A model reduction approach for inverse problems with operator valued data

Jürgen Dölz¹ · Herbert Egger² · Matthias Schlottbom³

Received: 22 December 2020 / Revised: 8 June 2021 / Accepted: 26 July 2021 /

Published online: 16 August 2021

© The Author(s) 2021

Abstract

We study the efficient numerical solution of linear inverse problems with operator valued data which arise, e.g., in seismic exploration, inverse scattering, or tomographic imaging. The high-dimensionality of the data space implies extremely high computational cost already for the evaluation of the forward operator which makes a numerical solution of the inverse problem, e.g., by iterative regularization methods, practically infeasible. To overcome this obstacle, we take advantage of the underlying tensor product structure of the problem and propose a strategy for constructing low-dimensional certified reduced order models of quasi-optimal rank for the forward operator which can be computed much more efficiently than the truncated singular value decomposition. A complete analysis of the proposed model reduction approach is given in a functional analytic setting and the efficient numerical construction of the reduced order models as well as of their application for the numerical solution of the inverse problem is discussed. In summary, the setup of a low-rank approximation can be achieved in an offline stage at essentially the same cost as a single evaluation of the forward operator, while the actual solution of the inverse problem in the online phase can be done with extremely high efficiency. The theoretical results are illustrated by application to a typical model problem in fluorescence optical tomography.

Mathematics Subject Classification 46N40 · 65J20 · 65N21

1 Introduction

We consider the numerical solution of linear inverse problems with operator valued data modeled by abstract operator equations

$$\mathcal{T}(c) = \mathcal{M}^\delta. \quad (1.1)$$

✉ Matthias Schlottbom
m.schlottbom@utwente.nl

Extended author information available on the last page of the article

Here $c \in \mathbb{X}$ is the quantity to be determined and we assume that $\mathcal{M}^\delta : \mathbb{Y} \rightarrow \mathbb{Z}'$, representing the possibly perturbed measurements, is a linear operator of Hilbert-Schmidt class between Hilbert spaces $\mathbb{Y}$ and $\mathbb{Z}'$, the dual of $\mathbb{Z}$. We further assume that the forward operator $\mathcal{T} : \mathbb{X} \rightarrow \mathbb{HS}(\mathbb{Y}, \mathbb{Z}')$ is linear and compact, and admits a factorization of the form

$$\mathcal{T}(c) = \mathcal{V}' \mathcal{D}(c) \mathcal{U}, \quad (1.2)$$

with $\mathcal{V}'$, $\mathcal{D}(c)$, and $\mathcal{U}$ again denoting appropriate linear operators. Problems of this kind arise in a variety of applications, e.g. in fluorescence tomography [1,30], inverse scattering [6,13], or source identification [17], but also as linearizations of related nonlinear inverse problems, see e.g., [8,34] or [23] and the references given there. In such applications, $\mathcal{U}$ typically models the propagation of excitation fields generated by the sources, $\mathcal{D}$ describes the interaction with the medium to be probed, and $\mathcal{V}'$ models the emitted fields which can be recorded by the detectors. In the following, we briefly outline our basic approach towards the numerical solution of (1.1)–(1.2) and report about related work in the literature.

1.1 Regularized inversion

By the particular functional analytic setting, the inverse problem (1.1)–(1.2) amounts to an ill-posed linear operator equation in Hilbert spaces and standard regularization theory can be applied for its stable solution [2,9]. Following standard arguments, we assume that $\mathcal{M}^\delta$ is a perturbed version of the exact data $\mathcal{M}$ and that a bound on the measurement noise

$$\|\mathcal{M} - \mathcal{M}^\delta\|_{\mathbb{HS}(\mathbb{Y}, \mathbb{Z}')} \leq \delta \quad (1.3)$$

is available. We further denote by $c^\dagger$ the minimum norm solution of (1.1) with $\mathcal{M}^\delta$ replaced by $\mathcal{M} = \mathcal{T}(c^\dagger)$. A stable approximation for the solution $c^\dagger$ can then be obtained by the regularized inversion of (1.1), e.g., by spectral regularization methods

$$c_\alpha^\delta = g_\alpha(\mathcal{T}^* \mathcal{T}) \mathcal{T}^* \mathcal{M}^\delta = \mathcal{T}^* g_\alpha(\mathcal{T} \mathcal{T}^*) \mathcal{M}^\delta. \quad (1.4)$$

Here $\mathcal{T}^* : \mathbb{HS}(\mathbb{Y}, \mathbb{Z}') \rightarrow \mathbb{X}$ denotes the adjoint of the operator $\mathcal{T}$ and $g_\alpha(\lambda)$ denotes a regularized approximation of $1/\lambda$, i.e., the filter function, satisfying some standard conditions; we refer to [2, Chapter 2] or [9, Chapter 4] for details and to [26] for generalizations. A typical example for the filter function is $g_\alpha(\lambda) = (\lambda + \alpha)^{-1}$, which leads to Tikhonov regularization $c_\alpha^\delta = (\mathcal{T}^* \mathcal{T} + \alpha \mathcal{I})^{-1} \mathcal{T}^* \mathcal{M}^\delta$, see [36]. Another filter function with certain optimality conditions results from truncated singular value decomposition and reads $g_\alpha(\lambda) = 1/\lambda$ if $\lambda \geq \alpha$ and $g_\alpha(\lambda) = 0$ otherwise.

Note that the choice of the regularization norm in (1.4) is incorporated implicitly in the definition of the function spaces and one can obtain convergence $c_\alpha^\delta \rightarrow c^\dagger$ of the regularized solutions to the minimum norm solution $c^\dagger$ in this norm if $\delta \rightarrow 0$ and $\alpha = \alpha(\delta, \mathcal{M}^\delta)$ is chosen appropriately; the rate of convergence will depend

on the properties of the filter function g_α , the parameter choice $\alpha(\delta, \mathcal{M}^\delta)$, and on the smoothness of the minimum norm solution $c^\dagger$; see [2,9,26] for details. For tomographic applications we have in mind, uniqueness results are usually available [18,28], i.e., $\mathcal{T}$ can be assumed to be injective, in which case $c^\dagger$ is independent of the regularization norm; see Remark 3.3 below. We will not step further into the analysis of regularization methods, but rather focus on their efficient numerical realization for problems with operator valued data.

For the actual computation of the regularized solution c_α^δ , a sufficiently accurate finite dimensional approximation of the operator $\mathcal{T}$ is required, which is usually obtained by some discretization procedure; in the language of model order reduction, this is called the *truth* or *high-fidelity approximation* [3,32]. In the following discussion, we will not distinguish between infinite dimensional operators and their truth approximations. We thus assume that $\dim(\mathbb{X}) = m$, $\dim(\mathbb{Y}) = k_Y$ and $\dim(\mathbb{Z}) = k_Z$. For ease of notation, we assume that $k_X = k_Y = k$ in the following. We may then identify c with a vector in $\mathbb{R}^m$, $\mathcal{M}^\delta$ with a matrix in $\mathbb{R}^{k \times k}$, and $\mathcal{T}$ with a third order tensor in $\mathbb{R}^{k \times k \times m}$ or a matrix in $\mathbb{R}^{k^2 \times m}$. In typical applications, like computerized tomography, the dimensions m and k are very large and one may assume that $k < m < k^2$; see [5,24]. The inverse problem (1.1) thus can be considered to be typically of *large scale* and *overdetermined*.

1.2 Model reduction and computational complexity

The high-dimensionality of the problem poses severe challenges for the numerical solution of the inverse problem (1.1)–(1.2) and different model reduction approaches have been proposed to reduce the computational complexity. We consider approximations

$$\mathcal{T}_N = \mathcal{Q}_N \mathcal{T}, \quad (1.5)$$

based on projection in data space, where $\mathcal{Q}_N$ is an orthogonal projection with finite rank N , which is the dimension of the range of $\mathcal{Q}_N$. Since $\mathcal{T}$ is assumed compact, we can always choose N sufficiently large such that

$$\|\mathcal{T}_N - \mathcal{T}\|_{\mathbb{X} \rightarrow \mathbb{HS}(\mathbb{Y}, \mathbb{Z})} \leq \delta, \quad (1.6)$$

and we may assume that typically $N \ll m, k$, where m, k are the dimensions of the truth approximation used for the computations.

For the stable and efficient numerical solution of the inverse problem (1.1)–(1.2), we may then consider the low-dimensional regularized approximation

$$c_{\alpha, N}^\delta = \mathcal{T}_N^* g_\alpha(\mathcal{T}_N \mathcal{T}_N^*) \mathcal{Q}_N \mathcal{M}^\delta. \quad (1.7)$$

As shown in [29], the low-rank approximation $c_{\alpha, N}^\delta$ defined in (1.7) has essentially the same quality as the infinite dimensional approximation c_α^δ , as long as the perturbation

bound (1.6) can be guaranteed. In the sequel, we therefore focus on the numerical realization of (1.7), which can be roughly divided into the following two stages:

- Setup of the approximations $\mathcal{Q}_N$, $\mathcal{T}_N^*$, and $\mathcal{T}_N \mathcal{T}_N^*$. This compute intensive part can be done in an *offline stage* and the constructed approximations can be used for repeated solution of the inverse problem (1.1) for multiple data.
- Computation of the regularized solution (1.7). This *online stage*, which is relevant for the actual solution of (1.1), comprises the following three steps:

Step	Computations	Complexity	Memory
Compression	$\mathcal{M}_N^\delta = \mathcal{Q}_N \mathcal{M}^\delta$	Nk^2	Nk^2
Analysis	$z_{\alpha,N}^\delta = g_\alpha(\mathcal{T}_N \mathcal{T}_N^*) \mathcal{M}_N^\delta$	N^2	N^2
Synthesis	$c_{\alpha,N}^\delta = \mathcal{T}_N^* z_{\alpha,N}^\delta$	Nm	Nm

Let us note that the analysis step is completely independent of the large system dimension k, m of the truth approximation and therefore the compression and synthesis step are the compute intensive parts in the online stage. If $k^2 > m > k$, which is the typical situation [5,24], the data compression turns out to be the most compute and memory expensive step. As we will explain below, the tensor product structure of (1.2) allows us to considerably reduce the memory cost in the compression step.

1.3 Low-rank approximations

A particular example of a low-rank approximation (1.5) is given by the truncated singular value decomposition $\mathcal{T}_{N^{\text{svd}}}$, for which $\mathcal{Q}_{N^{\text{svd}}}$ amounts to the orthogonal projection onto the N^{svd} -dimensional subspace of the left singular vectors corresponding to the largest singular values $\sigma_1 \geq \dots \geq \sigma_{N^{\text{svd}}}$ of the operator $\mathcal{T}$. By construction of the singular value decomposition, we have

$$\|\mathcal{T}_{N^{\text{svd}}} - \mathcal{T}\|_{\mathbb{X} \rightarrow \text{HS}(\mathbb{Y}, \mathbb{Z}')} = \sigma_{N^{\text{svd}}+1}, \quad (1.8)$$

which allows to guarantee the desired accuracy (1.6) by choosing $\sigma_{N^{\text{svd}}+1} \leq \delta < \sigma_{N^{\text{svd}}}$. From the Eckhard-Young-Mirsky theorem, we can conclude that $N = N^{\text{svd}}$ is the minimal rank of an operator $\mathcal{T}_N$ satisfying the perturbation bound (1.6), i.e., the truncated singular value decomposition certainly yields the best possible low-rank approximation with a given rank N .

Based on Fourier techniques, fast analytic singular value decompositions for linear operators arising in optical diffusion tomography have been constructed in [25] for problems with regular geometries and constant coefficients. In more general situations, the full assembly and decomposition of the operator $\mathcal{T}$ is, however, computationally prohibitive. Krylov subspace methods [16,35] and randomized algorithms [14,27] then provide alternatives that allow to construct approximate singular value decompositions using only a moderate number of evaluations of $\mathcal{T}$ and its adjoint $\mathcal{T}^*$. By

combining randomized singular value decompositions for subproblems associated to a single frequency in a recursive manner, approximate singular value decompositions for inverse medium problems have been constructed in [5].

In a recent work [24], motivated by [23] and [22], finite dimensional inverse scattering problems of the particular form

$$T(c) := V^\top D(c)U = M^\delta, \quad \text{with} \quad D(c) = \text{diag}(c) \quad (1.9)$$

are considered. Using the Khatri-Rao product $(A \odot B^\top)_{ij,l} = A_{i,l}B_{j,l}$ for matrices $A, B \in \mathbb{R}^{k \times m}$ with $ij = (k-1)i + j$ this problem can be cast into a linear system

$$(U^\top \odot V)c = \text{vec}(M^\delta), \quad (1.10)$$

where $\text{vec}(M) \in \mathbb{R}^{k^2}$ denotes the vectorization of the matrix M by columns. The Khatri-Rao product structure allows the efficient evaluation of $T^\top T$, required for the solution of the inverse problem, using pre-computed low-rank approximations for $UU^\top$ and $VV^\top$; we refer to [20] for a definition and properties of the Khatri-Rao product and a survey on tensor decompositions. Apart from the more restrictive assumptions on the problem structure, the computational cost of the reconstruction algorithms in [24] is still rather high, since the dimension m of the parameter c still appears in the system, which may be prohibitive for problems with distributed parameters.

Another popular strategy towards dimension reduction for inverse problems with multiple excitations consists in synthetically reducing the number of sources. Such *simultaneous* or *encoded sources* have been used, e.g., in geophysics [15,21] and tomography [37]; see [33] for further references. The systematic construction of low-rank approximations is investigated intensively also in the context of model order reduction; see [3,32] for a survey on results in this direction.

