This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Hannukainen, A.; Hyvönen, N.; Majander, H.; Tarvainen, T.

Efficient inclusion of total variation type priors in quantitative photoacoustic tomography

Published in:
SIAM JOURNAL ON IMAGING SCIENCES

DOI:
[10.1137/15M1051737](https://doi.org/10.1137/15M1051737)

Published: 16/08/2016

Document Version
Publisher's PDF, also known as Version of record

Please cite the original version:
Hannukainen, A., Hyvönen, N., Majander, H., & Tarvainen, T. (2016). Efficient inclusion of total variation type priors in quantitative photoacoustic tomography. *SIAM JOURNAL ON IMAGING SCIENCES*, 9(3), 1132-1153. <https://doi.org/10.1137/15M1051737>

Efficient Inclusion of Total Variation Type Priors in Quantitative Photoacoustic Tomography*

A. Hannukainen[†], N. Hyvönen[†], H. Majander[‡], and T. Tarvainen[§]

Abstract. Quantitative photoacoustic tomography is an emerging imaging technique aimed at estimating the distribution of optical parameters inside tissues from photoacoustic images, which are formed by combining optical information and ultrasonic propagation. This optical parameter estimation problem is ill-posed and needs to be approached within the framework of inverse problems. Photoacoustic images are three-dimensional and high-resolution. Furthermore, high-resolution reconstructions of the optical parameters are targeted. Therefore, in order to provide a practical method for quantitative photoacoustic tomography, the inversion algorithm needs to be able to perform successfully with problems of prominent size. In this work, an efficient approach for the inverse problem of quantitative photoacoustic tomography is proposed, assuming an edge-preferring prior for the optical parameters. The method is based on iteratively combining priorconditioned LSQR with a lagged diffusivity step and a linearization of the measurement model, with the needed multiplications by Jacobians performed in a matrix-free manner. The algorithm is tested with three-dimensional numerical simulations. The results show that the approach can be used to produce accurate and high quality estimates of absorption and diffusion in complex three-dimensional geometries with moderate computation time and cost.

Key words. quantitative photoacoustic tomography, priorconditioning, Perona–Malik, LSQR, matrix-free implementation, total variation

AMS subject classifications. 65N21, 35R30, 35Q60

DOI. 10.1137/15M1051737

1. Introduction. In *photoacoustic tomography* (PAT), high-contrast, high-resolution images of biological tissue are produced by utilizing the photoacoustic effect caused by an externally introduced light pulse. Absorption of the light pulse in a target generates an initial acoustic pressure distribution that is proportional to the absorbed energy density of the light. Due to the elastic nature of tissue, the initial pressure distribution propagates as an ultrasonic wave and can be measured on the surface of the tissue. These measurements are then

*Received by the editors December 8, 2015; accepted for publication (in revised form) May 16, 2016; published electronically August 16, 2016.

<http://www.siam.org/journals/siims/9-3/M105173.html>

Funding: The work of the second and third authors was supported by the Academy of Finland (decision 267789). The work of the fourth author was supported by the Academy of Finland (decisions 136220, 272803, and the Finnish Centre of Excellence in Inverse Problems Research 250215).

[†]Department of Mathematics and Systems Analysis, Aalto University, FI-00076 Aalto, Finland (antti.hannukainen@aalto.fi, nuutti.hyvonen@aalto.fi).

[‡]Department of Mathematics and Systems Analysis, Aalto University, FI-00076 Aalto, Finland, and École Polytechnique, Centre de Mathématiques Appliquées, Route de Saclay, 91128 Palaiseau Cedex, France (helle.majander@aalto.fi).

[§]Department of Applied Physics, University of Eastern Finland, FI-70211 Kuopio, Finland, and Department of Computer Science, University College London, WC1E 6BT London, UK (tanja.tarvainen@uef.fi).

used to reconstruct the initial pressure and form images of the target. Photoacoustic imaging combines the benefits of optical contrast and ultrasound propagation. The optical methods provide information about the distribution of chromophores which are light absorbing molecules within the tissue. The chromophores of interest are, e.g., haemoglobin, melanin, and various contrast agents. The ultrasonic waves carry this optical information directly to the surface with minimal scattering, thus retaining accurate spatial information as well. PAT has successfully been applied to the visualization of different structures in biological tissues such as human blood vessels, microvasculature of tumors, and cerebral cortex in small animals. For more information about PAT, see, e.g., [51, 29, 50, 6] and the references therein.

Quantitative photoacoustic tomography (QPAT) is a technique aimed at estimating the concentrations of the chromophores [12]. The inverse problem associated with QPAT is two-fold. First, the initial acoustic pressure distribution is estimated from the measured acoustic waves. This is an inverse initial value problem of acoustics and it has been studied extensively; see, e.g., [51, 27, 50] and the references therein. The second stage of QPAT consists of the optical inverse problem of determining the concentrations of chromophores. These concentrations can be reconstructed either by directly estimating them from photoacoustic images obtained at various wavelengths [13, 28, 5, 36] or by first recovering the absorption coefficients at different wavelengths and then calculating the concentrations based on the absorption spectra [39, 13, 5]. In order to obtain accurate estimates, scattering effects need to be taken into account [4, 12, 47, 38]. As an alternative to the two-step approach, the estimation of the optical parameters directly from photoacoustic time-series has also been considered [45, 46, 21, 17, 15, 35].

In this work, the optical inverse problem of QPAT is studied assuming the corresponding acoustic inverse problem has already been solved. The estimation of absorption and diffusion at a single wavelength of light is considered, but the extension to multiple wavelengths is straightforward. Furthermore, it is assumed that the photoacoustic efficiency, which connects the acoustic pressure with the absorbed optical energy density and can be identified with the Grüneisen parameter for an absorbing fluid, is known. For discussions about the estimation of the Grüneisen parameter simultaneously with the optical parameters, see, e.g., [44, 5, 31, 36, 1]. In the optical inverse problem of QPAT, two models of light propagation have been used: the *radiative transfer equation* (RTE) [47, 43, 30, 37, 21] and its *diffusion approximation* (DA) [19, 4, 44, 52, 47, 48, 31, 38, 40]. The forward model used here is based on the DA, although our approach could also be implemented with the RTE.

The optical inverse problem of QPAT is nonlinear and ill-posed. In order to overcome the ill-posedness, regularization or Bayesian methods need to be used [26]. In this work, a Bayesian approach with a Gaussian model for the measurement noise and an edge-preferring prior for the to-be-estimated optical parameters are employed. To be more precise, we consider total variation [41] type priors (particularly Perona–Malik [34]), which are edge-preserving and support piecewise constant images which consist of a few homogeneous levels. Total variation priors/regularization have previously been utilized in QPAT in, e.g., [18, 19, 4, 47] and a Mumford–Shah type approach in [7].

PAT images are *three-dimensional* (3D) and high-resolution, and hence they contain a significant amount of data. Furthermore, QPAT aims at high-resolution 3D reconstructions of the optical parameters. In consequence, a practical inversion method for QPAT must be able to

successfully tackle problems with tens, or even hundreds of thousands of data and unknowns. In this work, we propose an efficient algorithm for the optical inverse problem of QPAT, capable of handling edge-preferring, total variation type priors for the optical parameters. The approach is based on iteratively combining a lagged diffusivity step and a linearization of the measurement model of QPAT with priorconditioned LSQR. The algorithm is a modified version of the one introduced for inverse elliptic boundary value problems in [22, 24]; see also [2] for the original ideas behind the technique. In particular, to facilitate the treatment of far greater amounts of data compared to [22, 24], we implement a matrix-free technique for multiplying vectors by the Jacobian of the measurement map, rendering it possible to painlessly handle (full) Jacobians with, say, 10^5 rows and columns. This is one of the few studies where QPAT is investigated in 3D; for previous works, see [42, 31, 35]. In particular, [18, 19] studied a gradient-based (bound-constrained) split Bregman method for handling TV regularization in 3D QPAT.

The structure of this paper is as follows. The optical measurement model of QPAT is described in section 2. The Bayesian framework is introduced in section 3 and the algorithm itself in section 4. Section 5 tests the approach with 3D numerical simulations. The conclusions are drawn in section 6.

