner starting from the linearized equation (6). This derivation is shown in Appendix B and it takes advantage of the algebraic properties of K (see Remark 3). It is important to realize that, although (34) can be obtained from (6) by a series of linear transformations, the two equations are not equivalent in the following sense: if K is not invertible, then the pseudoinverse solutions of the two equations can be different. However, if K is invertible, then the two equations have the same unique solution.

Of course, iteration (32) is only a particular numerical method of solving (34) and not the most efficient one: conjugate-gradient descent is expected to provide better computational performance. However, consideration of DCTMC in the linear regime is not a vane or trivial exercise but is useful in several respects. First, it gives us an insight into the convergence properties of DCTMC. Second, it gives us an idea of how DCTMC iterations can be regularized. Indeed, so far we did not discuss regularization, apart from what is afforded by imposing physical constraints (Sec. 6.3). Convergence and regularization will be discussed in the remainder of this section.

It is obvious that the iterations converge to the fixed point provided that $|1 - w_n| < 1$ for all n , where w_n are the eigenvalues of W . Since W is Hermitian, all its eigenvalues

are real and therefore the convergence condition reads $0 < w_n < 2$. Under the same condition the inverse W^{-1} exists. We will now prove that

$$0 \leq w_n \leq 1 . \quad (35)$$

The fact that W is non-negative definite is obvious from (33). We will however make an additional step and recall that the columns of P_A are the singular vectors $|g_\mu^A\rangle$ for $\mu = 1, \dots, M_A$ and zeros otherwise while the columns of P_B are the singular vectors $|f_\nu^B\rangle$ for $\nu = 1, \dots, M_B$ and zeros otherwise. We then obtain in a straightforward manner:

$$W_{ij} = \sum_{\mu=1}^{M_A} \sum_{\nu=1}^{M_B} \langle i | g_\mu^A \rangle \langle g_\mu^A | j \rangle \langle j | f_\nu^B \rangle \langle f_\nu^B | i \rangle . \quad (36)$$

Let $|x\rangle$ be an arbitrary nonzero vector of length N_v and X be an $N_v \times N_v$ matrix with the elements of $|x\rangle$ on the diagonal and zeros elsewhere. Then

$$\langle x | W | x \rangle = \sum_{\mu=1}^{M_A} \sum_{\nu=1}^{M_B} |\langle g_\mu^A | X | f_\nu^B \rangle|^2 \geq 0 . \quad (37)$$

We therefore have proved that $w_n \geq 0$. Next, we use the orthonormality of each set of singular vectors to write the following identities

$$\langle x | x \rangle = \sum_{i=1}^{N_v} \langle i | X^* X | i \rangle = \sum_{\mu=1}^{N_v} \sum_{\nu=1}^{N_v} |\langle g_\mu^A | X | f_\nu^B \rangle|^2 . \quad (38)$$

Since $M_a, M_b \leq N_v$, $\langle x | W | x \rangle \leq \langle x | x \rangle$ and we have proved (35). The equality $\langle x | W | x \rangle = \langle x | x \rangle$ holds only in the case $M_A = M_B = N_v$, in which case $W = I$ and the iteration (32) trivially converges to its fixed point right upon making the initial guess. In this unrealistic case, all elements of the T-matrix are determined from the data and no iterations are needed.

We thus conclude that convergence can be slow in the case W has a small (or zero) eigenvalue. We can define the characteristic *overlap* of singular vectors related to detectors and sources as

$$\xi = \inf_{X \neq 0} \left\{ \frac{\sum_{\mu=1}^{M_A} \sum_{\nu=1}^{M_B} |\langle g_\mu^A | X | f_\nu^B \rangle|^2}{\sum_{\mu=1}^{N_v} \sum_{\nu=1}^{N_v} |\langle g_\mu^A | X | f_\nu^B \rangle|^2} \right\} . \quad (39)$$

The iterations (32) converge at least as fast as the power series $\sum_n (1 - \xi)^n$. If ξ is close to zero, the convergence can be slow. This observation gives us an idea of how the iterations can be regularized. This can be accomplished by replacing W by $W + \lambda^2 I$, where λ is a regularization parameter. As is shown in Appendix B, this procedure is equivalent to Tikhonov regularization of equation (B.8), which can be derived from (6) by several linear operations. However, the substitution $W \rightarrow W + \lambda^2 I$ is not equivalent to Tikhonov regularization of (6).