Let us note that in the context of inverse problems, the mapping $\mathcal{T}$ here has to be understood as a linear operator, i.e., a tensor of order 2, and the norms in which the approximation quality should be measured are prescribed by the functional-analytic setting; see (1.8). By the Eckhard-Young-Mirsky theorem, the truncated singular value decomposition therefore yields the optimal low-rank approximation, and the main aspect here is to compute the truncated singular value decomposition of the operator $\mathcal{T}$, or a sufficiently good approximation, with minimal effort. As we will see, this can be done on a rather general level, only using the particular form (1.2) and some abstract smoothness conditions on the involved operators.

1.4 Contributions and outline of the paper

The main scope of this paper is the systematic construction and analysis of low-rank approximations $\mathcal{T}_N$ for operators $\mathcal{T}$ of the particular form (1.2) with

- *Certified* approximation error bounds (1.6), and
- *Quasi-optimal* rank N comparable to that of the truncated singular value decomposition.

The stable solution of the inverse problem (1.1) can then be achieved by (1.7) in a highly efficient manner. The tensor-product structure (1.2) will further allow us to

- Set up $\mathcal{T}_N$ at essentially the same cost as a *single* evaluation of $\mathcal{T}(c)$;
- Compress the data $\mathcal{M}^\delta$ on the fly already during recording.

Before diving into the detailed analysis of our approach, let us briefly highlight the main underlying principles and key steps of the construction.

1.4.1 Sparse tensor product compression

Let $\mathcal{Q}_{K,\mathcal{U}}$, $\mathcal{Q}_{K,\mathcal{V}}$ denote orthogonal projections of rank K in the space $\mathbb{Y}$ of sources and the space $\mathbb{Z}$ of detectors, respectively. If $\dim(\mathbb{Y}) = \dim(\mathbb{Z}) = k$, then clearly $K \leq k$. One may also use different ranks $K_{\mathcal{U}}$, $K_{\mathcal{V}}$ for the two approximations, but for ease of notation we take $K = K_{\mathcal{U}} = K_{\mathcal{V}}$. In the language of [15,21], the columns of the operators $\mathcal{Q}_{K,\mathcal{U}}$ and $\mathcal{Q}_{K,\mathcal{V}}$ amount to *optimal sources* and *detectors*, and the appropriate choice of $\mathcal{Q}_{K,\mathcal{U}}$ and $\mathcal{Q}_{K,\mathcal{V}}$ is also related to *optimal experiment design* [31]. We define corresponding approximations

$$\mathcal{U}_K = \mathcal{U}\mathcal{Q}_{K,\mathcal{U}} \quad \text{and} \quad \mathcal{V}_K = \mathcal{V}\mathcal{Q}_{K,\mathcal{V}}$$

for the operators $\mathcal{U} : \mathbb{Y} \rightarrow \mathbb{U}$ and $\mathcal{V} : \mathbb{Z} \rightarrow \mathbb{V}$, each of rank K . Since $\mathcal{U}$ and $\mathcal{V}$ are compact operators, we may choose K large enough such that the error in these approximations is as small as desired. The resulting tensor product approximation

$$\mathcal{T}_{K,K}(c) = \mathcal{V}'_K \mathcal{D}(c) \mathcal{U}_K$$

may then be used as an approximation for $\mathcal{T}$. Following our notation, we may write $\mathcal{T}_{K,K} = \mathcal{Q}_{K,K} \mathcal{T}$, with $\mathcal{Q}_{K,K}$ denoting a tensor-product projection in data space. Unfortunately, the rank of $\mathcal{T}_{K,K}$ is in general K^2 , which turns out to be typically much larger than the optimal rank achievable by truncated singular value decomposition of the same accuracy. Instead of $\mathcal{T}_{K,K}$, we therefore consider a hyperbolic-cross approximation [7], which has the general form

$$\mathcal{T}_{\widehat{K}} = \mathcal{Q}_{\widehat{K}} \mathcal{T}_{K,K} = \mathcal{Q}_{\widehat{K}} \mathcal{T},$$

with an orthogonal projection $\mathcal{Q}_{\widehat{K}}$ onto the $\widehat{K}$ most significant components in the range of $\mathcal{T}_{K,K}$. We will show in detail how to construct sparse-tensor product approximations $\mathcal{T}_{\widehat{K}}$ of any desired accuracy, using knowledge about $\mathcal{U}_K$, $\mathcal{V}_K$ and $\mathcal{D}(c)$ only. Moreover, under some mild conditions on the operators $\mathcal{U}$, $\mathcal{V}$, we will see that $\widehat{K} \approx K$ is sufficient to guarantee essentially the same accuracy as the full tensor-product approximation. Let us further note that $\mathcal{T}_{\widehat{K}}$ does not have a tensor-product structure, but is formally based on a tensor-product approximation, which turns out to be advantageous when computing the projected data $\mathcal{M}_{\widehat{K}}^\delta = \mathcal{Q}_{\widehat{K}} \mathcal{M}^\delta$; see below.

1.4.2 Recompression

By truncated singular value decomposition, we can further reduce the rank of the sparse-tensor product approximation $\mathcal{T}_{\hat{K}}$ leading to a final approximation

$$\mathcal{T}_N = \mathcal{Q}_N \mathcal{T} \quad \text{with} \quad \mathcal{Q}_N = \mathcal{P}_N \mathcal{Q}_{\hat{K}} = \mathcal{P}_N \mathcal{Q}_{K,K}, \quad (1.11)$$

which can be shown to have essentially the rank $N \approx N^{\text{svd}}$ of the truncated singular value decomposition of $\mathcal{T}$ with the same accuracy; see Subsect. 2.4 for details. We thus obtain computable approximations $\mathcal{T}_N$ for $\mathcal{T}$ with essentially the same rank as the truncated singular value decomposition of similar accuracy. Only some mild assumptions on the mapping properties of the operators $\mathcal{U}$ and $\mathcal{V}$ are required to rigorously establish and guarantee the approximation property (1.6).

1.4.3 Summary of basic properties

It turns out that the proposed two-step construction of the approximation $\mathcal{T}_N$, which is based on the underlying tensor-product structure of the problem, has significant advantages compared to the truncated singular value decomposition in the setup, i.e., $\mathcal{T}_N$ can be computed at the computational cost of essentially *one single evaluation of the forward operator* $\mathcal{T}(c)$. Moreover, the underlying tensor-product structure also allows to compute the projection $\mathcal{M}_N^\delta = \mathcal{Q}_N \mathcal{M}^\delta$ in an efficient manner. By construction of the projection $\mathcal{Q}_N = \mathcal{P}_N \mathcal{Q}_{K,K}$, we have $\mathcal{M}_N^\delta = \mathcal{P}_N \mathcal{M}_{K,K}^\delta$ with pre-compressed data

$$\mathcal{M}_{K,K}^\delta = \mathcal{Q}_{K,K} \mathcal{M}^\delta = (\mathcal{Q}'_{K,\mathcal{V}} \mathcal{M}^\delta) \mathcal{Q}_{K,\mathcal{U}}$$

that can be computed by two separate projections $\mathcal{Q}_{K,\mathcal{U}}$, $\mathcal{Q}_{K,\mathcal{V}}$ of rank K in the spaces $\mathbb{Y}$ and $\mathbb{Z}$ of sources and detectors. Since the projection $\mathcal{Q}'_{K,\mathcal{V}}$ can already be applied during recording, simultaneous access to the full data $\mathcal{M}^\delta$ is never required. As a consequence, the memory cost of data recording and compression is thereby reduced to $3Kk + K^2$.

1.4.4 Outline

The remainder of the manuscript is organized as follows: In Sect. 2, we discuss in detail the construction of quasi-optimal low-rank approximations $\mathcal{T}_N$ for problems of the form (1.2) with guaranteed accuracy (1.6). To illustrate the applicability of our theoretical results, we discuss in Sect. 3 a particular example stemming from fluorescence diffuse optical tomography. An appropriate choice of function spaces allows us to verify all conditions required for the analysis of our approach. In Sect. 4, we report in detail about numerical tests, in which we demonstrate the computational efficiency of the model reduction approach and the resulting numerical solution of the inverse problems.

2 Analysis of the model reduction approach

We will start with introducing our basic notation and then provide a complete analysis of the data compression and model reduction approach outlined in the introduction.

2.1 Notation

Function spaces will be denoted by $\mathbb{A}, \mathbb{B}, \dots$ and assumed to be separable Hilbert spaces with scalar product $(\cdot, \cdot)_{\mathbb{A}}$ and norm $\|\cdot\|_{\mathbb{A}}$. By $\mathbb{A}'$ we denote the dual of $\mathbb{A}$, i.e., the space of bounded linear functionals on $\mathbb{A}$, and by $\langle a', a \rangle_{\mathbb{A}' \times \mathbb{A}}$ the corresponding duality product. Furthermore, $\mathcal{L}(\mathbb{A}, \mathbb{B})$ denotes the Banach space of linear operators $\mathcal{S} : \mathbb{A} \rightarrow \mathbb{B}$ with norm $\|\mathcal{S}\|_{\mathcal{L}(\mathbb{A}, \mathbb{B})} = \sup_{\|a\|_{\mathbb{A}}=1} \|\mathcal{S}a\|_{\mathbb{B}} < \infty$. We write $\mathcal{R}(\mathcal{S}) = \{\mathcal{S}a : a \in \mathbb{A}\}$ for the range of the operator $\mathcal{S}$ and define $\text{rank}(\mathcal{S}) = \dim(\mathcal{R}(\mathcal{S}))$. By $\mathcal{S}' : \mathbb{B}' \rightarrow \mathbb{A}'$ and $\mathcal{S}^* : \mathbb{B} \rightarrow \mathbb{A}$ we denote the dual and the adjoint of a bounded linear operator $\mathcal{S} \in \mathcal{L}(\mathbb{A}, \mathbb{B})$ defined, respectively, for all $a \in \mathbb{A}$, $b \in \mathbb{B}$, and $b' \in \mathbb{B}'$ by

$$\langle \mathcal{S}'b', a \rangle_{\mathbb{A}' \times \mathbb{A}} = \langle b', \mathcal{S}a \rangle_{\mathbb{B}' \times \mathbb{B}} \quad \text{and} \quad (\mathcal{S}^*b, a)_{\mathbb{B}} = (b, \mathcal{S}a)_{\mathbb{A}}. \quad (2.1)$$

The two operators $\mathcal{S}'$ and $\mathcal{S}^*$ are directly related by Riesz-isomorphisms.

Any operator $\mathcal{S}^{\delta} : \mathbb{A} \rightarrow \mathbb{B}$ with $\|\mathcal{S} - \mathcal{S}^{\delta}\|_{\mathcal{L}(\mathbb{A}, \mathbb{B})} \lesssim \delta$ will be called a δ -approximation for $\mathcal{S}$ in the following. Let us recall that any compact linear operator $\mathcal{S} : \mathbb{A} \rightarrow \mathbb{B}$ has a singular value decomposition, i.e., a countable system $\{(\sigma_k, a_k, b_k)\}_{k \geq 1}$ such that

$$\mathcal{S}a = \sum_{k \geq 1} (a, a_k)_{\mathbb{A}} \sigma_k b_k, \quad (2.2)$$

with singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq 0$ and $\{a_k : \sigma_k > 0\}$ and $\{b_k : \sigma_k > 0\}$ denoting orthonormal basis for $\mathcal{R}(\mathcal{S}^*) \subset \mathbb{A}$ and $\mathcal{R}(\mathcal{S}) \subset \mathbb{B}$, respectively. Also note that $\|\mathcal{S}\|_{\mathcal{L}(\mathbb{A}, \mathbb{B})} = \sigma_1$ and $\text{rank}(\mathcal{S}) = \sup\{k : \sigma_k > 0\}$. Moreover, by the Courant-Fisher min-max principle [12], also known as the Eckart-Young-Mirsky theorem, the k th singular value can be characterized by

$$\sigma_k = \min_{\mathbb{A}_{k-1}} \max_{a \in \mathbb{A}_{k-1}^{\perp}} \|\mathcal{S}a\|_{\mathbb{B}} / \|a\|_{\mathbb{A}}, \quad (2.3)$$

where $\mathbb{A}_{k-1}$ are the $(k-1)$ -dimensional subspaces of $\mathbb{A}$. Hence every linear compact operator $\mathcal{S} : \mathbb{A} \rightarrow \mathbb{B}$ can be approximated by truncated singular value decompositions

$$\mathcal{S}_K a = \sum_{k \leq K} (a, a_k)_{\mathbb{A}} \sigma_k b_k, \quad (2.4)$$

with error $\|\mathcal{S} - \mathcal{S}_K\|_{\mathcal{L}(\mathbb{A}, \mathbb{B})} = \sigma_{K+1}$ and the truncated singular value decomposition can be used to construct δ -approximations of minimal rank.

We further denote by $\text{HS}(\mathbb{A}, \mathbb{B}) \subset \mathcal{L}(\mathbb{A}, \mathbb{B})$ the Hilbert-Schmidt class of compact linear operators whose singular values are square summable. Note that $\text{HS}(\mathbb{A}, \mathbb{B})$ is a Hilbert space equipped with the scalar product $(\mathcal{S}, \mathcal{R})_{\text{HS}(\mathbb{A}, \mathbb{B})} = \sum_{k \geq 1} (\mathcal{S}a_k, \mathcal{R}a_k)_{\mathbb{B}}$,

where $\{a_k\}_{k \geq 1}$ is an orthonormal basis of $\mathbb{A}$. Moreover, the scalar product and the associated norm are independent of the choice of this basis. Let us mention the following elementary results, which will be used several times later on.