2. Measurement model. We assume the measurements are static in time and model the examined physical body as a bounded domain $\Omega \subset \mathbb{R}^3$ with a connected complement and a Lipschitz boundary. The domain Ω is assumed to be isotropic and the associated diffusion and absorption coefficients are denoted by $\kappa, \mu \in L_+^\infty(\Omega)$, respectively, with the definition

$$L_+^\infty(\Omega) = \{v \in L^\infty(\Omega) \mid \text{ess inf } v > 0\}$$

that takes into account the positivity of these physical quantities. In this work the DA of the RTE is used as the model for light transport. Compared to the full RTE, the DA generally allows faster reconstruction methods due to its simplicity. According to the DA, the photon fluence $\varphi \in H^1(\Omega)$ corresponding to the photon flux $\Phi \in L^2(\partial\Omega)$ through $\partial\Omega$ satisfies the elliptic Robin boundary value problem

$$(2.1) \quad \begin{cases} -\nabla \cdot (\kappa \nabla \varphi) + \mu \varphi = 0 & \text{in } \Omega, \\ \frac{1}{4}\varphi + \frac{1}{2}\nu \cdot \kappa \nabla \varphi = \Phi & \text{on } \partial\Omega, \end{cases}$$

where $\nu : \partial\Omega \rightarrow \mathbb{R}^3$ is the exterior unit normal of $\partial\Omega$ (cf. [20]). We assume the available photoacoustic measurements are noisy samples of the absorbed energy density $H \in L^2(\Omega)$ defined via

$$(2.2) \quad H = \mu \varphi,$$

which obviously depends on the optical parameters κ and μ .

Lemma 2.1. *The Fréchet derivative of the measurement map*

$$[L_+^\infty(\Omega)]^2 \ni (\kappa, \mu) \mapsto H \in L^2(\Omega)$$

at $(\kappa, \mu) \in [L_+^\infty(\Omega)]^2$ is given by the linear mapping

$$[L^\infty(\Omega)]^2 \ni (\vartheta, \theta) \mapsto \mu \varphi' + \theta \varphi \in L^2(\Omega)$$

where $\varphi = \varphi(\kappa, \mu) \in H^1(\Omega)$ is the unique solution of (2.1) and $\varphi' = (\varphi'(\kappa, \mu))(\vartheta, \theta) \in H^1(\Omega)$ that of the variational problem

$$(2.3) \quad \begin{aligned} \int_{\Omega} (\kappa \nabla \varphi' \cdot \nabla v + \mu \varphi' v) dx + \frac{1}{2} \int_{\partial\Omega} \varphi' v dS \\ = - \int_{\Omega} \vartheta \nabla \varphi \cdot \nabla v dx - \int_{\Omega} \theta \varphi v dx \end{aligned}$$

for all $v \in H^1(\Omega)$.

Proof. It is well known that the map

$$[L_+^\infty(\Omega)]^2 \ni (\kappa, \mu) \mapsto \varphi \in H^1(\Omega)$$

is Fréchet differentiable in $[L_+^\infty(\Omega)]^2$ and that the corresponding derivative is given by (cf., e.g., [14])

$$[L^\infty(\Omega)]^2 \ni (\vartheta, \theta) \mapsto \varphi' \in H^1(\Omega).$$

Since the bilinear map

$$L^\infty(\Omega) \times H^1(\Omega) \ni (\lambda, \psi) \mapsto \lambda \psi \in L^2(\Omega)$$

is obviously continuous, the claim follows from the product and the chain rules for Banach spaces. ■

3. Bayesian framework and the choice of prior. Let us next consider the discretized version of (2.1)–(2.2). To begin with, we express the diffusion κ and the absorption μ as exponential quantities,

$$(3.1) \quad \kappa(\tilde{\kappa}) = \kappa_0 \exp(\tilde{\kappa}) \quad \text{and} \quad \mu(\tilde{\mu}) = \mu_0 \exp(\tilde{\mu}),$$

where

$$(3.2) \quad \tilde{\kappa} = \sum_{n=1}^N \tilde{\kappa}_n \phi_n \quad \text{and} \quad \tilde{\mu} = \sum_{n=1}^N \tilde{\mu}_n \phi_n$$

are representations of $\tilde{\kappa}$ and $\tilde{\mu}$ with respect to (w.r.t.) a piecewise linear *finite element* (FE) basis $\{\phi_n\}_{n=1}^N \subset H^1(\Omega)$ corresponding to a chosen (tetrahedral) partition of Ω . If there is no possibility of confusion, we denote by $\tilde{\kappa}$ and $\tilde{\mu}$ both the corresponding vectors of coefficients, $\tilde{\kappa}, \tilde{\mu} \in \mathbb{R}^N$, and the functions defined in (3.2). The positive real numbers $\kappa_0, \mu_0 > 0$ in (3.1) are the constant diffusion and absorption levels that produce an energy density that is the most compatible with the available measurements (cf. Algorithm 1 in section 4). Notice that we have chosen the logarithms of the diffusion and absorption as the free variables since this automatically guarantees the positivity of the coefficient functions in (2.1).

We assume the available measurement is

$$(3.3) \quad \chi = h(\tilde{\kappa}, \tilde{\mu}) + \eta \in \mathbb{R}^N,$$

where $h : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}^N$ are the coefficients in an approximation of $H(\tilde{\kappa}, \tilde{\mu}) = \mu(\tilde{\mu})\varphi(\kappa(\tilde{\kappa}), \mu(\tilde{\mu}))$ w.r.t. the FE basis $\{\phi_n\}_{n=1}^N$,¹

$$(3.4) \quad H \approx \sum_{n=1}^N h_n \phi_n.$$

Moreover, $\eta \in \mathbb{R}^N$ is a realization of a normally distributed random variable with zero mean and a known, symmetric, and positive definite covariance matrix $\Gamma \in \mathbb{R}^{N \times N}$. It easily follows that the probability density of the measurements given the parameters is

$$p(\chi | \tilde{\kappa}, \tilde{\mu}) \propto \exp\left(-\frac{1}{2}(\chi - h(\tilde{\kappa}, \tilde{\mu}))^T \Gamma^{-1}(\chi - h(\tilde{\kappa}, \tilde{\mu}))\right),$$

where the constant of proportionality is independent of $\tilde{\kappa}$ and $\tilde{\mu}$.

The prior information that the optical properties of the examined body are approximately homogeneous apart from clearly distinguishable inhomogeneities is taken into account by equipping the logarithms of the diffusion and absorption with the prior densities

$$(3.5) \quad p(\tilde{\kappa}) \propto \exp(-a R(\tilde{\kappa})) \quad \text{and} \quad p(\tilde{\mu}) \propto \exp(-b R(\tilde{\mu}))$$

where $a, b > 0$ are free parameters and R is of the form

$$(3.6) \quad R(u) = \int_{\Omega} r(|\nabla u(x)|) dx,$$

with $r : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ being a suitable, continuously differentiable, monotonically increasing function (cf., e.g., [2]). All numerical examples presented in this work are based on a Perona–Malik prior, i.e.,

$$(3.7) \quad r(t) = \frac{1}{2} T^2 \log(1 + (t/T)^2),$$

where $T > 0$ is a small parameter controlling the size of detectable edges [34]. However, exactly the same reconstruction algorithm could as well be employed in the context of, e.g., (smoothened) TV or TV^q regularization by simply using another choice of r (cf. [2, 24]).