We can now introduce regularization of the general iterative algorithm of Sec. 5.3, which is applicable to the nonlinear case as well. Namely, for any matrix X , we replace the linear transformation $(P_A P_A^*)X(P_B P_B^*)$ by $(P_A P_A^*)X(P_B P_B^*) + \lambda^2 \mathcal{D}[X]$. This entails the following modified Step 5 of the algorithm:

$$\begin{aligned} 5: T_{k+1} &= T'_k - \lambda^2 \mathcal{D}[T'_k] + T_{\text{exp}} - P_A (P_A^* T'_k P_B) P_B^* \quad , \\ \Lambda_{k+1} &= \Lambda'_k - \lambda^2 \mathcal{D}[\Lambda'_k] + \Lambda_{\text{exp}} - Q_A (P_A^* T'_k P_B) P_B^* \quad . \end{aligned}$$

8. Discussion

This paper contains a mathematical description of a novel method for solving nonlinear inverse scattering problems. The method is based on iterative completion of the unknown entries of the T-matrix and we refer to it as to the data-compatible T-matrix completion (DCTMC). It should be emphasized that the constraint that we apply to the T-matrix (namely, that it corresponds to a nearly diagonal interaction matrix V) is not the same as in the conventional formulation of the matrix completion problem. The method developed in this paper is well suited for overdetermined ISPs in which the number of volume voxels is not too large (e.g., $\lesssim 10^4$) or the target is sparse. The size of the data set is not a limiting factor for this method, unlike in many traditional approaches to the same problem.

In the case of ill-posed ISPs, regularization plays the key role. One should not expect to recover a reasonable image without some form of regularization. DCTMC allows for two types of regularization: (i) by imposing physical constraints and (ii) by regularizing the matrix $W_{ij} = (P_A P_A^*)_{ij} (P_B P_B^*)_{ji}$. In the linear regime, the approach (ii) corresponds to Tikhonov regularization of the linearized equation $\Theta U K |v\rangle = \Theta U |\phi\rangle$, which is obtained from Eq. (6) by multiplying the latter by ΘU from the left; the unitary matrix U and the matrix of diagonal scaling Θ are defined in Appendix B. In the nonlinear regime, the approach (ii) is somewhat *ad hoc* and its applicability requires additional research.

A potential advantage of DCTMC is its computational efficiency in solving strongly overdetermined inverse problems with large data sets. This advance is obtained by exploiting the algebraic structure of the ISP rather than stating it in the conventional generic form $F[v] = 0$ (or $Kv = 0$ in the linear approximation) where $F[\cdot]$ is a general nonlinear functional, K is its linear approximation and v is the vector consisting of the diagonal elements of V . In our previous work, we have shown that strongly overdetermined data sets are required to obtain the optimal image resolution [28, 29]. Specifically, the fundamental limit of *lateral* resolution of diffuse optical tomography (DOT) in the slab geometry is equal to the step h of the mesh of sources and detectors on each face of the slab, provided that the imaging windows are significantly larger than the field of view of the reconstruction. Thus, to reconstruct a lateral central cross section of the slab on a 100×100 mesh of step h , one needs about 300^2 sources on one side and the same number of detectors on the other side of the slab. This translates into $\sim 10^{10}$ data points. This result is not specific to DOT but also holds for diffraction

tomography that we consider numerically in [27]. However, the fundamental resolution limit is not always achievable. If the inverse problem is ill-posed or noise is present in the data, the theoretical limit of resolution can not be achieved and smaller data sets can suffice to obtain the optimal image quality. Therefore, the optimal size of the data set is determined by a complex interplay between the ill-posedness of the inverse problem and the statistical properties of the noise. Experimental DOT reconstruction with large data sets was demonstrated in [30], experimental determination of the optimal size of the data set was in [30,31] and further insights on selection of useful data points in the presence of strong nonlinearity were provided in [32]. However, all references just quoted used linearized image reconstruction. DCTMC allows one to overcome this limitation and obtain nonlinear reconstructions with very large data sets.