Lemma 2.1 (a) Let $S \in \mathcal{HS}(\mathbb{A}, \mathbb{B})$. Then there exists a sequence $\{S_K\}_{K \in \mathbb{N}}$ of linear operators of rank K , such that $\|S - S_K\|_{\mathcal{L}(\mathbb{A}, \mathbb{B})} \lesssim K^{-1/2}$.

(b) Let $S : \mathbb{A} \rightarrow \mathbb{B}$, $\mathcal{R} : \mathbb{B} \rightarrow \mathbb{C}$ be two linear bounded operators and at least one of them Hilbert-Schmidt. Then the composition $\mathcal{R}S : \mathbb{A} \rightarrow \mathbb{C}$ is Hilbert-Schmidt and

$$\begin{aligned}\|\mathcal{R}S\|_{\mathcal{HS}(\mathbb{A}, \mathbb{C})} &\leq \|\mathcal{R}\|_{\mathcal{L}(\mathbb{B}, \mathbb{C})} \|S\|_{\mathcal{HS}(\mathbb{A}, \mathbb{B})}, \quad \text{or} \\ \|\mathcal{R}S\|_{\mathcal{HS}(\mathbb{A}, \mathbb{C})} &\leq \|\mathcal{R}\|_{\mathcal{HS}(\mathbb{B}, \mathbb{C})} \|S\|_{\mathcal{L}(\mathbb{A}, \mathbb{B})}.\end{aligned}$$

Here and below, we use $a \lesssim b$ to express $a \leq Cb$ with some generic constant C that is independent of the relevant context, and we write $a \simeq b$ when $a \lesssim b$ and $b \lesssim a$.

For convenience of the reader, we provide a short proof of these assertions.

Proof The assumption $S \in \mathcal{HS}(\mathbb{A}, \mathbb{B})$ implies that S is compact with square summable singular values, and hence $\sigma_{K,S} \lesssim K^{-1/2}$. The truncated singular value decomposition S_K then satisfies $\|S - S_K\|_{\mathcal{L}(\mathbb{A}, \mathbb{B})} \leq \sigma_{K,S} \lesssim K^{-1/2}$ which yields (a). After choosing an orthonormal basis $\{a_k\}_{k \geq 1} \subset \mathbb{A}$, we can write

$$\begin{aligned}\|\mathcal{R}S\|_{\mathcal{HS}(\mathbb{A}, \mathbb{C})}^2 &= \sum_k \|\mathcal{R}Sa_k\|_{\mathbb{C}}^2 \\ &\leq \|\mathcal{R}\|_{\mathcal{L}(\mathbb{B}, \mathbb{C})}^2 \sum_{k \geq 1} \|Sa_k\|_{\mathbb{B}}^2 = \|\mathcal{R}\|_{\mathcal{L}(\mathbb{B}, \mathbb{C})}^2 \|S\|_{\mathcal{HS}(\mathbb{A}, \mathbb{B})}^2\end{aligned}$$

which implies the first inequality of assertion (b). The second inequality follows from the same arguments applied to the adjoint $(\mathcal{R}S)^* = S^* \mathcal{R}^*$ and noting that the respective norms of an operator and its adjoint are the same. $\square$

2.2 Preliminaries and basic assumptions

We now introduce in more detail the functional analytic setting for the inverse problem (1.1) used for our considerations. We assume that the operators $\mathcal{U}$, $\mathcal{V}$, $\mathcal{D}$ appearing in definition (1.2) satisfy

Assumption 2.2 Let $\mathcal{U} \in \mathcal{HS}(\mathbb{Y}, \mathbb{U})$, $\mathcal{V} \in \mathcal{HS}(\mathbb{Z}, \mathbb{V})$, and $\mathcal{D} \in \mathcal{L}(\mathbb{X}, \mathcal{L}(\mathbb{U}, \mathbb{V}'))$.

Following our convention, all function spaces appearing in these conditions, except the space $\mathcal{L}(\cdot, \cdot)$, are separable Hilbert spaces. We can now prove the following assertions.

Lemma 2.3 Let Assumption 2.2 be valid. Then $\mathcal{T}(c) = \mathcal{V}' \mathcal{D}(c) \mathcal{U}$ defines a bounded linear operator $\mathcal{T} : \mathbb{X} \rightarrow \mathcal{HS}(\mathbb{Y}, \mathbb{Z}')$ and, additionally, $\mathcal{T}$ is compact.

Proof Linearity of $\mathcal{T}$ is clear by construction and the linearity of $\mathcal{U}$, $\mathcal{V}$, and $\mathcal{D}$. Now let $\{y_k\}_{k \geq 1}$ denote an orthonormal basis of $\mathbb{Y}$ and let $c \in \mathbb{X}$ be arbitrary. Then

$$\begin{aligned} \|\mathcal{T}(c)\|_{\mathbb{HS}(\mathbb{Y}, \mathbb{Z}')}^2 &= \sum_{k \geq 1} \|\mathcal{V}' \mathcal{D}(c) \mathcal{U} y_k\|_{\mathbb{Z}'}^2 \leq \|\mathcal{V}' \mathcal{D}(c)\|_{\mathcal{L}(\mathbb{U}, \mathbb{Z}')}^2 \sum_{k \geq 1} \|\mathcal{U} y_k\|_{\mathbb{U}}^2 \\ &\leq \|\mathcal{V}'\|_{\mathcal{L}(\mathbb{V}', \mathbb{Z}')}^2 \|\mathcal{D}\|_{\mathcal{L}(\mathbb{X} \rightarrow \mathcal{L}(\mathbb{U}, \mathbb{V}'))}^2 \|c\|_{\mathbb{X}}^2 \|\mathcal{U}\|_{\mathbb{HS}(\mathbb{Y}, \mathbb{U})}^2, \end{aligned}$$

where we used Lemma 2.1 in the second step, and the boundedness of the operators in the third. Since $\|\mathcal{V}'\|_{\mathcal{L}(\mathbb{V}', \mathbb{Z}')} = \|\mathcal{V}\|_{\mathcal{L}(\mathbb{Z}, \mathbb{V})} \leq \|\mathcal{V}\|_{\mathbb{HS}(\mathbb{Z}, \mathbb{V})}$, we obtain

$$\|\mathcal{T}(c)\|_{\mathbb{HS}(\mathbb{Y}, \mathbb{Z}')} \leq \|\mathcal{U}\|_{\mathbb{HS}(\mathbb{Y}, \mathbb{U})} \|\mathcal{V}\|_{\mathbb{HS}(\mathbb{Z}, \mathbb{V})} \|\mathcal{D}\|_{\mathcal{L}(\mathbb{X}, \mathcal{L}(\mathbb{U}, \mathbb{V}'))} \|c\|_{\mathbb{X}}$$

for all $c \in \mathbb{X}$, which shows that $\mathcal{T}$ is bounded. Using Lemma 2.1(a), we can further approximate $\mathcal{U}$ and $\mathcal{V}$ by operators $\mathcal{U}_K$, $\mathcal{V}_K$ of rank K , such that

$$\|\mathcal{U} - \mathcal{U}_K\|_{\mathcal{L}(\mathbb{Y}, \mathbb{U})} \lesssim K^{-1/2} \quad \text{and} \quad \|\mathcal{V} - \mathcal{V}_K\|_{\mathcal{L}(\mathbb{Z}, \mathbb{V})} \lesssim K^{-1/2}, \quad (2.5)$$

and we can define an operator $\mathcal{T}_{K,K} : \mathbb{X} \rightarrow \mathbb{HS}(\mathbb{Y}, \mathbb{Z}')$ by $\mathcal{T}_{K,K}(c) = \mathcal{V}'_K \mathcal{D}(c) \mathcal{U}_K$, which defines an approximation of $\mathcal{T}$ of rank K^2 . From Lemma 2.1(b), we infer that

$$\begin{aligned} \|\mathcal{T} - \mathcal{T}_{K,K}\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))} &= \sup_{\|c\|_{\mathbb{X}}=1} \|\mathcal{V}' \mathcal{D}(c) \mathcal{U} - \mathcal{V}'_K \mathcal{D}(c) \mathcal{U}_K\|_{\mathbb{HS}(\mathbb{Y}, \mathbb{Z}')} \\ &\leq (\|\mathcal{V}' - \mathcal{V}'_K\|_{\mathcal{L}(\mathbb{V}', \mathbb{Z}')} \|\mathcal{U}\|_{\mathbb{HS}(\mathbb{Y}, \mathbb{U})} \\ &\quad + \|\mathcal{V}'\|_{\mathbb{HS}(\mathbb{V}', \mathbb{Z}')} \|\mathcal{U} - \mathcal{U}_K\|_{\mathcal{L}(\mathbb{Y}, \mathbb{U})}) \|\mathcal{D}\|_{\mathcal{L}(\mathbb{X}, \mathcal{L}(\mathbb{U}, \mathbb{V}'))}. \end{aligned}$$

Using Assumption 2.2 and the bounds (2.5), we thus conclude that $\mathcal{T}$ can be approximated uniformly by finite-rank operators, and hence $\mathcal{T}$ is compact. $\square$

2.3 Sparse tensor product approximation

As direct consequence of the arguments used in the previous result, we obtain the following preliminary approximation result.

Lemma 2.4 *Let Assumption 2.2 hold. Then for any $\delta > 0$ there exists $K \in \mathbb{N}$ with $K \lesssim \delta^{-2}$ and rank K approximations $\mathcal{U}_K = \mathcal{U} \mathcal{Q}_{K,\mathcal{U}}$ and $\mathcal{V}_K = \mathcal{V} \mathcal{Q}_{K,\mathcal{V}}$ such that*

$$\|\mathcal{U} - \mathcal{U}_K\|_{\mathcal{L}(\mathbb{Y}, \mathbb{U})} \leq \delta \quad \text{and} \quad \|\mathcal{V} - \mathcal{V}_K\|_{\mathcal{L}(\mathbb{Z}, \mathbb{V})} \leq \delta. \quad (2.6)$$

Here, $\mathcal{Q}_{K,\mathcal{U}}$ and $\mathcal{Q}_{K,\mathcal{V}}$ are orthogonal projections on $\mathbb{Y}$ and $\mathbb{Z}$, respectively. Furthermore, the operator $\mathcal{T}_{K,K}$ defined by $\mathcal{T}_{K,K}(c) = \mathcal{V}'_K \mathcal{D}(c) \mathcal{U}_K$ has rank K^2 and satisfies

$$\|\mathcal{T} - \mathcal{T}_{K,K}\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))} \lesssim \delta. \quad (2.7)$$

If the singular values of $\mathcal{U}$ and $\mathcal{V}$ satisfy $\sigma_{k,\mathcal{U}}, \sigma_{k,\mathcal{V}} \lesssim k^{-\alpha}$ for some $\alpha > 1/2$, then the assertions hold with $K \simeq \delta^{-1/\alpha}$, and consequently $\text{rank}(\mathcal{T}_{K,K}) \lesssim \delta^{-2/\alpha}$.

Note that different ranks $K_{\mathcal{U}}$, $K_{\mathcal{V}}$ could be chosen for the approximations of $\mathcal{U}$ and $\mathcal{V}$, but for ease of notation, we assume $K_{\mathcal{U}} = K_{\mathcal{V}} = K$. Since $\mathcal{U}$ and $\mathcal{V}$ are Hilbert-Schmidt, we know that $\sigma_{k,\mathcal{U}}, \sigma_{k,\mathcal{V}} \lesssim k^{-\alpha}$ with $\alpha \geq 1/2$, and, thus, the decay assumption on the singular values are not very restrictive. In general, $\text{rank } \mathcal{T}_{K,K} = K^2$ may however be substantially larger than the optimal rank N^{svd} of the truncated singular value decomposition satisfying a similar perturbation bound. We now show that based on the approximations $\mathcal{U}_K, \mathcal{V}_K$, and assuming sufficient decay of the singular values $\sigma_{k,\mathcal{U}}, \sigma_{k,\mathcal{V}}$, one can construct a δ -approximation $\mathcal{T}_{\widehat{K}}$ for $\mathcal{T}$ of rank $\widehat{K} \ll K^2$.