Under the assumption that $\tilde{\kappa}$ and $\tilde{\mu}$ are independent, the Bayes' formula yields

$$\begin{aligned} p(\tilde{\kappa}, \tilde{\mu} | \chi) &\propto p(\chi | \tilde{\kappa}, \tilde{\mu}) p(\tilde{\kappa}) p(\tilde{\mu}) \\ &\propto \exp\left(-\frac{1}{2}(\chi - h(\tilde{\kappa}, \tilde{\mu}))^T \Gamma^{-1}(\chi - h(\tilde{\kappa}, \tilde{\mu})) - a R(\tilde{\kappa}) - b R(\tilde{\mu})\right), \end{aligned}$$

¹In our numerical experiments, the actual measurement for given target optical coefficients—which are *not* originally represented in the basis $\{\phi_n\}_{n=1}^N$ —are simulated by approximating H in a piecewise linear FEM basis on a denser simulation grid, interpolating the obtained node values onto the mesh corresponding to $\{\phi_n\}_{n=1}^N$, and adding noise. In particular, the model (3.3) is clearly inexact as it does not account for the interpolation step in the data simulation; this corresponds to avoiding an obvious inverse crime.

where the constants of proportionality do not depend on $\tilde{\kappa}$ and $\tilde{\mu}$. The algorithm described in the following section seeks an (approximate) maximum a posteriori (MAP) estimate, i.e., the maximizer of this posterior, or, equivalently, tries to approximate the minimizer for the functional

$$(3.8) \quad F(\tilde{\kappa}, \tilde{\mu}) := \frac{1}{2}(\chi - h(\tilde{\kappa}, \tilde{\mu}))^T \Gamma^{-1}(\chi - h(\tilde{\kappa}, \tilde{\mu})) + a R(\tilde{\kappa}) + b R(\tilde{\mu}).$$

In the following, we denote $\beta = [\tilde{\kappa}^T, \tilde{\mu}^T]^T \in \mathbb{R}^{2N}$ and write occasionally $h(\beta)$ and $F(\beta)$ to shorten the notation.

Remark 3.1. It is arguably unrealistic to assume that $\tilde{\kappa}$ and $\tilde{\mu}$ are independent because both of them are prone to change at interfaces between different tissues. However, since assuming a dependence between the two parameters would make the setting less general as well as less ill-posed, we leave considerations on introducing a suitable joint prior for future studies.

4. The algorithm. In this section we briefly introduce our method for minimizing (3.8); for more information, see [2, 22, 24]. The algorithm is only described here for a single boundary photon flux Φ , but it trivially generalizes to the case of multiple illuminations.

The basic version of the iterative algorithm starts at the initial guess $\beta^{(0)} = [\tilde{\kappa}_{\text{init}}^T, \tilde{\mu}_{\text{init}}^T]^T$, where $\tilde{\kappa}_{\text{init}} = 0 \in \mathbb{R}^N$ corresponds to the constant diffusion κ_0 (cf. (3.1)) and $\tilde{\mu}_{\text{init}} = \log(\chi/(\varphi_0 \mu_0))$, with $\varphi_0 = \varphi(\kappa_0, \mu_0)$ being the fluence corresponding to the homogeneous parameter values κ_0 and μ_0 . Note that the choice of $\tilde{\mu}_{\text{init}}$ is motivated by (2.2) and (3.1). Linearizing $h(\beta)$ around $\beta^{(l)}$ in (3.8) results in a new functional

$$(4.1) \quad F^{(l)}(\beta) := \frac{1}{2}(y^{(l)} - J^{(l)}\beta)^T \Gamma^{-1}(y^{(l)} - J^{(l)}\beta) + a R(\tilde{\kappa}) + b R(\tilde{\mu}),$$

where the matrix $J^{(l)} \in \mathbb{R}^{N \times 2N}$ is the Jacobian of the map $\beta \mapsto h(\beta)$ evaluated at $\beta^{(l)}$ and

$$y^{(l)} = \chi - h(\beta^{(l)}) + J^{(l)}\beta^{(l)} \in \mathbb{R}^N.$$

For a given illumination $\Phi \in L^2(\Omega)$, the coefficients $h(\beta^{(l)})$ in (3.4) are solved from (2.1)–(2.2) with the optical parameters defined via (3.1) by the *finite element method* (FEM) employing the aforementioned piecewise linear basis functions $\{\phi_n\}_{n=1}^N$. The approximation of the elements in $J^{(l)}$ is based on Lemma 2.1 and the chain rule; the details are given in Appendix A. In particular, the Jacobians are not formed explicitly, but the needed matrix-vector multiplications are performed row by row in a matrix-free manner. This enables painless handling of (full) Jacobians with tens of thousands of rows and columns in our numerical experiments. Details about the utilized FE meshes can be found in section 5.

Taking the gradient of (4.1), one obtains the necessary condition for a minimizer, that is,

$$(4.2) \quad (J^{(l)})^T \Gamma^{-1} J^{(l)} \beta + a \begin{bmatrix} (\nabla R)(\tilde{\kappa}) \\ 0 \end{bmatrix} + b \begin{bmatrix} 0 \\ (\nabla R)(\tilde{\mu}) \end{bmatrix} = (J^{(l)})^T \Gamma^{-1} y^{(l)},$$

where $0 \in \mathbb{R}^N$. The gradient $\nabla R : \mathbb{R}^N \rightarrow \mathbb{R}^N$ can be given as [2]

$$(\nabla R)(u) = M(u)u,$$

where $M \in \mathbb{R}^{N \times N}$ is the FEM system matrix in the basis $\{\phi_n\}_{n=1}^N$ for the elliptic partial differential operator

$$(4.3) \quad -\nabla \cdot c_u \nabla$$

with a natural boundary condition on $\partial\Omega$ and the positive-valued diffusion coefficient

$$c_u : x \mapsto \frac{r'(|\nabla u(x)|)}{|\nabla u(x)|}, \quad \Omega \rightarrow \mathbb{R}_+.$$

It follows easily that M is positive semidefinite with the one-dimensional kernel $\text{Ker}(M) = \text{span}\{[1, \dots, 1]^T\}$.

We rewrite (4.2) as

$$(4.4) \quad \left((J^{(l)})^T \Gamma^{-1} J^{(l)} + \begin{bmatrix} a M(\tilde{\kappa}) & 0 \\ 0 & b M(\tilde{\mu}) \end{bmatrix} \right) \beta = (J^{(l)})^T \Gamma^{-1} y^{(l)},$$

and get rid of its nonlinearity w.r.t. $\beta = [\tilde{\kappa}^T, \tilde{\mu}^T]^T$ by substituting $M(\tilde{\kappa}^{(l)})$ and $M(\tilde{\mu}^{(l)})$ for $M(\tilde{\kappa})$ and $M(\tilde{\mu})$, respectively. This corresponds to a single lagged diffusivity step [49]. Denoting a Cholesky factor of Γ^{-1} by $\Gamma^{-1/2}$ and setting

$$(4.5) \quad A = \Gamma^{-1/2} J^{(l)}, \quad M = \begin{bmatrix} M(\tilde{\kappa}^{(l)}) & 0 \\ 0 & \frac{b}{a} M(\tilde{\mu}^{(l)}) \end{bmatrix}, \quad \tilde{y} = \Gamma^{-1/2} y^{(l)},$$

we finally arrive at the equation

$$(4.6) \quad (A^T A + aM) \beta = A^T \tilde{y}, \quad a > 0,$$

from which $\beta^{(l+1)}$ is to be solved.

Solving (4.6) is equivalent to determining the MAP or *conditional mean* (CM) estimate for the linear model

$$(4.7) \quad A\beta = \tilde{y}$$

assuming a suitable additive Gaussian measurement noise model and for β a zero-mean (improper) Gaussian prior with a scaled version of M as the inverse covariance matrix. In our setting, the leading idea of priorconditioning [8, 9, 10, 11]² is to include the prior information in M directly in the Krylov subspace structure produced by LSQR [32, 33]. As M is only positive semidefinite with a nontrivial kernel, we approximate it in (4.6) by the positive definite matrix

$$(4.8) \quad M_\delta = M + \delta I$$

where $\delta > 0$ is a small positive constant, that is, we consider

$$(4.9) \quad (A^T A + aM_\delta) \beta = A^T \tilde{y}, \quad a > 0,$$

²Priorconditioning is related to transforming a Tikhonov functional into the standard form [16, 23, 25].

in place of (4.6). In terms of the MAP estimate for (4.7), this corresponds to assuming that the inverse covariance matrix of the Gaussian prior is proportional to M_δ instead of M , making the prior proper.