Although the main goal of DCTMC is to solve nonlinear problems, we note that it has conceptual similarity with the image reconstruction methods of references [29,33–35] that were developed for solving linear ISPs. The similarity lies in exploring the algebraic structure of the matrix K , which is obtained as a product of two unperturbed Green's functions. In the nonlinear case, we use this algebraic structure to obtain an expression of the form ATB , where the T -matrix T is viewed as the fundamental unknown. This replaces the traditional approach in which one writes $ATB = F[V]$, where $F[V]$ is a nonlinear functional and the strictly diagonal (by definition) matrix V is viewed as the fundamental unknown. DCTMC also bears conceptual similarity to the methods of reference [14] in which the notion of the fundamental unknown is also expanded to include the internal fields or the complete Green's function G .

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 US National Science Foundation under Grant DMS 1115616 and by the US National Institutes of health under Grant P41 RR002305. The authors are grateful to J.C.Schotland and A.Yodh for very useful discussions.

References

- [1] Snieder R 1998 *Inverse Problems* **14** 387404
- [2] Seo J K and Woo E J 2012 *Nonlinear Inverse Problems in Imaging* (Wiley)
- [3] 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
- [4] Arridge S R and Schotland J C 2009 *Inverse Problems* **25** 123010
- [5] Bronstein M M, Bronstein A M, Zibulevski M and Azhari H 2002 *IEEE Trans. Med. Imag.* **21** 1395–1401
- [6] Devaney A J 1984 *IEEE Trans. Geosci. Remote Sensing* **GE-22** 3–13
- [7] Berryman J G and Kohn R V 1990 *Phys. Rev. Lett.* **65** 325–328
- [8] Isaacson D, Mueller J L, Newell J C and Siltanen S 2004 *IEEE Trans. Med. Imag.* **23** 821–828

- [9] Arridge A, Moskow S and Schotland J C 2012 *Inverse Problems* **28** 035003
- [10] Carney P S, Frazin R A, Bozhevolnyi S I, Volkov V S, Boltasseva A and Schotland J C 2004 *Phys. Rev. Lett.* **92** 163903
- [11] Belkebir K, Chaumet P C and Sentenac A 2005 *J. Opt. Soc. Am. A* **22** 1889–1897
- [12] Bao G and Li P 2007 *Opt. Lett.* **32** 1465–1467
- [13] Belkebir K, Chaumet P C and Sentenac A 2006 *J. Opt. Soc. Am. A* **23** 586–595
- [14] Mudry E, Chaumet P C, Belkebir K and Sentenac A 2012 *Inverse Problems* **28** 065007
- [15] Jakobsen M 2012 *Stud. Geophys. Geod.* **56** 1
- [16] Jakobsen M and Ursin B 2015 *J. Geophys. Eng.* **12** 400
- [17] Newton R G 1966 *Scattering theory of waves and particles* (New York: McGraw-Hill)
- [18] Colton D and Kress R 1998 *Inverse Acoustic and Electromagnetic Scattering Theory (Applied Mathematical Sciences vol 93)* (Berlin: Springer)
- [19] Aster R C, Borchers B and Thurber C H 2005 *Parameter Estimation and Inverse Problems* (Amsterdam: Elsevier)
- [20] Engl H W and Kugler P 2005 *Multidisciplinary Methods for Analysis Optimization and Control of Complex Systems Mathematics in Industry vol 6* (Springer) chap Nonlinear inverse problems: Theoretical aspects and some industrial applications, pp 3–47
- [21] Markel V A, O’Sullivan J A and Schotland J C 2003 *J. Opt. Soc. Am. A* **20** 903–912
- [22] Moskow S and Schotland J 2008 *Inverse Problems* **24** 065005
- [23] Moskow M and Schotland J 2009 *Inverse Problems* **25** 095007
- [24] Watzenig D 2007 *e & i Elektrotechnik und Informationstechnik* vol 124 (Springer) chap Bayesian inference for inverse problems - statistical inversion, pp 240–247
- [25] Stolt R H and Jacobs B 1980 Inversion of seismic data in a laterally heterogeneous medium Tech. Rep. 24 Stanford Exploration Project (SEP)
- [26] Kouri D J and Vijay A 2003 *Phys. Rev. E* **67** 046614
- [27] Levinson H W and Markel V A 2014 *Inverse Problems* Submitted
- [28] Markel V A and Schotland J C 2002 *Appl. Phys. Lett.* **81** 1180–1182
- [29] Markel V A and Schotland J C 2004 *Phys. Rev. E* **70** 056616
- [30] Wang Z M, Panasyuk G Y, Markel V A and Schotland J C 2005 *Opt. Lett.* **30** 3338–3340
- [31] Konecky S D, Panasyuk G Y, Lee K, Markel V, Yodh A and Schotland J C 2008 *Opt. Expr.* **16** 5048–5060
- [32] 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
- [33] Schotland J C 1997 *J. Opt. Soc. Am. A* **14** 275–279
- [34] Schotland J C and Markel V A 2001 *J. Opt. Soc. Am. A* **18** 2767–2777
- [35] Markel V A, Mital V and Schotland J C 2003 *J. Opt. Soc. Am. A* **20** 890–902