Lemma 2.5 *Let $\sigma_{k,\mathcal{U}} \lesssim k^{-\beta}$ and $\sigma_{k,\mathcal{V}} \lesssim k^{-\alpha}$ (or $\sigma_{k,\mathcal{U}} \lesssim k^{-\alpha}$ and $\sigma_{k,\mathcal{V}} \lesssim k^{-\beta}$) for some $\beta > 1/2$ and $\alpha > \beta + 1/2$. Then $\sigma_{k,\mathcal{T}} \lesssim k^{-\beta}$, and for any $\delta > 0$, we can find $\widehat{K} \in \mathbb{N}$ with $\widehat{K} \lesssim \delta^{-1/\beta}$ and an approximation $\mathcal{T}_{\widehat{K}} = \mathcal{Q}_{\widehat{K}} \mathcal{T}$ of rank $\widehat{K}$, such that*

$$\|\mathcal{T} - \mathcal{T}_{\widehat{K}}\|_{\mathcal{L}(\mathbb{X}, \text{HS}(\mathbb{V}, \mathbb{Z}'))} \lesssim \delta. \quad (2.8)$$

Proof Let $\{\sigma_{k,*}, a_{k,*}, b_{k,*}\}$ denote the singular systems for $\mathcal{U}$ and $\mathcal{V}'$, respectively. We now show that the hyperbolic cross approximation [7]

$$\mathcal{T}_{\widehat{K}}(c) = \sum_{k \geq 1} \sum_{\ell=1}^{L_k} \sigma_{\ell,\mathcal{U}} \sigma_{k,\mathcal{V}'}(\cdot, a_{\ell,\mathcal{U}})_{\mathbb{Y}} \langle \mathcal{D}(c) b_{\ell,\mathcal{U}}, a_{k,\mathcal{V}'} \rangle_{\mathbb{V}' \times \mathbb{V}} b_{k,\mathcal{V}'},$$

with the choice $L_k = \lfloor \widehat{K}/k^{1+\epsilon} \rfloor$, $\widehat{K} \simeq \delta^{-1/\beta}$, and $\epsilon = (\alpha - \beta - 1/2)/(2\beta) > 0$ has the required properties. By counting, one can verify that $\text{rank}(\mathcal{T}_{\widehat{K}}) \lesssim \sum_{k \geq 1} L_k \lesssim \widehat{K}$, since by construction $L_k \simeq \widehat{K}/k^{1+\epsilon}$ is summable. Furthermore, we can bound

$$\begin{aligned} \|\mathcal{T}(c) - \mathcal{T}_{\widehat{K}}(c)\|_{\text{HS}(\mathbb{U}, \mathbb{V}')}^2 &= \sum_{m \geq 1} \|(\mathcal{T}(c) - \mathcal{T}_{\widehat{K}}(c)) a_{m,\mathcal{U}}\|_{\mathbb{V}'}^2 \\ &= \sum_{k \geq 1} \sigma_{k,\mathcal{V}'}^2 \left| \sum_{\ell \geq L_k+1} \sigma_{\ell,\mathcal{U}} \langle \mathcal{D}(c) b_{\ell,\mathcal{U}}, a_{k,\mathcal{V}'} \rangle_{\mathbb{V}' \times \mathbb{V}} \right|^2 \\ &\leq \sum_{k \geq 1} \sigma_{k,\mathcal{V}'}^2 \sigma_{L_k}^2 \|\mathcal{D}(c)\|_{\mathcal{L}(\mathbb{U}, \mathbb{V})}^2 \|a_{k,\mathcal{V}'}\|_{\mathbb{V}'}^2. \end{aligned}$$

By observing that $\|a_{k,\mathcal{V}'}\|_{\mathbb{V}'} = 1$, $\|\mathcal{D}(c)\|_{\mathcal{L}(\mathbb{U}, \mathbb{V})} \lesssim \|c\|_{\mathbb{X}}$, and $\sigma_{k,\mathcal{V}'} = \sigma_{k,\mathcal{V}}$ and by using the decay properties of the singular values, we obtain

$$\|\mathcal{T} - \mathcal{T}_{\widehat{K}}\|_{\mathcal{L}(\mathbb{X}, \text{HS}(\mathbb{U}, \mathbb{V}'))}^2 \lesssim \sum_{k \geq 1} k^{-2\alpha+2\beta(1+\epsilon)} \widehat{K}^{-2\beta} \lesssim \delta^2.$$

In the last step, we used the fact that $-2\alpha + 2\beta(1 + \epsilon) < -1$ and $\widehat{K} \simeq \delta^{-1/\beta}$, which follows immediately from the construction. $\square$

Remark 2.6 Comparing the results of Lemma 2.4 and 2.5, we expect to obtain a tensor product approximation $\mathcal{T}_{K,K}$ of rank $K^2 \simeq \delta^{-2/\alpha}$ while the hyperbolic cross approximation $\mathcal{T}_{\widehat{K}}$ and consequently also the truncated singular value decomposition of the same accuracy only have rank $\widehat{K} \lesssim \delta^{-1/(\alpha-1/2-\epsilon)}$, with $\epsilon = \alpha - \beta - 1/2 > 0$, which may be substantially smaller for $\alpha > 1$. Note that, like $\mathcal{T}_{K,K}$, the sparse tensor product approximation $\mathcal{T}_{\widehat{K}}$ can be constructed directly from the low-rank approximations $\mathcal{U}_K$ and $\mathcal{V}_K$.

2.4 Quasi-optimal low-rank approximation

Let $P_{N^{\text{svd}}}$ denote the orthogonal projection onto the space spanned by the left singular vectors of $\mathcal{T}$ corresponding to the first N^{svd} singular values, and let N^{svd} be chosen such that

$$\|P_{N^{\text{svd}}} \mathcal{T} - \mathcal{T}\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))} \leq \delta.$$

We now show that a compression of any δ -approximation $\mathcal{T}^\delta$, e.g. the tensor product approximation $\mathcal{T}_{K,K}$ or its hyperbolic-cross approximation $\mathcal{T}_{\hat{K}}$, allows to construct another δ -approximation $\mathcal{T}_{N^\delta}^\delta = P_{N^\delta}^\delta \mathcal{T}^\delta$ for $\mathcal{T}$ with quasi-optimal rank $N^\delta \leq N^{\text{svd}}$.

Lemma 2.7 *Let $\delta > 0$ and let $\mathcal{T}^\delta : \mathbb{X} \rightarrow \mathbb{HS}(\mathbb{Y}, \mathbb{Z}')$ be a linear compact operator such that $\|\mathcal{T}^\delta - \mathcal{T}\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))} \leq C\delta$ for some $C > 0$. Let $\mathcal{P}_{N^\delta}^\delta \mathcal{T}^\delta$ denote the truncated singular value decomposition of $\mathcal{T}^\delta$ with minimal rank N^δ such that*

$$\|\mathcal{T}^\delta - \mathcal{P}_{N^\delta}^\delta \mathcal{T}^\delta\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))} \leq (C + 1)\delta. \quad (2.9)$$

Then $N^\delta \leq N^{\text{svd}}$ and

$$\|\mathcal{T} - \mathcal{P}_{N^\delta}^\delta \mathcal{T}^\delta\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))} \leq (2C + 1)\delta,$$

i.e., $\mathcal{P}_{N^\delta}^\delta \mathcal{T}^\delta$ is a δ -approximation for $\mathcal{T}$ with quasi-optimal rank.

Proof *Step 1.* We start by recalling a well-known perturbation result for singular values [19], i.e., we show that for each $k \in \mathbb{N}$ one has

$$\sigma_k - C\delta \leq \sigma_k^\delta \leq \sigma_k + C\delta, \quad (2.10)$$

where $\{\sigma_k\}$ and $\{\sigma_k^\delta\}$ denote the singular values of $\mathcal{T}$ and $\mathcal{T}^\delta$, respectively. Let us abbreviate $\|\cdot\| = \|\cdot\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))}$, choose $\varepsilon > 0$, and let $\mathcal{P}_M^{\text{svd}} \mathcal{T}$ denote the truncated singular value decomposition of $\mathcal{T}$ with optimal rank such that $\|\mathcal{P}_M^{\text{svd}} \mathcal{T} - \mathcal{T}\| < \varepsilon$. Using the optimality of M and the non-expansiveness of the projection, we estimate

$$\begin{aligned} \|(\mathcal{I} - \mathcal{P}_M^\delta) \mathcal{T}^\delta\| &\leq \|(\mathcal{I} - \mathcal{P}_M^{\text{svd}}) \mathcal{T}^\delta\| \\ &\leq \|(\mathcal{I} - \mathcal{P}_M^{\text{svd}}) \mathcal{T}\| + \|(\mathcal{I} - \mathcal{P}_M^{\text{svd}})(\mathcal{T} - \mathcal{T}^\delta)\| \leq \varepsilon + C\delta. \end{aligned}$$

For $\varepsilon = \sigma_{k+1}$, we have $M = M(\varepsilon) = k$, and we conclude that $\|(\mathcal{I} - \mathcal{P}_M^\delta) \mathcal{T}^\delta\| = \sigma_{M+1}^\delta$ and $\sigma_{k+1}^\delta \leq \sigma_{k+1} + C\delta$. The second inequality in (2.10) follows by considering $\mathcal{T}$ as a δ -approximation for the operator $\mathcal{T}^\delta$.

Step 2. Now let N^δ be as in the statement of the lemma and $N^{\text{svd}} = M(\delta)$ as defined in Step 1. Then (2.10) implies that $\sigma_{N^{\text{svd}}+1}^\delta \leq \sigma_{N^{\text{svd}}+1} + C\delta \leq (C + 1)\delta$. Optimality of N^δ then implies $\sigma_{N^{\text{svd}}+1}^\delta \leq \sigma_{N^\delta+1}^\delta \leq (C + 1)\delta < \sigma_{N^\delta}^\delta$, and from the monotonicity of the singular values, we conclude that $N^\delta \leq N^{\text{svd}}$. Furthermore,

$$\|\mathcal{P}_{N^\delta}^\delta \mathcal{T}^\delta - \mathcal{T}\| \leq \|\mathcal{P}_{N^\delta}^\delta \mathcal{T}^\delta - \mathcal{T}^\delta\| + \|\mathcal{T}^\delta - \mathcal{T}\| \leq (2C + 1)\delta,$$

i.e., $\mathcal{P}_{N^\delta}^\delta \mathcal{T}^\delta$ is a δ -approximation for $\mathcal{T}$ with quasi-optimal rank $N^\delta \leq N^{\text{svd}}$. $\square$

We now apply the previous lemma with $\mathcal{T}^\delta = \mathcal{T}_{\hat{K}}$ and $\mathcal{P}_N^\delta = \mathcal{P}_N$ the projection of the corresponding truncated singular value decomposition, which leads to

$$\mathcal{T}_N = \mathcal{P}_N \mathcal{T}_{\hat{K}} = \mathcal{P}_N (\mathcal{Q}_{\hat{K}} \mathcal{T}). \quad (2.11)$$

By construction this is a δ -approximation for $\mathcal{T}$ of quasi-optimal rank $N \leq N^{\text{svd}}$.

2.5 Summary

Let us briefly summarize the main observations and results of this section. Based on δ -approximations $\mathcal{U}_K$, $\mathcal{V}_K$ for the operators $\mathcal{U}$ and $\mathcal{V}$, we can construct a low-rank δ -approximation $\mathcal{T}_{\hat{K}} = \mathcal{Q}_{\hat{K}} \mathcal{T}$ for $\mathcal{T}$ via hyperbolic-cross approximation. We observed that in typical situations $\mathcal{T}_{\hat{K}}$ is of much lower rank than the corresponding full tensor product approximations $\mathcal{T}_{K,K}$, whose assembly can be completely avoided. By truncated singular value decomposition of $\mathcal{T}_{\hat{K}}$, we obtained another δ -approximation $\mathcal{T}_N = \mathcal{P}_N \mathcal{T}_{\hat{K}} = \mathcal{P}_N (\mathcal{Q}_{\hat{K}} \mathcal{T})$ with quasi-optimal rank $N \leq N^{\text{svd}}$. In summary, we thus efficiently computed a low-rank approximation $\mathcal{T}_N$ for $\mathcal{T}$ with similar rank and approximation properties as the truncated singular value decomposition.

The analysis in this section was done in abstract function spaces and applies verbatim to infinite-dimensional operators as well as to their finite-dimensional (truth) approximations obtained after discretization. As a consequence, the computational results, e.g., the ranks K and N of the approximations, can be expected to be essentially independent of the actual truth approximation used for computations. In the language of model-reduction, the low-rank approximation $\mathcal{T}_N$ is a certified reduced order model.

3 Fluorescence optical tomography

In order to illustrate the viability of the theoretical results derived in the previous section, we now consider in some detail a typical application arising in medical imaging.

3.1 Model equations

Fluorescence optical tomography aims at retrieving information about the concentration c of a fluorophore inside an object by illuminating this object from outside with near infrared light and measuring the light reemitted by the fluorophores at a different wavelength. The distribution $u_x = u_x(q_x)$ of the light intensity inside the object generated by a source q_x at the boundary is described by

$$-\nabla \cdot (\kappa_x \nabla u_x) + \mu_x u_x = 0, \quad \text{in } \Omega, \quad (3.1)$$

$$\kappa_x \partial_n u_x + \rho_x u_x = q_x, \quad \text{on } \partial\Omega. \quad (3.2)$$

We assume that $\Omega \subset \mathbb{R}^d$, $d = 2, 3$, is a bounded domain with smooth boundary enclosing the object under consideration. The light intensity $u_m = u_m(u_x, c)$ emitted by the fluorophores is described by a similar equation

$$-\nabla \cdot (\kappa_m \nabla u_m) + \mu_m u_m = c u_x, \quad \text{in } \Omega, \quad (3.3)$$

$$\kappa_m \partial_n u_m + \rho_m u_m = 0, \quad \text{on } \partial\Omega. \quad (3.4)$$

The model parameters κ_i , μ_i , and ρ_i , $i = x, m$, characterize the optical properties of the medium at excitation and emission wavelength; we assume these parameters to be known, e.g., determined by independent measurements [1]. As shown in [8], the above linear model, which can be interpreted as a Born approximation or linearization, is a valid approximation for moderate fluorophore concentrations.