We formally introduce a (Cholesky) factorization $M_\delta = L^T L$, but emphasize that such is not actually needed in the final algorithm because we resort to a version of LSQR that is compatible with symmetric preconditioning [2]. Subsequently, (4.9) is multiplied from the left by $(L^{-1})^T$ and a is chosen to be zero, which altogether leads to the ill-posed linear equation

$$(4.10) \quad (L^{-1})^T A^T A L^{-1} \tilde{\beta} = (L^{-1})^T A^T \tilde{y},$$

where $\tilde{\beta} = L\beta$. We solve (4.10) by combining LSQR [2] with an early stopping rule; loosely speaking, the regularization provided by $a > 0$ is replaced with the early stopping of a Krylov subspace method. Each round of LSQR includes one multiplication with M_δ^{-1} , which is not overly expensive as M_δ results from a discretization of an elliptic partial differential equation and is, in particular, sparse. If one starts the LSQR iteration from $\tilde{\beta} = 0$, it is easy to see that the approximate solution is in the range of M_δ^{-1} regardless of the number of iterations; see [2, 22] for more details. As M_δ^{-1} is proportional to the (fictive) prior covariance matrix for (4.9), this means that the prior information in M_δ —originating from the previous iterates $\tilde{\kappa}^{(l)}$ and $\tilde{\mu}^{(l)}$ of the outer loop (cf. (4.5) and (4.8))—is indeed directly included in the candidate solutions for (4.10) produced by the LSQR sequence.

The LSQR iteration is terminated when the residual for (4.7) ceases to decrease substantially: We monitor the relative reduction of the residual over a window of m_0 LSQR steps,

$$(4.11) \quad r_m = 1 - \frac{|A\beta_m - \tilde{y}|}{|A\beta_{m-m_0} - \tilde{y}|}, \quad m > m_0,$$

where β_m is the m th element in the LSQR sequence. Once $r_m \leq \tau$, for some user specified $\tau > 0$ and $m_0 \in \mathbb{N}$, the latest iterate β_m is named $\beta^{(l+1)}$ and one proceeds to the next linearization of the measurement model (cf. (4.1)).

Including an overall stopping criterion based on tracking the decrease of the residual for the original nonlinear measurement model, our reconstruction algorithm is altogether as follows:

Algorithm 1. Select $T > 0$, the ratio b/a for (3.5), the parameters $m_0 \in \mathbb{N}$ and $\tau > 0$ related to (4.11), and $\delta > 0$. Let $\mathbf{1} = [1, \dots, 1]^T \in \mathbb{R}^N$ and determine (κ_0, μ_0) as the minimizing pair for

$$|\Gamma^{-1/2}(\chi - h(\log(\kappa)\mathbf{1}, \log(\mu)\mathbf{1}))|$$

over $(\kappa, \mu) \in \mathbb{R}_+^2$. Solve $\varphi_0 = \varphi(\kappa_0\mathbf{1}, \mu_0\mathbf{1})$ and initialize $\tilde{\kappa}_{\text{init}} = 0 \in \mathbb{R}^N$, $\tilde{\mu}_{\text{init}} = \log(\chi/(\varphi_0\mu_0))$. Set $\beta^{(0)} = [\tilde{\kappa}_{\text{init}}^T, \tilde{\mu}_{\text{init}}^T]^T$ and $l = 0$.

1. Let $h = h(\beta^{(l)})$, $J = J^{(l)}$, and $\tilde{y} = \Gamma^{-1/2}(\chi - h + J\beta^{(l)})$. (Recall that J is not built explicitly, but the corresponding matrix-vector multiplications are performed in a matrix-free manner as explained in Appendix A.)
2. Build $M(\tilde{\kappa}^{(l)})$ and $M(\tilde{\mu}^{(l)})$ as finite element discretizations of (4.3) and form $M_\delta = L^T L$ according to (4.5) and (4.8).

3. Apply the LSQR algorithm of [2] to

$$(L^{-1})^T J^T (\Gamma^{-1/2})^T \Gamma^{-1/2} J L^{-1} \tilde{\beta} = (L^{-1})^T J^T (\Gamma^{-1/2})^T \tilde{y}, \quad \beta = L^{-1} \tilde{\beta},$$

starting from $\tilde{\beta} = 0$. Terminate the iteration when $r_m \leq \tau$ (cf. (4.11)) and denote the corresponding solution $\beta^{(l+1)}$.

4. If the nonlinear residual corresponding to (3.3) has not decreased, i.e.,

$$|\Gamma^{-1/2}(\chi - h(\beta^{(l+1)}))| \geq |\Gamma^{-1/2}(\chi - h)|,$$

substitute the previous parameter vector $\beta^{(l)} = [(\tilde{\kappa}^{(l)})^T, (\tilde{\mu}^{(l)})^T]^T$ in (3.1) and declare the resulting κ and μ the reconstruction. Otherwise, set $l \leftarrow l+1$ and return to step 1.

The performance of Algorithm 1 is relatively insensitive to the choice of the free parameters $T, \delta > 0$, which are set to $T = 5 \cdot 10^{-3}$ and $\delta = 10^{-6}$ in our numerical experiments. Moreover, we choose $b/a = 1$, which means that we assume as strong priors for $\tilde{\kappa} = \log(\kappa/\kappa_0)$ and $\tilde{\mu} = \log(\mu/\mu_0)$; cf. (3.5). The choice of m_0 and τ is a more delicate issue: via trial and error, we ended up setting $m_0 = 10$ and $\tau = 10^{-2}$, that is, all LSQR iterations in our numerical tests are terminated when the residual for (4.7) decreases less than one percent over ten steps. These are certainly not optimal values for m_0 and τ , but they seem rather generic and result in adequate reconstructions.

With K illuminations $\Phi^{(1)}, \dots, \Phi^{(K)} \in L^2(\partial\Omega)$, the number of measurements increases from N to KN , which in turn results in larger Jacobians and slower computations. Be that as it may, Algorithm 1 trivially generalizes to such a setting: one just needs to stack the individual measurements in a vector of length KN , form the corresponding covariance matrix for the measurement noise as a block diagonal matrix of the original covariances, and build the “total Jacobian” by piling the “sub-Jacobians” on top of each other (in a matrix-free manner). The only essential change concerns the initialization of the absorption coefficient: we form $\exp(\tilde{\mu}_{\text{init}}^{(k)})$, $k = 1, \dots, K$, separately for each illumination as described in Algorithm 1 and subsequently choose the exponential of the actual initial guess $\exp(\tilde{\mu}_{\text{init}})$ to be their average.

Remark 4.1. The early stopping of LSQR in step 3 of Algorithm 1 has previously been successfully implemented in the contexts of electrical impedance and optical tomography by resorting to the Morozov discrepancy principle [24, 22]. In our setting, this would lead to monitoring when the residual for (4.7) falls below the (whitened) noise level

$$\sqrt{\mathbb{E}(|\Gamma^{-1/2}\eta|^2)} = \sqrt{KN}.$$

However, according to our experience, such an approach does not work, in general, for QPAT based on *simulated* interior measurements without committing an inverse crime or using a dubiously large “fudge factor” to scale the noise level: The discrepancy associated to the interpolation of the target energy density H , which is computed on a denser simulation grid (cf. (3.4)), onto the FEM mesh employed in Algorithm 1 is easily of the same order as the error corresponding to the one percent of artificial noise that is added to the data in the numerical experiments of section 5. This effect is particularly pronounced if the target absorption and diffusion contain jumps that cause quick deviations in H and, therefore, also lead to larger interpolation errors.

Figure 1. Test 1. Left: the target absorption. Right: the target diffusion. The values in the intervals $[\mu_{bg} \pm 0.001]$ and $[\kappa_{bg} \pm 0.01]$ indicated in the colorbars are transparent in the respective images.

5. Numerical experiments. To demonstrate the performance of our algorithm, we present two numerical experiments. The first one considers a complicated target and several illuminations. The second test studies how the number and directions of illuminations affect the quality of reconstructions for a somewhat simpler phantom.