Appendix A. Linearizing approximations

In this Appendix we state the linearizing approximations without derivation or analysis as this is outside of the scope of this paper. We only note that first Born approximation is obtained by retaining the first-order term in the power-series expansion of the complete Green’s function G ; first Rytov approximation is obtained by retaining the first-order term in the cumulant expansion of $\log(G)$ and the mean-field approximation is obtained by using the first-order approximant in the continued-fraction expansion of G . More details are given in [29]. Only first Born approximation can be stated in abstract form while the other two approximations involve entry-wise expressions. Correspondingly, the accuracy of these two approximations depends on the matrix representation while the accuracy of first Born approximation is representation-independent.

(i) *First Born approximation.* The simplest approach to linearization is the first Born approximation, according to which

$$A(I - V\Gamma)^{-1}VB \approx AVB . \quad (\text{A.1})$$

Obviously, the first Born approximation is valid if $\|V\Gamma\|_2 \ll 1$. By substituting the above approximation into the left-hand side of (4), we obtain (5) in which $\Psi = \Phi$. Therefore, the data transformation in the case of first Born approximation is trivial.

(ii) *First Rytov approximation.* The first Rytov approximation can be stated as

$$[A(I - V\Gamma)^{-1}VB]_{ij} \approx C_{ij} \left\{ \exp \left[\frac{(AVB)_{ij}}{C_{ij}} \right] - 1 \right\} , \quad (\text{A.2})$$

where $C_{ij} = G_0(\mathbf{r}_i, \mathbf{r}_j)$ and $\mathbf{r}_i \in \Sigma_d$, $\mathbf{r}_j \in \Sigma_s$. Thus we have encountered yet another restriction of $G_0(\mathbf{r}, \mathbf{r}')$. Obviously, with this restriction used, $G_0(\mathbf{r}, \mathbf{r}')$ is the *direct* (unscattered) field that would have been produced by a source located at \mathbf{r} and measured by a detector at \mathbf{r}' in the case $V = 0$. If we substitute the approximation (A.2) into the left-hand side of (4) and introduce the data transformation

$$\Psi_{ij} = C_{ij} \log(1 + \Phi_{ij}/C_{ij}) , \quad (\text{A.3})$$

we would arrive again at (5). Therefore, (A.3) defines the data transformation of first Rytov approximation.

(iii) *Mean-field approximation.* The mean-field approximation is

$$[A(I - V\Gamma)^{-1}VB]_{ij} \approx \frac{(AVB)_{ij}}{1 - (AVB)_{ij}/C_{ij}} \quad (\text{A.4})$$

The data transformation of the mean-field approximation has the form of element-wise harmonic average and reads

$$\Psi_{ij} = \frac{1}{1/\Phi_{ij} + 1/C_{ij}} . \quad (\text{A.5})$$

Appendix B. Derivation of the equation $W|v\rangle = |v_{\text{exp}}\rangle$ [Eq. (34)] from (6).