3.2 Forward operator

The forward problem in fluorescence optical tomography models an experiment in which the emitted light resulting from excitation with a known source and after interaction with a given fluorophore concentration is measured at the boundary. The measurable quantity is the outward photon flux, which is proportional to u_m ; see [1] for details. The potential data for a single excitation with source q_x measured by a detector with characteristic q_m can be described by

$$\langle \mathcal{T}(c) q_x, q_m \rangle = \int_{\partial\Omega} u_m q_m ds(x), \quad (3.5)$$

where u_m and u_x are determined by the boundary value problems (3.1)–(3.4). The inverse problem finally consists of determining the concentration c of the fluorophore marker from measurements $\langle \mathcal{T}(c) q_x, q_m \rangle$ for multiple excitations q_x and detectors q_m .

We now illustrate that fluorescence optical tomography perfectly fits into the abstract setting of Section 2. Let us begin with defining the excitation operator

$$\mathcal{U} : H^1(\partial\Omega) \rightarrow H^1(\Omega), \quad q_x \mapsto \mathcal{U}q_x := u_x, \quad (3.6)$$

which maps a source q_x to the corresponding weak solution u_x of (3.1)–(3.2). The interaction with the fluorophore can be described by the multiplication operator

$$\mathcal{D} : L^2(\Omega) \rightarrow \mathcal{L}(H^1(\Omega), H^1(\Omega)'), \quad \mathcal{D}(c)u = cu. \quad (3.7)$$

In dimension $d \leq 3$, the product cu of two functions $c \in L^2(\Omega)$ and $u \in H^1(\Omega)$, lies in $L^{3/2}(\Omega)$ and can thus be interpreted as a bounded linear functional on $H^1(\Omega)$; this shows that $\mathcal{D}$ is a bounded linear operator. We further introduce the linear operator

$$\mathcal{V} : H^1(\partial\Omega) \rightarrow H^1(\Omega), \quad q_m \mapsto \mathcal{V}q_m := v_m, \quad (3.8)$$

which maps q_m to the weak solution v_m of the adjoint emission problem

$$-\nabla \cdot (\kappa_m \nabla v_m) + \mu_m v_m = 0, \quad \text{in } \Omega, \quad (3.9)$$

$$\kappa_m \partial_n v_m + \rho_m v_m = q_m, \quad \text{on } \partial\Omega. \quad (3.10)$$

One can verify that $\mathcal{V}$ is the dual of the solution operator $u_m|_{\partial\Omega} = \mathcal{V}'\mathcal{D}(c)u_x$ of the system (3.3)–(3.4); see [8] for details. Hence we may express the forward operator as

$$\mathcal{T}(c) = \mathcal{V}'\mathcal{D}(c)\mathcal{U}. \quad (3.11)$$

As function spaces we choose $\mathbb{U} = \mathbb{V} = H^1(\Omega)$, $\mathbb{Y} = \mathbb{Z} = H^1(\partial\Omega)$, and $\mathbb{X} = L^2(\Omega)$.

In order to apply the results of Section 2, it remains to verify Assumption 2.2. We already showed that $\mathcal{D} \in \mathcal{L}(\mathbb{X}, \mathcal{L}(\mathbb{U}, \mathbb{V}'))$ is a bounded linear operator. The following assertion states that also the remaining conditions on $\mathcal{U}$ and $\mathcal{V}$ hold true.

Lemma 3.1 *The operators $\mathcal{U}$ and $\mathcal{V}$ defined in (3.6) and (3.8) are Hilbert-Schmidt and their singular values decay like $\sigma_{k,\mathcal{U}} \lesssim k^{-3/(2d-2)}$ and $\sigma_{k,\mathcal{V}} \lesssim k^{-3/(2d-2)}$.*

Proof The Hilbert-Schmidt property follows immediately from the decay behavior of the singular values. Let Ω_h be a quasi-uniform triangulation of the domain Ω of mesh size h and $\partial\Omega_h$ be the induced segmentation of the boundary $\partial\Omega$. Further, let $\mathbb{Y}_h = P_1(\partial\Omega_h) \cap H^1(\partial\Omega) \subset H^{-1/2}(\partial\Omega)$ be the space of piecewise linear finite elements on $\partial\Omega_h$. Let $\mathcal{Q}_h$ be the L^2 -orthogonal projection onto $\mathbb{Y}_h$ and $q \in H^1(\partial\Omega)$ arbitrary. Then standard approximation error estimates, see e.g., [4], yield

$$\|q - \mathcal{Q}_h q\|_{H^{-1/2}(\partial\Omega)} \lesssim h^{3/2} \|q\|_{H^1(\partial\Omega)}.$$

A-priori estimates for elliptic PDEs yield $\|\mathcal{U}q\|_{H^1(\Omega)} \lesssim \|q\|_{H^{-1/2}(\partial\Omega)}$, and hence $\mathcal{U}$ can be continuously extended to an operator on $H^{-1/2}(\partial\Omega)$; see e.g. [10]. This yields

$$\|\mathcal{U} - \mathcal{U}\mathcal{Q}_h\|_{\mathcal{L}(H^1(\partial\Omega), H^1(\Omega))} \lesssim h^{3/2} \lesssim k^{-3/(2d-2)},$$

where $k = \dim(\mathbb{Y}_h) = \text{rank}(\mathcal{Q}_h) \simeq h^{-(d-1)}$ is the dimension of the space $\mathbb{Y}_h$. From the min-max characterization of the singular values (2.3), we may therefore conclude that $\sigma_{k,\mathcal{U}} \lesssim k^{-3/(2d-2)}$ as required. The result for $\sigma_{k,\mathcal{V}}$ follows in the same way. $\square$

Remark 3.2 If prior knowledge $\text{supp}(c) \subset \Omega$ on the support of the fluorophore concentration is available, which is frequently encountered in practice, elliptic regularity [10] implies exponential decay of the singular values $\sigma_{k,\mathcal{U}}$ and $\sigma_{k,\mathcal{V}}$. In such a situation, the ranks K and N in Lemma 2.4 and 2.7 will depend only logarithmically on the noise level δ , and an accurate approximation $\mathcal{T}_N$ of very low rank can be found.

Remark 3.3 If c_1 and c_2 are two fluorophore concentrations leading to the same measurements, that is $\mathcal{T}(c_1) = \mathcal{T}(c_2)$, then the factorization (3.11) shows that

$$\int_{\Omega} (c_1 - c_2) u_x v_m \, dx = 0 \quad (3.12)$$

Fig. 1 Left: computational domain and coarsest mesh used for our computations. Middle: minimum norm solution $c^\dagger$. Right: reconstructed fluorophore concentration c_α^δ for $\delta = 10^{-5}$

for all possible excitation fields u_x satisfying (3.1) and adjoint emission fields v_m satisfying (3.9). Under certain regularity conditions, density results for the set of products $\{u_x v_m\}$ imply $c_1 = c_2$, see [18, Chapter 5] for precise statements. Hence, in such a situation, the solution $c^\dagger$ of (1.1) is unique.

4 Algorithmic realization and complexity estimates

We will now discuss in detail the implementation of the model reduction approach presented in Sect. 2 for the fluorescence optical tomography problem and demonstrate its viability by some preliminary considerations. For ease of presentation, we consider a simple two-dimensional test problem. Our observations, however, carry over almost verbatim also to three dimensional problems of similar dimensions.

4.1 Problem setup

For the discretization of (3.1)–(3.2) and (3.9)–(3.10), we use a standard finite element method with continuous piecewise linear polynomials. The computational meshes used for the truth approximations are obtained by successive uniform refinement of the initial mesh, leading to quasi-uniform conforming triangulations Ω_h of the domain Ω with $h > 0$ denoting the mesh size. Thus, the corresponding spaces $\mathbb{U}_h, \mathbb{V}_h \subset H^1(\Omega)$ then have dimension $m \simeq h^{-d}$ each. For our test problem, we have $d = 2$, since we consider a two-dimensional setting. We choose the same finite element space $\mathbb{X}_h$ also for the approximation of the concentration c . The sources q_x, q_m for the forward and the adjoint problem are approximated by piecewise linear functions on the boundary of the same mesh Ω_h ; hence $\mathbb{Y}_h, \mathbb{Z}_h \subset H^1(\partial\Omega)$ have dimension $k \simeq h^{d-1}$. All approximation spaces are equipped with the topologies induced by their infinite dimensional counterparts. Standard error estimates allow to quantify the discretization errors in the resulting truth approximation of the forward operator and to establish the δ -approximation property for h small enough. The error introduced by the discretization can therefore be assumed to be negligible.

A sketch of the domain $\Omega \subset \mathbb{R}^2$ and the coarsest mesh Ω_h used for our computations as well as the parameter $c^\dagger$ to be identified are depicted in Fig. 1.

Table 1 Dimensions $m = \dim(\mathbb{X}_h) = \dim(\mathbb{U}_h) = \dim(\mathbb{V}_h)$ and $k = \dim(\mathbb{Y}_h) = \dim(\mathbb{Z}_h)$ of the relevant function spaces and discretization errors $de_h = \|\mathcal{T}_h - \mathcal{T}\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))}$ between approximation $\mathcal{T}_h$ on refinement level ref and the truth approximation $\mathcal{T}$ on level ref = 5. The norm of the forward operator is $\|\mathcal{T}_h\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))} = 0.2804$ on all mesh levels

Ref	0	1	2	3	4	5
m	993	3 881	15 345	61 025	243 393	927 161
k	88	176	352	704	1 408	2 816
de_h	$6.31 \cdot 10^{-4}$	$1.69 \cdot 10^{-4}$	$4.34 \cdot 10^{-5}$	$1.10 \cdot 10^{-5}$	$2.75 \cdot 10^{-6}$	—

The characteristic dimension of the relevant function spaces after discretization can be deduced from Table 1. The numbers in the table also illustrate the assumption $k < m < k^2$ and that the discretized inverse problem is formally overdetermined.

Note that an approximation on level ref ≥ 4 is required to guarantee a discretization error of $de_h = \|\mathcal{T}_h - \mathcal{T}\|_{\mathcal{L}(\mathbb{X}, \mathbb{HS}(\mathbb{Y}, \mathbb{Z}'))} \leq 10^{-5}$. Further observe that for the finest mesh level ref = 5, the discretized forward operator amounts to a linear mapping $\mathbb{T}$ from $\mathbb{R}^{927\,161}$ to $\mathbb{R}^{2\,816 \times 2\,816}$. The storage of the matrix $\mathbf{A}$ representing the forward operator would require approximately 56TB of memory and even one single evaluation of $\mathbb{T}(\mathbf{c})$ via the matrix product $\mathbf{A}\mathbf{c}$ would require approximately 7Tflops. It should be clear that more sophisticated algorithms are required to make even the evaluation of the forward operator feasible.

4.2 Truth approximation

Let us briefly discuss in a bit more detail the algebraic structure of the resulting problems arising in the truth approximation. Under the considered setup, the finite element approximation of problem (3.1)–(3.2) leads to the linear system

$$(\mathbf{K}_x + \mathbf{M}_x + \mathbf{R}_x) \mathbf{U} = \mathbf{E}_x \mathbf{Q}_x. \quad (4.1)$$

Here $\mathbf{K}_x, \mathbf{M}_x \in \mathbb{R}^{m \times m}$ are the stiffness and mass matrices with coefficients κ_x, μ_x , and the matrices $\mathbf{R}_x \in \mathbb{R}^{m \times m}$, $\mathbf{E}_x \in \mathbb{R}^{m \times k}$ stem from the discretization of the boundary conditions. The columns of regular $\mathbf{Q}_x \in \mathbb{R}^{k \times k}$ represent the individual independent sources in the basis of $\mathbb{Y}_h$. Any excitation generated by a source in $\mathbb{Y}_h$ can thus be expressed as a linear combination of columns of the excitation matrix $\mathbf{U} \in \mathbb{R}^{m \times k}$, which serves as a discrete counterpart of the operator $\mathcal{U}$. In a similar manner, the discretization of the adjoint problem (3.9)–(3.10) leads to

$$(\mathbf{K}_m + \mathbf{M}_m + \mathbf{R}_m) \mathbf{V} = \mathbf{E}_m \mathbf{Q}_m, \quad (4.2)$$

whose solution matrix $\mathbf{V} \in \mathbb{R}^{m \times k}$ can be interpreted as the discrete counterpart of the operator $\mathcal{V}$. The system matrices $\mathbf{K}_m, \mathbf{M}_m, \mathbf{R}_m$, and $\mathbf{E}_m$ have a similar meaning as above, and the columns of $\mathbf{Q}_m$ represent the individual detector characteristics. Recall from Subsection 1.1 that $k = k_{\mathbb{Y}} = k_{\mathbb{Z}}$ only to simplify exposition. The algebraic form of

the truth approximation finally reads

$$\mathbf{T}(\mathbf{c}) = \mathbf{V}^\top \mathbf{D}(\mathbf{c}) \mathbf{U}, \quad (4.3)$$

where $\mathbf{D}(\mathbf{c}) \in \mathbb{R}^{m \times m}$ is the matrix representation of the finite element approximation for the operator $\mathcal{D}(c_h)$ with $\mathbf{c} \in \mathbb{R}^m$ denoting the coordinates of the function $c_h \in \mathbb{X}_h$. The discrete measurement $M_{ij} = (\mathbf{V}^\top \mathbf{D}(\mathbf{c}) \mathbf{U})_{ij} = \mathbf{V}(:, i)^\top \mathbf{D}(\mathbf{c}) \mathbf{U}(:, j)$ then approximates the data taken by the i th detector for excitation with the j th source.