Test 1. We first examine the cylindrical body with constant background properties and embedded inhomogeneities visualized in Figure 1. The shapes of the target inclusions are described in Table 1 and the corresponding constant values of the optical parameters are listed in Table 2. In particular, the target diffusion corresponds to a reduced scattering coefficient that varies in the interval $[0.5, 6.7] \text{ mm}^{-1}$, which means that our values for the optical parameters could mimic, e.g., those in breast tissue [3]. We illuminate the object in turns with $K = 4$ photon fluxes that penetrate the boundary through rectangular regions located symmetrically around the curved side of the cylinder. The illuminations are homogeneous along the axis of the cylinder, their central polar angles are $\theta_0^{(k)} = (k-1)\pi/2$, $k = 1, \dots, 4$, and the angular width of each illumination is $\pi/4$. The amplitude of the input flux $\Phi^{(k)} : \partial\Omega \rightarrow \mathbb{R}$, $k = 1, \dots, 4$, is modeled as

$$\Phi^{(k)}(\theta) = \begin{cases} \cos(4(\theta - \theta_0^{(k)})) & \text{if } \theta \in [\theta_0^{(k)} - \frac{\pi}{8}, \theta_0^{(k)} + \frac{\pi}{8}], \\ 0 & \text{otherwise,} \end{cases}$$

where θ is the polar angle w.r.t. the axis of the cylindrical domain.

To simulate the data for the numerical experiments, we first solve for each illumination $\Phi = \Phi^{(k)}$ the photon fluence $\varphi^{(k)}$ from (2.1) by using a fine FE mesh with $N_f = 130091$ nodes and 690905 tetrahedrons, and then calculate the corresponding node values of the absorbed energy density $h_f^{(k)} \in \mathbb{R}^{N_f}$ according to (2.2). Next, in order to avoid an obvious inverse crime, we project the absorbed energy density onto a coarser FE mesh ($N = 51794$ nodes and 260216 tetrahedrons) which is also used in the reconstruction algorithm. In other words, the

Table 1

Test 1. Geometrical specification of the inclusions in the target absorption and diffusion illustrated in Figure 1. The unit of length is mm.

Absorption	Diffusion
rectangles: size $4 \times 6 \times 4$	(1) cylinder: radius 1, center $(0, y, 0)$, $y \in [-20, 20]$
(1) center $(0, -11, 0)$	
(2) center $(0, 0, 0)$	
(3) center $(0, 11, 0)$	cubes: size 4^3 , center $(\rho \cos \theta, y, \rho \sin \theta)$, $\rho = 5.5$
helical cylinders: radius 1, center $(\rho \cos \theta, y, \rho \sin \theta)$, $\rho = 5.5$, $y \in [-16, 16]$	(2) $\theta = \pi/6$, $y = -15$
(4) $\theta \in [\pi/6, 11\pi/6]$	(3) $\theta = 3\pi/6$, $y = -9$
(5) $\theta \in [7\pi/6, 17\pi/6]$	(4) $\theta = 5\pi/6$, $y = -3$
	(5) $\theta = 7\pi/6$, $y = 3$
	(6) $\theta = 9\pi/6$, $y = 9$
	(7) $\theta = 11\pi/6$, $y = 15$

Figure 2. Test 1. Visualization of the measurement. Two slices of each $h^{(k)}$, $k = 1, \dots, 4$, are plotted: one at the position $y = 0$, the other at $z = 0$ for $k = 1, 3$ (left) and $x = 0$ for $k = 2, 4$ (right) depending on the direction of the illumination.

(noiseless) measurement, which is illustrated in Figure 2, consists of the nodal values

$$h^{(k)} = Ph_f^{(k)} \in \mathbb{R}^N, \quad k = 1, \dots, 4,$$

where the matrix $P \in \mathbb{R}^{N \times N_f}$ describes the linear interpolation from the fine FE mesh onto the coarse one. Finally, we corrupt the measurement with 1% of Gaussian noise, that is, we end up with the data $\chi = \{\chi^{(k)}\}_{k=1,\dots,4} \subset \mathbb{R}^N$, where

$$\chi_i^{(k)} = h_i^{(k)} + \eta_i^{(k)}$$

and $\eta_i^{(k)} \sim \mathcal{N}(0, (0.01|h_i^{(k)}|)^2)$ for $i = 1, \dots, N$ and $k = 1, \dots, K$.

In the initialization phase of Algorithm 1, we choose the free parameters as described at the end of section 4,

$$T = 5 \cdot 10^{-3}, \quad b/a = 1, \quad m_0 = 10, \quad \tau = 10^{-2}, \quad \delta = 10^{-6},$$

and obtain the approximate background values $\kappa_0 = 0.26 \text{ mm}$ and $\mu_0 = 0.0087 \text{ mm}^{-1}$ for the optical parameters. To construct the initial guess for the absorption shown in Figure 3, we

Figure 3. Test 1. The initial guess μ_{init} for the absorption. Left: a 3D-visualization; the values in the interval $[\mu_0 \pm 0.004]$ indicated in the colorbar are transparent. Right: slices at the positions $y = -11, 0, 11$ (top row) and $x = 0$ (bottom row).

compute $\varphi_0^{(k)} = \varphi^{(k)}(\kappa_0 \mathbf{1}, \mu_0 \mathbf{1})$ for $k = 1, \dots, 4$ and set $\tilde{\mu}_{\text{init}} = \log(\mu_{\text{init}}/\mu_0)$, with

$$\mu_{\text{init}} = \frac{1}{4} \sum_{k=1}^4 \frac{\chi^{(k)}}{\varphi_0^{(k)}}.$$

According to Algorithm 1, the initial guess for the logarithm of the diffusion is $\tilde{\kappa}_{\text{init}} = 0 \in \mathbb{R}^N$, which corresponds to the homogeneous $\kappa_{\text{init}} = \kappa_0 \mathbf{1} \in \mathbb{R}^N$ estimate. However, starting from a homogeneous κ_{init} leads to somewhat slow convergence during the first rounds of Algorithm 1. To avoid this, we employ the (already quite reasonable) initial guess μ_{init} for the absorption to run a single LSQR iteration only for the diffusion. To be more precise, we include a “zeroth step” in Algorithm 1 in between the initialization and step 1:

0. Set $h = h(\tilde{\kappa}_{\text{init}}, \tilde{\mu}_{\text{init}})$. Form the Jacobian $J = J_{\tilde{\kappa}}(\tilde{\kappa}_{\text{init}}, \tilde{\mu}_{\text{init}})$ (w.r.t. $\tilde{\kappa}$ evaluated at $(\tilde{\kappa}_{\text{init}}, \tilde{\mu}_{\text{init}})$), set $\hat{y} = \Gamma^{-1/2}(\chi - h + J\tilde{\kappa}_{\text{init}})$, and build the matrix $M_{\delta}(\tilde{\kappa}_{\text{init}}) = L^T L$. Apply the LSQR algorithm to solve

$$(L^{-1})^T J^T (\Gamma^{-1/2})^T \Gamma^{-1/2} J L^{-1} \hat{\kappa} = (L^{-1})^T J^T (\Gamma^{-1/2})^T \hat{y}, \quad \tilde{\kappa} = L^{-1} \hat{\kappa}$$

stating from $\hat{\kappa} = 0$ and terminating when $r_m \leq \tau$. (Re)define the initial guess $\tilde{\kappa}_{\text{init}}$ to be the corresponding solution. Reset $\beta^{(0)} = [\tilde{\kappa}_{\text{init}}^T, \tilde{\mu}_{\text{init}}^T]^T$ and continue to step 1.

The resulting diffusion estimate $\kappa_{\text{init}} = \kappa_0 \exp(\tilde{\kappa}_{\text{init}})$ is illustrated in Figure 4. It is not as accurate as the initial guess for the absorption, but inclusions have already started to form at the correct positions.

The algorithm then proceeds in the standard way, i.e., with the simultaneous reconstruction of the absorption and the diffusion; see steps 1–4 in Algorithm 1. The residuals for both the outer and the inner iteration are depicted in Figure 5. The algorithm terminates after five linearizations (including step 0), meaning that the reconstruction, presented in Figure 6, is a result of four linearizations of the measurement model (cf. step 4 of Algorithm 1). The reconstructions of both optical parameters are in accordance with the prior information: there are well localized inclusions in an approximately constant background. Moreover, the

Figure 4. Test 1. The (refined) initial guess κ_{init} for the diffusion. Left: a 3D-visualization; the values in the interval $[\kappa_0 \pm 0.075]$ indicated in the colorbar are transparent. Right: slices at the positions $y = -15, -9, -3$ (top row), $y = 3, 9, 15$ (middle row), and $x = 0$ (bottom row).