Consider the linearized equation (6) in the first Born approximation, that is, with the trivial data transformation $\Psi = \Phi$ or equivalently $|\psi\rangle = |\phi\rangle$. We recall that $K_{(mn),j} = A_{mj}B_{jn}$ and use (12) for A_{mj} and B_{jn} to obtain the following result for the elements of K :

$$K_{(mn),j} = \sum_{\mu=1}^{N_d} \sum_{\nu=1}^{N_s} \sigma_{\mu}^A \sigma_{\nu}^B \langle m | f_{\mu}^A \rangle \langle g_{\mu}^A | j \rangle \langle j | f_{\nu}^B \rangle \langle g_{\nu}^B | n \rangle . \quad (\text{B.1})$$

Now define the unitary matrix U with the elements

$$U_{(\mu\nu), (mn)} = \langle f_{\mu}^A | m \rangle \langle n | g_{\nu}^B \rangle , \quad 1 \leq \mu, m \leq N_d , \quad 1 \leq \nu, n \leq N_s \quad (\text{B.2})$$

and multiply (6) by U from the left:

$$(UK)|v\rangle = U|\phi\rangle . \quad (\text{B.3})$$

We emphasize that multiplication of any linear equation by a unitary matrix does not change its Tikhonov-regularized pseudoinverse solution. This follows immediately from $(UK)^*(UK) = K^*K$ and $(UK)^*U = K^*$. Now we use the equalities

$$(UK)_{(\mu\nu),j} = \sigma_\mu^A \sigma_\nu^B \langle g_\mu^A | j \rangle \langle j | f_\nu^B \rangle , \quad \langle (\mu\nu) | U | \phi \rangle = \tilde{\Phi}_{\mu\nu} \quad (\text{B.4})$$

to write (B.3) as

$$\sigma_\mu^A \sigma_\nu^B \sum_{j=1}^{N_v} \langle g_\mu^A | j \rangle \langle j | f_\nu^B \rangle \langle j | v \rangle = \tilde{\Phi}_{\mu\nu} , \quad 1 \leq \mu \leq N_d , \quad 1 \leq \nu \leq N_s . \quad (\text{B.5})$$

Next we observe that the set (B.5) may contain some equations in which all coefficients are zero or very small. These equations can be safely discarded and this operation still does not affect the pseudo-inverse. Therefore we obtain

$$\sigma_\mu^A \sigma_\nu^B \sum_{j=1}^{N_v} \langle g_\mu^A | j \rangle \langle j | f_\nu^B \rangle \langle j | v \rangle = \tilde{\Phi}_{\mu\nu} , \quad 1 \leq \mu \leq M_A , \quad 1 \leq \nu \leq M_B , \quad (\text{B.6})$$

where M_A and M_B are the dimensions of the shaded rectangle in Fig. 3, which is the region where the inequality $\sigma_\mu^A \sigma_\nu^B > \epsilon^2$ holds and ϵ is the small positive constant introduced in (16). Note that (B.6) is in all respects equivalent to (6).

At this point however, we make a transformation that will, in fact, change the equation. Namely, we divide (B.6) by the factor $\sigma_\mu^A \sigma_\nu^B$, which is larger than ϵ^2 for all equations included in (B.6). In computational linear algebra, this operation is known as preconditioning by diagonal scaling. We thus obtain

$$\sum_{j=1}^{N_v} \langle g_\mu^A | j \rangle \langle j | f_\nu^B \rangle \langle j | v \rangle = \left(\tilde{T}_{\text{exp}} \right)_{\mu\nu} , \quad 1 \leq \mu \leq M_A , \quad 1 \leq \nu \leq M_B , \quad (\text{B.7})$$

where we have used the definition (18) of \tilde{T}_{exp} . The diagonal scaling that was applied to obtain (B.7) is invertible. Therefore, if (6) has a solution in the ordinary sense, then (B.7) has the same solution. However, if (6) is not invertible, then the two equations have different pseudoinverse solutions that can not be related to each other by a simple transformation. In this sense, the two equations (6) and (B.7) are no longer equivalent.