From the particular form (4.3) of the forward operator, one can see that only the matrices $\mathbf{U}$, $\mathbf{V}$, and a routine for the application of $\mathbf{D}(\mathbf{c})$ are required to evaluate $\mathbf{T}(\mathbf{c})$. In our example, the application of $\mathbf{D}(\mathbf{c})$ amounts to the multiplication by a diagonal matrix, which leads to a complexity of $O(k^2 m)$ flops for the application and a memory requirement of $O(2km + m)$ bytes for the forward operator. Note that the matrices $\mathbf{U}$ and $\mathbf{V}$ can now be stored on a standard workstation, while the application of the forward operator is still too compute intensive to be useful for the efficient solution of the inverse problem under consideration.

Remark 4.1 Let $\mathbf{SY}$, $\mathbf{SZ}$ be the matrix representation of the $H^1(\partial\Omega)$ inner products for the spaces $\mathbb{Y}_h$, $\mathbb{Z}_h$. Furthermore, let $\mathbf{AY}$, $\mathbf{AZ} \in \mathbb{R}^{k \times k}$ be orthogonal with respect to $\mathbf{SY}$ and $\mathbf{SZ}$, i.e., $\mathbf{AY}^\top * \mathbf{SY} * \mathbf{AY} = \mathbf{I}$ and $\mathbf{AZ}^\top * \mathbf{SZ} * \mathbf{AZ} = \mathbf{I}$. Then $\mathbf{AY} = \mathbf{Qx} * \mathbf{Ax}$ and $\mathbf{AZ} = \mathbf{Qm} * \mathbf{Am}$ with $\mathbf{Ax} = \mathbf{Qx}^{-1} * \mathbf{AY}$ and $\mathbf{Am} = \mathbf{Qm}^{-1} * \mathbf{AZ}$, and the Hilbert-Schmidt norm of the measurement matrix $\mathbf{M} = \mathbf{T}(\mathbf{c})$ can be expressed by the Frobenius norm

$$\|\mathbf{M}\|_{\text{HS}} := \|\mathbf{Am}^\top * \mathbf{M} * \mathbf{Ax}\|_{\mathbb{F}}.$$

Note that we simply have $\|\mathbf{M}\|_{\text{HS}} := \|\mathbf{M}\|_{\mathbb{F}}$ if the columns of $\mathbf{Qx}$, $\mathbf{Qm}$ are chosen orthonormal with respect to the $\mathbf{SY}$ and $\mathbf{SZ}$ inner products right from the beginning. We will use this fact in our numerical tests below.

Remark 4.2 Using multigrid solvers, the matrices $\mathbf{U}$ and $\mathbf{V}$ can be computed in $O(mk)$ operations [11], which is, at least asymptotically, negligible compared to the application of $\mathbf{T}(\mathbf{c})$ in tensor product form. In our computational tests, we utilize sparse direct solvers for the computation of $\mathbf{U}$ and $\mathbf{V}$, for which the computational cost is $O(m^{3/2} + km \log(m))$. Since $m^{3/2} \leq mk$ and $\log(m) \leq k$ in our two-dimensional setting, this is still of lower complexity than even a single evaluation of $\mathbf{T}(\mathbf{c})$.

4.3 Model reduction—offline phase

With $\mathbf{U}$ and $\mathbf{V}$ obtained, we are now in the position to compute our reduced order model.

4.3.1 Orthonormalization

Let $\mathbf{SX}$, $\mathbf{SY}$, $\mathbf{SZ}$ be the Gram matrices representing the scalar products of the function spaces $\mathbb{X}_h$, $\mathbb{Y}_h$, $\mathbb{Z}_h$. As a next step, we compute the approximations for the singular value decompositions of the excitation and emission operators. For this, we recall that

the right singular vectors of an operator $\mathcal{U}$ correspond to the eigenvectors of $\mathcal{U}^* \mathcal{U}$. The singular value decompositions for the matrices $\mathbf{U}$ and $\mathbf{V}$ can thus be computed by the generalized eigenvalue decompositions

$$\begin{aligned} [\mathbf{A}\mathbf{x}, \mathbf{D}\mathbf{x}] &= \text{eigs}(\mathbf{U}' * \mathbf{S}\mathbf{X} * \mathbf{U}, \mathbf{S}\mathbf{Y}); \\ [\mathbf{A}\mathbf{m}, \mathbf{D}\mathbf{m}] &= \text{eigs}(\mathbf{V}' * \mathbf{S}\mathbf{X} * \mathbf{V}, \mathbf{S}\mathbf{Z}); \end{aligned}$$

Note that some slight modifications would be required here, if the source and detector matrices $\mathbf{Q}\mathbf{x}$ and $\mathbf{Q}\mathbf{m}$ would not be chosen as the identity matrices. The columns of $\mathbf{A}\mathbf{x}$ and $\mathbf{A}\mathbf{m}$ are orthogonal with respect to the $\mathbf{S}\mathbf{Y}$ and $\mathbf{S}\mathbf{Z}$ scalar product and thus define bases of the discrete source and detector spaces. After appropriate scaling, the columns can be assumed to be normalized such that $\mathbf{A}\mathbf{x}' * \mathbf{S}\mathbf{Y} * \mathbf{A}\mathbf{x}$ and $\mathbf{A}\mathbf{m}' * \mathbf{S}\mathbf{Z} * \mathbf{A}\mathbf{m}$ equal the identity matrix. To simplify the subsequent discussion, we change the definition of the sources and detectors as well as of the excitation and emission matrices, and redefine the forward operator according to

$$\begin{aligned} \mathbf{Q}\mathbf{x} &= \mathbf{Q}\mathbf{x} * \mathbf{A}\mathbf{x}; & \mathbf{U} &= \mathbf{U} * \mathbf{A}\mathbf{x}; \\ \mathbf{Q}\mathbf{m} &= \mathbf{Q}\mathbf{m} * \mathbf{A}\mathbf{m}; & \mathbf{V} &= \mathbf{V} * \mathbf{A}\mathbf{m}; \\ \mathbf{T} &= \mathcal{A}(\mathbf{c}) \quad \mathbf{V}' * \mathbf{D}(\mathbf{c}) * \mathbf{U}; \end{aligned}$$

The columns of $\mathbf{Q}\mathbf{x}$ and $\mathbf{Q}\mathbf{m}$ are now orthogonal with respect to the $\mathbf{S}\mathbf{Y}$ and $\mathbf{S}\mathbf{Z}$ scalar products, and as a consequence, the Hilbert-Schmidt norm in the measurement space amounts to the Frobenius norm of $\mathbf{M} = \mathbf{T}(\mathbf{c})$; see Remark 4.1 for details.

The memory cost for storing $\mathbf{A}\mathbf{Y} = \mathbf{U}' * \mathbf{S}\mathbf{X} * \mathbf{U}$ and $\mathbf{A}\mathbf{Z} = \mathbf{V}' * \mathbf{S}\mathbf{X} * \mathbf{V}$ amounts to $O(k^2)$ bytes, while the computation of the matrix products requires $O(k^2 m)$ flops. The complexity for the eigenvalue decompositions finally is $O(k^3)$ flops. Note that the setup of the matrices $\mathbf{A}\mathbf{Y}$ and $\mathbf{A}\mathbf{Z}$ has the same cost as a single evaluation $\mathbf{T}(\mathbf{c})$ of the forward operator. The additional memory required for storing the $k \times k$ matrices $\mathbf{A}\mathbf{Y}$ and $\mathbf{A}\mathbf{Z}$ is negligible.

4.3.2 Low-rank approximations for $\mathbf{U}$ and $\mathbf{V}$

The eigenvalues computed in the decompositions above correspond to the square of the singular values of $\mathbf{U}$ and $\mathbf{V}$. We here allow for different ranks in the approximation and define truncation indices

$$\begin{aligned} \mathbf{d}\mathbf{x} &= \text{diag}(\mathbf{D}\mathbf{x}); & \mathbf{x}\mathbf{K}\mathbf{K} &= \text{find}(\mathbf{d}\mathbf{x} > \delta^2); & \mathbf{x}\mathbf{K} &= \text{length}(\mathbf{x}\mathbf{K}\mathbf{K}); \\ \mathbf{d}\mathbf{m} &= \text{diag}(\mathbf{D}\mathbf{m}); & \mathbf{m}\mathbf{K}\mathbf{K} &= \text{find}(\mathbf{d}\mathbf{m} > \delta^2); & \mathbf{m}\mathbf{K} &= \text{length}(\mathbf{m}\mathbf{K}\mathbf{K}); \end{aligned}$$

We could further set $\mathbf{K} = \max(\mathbf{x}\mathbf{K}, \mathbf{m}\mathbf{K})$ to stay exactly with the notation used in Sect. 2, but we stress that our implementation is in full generality. The low-rank approximations for $\mathbf{U}$ and $\mathbf{V}$ and the resulting tensor product approximation of the forward operator $\mathbf{T}(\mathbf{c})$ are then given by

$$\begin{aligned} \mathbf{Q}\mathbf{x}\mathbf{K} &= \mathbf{Q}\mathbf{x}(:, \mathbf{x}\mathbf{K}\mathbf{K}); & \mathbf{U}\mathbf{K} &= \mathbf{U}(:, \mathbf{x}\mathbf{K}\mathbf{K}); \\ \mathbf{Q}\mathbf{m}\mathbf{K} &= \mathbf{Q}\mathbf{m}(:, \mathbf{m}\mathbf{K}\mathbf{K}); & \mathbf{V}\mathbf{K} &= \mathbf{V}(:, \mathbf{m}\mathbf{K}\mathbf{K}); \\ \mathbf{T}\mathbf{K}\mathbf{K} &= \mathcal{A}(\mathbf{c}) \quad \mathbf{V}\mathbf{K}' * \mathbf{D}(\mathbf{c}) * \mathbf{U}\mathbf{K}; \end{aligned}$$

Observe that the measurements $\mathbf{M}\mathbf{K}\mathbf{K} = \mathbf{T}\mathbf{K}\mathbf{K}(\mathbf{c})$ obtained by this approximation correspond to a sub-block of the full measurements, i.e., $\mathbf{M}\mathbf{K}\mathbf{K} = \mathbf{M}(\mathbf{m}\mathbf{K}\mathbf{K}, \mathbf{x}\mathbf{K}\mathbf{K})$.

Table 2 Memory and computation cost for storing and applying the tensor product approximation TKK of rank $K^2 = 10^4$ and the corresponding hyperbolic cross approximation TK of rank $2K = 200$. The theoretical memory and computation cost is given by $\text{mem}(\text{TKK}) = \text{mem}(\text{TK}) = 16Km$ bytes, $\text{ops}(\text{TKK}(c)) = Km + K^2m$ flops, and $\text{ops}(\text{TK}(c)) = 2Km$ flops, respectively

Ref	0	1	2	3	4	5
<i>Mem (GB)</i>	0.001	0.006	0.023	0.009	0.036	1.448
<i>Ops (TKK(c), Gflops)</i>	0.009	0.037	0.144	0.574	2.289	9.144
<i>Ops (TK(c), Gflops)</i>	0.000	0.001	0.003	0.011	0.045	0.181

4.3.3 Hyperbolic cross approximation

The proof of Lemma 2.5 shows that we may replace the tensor product operator $\text{TKK}(c)$ by the hyperbolic cross approximation $\text{TK}(c)$, which takes into account only the entries $M(k, l) = \text{MKK}(k, l)$ of the measurements $M = T(c)$ for indices $k \cdot l \leq N \lesssim \delta^{-\beta}$. In our computations, we actually replace N by K , i.e., we utilize the hyperbolic cross approximation $\text{TK}(c)$ of $\text{TKK}(c)$. The assembly of the matrix representation AK for $\text{TK}(c) = AK * c$ then reads

```

m=0;
for k=1:K
    for l=1:floor(K/k)
        m=m+1;
        AK(m,:) = (VK(:,k))' .* UK(:,l))' * DD;
    end
end
end

```

Here, DD is a diagonal matrix representing the numerical integration on the computational domain. Let us note that AK and thus also the operator $\text{TK}(c)$ do not have a tensor product structure any more; therefore the measurements $MK = \text{TK}(c)$ are stored as a column vector rather than a matrix. The norm in the reduced measurement space then is the Euclidean norm for vectors. Also note that the construction of AK and TK only requires access to the matrices UK and VK defining the operator TKK .

The tensor product approximation $\text{TKK}(c) = VK' * D(c) * UK$ requires only sub-blocks UK , VK of the excitation and emission matrices U , V and, therefore, no additional memory cost arises in setting up this approximation. To achieve a δ -approximation with $\delta = 10^{-3}$, for instance, we expect to require approximately $K = 100$ singular components of U and V ; see Lemma 3.1 for details. The tensor product approximation will then have rank $K^2 = 10^4$. For the hyperbolic cross approximation TK , we however expect to require only approximately $2K = 200$ components of the tensor product approximation TKK ; compare with Lemma 2.5. In Table 2 we summarize the expected memory and computational cost for the corresponding approximations.

Note that the tensor product structure allows to store the tensor product approximation TKK as efficiently as the hyperbolic cross approximation TK . The application of the latter is, however, substantially more efficient.