Figure 5. Test 1. The evolution of the residuals after each iteration step. During step 0, only the diffusion is estimated. The dashed lines indicate the theoretical (whitened) noise level $\sqrt{KN}$. Left: The nonlinear residuals $|\Gamma^{-1/2}(\chi - h(\beta))|$ correspond to the outer iteration. Right: The LSQR residuals $|A\beta - \tilde{y}|$ correspond to the inner iteration on a logarithmic scale.

inclusions lie at approximately correct positions. On the negative side, the reconstruction of the diffusion exhibits some instability near the object boundary, which causes very high jumps at a few isolated nodes: the highest and the lowest point values are $\max(\kappa) = 8.47$ and $\min(\kappa) = 0.0086$, respectively. However, discarding the nodes at the boundary, one gets the reasonable extremal values $\max(\kappa|_{\Omega \setminus \partial\Omega}) = 0.78$ and $\min(\kappa|_{\Omega \setminus \partial\Omega}) = 0.03$ (which are used

Figure 6. Test 1. The reconstructions of the absorption (above) and the diffusion (below). Left: 3D-visualizations; the values in the intervals $[\mu_0 \pm 0.0025]$ and $[\kappa_0 \pm 0.1]$ indicated in the colorbars are transparent in the respective images. Right: slices of the reconstructed absorption at the positions $y = -11, 0, 11$ and $x = 0$ and of the reconstructed diffusion at the positions $y = -15, -9, -3, 3, 9, 15$ and $x = 0$.

as the limits for the diffusion plot in Figure 6). In addition, in spite of the aforementioned outliers in the reconstructed diffusion, the mean values over the background and the inclusions presented in Table 2 are approximately correct for both parameters. Even though the reconstructed means are not exactly the same as the target values, they are very close to the corresponding mean values of the interpolated parameters $P\mu_{\text{target}}$ and $P\kappa_{\text{target}}$ (except for the “most difficult” diffusive inclusion (1) lying along the axis of Ω), which is arguably the best one can expect to achieve.

As mentioned after Algorithm 1, the reconstructions seem insensitive to the choice of the threshold parameter $T > 0$. Considering significantly higher noise levels (say, 10%) clearly deteriorates the quality of the reconstructions, with the diffusion coefficient exhibiting a higher

Table 2

Test 1. The mean values of the absorption and the diffusion in the target illustrated in Figure 1, in the target interpolated onto the sparser grid, and in the reconstruction shown in Figure 6. The mean values are taken over the correct supports of the inclusions listed in Table 1.

Absorption	Mean values (mm^{-1})			Diffusion	Mean values (mm)		
	μ_{target}	$P\mu_{\text{target}}$	μ_{rec}		κ_{target}	$P\kappa_{\text{target}}$	κ_{rec}
bg	0.01	0.0101	0.00996	bg	0.3	0.300	0.304
(1)	0.05	0.0459	0.0456	(1)	0.05	0.0874	0.0708
(2)	0.02	0.0189	0.0188	(2)	0.05	0.0668	0.0628
(3)	0.002	0.00266	0.00263	(3)	0.15	0.163	0.165
(4)	0.05	0.0432	0.0425	(4)	0.6	0.579	0.595
(5)	0.002	0.00338	0.00335	(5)	0.05	0.0792	0.0764
				(6)	0.15	0.169	0.174
				(7)	0.6	0.573	0.576

level of instability. However, e.g., doubling the noise level does not considerably alter the performance of the algorithm. Moderate changes in the density of the reconstruction mesh mainly affect the reconstruction process via the computation time.

The running time of Algorithm 1 for this experiment was approximately 12 minutes with a MATLAB (2014a) implementation on a laptop with 16 GB RAM and an Intel Core i7-4600U CPU having clock speed 2.10 GHz.

Test 2. In practical applications it is not always possible to illuminate the target from several different directions. We next study the effect of the number and position of illuminations on the cubical target of size 11^3 mm^3 visualized in Figure 7a. The absorption is composed of a homogeneous background $\mu_{\text{bg}} = 0.015 \text{ mm}^{-1}$ with two embedded inhomogeneities: a cross-shaped inclusion lying along the plane $z = x$ with absorption 0.01 mm^{-1} and an origin-centered spherical shell with outer radius 5 mm, inner radius 4 mm and the absorption level 0.02 mm^{-1} (except at the intersection with the cross). The background diffusion level is $\kappa_{\text{bg}} = 0.3 \text{ mm}$ and there are also two diffusive inhomogeneities: a cross-shaped inclusion lying along the plane $z = -x$ with diffusion 0.4 mm and a ball centered at the origin with radius 3 mm and diffusion 0.2 mm (except at the intersection with the cross).

To begin with, we illuminate the object through its bottom face with only one photon flux $\Phi = \Phi^{\text{btm}} : \partial\Omega \rightarrow \mathbb{R}$ that is modeled as the characteristic function

$$\Phi^{\text{btm}}(x) = \begin{cases} 1 & \text{if } x \in \partial\Omega_{\text{btm}}, \\ 0 & \text{otherwise,} \end{cases}$$

where $\partial\Omega_{\text{btm}} = \{(x, y, z) \in \mathbb{R}^3 : |x| \leq 5.5, |y| \leq 5.5, z = -5.5\}$. The simulation of the data follows the same steps as in the previous example; in particular, the level of additive noise is still one percent. (In this case, the fine FE mesh has $N_f = 133649$ nodes and 752914 tetrahedrons whereas the coarse one has $N = 54721$ nodes and 295176 tetrahedrons.)

Figure 7. Test 2. In the 3D-visualizations the values in the intervals indicated in the colorbars are transparent in the respective images. The slices are taken for both parameters along the planes $z = x$ (left) and $z = -x$ (right).

Using the same free parameters and forming the initial guesses $\tilde{\kappa}_{\text{init}}$ and $\tilde{\mu}_{\text{init}}$ as in the first example, we get the approximate background levels $\kappa_0 = 0.29$ and $\mu_0 = 0.015$ as well as the final reconstruction presented in Figure 7b. The reconstruction of the absorption is reasonable, although it includes shadows of the diffusive inclusions. On the other hand, the diffusion produced by Algorithm 1 is almost constant, making it practically useless. This reconstruction, resulting in a (nonlinear) residual clearly under the noise level $\sqrt{N}$, was obtained after only two steps of LSQR already in step 0 of the algorithm. The reconstructions one gets by starting with $\tilde{\kappa}_{\text{init}} = 0$, i.e., without step 0, or even by setting $\tilde{\mu}_{\text{init}} = 0$ are essentially the same. In addition, changing the direction of the illumination does not seem to affect the results. We conclude that when using only one illumination, the absorption can explain the measurement (almost) completely, and hence one cannot simultaneously reconstruct the diffusion. This is not very surprising as it is well known that the optical inverse problem of QPAT is nonunique for one illumination; see [4, 13, 31, 37, 44]. However, to add an extra twist, [7] recently presented a uniqueness result for *piecewise constant* absorption and diffusion with only one illumination.

Let us then complement the measurement corresponding to the flux Φ^{btm} with a second one. To evaluate the effect of the chosen illumination directions, we consider two cases: illuminations through adjacent faces and through opposite sides. To study the first case, we combine Φ^{btm} with Φ^{rgt} , i.e., a homogeneous flux through the right-hand face of the cube. In the second case, Φ^{btm} is accompanied by Φ^{top} , i.e., a homogeneous flux through the top facet. Both Φ^{rgt} and $\Phi^{\text{top}} : \partial\Omega \rightarrow \mathbb{R}$ are modeled by the appropriate characteristic functions on $\partial\Omega$. The reconstruction algorithm is run for the two cases with the same free parameters as previously. As a result, we get the same approximate background values as with one illumination, but the final reconstructions, presented in Figures 7c–7d, are significantly better: two illuminations result in decent reconstructions of both optical parameters. It has to be noted, however, that the directions of the two employed photon fluxes play a significant role: for the illuminations through adjacent faces, the reconstructed diffusion is inaccurate close to the opposite edge, whereas the opposite illuminations lead to good quality reconstructions of both parameters in the whole domain.