To obtain (34), we observe that (B.7) can be written as

$$Q|v\rangle = |\tau\rangle , \quad (\text{B.8})$$

where $Q_{(\mu\nu),j} = \langle g_\mu^A | j \rangle \langle j | f_\nu^B \rangle$ and $\langle (\mu\nu) | \tau \rangle = \left(\tilde{T}_{\text{exp}} \right)_{\mu\nu}$. We then multiply (B.6) by Q^* from the left and obtain

$$(Q^*Q)_{ij} = \sum_{\mu=1}^{M_A} \sum_{\nu=1}^{M_B} \langle f_\nu^B | i \rangle \langle i | g_\mu^A \rangle \langle g_\mu^A | j \rangle \langle j | f_\nu^B \rangle = W_{ij} , \quad (\text{B.9})$$

where $W = Q^*Q$ is the same matrix as in (33) (compare (B.9) to (36)). In a similar manner, we obtain

$$\langle i|Q^*|\tau\rangle = \sum_{\mu_1=1}^{M_A} \sum_{\nu=1}^{M_B} Q_{i,(\mu\nu)}^* \left(\tilde{T}_{\text{exp}} \right)_{\mu\nu} = (P_A \tilde{T}_{\text{exp}} P_B^*)_{ii} = (T_{\text{exp}})_{ii} = \langle i|v_{\text{exp}}\rangle. \quad (\text{B.10})$$

Therefore, $Q^*|\tau\rangle = |v_{\text{exp}}\rangle$. We thus conclude that (34) is obtained from (B.8) by multiplying both sides with Q^* . Moreover, the substitution $W \rightarrow W + \lambda^2 I$ is equivalent to Tikhonov-regularization of (B.8).

Finally, we can state the formal relation between K and W in the following form:

$$W = (\Theta UK)^*(\Theta UK) = K^* U^{-1} \Theta^2 UK, \quad (\text{B.11})$$

where U is the unitary matrix defined in (B.2) and Θ is the diagonal conditioning matrix containing the quantities $1/(\sigma_\mu^A \sigma_\nu^B)$ for $1 \leq \mu \leq M_A$ and $1 \leq \nu \leq M_B$ and zeros otherwise.

Appendix C. Definitions and properties of several functionals used in this paper

This appendix gives a summary of the functionals used in this paper. In the table, \mathcal{F} refers to any of the functionals \mathcal{T} , \mathcal{R} , \mathcal{D} , \mathcal{M} and \mathcal{N} and \mathcal{O} .

\mathcal{F}	Entry-wise?	$Y = \mathcal{F}[X]$	Invertible?	$X = \mathcal{F}^{-1}[Y]$
\mathcal{T}	No	$Y = (I - X\Gamma)^{-1}X$	Sometimes	$X = (I + Y\Gamma)^{-1}Y$
\mathcal{R}	No	$Y = R_A^* X R_B$	Yes	$X = R_A Y R_B^*$
\mathcal{D}	Yes	$Y_{ij} = X_{ij} \delta_{ij}$	No	N/A
\mathcal{M}	Yes	$\tilde{Y}_{\mu\nu} = \begin{cases} 0, & \sigma_\mu^A \sigma_\nu^B > \epsilon^2 \\ \tilde{X}_{\mu\nu}, & \text{otherwise} \end{cases}$	No	N/A
\mathcal{N}	Yes	$\tilde{Y}_{\mu\nu} = \begin{cases} \tilde{X}_{\mu\nu}, & \sigma_\mu^A \sigma_\nu^B > \epsilon^2 \\ 0, & \text{otherwise} \end{cases}$	No	N/A
\mathcal{O}	Yes	$\tilde{Y} = \mathcal{M}[\tilde{X}] + \tilde{T}_{\text{exp}}$ $= \tilde{X} - \mathcal{N}[\tilde{X}] + \tilde{T}_{\text{exp}}$	No	N/A