Table 3 Complexity for computing $BKK = AKK*(DD \setminus AKK')$ and $BK = AK*(DD \setminus AK')$. The estimates are $mem(BKK) = 8K^4$ bytes, $ops(BKK) = mK^4$ flops and $mem(BK) = 32K^2$ bytes, $ops(BK) = 2K$ flops

Ref	0	1	2	3	4	5
$Ops(BKK, \text{Gflops})$	92.480	361.44	1429.1	5683.4	22667	90540
$Ops(BK, \text{Gflops})$	0.000	0.001	0.003	0.011	0.045	0.181

4.3.4 Final recompression

The last step in our model reduction approach consists in a further compression of the hyperbolic cross approximation TK of the tensor product operator TKK ; cf. Subsection 2.4 for details. This can be realized by

```
AKt=DD\AK';
AKAKt=AK*AKt;
[VA, DA]=eigs(AKAKt, NK);
```

where $NK = \text{size}(AK, 1)$ is the number of terms used for the hyperbolic cross approximation. The recompression then consists of selecting the largest entries, i.e.,

```
N=find(diag(DA)>delta^2);
AN=AK*VA(:, N); ANt=DD\AN';
```

Matrix representations for the projection operators $\mathcal{Q}_{K,K}$, $\mathcal{P}_K$, and $\mathcal{P}_N$ corresponding to the tensor product, the hyperbolic cross, and the final approximation, can be assembled easily from the eigenvectors computed during the construction.

As before, the compression is based on singular value decompositions of operators via the solution of generalized eigenvalue problems for the matrices $BKK = AKK*(DD/AKK')$ respectively $BK = AK*(DD/AK')$, where AKK and AK are the matrix representations for the operators $TKK(c)$ and $TK(c)$ respectively. The computational cost for the assembly of BKK and BK is listed in Table 3. For an evaluation of the computational complexity, we again assume that TKK has rank $K^2 = 10^4$ and that TK is of rank $2K = 200$.

Let us note that the required memory for storing BKK and BK is independent of the mesh size; for the setting considered here, it is given by $mem(BKK) = 745\text{MB}$ and $mem(BK) = 0.3\text{MB}$. Assuming that an eigenvalue decomposition of an $n \times n$ matrix needs roughly $ops(eig) = 50n^3$ operations, we obtain $ops(eig(BKK)) = 46566 \text{ Gflops}$ and $ops(eig(BK)) = 0.373 \text{ Gflops}$. Even if a computationally more efficient low-rank approximation [14,35] for the tensor product operator TKK would be used, the evaluation of $TKK(c)$ remains rather expensive; see Table 2 for details. Therefore, the tensor product approximation TKK is not really useful for the computation of low-rank approximation on large computational meshes. As shown in Sect. 2, a quasi-optimal approximation TN can be computed also by truncation of the singular value decomposition of the hyperbolic cross approximation TK , which does not require any additional computations.

4.4 Online phase

After the construction of the low-rank approximation $\mathcal{T}\mathcal{N}(\mathbf{c})$ as outlined above, the actual solution of the inverse problem consists of three basic steps; see Sect. 1 for a brief explanation. The first step is the *data compression* which can be expressed as $\mathbf{M}\mathbf{N} = \mathbf{P}\mathbf{N} * \forall \mathbf{c} \in (\mathbf{M}\mathbf{K}\mathbf{K})$ with $\mathbf{M}\mathbf{K}\mathbf{K} = (\mathbf{Q}\mathbf{m}\mathbf{K}' * \mathbf{M}) * \mathbf{Q}\mathbf{x}\mathbf{K}$. Here we make explicit use of the tensor product structure, which allows us to efficiently compress the data already during recording. After this, only the second projection $\mathbf{P}\mathbf{N}$ has to be applied. The additional memory required for computing $\mathbf{M}\mathbf{K}\mathbf{K}$ is $O(k\mathbf{K})$ bytes for each, $\mathbf{Q}\mathbf{x}\mathbf{K}$, $\mathbf{Q}\mathbf{m}\mathbf{K}$, and $\mathbf{Q}\mathbf{m}\mathbf{K}' * \mathbf{M}$ and thus negligible. Note that storing the full data $\mathbf{M}$ requires $O(k^2)$ bytes which is substantially higher. The computational cost of the data compression step is $O(\mathbf{K}^2 k + k^2 \mathbf{K})$. As mentioned before, the data can be partially compressed already during recording, such that access to the full data is actually never required.

The final compression $\mathbf{M}\mathbf{N} = \mathbf{P}\mathbf{N} * \mathbf{M}\mathbf{K}\mathbf{K}$ is independent of the system dimensions m , k and its computational cost is therefore negligible. The same applies for the solution of the regularized inverse problem $\mathbf{z}\mathbf{a}\mathbf{d}\mathbf{N} = (\mathbf{A}\mathbf{N}\mathbf{A}\mathbf{T} + \alpha \mathbf{I}) / \mathbf{M}\mathbf{N}$, which is the second step in the online phase and only depends on the dimension $\mathbf{N}$ of the reduced model.

The synthesis of the solution according to (1.7) can finally be realized by simple multiplication $\mathbf{c}\mathbf{a}\mathbf{d}\mathbf{N} = \mathbf{A}\mathbf{N}\mathbf{T} * \mathbf{z}\mathbf{a}\mathbf{d}\mathbf{N}$, where $\mathbf{A}\mathbf{N}\mathbf{T}$ denotes the matrix representation of the adjoint of the fully reduced forward operator $\mathcal{T}\mathcal{N}$. The additional memory required for $\mathbf{A}\mathbf{N}\mathbf{T}$ is $O(m\mathbf{N})$ bytes, whereas the computation of $\mathbf{c}\mathbf{a}\mathbf{d}\mathbf{N}$ can be accomplished in $O(m\mathbf{N})$ flops. Thus, as claimed in the introduction, the most compute intensive part of the online phase is the data compression, even if the tensor product structure is utilized to compress the data already during recording.

5 Computational results

We now illustrate the practical performance of our model reduction approach for the test problem introduced in the previous section. For comparison, we also report on corresponding results for traditional iterative methods for solving the inverse problem (1.1)–(1.2), as well as for methods based on a tensor-product approximation. A snapshot of the geometry, the exact solution, and a typical reconstruction is depicted in Fig. 1.

For our numerical tests, the model parameters are set to $\kappa_x = 1$, $\mu_x = 0.2$, $\rho_x = 10$ and $\kappa_m = 2$, $\mu_m = 0.1$, and $\rho_m = 10$. We further assume prior knowledge that $\mathbf{c}$ is supported in a circle of radius 0.9, i.e., the distance of its support to the boundary $\partial\Omega$ is at least 0.1. The singular values of the operators $\mathcal{U}$ and $\mathcal{V}$ as well as of $\mathcal{T}$ can thus be assumed to decay exponentially. The noise level in (1.3) is set to $\delta = 10^{-5}$ and the regularization parameter α is chosen from $\{10^{-n}\}$ via the discrepancy principle. In all our computations, this led to $\alpha = 10^{-8}$ which complies to the theoretical prediction for exponentially ill-posed problems [9].

Let us note that, in order to ensure sufficient accuracy $\|\mathcal{T}_h - \mathcal{T}\|_{\mathcal{L}(\mathbf{X}; \mathbb{H}\mathbf{S}(\mathbf{Y}, \mathbf{Z}'))} \leq \delta$ of the truth approximation, one should utilize a discretization of the forward operator

Table 4 Computation times (sec) for the individual steps in the problem setup phase

Ref	0	1	2	3	4	5
Initialization of U, V	0.01	0.06	0.34	3.45	47.74	612.63
Setup of $U' * DX * U, V' * DX * V$	0.00	0.01	0.16	2.24	30.68	460.13
Eigenvalue decompositions	0.01	0.03	0.24	2.15	33.44	470.88
Orthogonalization of U, V	0.00	0.02	0.09	1.70	25.12	363.64

Table 5 Truncation indices $\times K$ and mK guaranteeing

$\ U - UK\ \leq \delta$ and	$\times K$	88	139	163	179	185	187
$\ V - VK\ \leq \delta$ with $\delta = 10^{-5}$	mK	88	121	133	137	137	137

on mesh level $\text{ref} \geq 4$ for the reconstruction; see Table 1 and Sect. 1. For evaluation of computational performance, we however also report about results on coarser meshes.

All computations are performed on a workstation with Intel(R) Xeon(R) Gold 6130 CPU @ 2.10GHz and 768GB of memory. In our tests we use only a single core of the processor and an implementation in MATLAB 9.6.0.

5.1 Problem initialization

This step consists of setting up the excitation and emission matrices U, V , which are required for the efficient evaluation of $T(c)$. In Table 4, we also report about the singular value decomposition of U and V , by which we orthogonalize the sources and detectors; as mentioned in Subsect. 4.3, this is required for computation of the Hilbert-Schmidt norm $\|\mathcal{M}^\delta\|_{\text{HS}(\mathbb{Y}, \mathbb{Z})}$ of the measurement operator.

While the theoretical complexity of the first and third step is somewhat smaller than that of the second and fourth step, the overall computation times for the individual steps in the setup phase are comparable.

Note that the computations reported in Table 4 are required for the solution of the inverse problem (1.1), independent of the particular solution strategy, and they can be performed in a pre-processing step.

5.2 Model reduction—offline phase

The singular values computed in the decompositions of U and V allow to determine the truncation indices $\times K$ and mK used to define the δ -approximations $UK = U(:, \times KK)$ and $VK = V(:, mKK)$. The values of $\times K$ and mK obtained in our numerical tests are depicted in Table 5.

On the coarsest mesh, the number of possible excitations and detectors is limited by the number of boundary vertices, but otherwise, the number of truncation indices $\times K$ and mK are almost independent of the truth approximation. This can be expected since the eigenvalues converge with increasing refinement of the mesh.

Table 6 Computation times (sec) for a single evaluation of $T(c)$ and $TKK(c)$

Ref	0	1	2	3	4	5
$T(c)$	0.00	0.01	0.07	1.03	13.99	214.20
$TKK(c)$	0.00	0.01	0.02	0.11	0.52	2.20

Table 7 Computation times (sec) for singular value decompositions of $T(c)$ and $TKK(c)$ and truncation indices N leading to corresponding δ -approximations

Ref	0	1	2	3	4	5
$\text{svd}(T)$	6.46	28.23	284.33	–	–	–
$N(T)$	231	303	473	–	–	–
$\text{svd}(TKK)$	6.45	15.05	48.40	248.42	994.66	–
$N(TKK)$	231	276	296	310	314	–

5.2.1 Forward evaluation

As outlined in Subsect. 4.3, the full operator and its tensor product approximation can now be simply defined by $T = \mathcal{A}(c) \quad V' * D(c) * U$ and $TKK = \mathcal{A}(c) \quad VK' * D(c) * UK$. In Table 6, we report about the computation times for a single evaluation of these operators.

As can be seen, even the problem adapted evaluation of the full operator $T(c)$ becomes practically useless for the solution of the inverse problem (1.1). The tensor product approximation $TKK(c)$, which is the underlying approximation for methods based on optimal sources [21] or based on the Kathri-Rhao product [24], seems somewhat better suited but, as we will see below, may still be not appropriate for the efficient solution of the inverse problem.

5.2.2 Truncated singular value decomposition

As a theoretical reference for model-reduction, we consider the low-rank approximation of $T(c)$ and $TKK(c)$ by truncated singular value decomposition, which can be computed via eigenvalue decompositions for the symmetric operators $T(Tt(M))$ and $TKK(TKKt(MKK))$. The latter can be computed numerically by the `eigs` routine of MATLAB in a matrix-free way, i.e., only requiring the application of the operators $T(c)$, $TKK(c)$ and their adjoints $Tt(M)$, $TKKt(M)$. The sum of xK and mK specifies the maximal number of eigenvalues to be considered by the algorithm. In Table 7, we display the computation times for eigenvalue solvers and the number N of relevant eigenvalues required to obtain a δ -approximation.

The computation times for the decomposition of the full operator T increase roughly by a factor of 8 per refinement, while those for the tensor product approximation only increase by a factor of 4. Computations taking longer than 1000sec were not conducted. Due to the substantially smaller rank, the evaluation $TN(c)$ of the low-rank approximations resulting from one of the singular value decompositions above is faster by a factor of more than 100 compared to that of the tensor product approximation $TKK(c)$, and even on the finest mesh only takes about 0.01sec. Let us recall that a discretization at mesh level $\text{ref} \geq 4$ is required to guarantee sufficient approximation

Table 8 Computation times (sec) for construction of the hyperbolic cross approximation $\text{TK}(\mathbf{c})$ and its singular value decomposition used for constructing the final approximation $\text{TN}(\mathbf{c})$ with rank $N(\text{TK})$

Ref	0	1	2	3	4	5
setup of TK, $\text{TKTK}^\dagger$	0.01	0.013	1.38	6.31	30.48	140.32
$\text{svd}(\text{TK})$	0.08	0.87	2.57	3.51	3.87	4.03
$\text{rank}(\text{TK})$	403	933	1 725	1 867	1 905	1 917
$\text{rank}(\text{TN})$	166	266	391	396	401	403

$\|\mathcal{T}_h - \mathcal{T}\|_{\mathcal{L}(\mathbb{X}; \text{HS}(\mathbb{Y}, \mathbb{Z}))} \leq \delta$ of the truth approximation $\mathcal{T}_h$ used for the solution of the inverse problem.

5.2.3 Setup of reduced order model

As described in Section 2, we can utilize the hyperbolic cross approximation TK instead of the full tensor product approximation TKK without losing the δ -approximation property. In Table 8, we summarize the computation times for assembling the hyperbolic cross approximation TK and the subsequent singular value decomposition used in the final recompression step.

Note that the setup cost for the hyperbolic cross approximation increases roughly by a factor of 4 for each refinement, while the subsequent singular value decomposition and the ranks are essentially independent of the mesh level.