Finally, we add a third measurement to the two-illumination settings considered above. The additional measurement is induced by a homogeneous photon flux through the back of the cube, with its amplitude $\Phi^{\text{bck}} : \partial\Omega \rightarrow \mathbb{R}$ modeled once again by the appropriate characteristic function. The triplet Φ^{btm} , Φ^{rgt} , and Φ^{bck} corresponds to illuminations around one corner of the cube, while the facetwise fluxes Φ^{btm} , Φ^{top} , and Φ^{bck} for the other three-illumination test cover the boundary of the cube as evenly as possible. The resulting reconstructions are presented in Figures 7e–7f. Compared to the two-illumination cases, the reconstructions of the diffusion are now somewhat sharper. However, as with two adjacent illuminations, when the three photon fluxes are supported around one corner, the reconstruction of the diffusion close to the opposite corner is far from satisfactory. For three illuminations, there should be no problem with the (theoretical) uniqueness, so it seems that the physics of the measurement setting limits the reconstruction quality: the light does not penetrate deep enough into the object in order to provide reliable information about the diffusion near the opposite corner, where the reconstruction mainly reflects the prior information. On the other hand, complementing the opposite illuminations with a third measurement leads to slightly

more detailed reconstructions, but the difference between the two- and three-illumination cases is not substantial (cf. Figures 7d and 7f).

The reconstructions corresponding to two illuminations shown in Figures 7c–7d resulted from three linearizations of the measurement model in Algorithm 1 (including step 0), while it took four linearizations to produce the three-illumination reconstructions in Figures 7e–7f. The corresponding computation times were about five and ten minutes, respectively, with the same hardware as in Test 1.

6. Concluding remarks. In this work, the inverse problem of QPAT was investigated. The aim of QPAT is to produce high-resolution reconstructions of the optical parameters of interest from given 3D high-resolution PAT images. A computationally efficient algorithm for the inverse problem of QPAT was introduced by combining priorconditioned LSQR, a lagged diffusivity step, and linearizations of the measurement model in two nested iterations. To facilitate handling a high number of measurements and unknowns, all multiplications by Jacobians in the algorithm were implemented in a matrix-free manner. The numerical studies exclusively employed a Perona–Malik prior, although other types of edge-preserving priors can as easily be utilized in the described approach. The proposed reconstruction algorithm was tested with 3D numerical simulations. The results demonstrate that the approach is capable of producing accurate and good quality estimates of the absorption and diffusion in complex 3D geometries in a reasonable computation time, even when the tests are run on a standard laptop. This suggests that QPAT can be developed into a practical method without compromising the accuracy of the estimates or the good resolution the method can potentially provide.

Appendix A. Matrix-free approximation of the Jacobians. The aim of this appendix is to explain how to efficiently approximate multiplications by the Jacobian $J_{\tilde{\kappa}, \tilde{\mu}}^h$ of the map

$$\mathbb{R}^{2N} \ni (\tilde{\kappa}, \tilde{\mu}) \mapsto h(\kappa(\tilde{\kappa}), \mu(\tilde{\mu})) \in \mathbb{R}^N,$$

where h is the discretization of the measurement $H(\kappa(\tilde{\kappa}), \mu(\tilde{\mu})) = \mu(\tilde{\mu})\varphi(\kappa(\tilde{\kappa}), \mu(\tilde{\mu}))$ (cf. (3.4)) and

$$\kappa(\tilde{\kappa}) = \kappa_0 \exp(\tilde{\kappa}) \quad \text{and} \quad \mu(\tilde{\mu}) = \mu_0 \exp(\tilde{\mu})$$

are the elementwise transformations of the parameters introduced in (3.1). As in sections 3 and 4, we only consider here the case of one illumination, which corresponds to the photon flux Φ ; cf. (2.1). In the case of multiple illuminations, more bookkeeping of the indices is required, but otherwise the described methodology works as for a single flux of photons.

First of all, since $\kappa_n = \kappa(\tilde{\kappa}_n)$ and $\mu_n = \mu(\tilde{\mu}_n)$ for all $n = 1, \dots, N$, we get the Jacobian $J_{\tilde{\kappa}, \tilde{\mu}}^h = [J_{\tilde{\kappa}}^h, J_{\tilde{\mu}}^h] \in \mathbb{R}^{N \times 2N}$ by the chain rule

$$\begin{aligned} J_{\tilde{\kappa}}^h &= \left[\frac{\partial h_i}{\partial \tilde{\kappa}_j} \right]_{i,j=1}^N = \left[\frac{\partial h_i}{\partial \kappa_j} \frac{d\kappa_j}{d\tilde{\kappa}_j} \right]_{i,j=1}^N = \left[\frac{\partial h_i}{\partial \kappa_j} \kappa_j \right]_{i,j=1}^N = J_{\kappa}^h \text{diag}(\kappa), \\ J_{\tilde{\mu}}^h &= \left[\frac{\partial h_i}{\partial \tilde{\mu}_j} \right]_{i,j=1}^N = \left[\frac{\partial h_i}{\partial \mu_j} \frac{d\mu_j}{d\tilde{\mu}_j} \right]_{i,j=1}^N = \left[\frac{\partial h_i}{\partial \mu_j} \mu_j \right]_{i,j=1}^N = J_{\mu}^h \text{diag}(\mu). \end{aligned}$$

In particular, operating with $J_{\tilde{\kappa}, \tilde{\mu}}^h$ boils, in essence, down to multiplying with $J_{\kappa, \mu}^h = [J_{\kappa}^h, J_{\mu}^h]$. According to Lemma 2.1, the Fréchet derivative of H is

$$(H'(\kappa, \mu))(\vartheta, \theta) = \mu\varphi' + \theta\varphi,$$

where $\varphi(\kappa, \mu)$ is the solution of (2.1) and $(\varphi'(\kappa, \mu))(\vartheta, \theta)$ is the solution of the variational problem (2.3). To approximate $J_{\kappa, \mu}^h$ at the optical parameters $\kappa(\tilde{\kappa})$ and $\mu(\tilde{\mu})$ corresponding to given $\tilde{\kappa}$ and $\tilde{\mu}$, we use the FEM with the piecewise linear basis $\{\phi_n\}_{n=1}^N$ appearing in (3.2) to evaluate $H'(\kappa, \mu)$, letting the perturbations ϑ and θ run in turns through $\{\phi_n\}_{n=1}^N$. To be more precise, denoting the node values of the FEM approximations for $\varphi(\kappa, \mu)$, $(\varphi'(\kappa, \mu))(\vartheta, \theta)$, and θ by $\{\varphi_n\}$, $\{\varphi'_n(\vartheta, \theta)\}$, and $\{\theta_n\}$, respectively, we approximate

$$(H'(\kappa, \mu))(\vartheta, \theta) \approx \sum_{n=1}^N (h'_n(\kappa, \mu))(\vartheta, \theta) \phi_n,$$

where

$$(h'_n(\kappa, \mu))(\vartheta, \theta) = \mu_n \varphi'_n(\vartheta, \theta) + \theta_n \varphi_n, \quad n = 1, \dots, N.$$

Accordingly, the (approximate) Jacobian $J_{\kappa, \mu}^h = [J_{\kappa}^h, J_{\mu}^h]$ is formed as

$$\begin{aligned} J_{\kappa}^h &= [(h'_i(\kappa, \mu))(\phi_j, 0)]_{i,j=1}^N = \text{diag}(\mu) J_{\kappa}^{\varphi}, \\ J_{\mu}^h &= [(h'_i(\kappa, \mu))(0, \phi_j)]_{i,j=1}^N = \text{diag}(\mu) J_{\mu}^{\varphi} + \text{diag}(\varphi). \end{aligned}$$

We still need to consider how to handle multiplications by J_{κ}^{φ} and J_{μ}^{φ} matrix-freely.