Due to the moderate rank $K = \text{rank}(\text{TK})$ of the hyperbolic cross approximation, it pays off to compute the matrix approximation of $\text{TKTK}^\dagger$ and to use it for the subsequent eigenvalue decomposition. As can be seen from Table 8, the recompression step allows to reduce the rank by another factor of about 5. As predicted by our theoretical investigations, the rank of the final approximation TN is comparable to that of the truncated singular value decomposition of the full operator T or its tensor product approximation TKK; cf. Table 7. The use of the hyperbolic cross approximation TK instead of the full operator or its tensor product approximation however allows to speed up the computation of the final low-rank approximation TN substantially. Again, the rank of the approximation becomes essentially independent of the mesh after some initial refinements, reflecting the mesh-independence of our approach.

5.3 Solution of inverse problem – online phase

We now turn to the online phase of the solution process. Iterative methods are used for the solution of the inverse problem with the full operator T and its tensor product approximation TKK. As mentioned before, we choose a regularization parameter $\alpha = 10^{-8}$, which was determined by the discrepancy principle. For the computation of the regularized solution (1.7) with full operator $\text{T}(\mathbf{c})$ and the tensor product approximation $\text{TKK}(\mathbf{c})$, we use MATLAB's `pcg` routine with tolerance set to $\text{tol} = \alpha\delta^2$. Since the rank of the final reduced order model TN is rather small, we can use a direct solution of (1.7) by MATLAB's `backslash` operator in that case.

Table 9 Computation times (sec) for the solution of the inverse problem via Tikhonov regularization. Iterative methods are utilized for the solution of (1.4) in the first two cases working with operators $\mathbb{T}$ and $\mathbb{TKK}$, while a direct solver is used for the low-rank approximation $\mathbb{TN}$

Ref	0	1	2	3	4	5
$\mathbb{T}$	1.24	13.91	320.73	—	—	—
$\mathbb{TKK}$	1.22	10.07	65.02	382.76	—	—
$\mathbb{TN}$	0.01	0.01	0.03	0.13	0.52	1.94

Table 10 Computation times (sec) for the individual steps of the online phase for inversion with reduced order model $\mathcal{T}_N$ of quasi-optimal rank

Ref	0	1	2	3	4	5
Data compression	0.001	0.005	0.028	0.114	0.457	1.831
Regularized normal equations	0.002	0.001	0.002	0.003	0.003	0.003
Synthesis	0.001	0.001	0.004	0.015	0.061	0.107

In Table 9, we display the online solution times and the error $\text{err} = \|c_\alpha^\delta - c^\dagger\|$ obtained for the final iterate.

Approximately 1 800 iterations are required for the iterative solution of (1.4) with the full operator $\mathbb{T}$ and the tensor-product approximation $\mathbb{TKK}$ on all mesh levels, which again illustrates the mesh-independence of the algorithms. Note that even for the tensor product approximation, the iterative solution on fine meshes becomes practically infeasible, while for the low-rank approximation $\mathcal{T}_N$ of quasi-optimal rank, the inverse problem solution remains extremely efficient up the finest mesh.

In Table 10 we discuss in more detail the computation times for the individual steps in (1.7), namely the *data compression*, the *solution of the regularized normal equations*, and the *synthesis of the reconstruction*.

Similar online computation times are also obtained for the low-rank approximation computed by truncated singular value decomposition of the full operator $\mathcal{T}$, since its rank and approximation properties are very similar to that of the approximation constructed by our approach.

As announced in the introduction and predicted by our complexity estimates, the data compression step becomes the most compute-intensive task in the online solution via the low-rank reduced order model $\mathbb{TN}$. While the data compression and synthesis step depend on the dimension of the truth approximation, the solution of the regularized normal equations becomes completely independent of the computational mesh. Also observe that the quality of the reconstruction is not degraded by the use of a low-rank approximation in the solution process. Overall, we thus obtained an extremely efficient, stable, and accurate reconstruction for fluorescence tomography.

6 Summary

A novel approach towards the systematic construction of approximations for high dimensional linear inverse problems with operator valued data was proposed yielding certified reduced order models of quasi-optimal rank. The approach was fully analyzed in a functional analytic setting and the theoretical results were illustrated by an application to fluorescence optical tomography. The main advantages of our approach, compared to more conventional low-rank approximations, like truncated singular value decomposition, lies in a vastly improved setup time and the possibility to partially compress the data already during recording. In particular, the computational effort of setting up the reduced order model $\mathcal{T}_N$ is comparable to that of *one single* evaluation $\mathcal{T}(c)$ of the forward operator. Due to the underlying tensor-product approximation, access to the full data $\mathcal{M}^\delta$ is not required.

The most compute intensive part in the offline phase consists in the setup of the discrete representations for $\mathcal{U}$ and $\mathcal{V}$ as well as their eigenvalue decomposition. A closer investigation and the use of parallel computation could certainly further improve the computation times for this step. Further acceleration of the data compression and synthesis step could probably be achieved by using computer graphics hardware. The low-dimensional reduced order models obtained in this paper may also serve as preconditioners for the iterative solution of related nonlinear inverse problems, which would substantially increase the field of potential applications.

Acknowledgements The work of the second author was supported by the German Research Foundation (DFG) via grants TRR 146 C3 and TRR 154 C04 and via the “Center for Computational Engineering” at TU Darmstadt.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Arridge, S.R., Schotland, J.C.: Optical tomography: forward and inverse problems. *Inverse Problems* **25**, 123010 (2009). <https://doi.org/10.1088/0266-5611/25/12/123010>
2. Bakushinsky, A.B., Kokurin, M.Y.: *Iterative Methods for Approximate Solution of Inverse Problems*. Springer, Dordrecht (2004). <https://doi.org/10.1007/978-1-4020-3122-9>
3. Benner, P., Gugercin, S., Willcox, K.: A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Rev* **57**, 483–531 (2015). <https://doi.org/10.1137/130932715>
4. Brenner, S.C., Scott, L.R.: *The Mathematical Theory of Finite Element Methods*. Springer, New York (2008). <https://doi.org/10.1007/978-0-387-75934-0>
5. Chaillat, S., Biros, G.: FaIMS: a fast algorithm for the inverse medium problem with multiple frequencies and multiple sources for the scalar Helmholtz equation. *J. Comput. Phys.* **231**, 4403–4421 (2012). <https://doi.org/10.1016/j.jcp.2012.02.006>
6. Colton, D., Kress, R.: *Inverse Acoustic and Electromagnetic Scattering Theory*, 4th edn. Springer, Cham (2019). <https://doi.org/10.1007/978-3-030-30351-8>

7. Dng, D., Temlyakov, V., Ullrich, T.: Hyperbolic Cross Approximation. Birkhuser/Springer, Cham (2018)
8. Egger, H., Freiberger, M., Schlottbom, M.: On forward and inverse models in fluorescence diffuse optical tomography. *Inverse Problems Imag.* **4**, 411–427 (2010). <https://doi.org/10.3934/ipi.2010.4.411>
9. Engl, H.W., Hanke, M., Neubauer, A.: Regularization of Inverse Problems. Kluwer Academic Publishers Group, Dordrecht (1996)
10. L. C. Evans, *Partial differential equations*. 2nd ed., AMS, 2010
11. Freiberger, M., Egger, H., Liebmnn, M., Scharfetter, H.: High-performance image reconstruction in fluorescence tomography on desktop computers and graphics hardware. *Biomed. Opt. Express* **2**, 3207–3222 (2011). <https://doi.org/10.1364/BOE.2.003207>
12. Golub, G.H., Van Loan, C.F.: Matrix Computations, 4th edn. Johns Hopkins University Press, Baltimore, MD (2013)
13. Grinberg, N., Kirsch, A.: The Factorization Method for Inverse Problems of Oxford Lecture Series in Mathematics and its Applications. Oxford University Press, New York (2008)
14. Halko, N., Martinsson, P.G., Tropp, J.A.: Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev* **53**, 217–288 (2011). <https://doi.org/10.1137/090771806>
15. Herrmann, F.J., Erlangga, Y.A., Lin, T.T.: Compressive simultaneous full-waveform simulation. *Geophysics* **74**, A35–A40 (2009). <https://doi.org/10.1190/1.3115122>
16. Hochstenbach, M.E.: A Jacobi-Davidson type SVD method. *SIAM J. Sci. Comput.* **23**, 606–628 (2001). <https://doi.org/10.1137/S1064827500372973>
17. Hohage, T., Raumer, H.-G., Spehr, C.: Uniqueness of an inverse source problem in experimental aeroacoustics. *Inverse Problems* **36**, 075012 (2020)
18. Isakov, V.: Inverse Problems for Partial Differential Equations, 3rd edn. Springer, Cham (2017)
19. Kato, T.: Perturbation theory for linear operators. Springer, Berlin Heidelberg (1966). <https://doi.org/10.1007/978-3-662-12678-3>
20. Kolda, T.G., Bader, B.W.: Tensor decompositions and applications. *SIAM Rev* **51**, 455–500 (2009). <https://doi.org/10.1137/07070111X>
21. Krebs, J.R., Anderson, J.E., Hinkley, D., Neelamani, R., Lee, S., Baumstein, A., Lacasse, M.-D.: Fast full-wavefield seismic inversion using encoded sources. *Geophysics* **74**, WCC177–WCC188 (2009). <https://doi.org/10.1190/1.3230502>
22. Lev-Ari, H.: Efficient solution of linear matrix equations with applications to multistaging antenna array processing. *Commun Inf Syst* **5**, 123–130 (2005). <https://doi.org/10.4310/CIS.2005.v5.n1.a5>
23. Levinson, H.W., Markel, V.A.: Solution of the nonlinear inverse scattering problem by *t*-matrix completion. i. theory. *Phys. Rev. E* **94**, 043317 (2016). <https://doi.org/10.1103/PhysRevE.94.043317>
24. Markel, V.A., Levinson, H., Schotland, J.C.: Fast linear inversion for highly overdetermined inverse scattering problems. *Inverse Problems* **35**, 124002 (2019). <https://doi.org/10.1088/1361-6420/ab44e7>
25. Markel, V.A., Mital, V., Schotland, J.C.: Inverse problem in optical diffusion tomography iii inversion formulas and singular-value decomposition. *J. Opt. Soc. Am. A* **20**, 890 (2003). <https://doi.org/10.1364/josaa.20.000890>
26. Math, P., Pereverzev, S.: Geometry of linear ill-posed problems in variable hilbert scales. *Inverse Problems* **19**, 789–803 (2003). <https://doi.org/10.1088/0266-5611/19/3/319>
27. C. Musco and C. Musco, *Randomized block Krylov methods for stronger and faster approximate singular value decomposition*, in Advances in Neural Information Processing Systems 28, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, eds., Curran Associates, Inc., 2015, pp. 1396–1404
28. F. Natterer, The mathematics of computerized tomography, vol. 32 of Classics in Applied Mathematics, SIAM, Philadelphia, PA, (2001)
29. Neubauer, A.: An a posteriori parameter choice for Tikhonov regularization in the presence of modeling error. *Appl. Numer. Math.* **4**, 507–519 (1988). [https://doi.org/10.1016/0168-9274\(88\)90013-x](https://doi.org/10.1016/0168-9274(88)90013-x)
30. Ntziachristos, V.: Fluorescence molecular imaging. *Annual Rev. Biomed. Eng.* **8**, 1–33 (2006). <https://doi.org/10.1146/annurev.bioeng.8.061505.095831>
31. Pukelsheim, F.: Optimal Design of Experiments (Classics in Applied Mathematics) (Classics in Applied Mathematics, 50). Society for Industrial and Applied Mathematics, USA (2006)
32. Quarteroni, A., Manzoni, A., Negri, F.: Reduced basis methods for partial differential equations. Springer, Cham (2016). <https://doi.org/10.1007/978-3-319-15431-2>

33. Roosta-Khorasani, F., van den Doel, K., Ascher, U.: Stochastic algorithms for inverse problems involving PDEs and many measurements. *SIAM J. Sci. Comput.* **36**, S3–S22 (2014). <https://doi.org/10.1137/130922756>
34. Somersalo, E., Isaacson, D., Cheney, M.: A linearized inverse boundary value problem for Maxwell's equations. *J. Comput. Appl. Math.* **42**, 123–136 (1992). [https://doi.org/10.1016/0377-0427\(92\)90167-V](https://doi.org/10.1016/0377-0427(92)90167-V)
35. Stoll, M.: A Krylov-Schur approach to the truncated SVD. *Linear Algebra Appl.* **436**, 2795–2806 (2012). <https://doi.org/10.1016/j.laa.2011.07.022>
36. Tikhonov, A., Arsenin, V.: *Solutions of Ill-posed Problems*. Wiley, Hoboken (1977)
37. van den Doel, K., Ascher, U.M.: Adaptive and stochastic algorithms for electrical impedance tomography and DC resistivity problems with piecewise constant solutions and many measurements. *SIAM J. Sci. Comput.* **34**, A185–A205 (2012). <https://doi.org/10.1137/110826692>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Jürgen Dölz¹ · Herbert Egger² · Matthias Schlottbom³

Jürgen Dölz
doelz@ins.uni-bonn.de

Herbert Egger
herbert.egger@ricam.oeaw.ac.at

- ¹ Institute for Numerical Simulation, University of Bonn, Friedrich-Hirzebruch-Allee 7, 53115 Bonn, Germany
- ² Institute for Computational Mathematics, Johannes Kepler University Linz and Johann Radon Institute for Computational and Applied Mathematics, Austrian Academy of Sciences, Altenbergerstr. 69, 4040 Linz, Austria
- ³ Department of Applied Mathematics, University of Twente, P.O. Box 217, 7500, AE Enschede, The Netherlands