Based on the weak formulation of (2.1),

$$\int_{\Omega} (\kappa \nabla \varphi \cdot \nabla v + \mu \varphi v) dx + \frac{1}{2} \int_{\partial\Omega} \varphi v dS = 2 \int_{\partial\Omega} \Phi v dS \quad \text{for all } v \in H^1(\Omega),$$

we first form the FEM system matrix $K = K(\kappa, \mu) \in \mathbb{R}^{N \times N}$ and load vector $f = f(\Phi) \in \mathbb{R}^N$ corresponding to the basis $\{\phi_n\}_{n=1}^N$, and solve the node values $\varphi \in \mathbb{R}^N$ from the equation

$$K\varphi = f.$$

Next, we approximate the derivative $(\varphi'(\kappa, \mu))(\vartheta, \theta)$ based on the variational problem (2.3). Since the left-hand side of (2.3) is identical to that in the weak formulation of (2.1), the FEM system matrix $K \in \mathbb{R}^{N \times N}$ stays the same. Letting the perturbations on the right-hand side of (2.3) run in turns through $\{\phi_n\}_{n=1}^N$, we get the load matrix $G(\varphi) \in \mathbb{R}^{N \times 2N}$ and the equation

$$K J_{\kappa, \mu}^{\varphi} = G \quad \text{i.e.,} \quad K [J_{\kappa}^{\varphi}, J_{\mu}^{\varphi}] = [G^{(1)}, G^{(2)}],$$

where

$$G_{i,j}^{(1)} = - \int_{\Omega} \phi_j \nabla \varphi \cdot \nabla \phi_i dx, \quad G_{i,j}^{(2)} = - \int_{\Omega} \phi_j \varphi \phi_i dx,$$

for $i, j = 1, \dots, N$.

Finally, instead of solving the huge system to form the Jacobian explicitly, we use the matrices K and G to implicitly operate on a given vector: multiplying $s = [(s^{(1)})^T, (s^{(2)})^T]^T \in \mathbb{R}^{2N}$ by $J_{\tilde{\kappa}, \tilde{\mu}}^h \in \mathbb{R}^{N \times 2N}$ gives

$$\begin{aligned} J_{\tilde{\kappa}, \tilde{\mu}}^h s &= J_{\tilde{\kappa}}^h s^{(1)} + J_{\tilde{\mu}}^h s^{(2)} \\ &= \text{diag}(\mu) K^{-1} G^{(1)} \text{diag}(\kappa) s^{(1)} \\ &\quad + \left(\text{diag}(\mu) K^{-1} G^{(2)} + \text{diag}(\varphi) \right) \text{diag}(\mu) s^{(2)} \\ &= \text{diag}(\mu) K^{-1} \left(G^{(1)} \text{diag}(\kappa) s^{(1)} + G^{(2)} \text{diag}(\mu) s^{(2)} \right) \\ &\quad + \text{diag}(\varphi) \text{diag}(\mu) s^{(2)}. \end{aligned}$$

We also need to be able to multiply a given vector $t \in \mathbb{R}^N$ with the transpose $(J_{\tilde{\kappa}, \tilde{\mu}}^h)^T \in \mathbb{R}^{2N \times N}$. Since matrices K and $G^{(2)}$ are symmetric, we get

$$\begin{aligned} J_{\tilde{\kappa}, \tilde{\mu}}^{hT} t &= \begin{bmatrix} J_{\tilde{\kappa}}^{hT} t \\ J_{\tilde{\mu}}^{hT} t \end{bmatrix} = \begin{bmatrix} \text{diag}(\kappa) (G^{(1)})^T K^{-1} \text{diag}(\mu) t \\ \text{diag}(\mu) (G^{(2)} K^{-1} \text{diag}(\mu) + \text{diag}(\varphi)) t \end{bmatrix} \\ &= \begin{bmatrix} \text{diag}(\kappa) & 0 \\ 0 & \text{diag}(\mu) \end{bmatrix} \left(\begin{bmatrix} (G^{(1)})^T \\ G^{(2)} \end{bmatrix} K^{-1} \text{diag}(\mu) t + \begin{bmatrix} 0 \\ \text{diag}(\varphi) t \end{bmatrix} \right). \end{aligned}$$

Numerically speaking, multiplying a vector by the Jacobian $J_{\tilde{\kappa}, \tilde{\mu}}^h$ or its transpose is relatively cheap: one need only perform elementwise multiplications of vectors and matrix-vector multiplications with sparse matrices, and to operate with the inverse of the FEM system matrix K on a vector. However, when the number of degrees of freedom and/or illuminations is very high, the sizes of the matrices $G^{(1)}$ and $G^{(2)}$ may become impractically large. In this case, one can continue the analysis to avoid forming the matrices $G^{(1)}$ and $G^{(2)}$ altogether, but the details are omitted here for brevity.

REFERENCES

- [1] G. S. ALBERTI AND H. AMMARI, *Disjoint sparsity for signal separation and applications to hybrid inverse problems in medical imaging*, Appl. Comput. Harmon. Anal., (2015), in press.
- [2] S. ARRIDGE, M. BETCKE, AND L. HARHANEN, *Iterated preconditioned LSQR method for inverse problems on unstructured grids*, Inverse Problems, 30 (2014), 075009.
- [3] S. R. ARRIDGE, *Optical tomography in medical imaging*, Inverse Problems, 15 (1999), pp. R41–R93.
- [4] G. BAL AND K. REN, *Multi-source quantitative photoacoustic tomography in a diffusive regime*, Inverse Problems, 27 (2011), p. 075003.
- [5] G. BAL AND K. REN, *On multi-spectral quantitative photoacoustic tomography in a diffusive regime*, Inverse Problems, 28 (2012), 025010.
- [6] P. BEARD, *Biomedical photoacoustic imaging*, Interface Focus, 1 (2011), pp. 602–631.
- [7] E. BERETTA, M. MUSZKIETA, W. NAETAR, AND O. SCHERZER, *A Variational Method for Quantitative Photoacoustic Tomography with Piecewise Constant Coefficients*, <http://arxiv.org/abs/1509.04982>, 2015.
- [8] D. CALVETTI, *Preconditioned iterative methods for linear discrete ill-posed problems from a Bayesian inversion perspective*, J. Comput. Appl. Math., 198 (2007), pp. 378–395.

- [9] D. CALVETTI, D. MCGIVNEY, AND E. SOMERSALO, *Left and right preconditioning for electrical impedance tomography with structural information*, Inverse Problems, 28 (2012), 055015.
- [10] D. CALVETTI AND E. SOMERSALO, *Priorconditioners for linear systems*, Inverse Problems, 21 (2005), pp. 1397–1418.
- [11] D. CALVETTI AND E. SOMERSALO, *Hypermodels in the Bayesian imaging framework*, Inverse Problems, 24 (2008), 034013.
- [12] B. COX, J. G. LAUFER, S. R. ARRIDGE, AND P. C. BEARD, *Quantitative spectroscopic photoacoustic imaging: A review*, J. Biomed. Opt., 17 (2012), 061202.
- [13] B. T. COX, S. R., ARRIDGE, AND P. C. BEARD, *Estimating chromophore distributions from multiwavelength photoacoustic images*, J. Opt. Soc. Am. A, 26 (2009), pp. 443–455.
- [14] T. DIERKES, O. DORN, F. NATTERER, V. PALAMODOV, AND H. SIELSCHOTT, *Fréchet derivatives for some bilinear inverse problems*, SIAM J. Appl. Math., 62 (2002), pp. 2092–2113, doi:10.1137/S0036139901386375.
- [15] T. DING, K. REN, AND S. VALLÉLIAN, *A one-step reconstruction algorithm for quantitative photoacoustic imaging*, Inverse Problems, 31 (2015), 095005.
- [16] L. ELDÉN, *A weighted pseudoinverse, generalized singular values, and constrained least squares problem*, BIT, 22 (1982), pp. 487–502.
- [17] H. GAO, J. FENG, AND L. SONG, *Limited-view multi-source quantitative photoacoustic tomography*, Inverse Problems, 31 (2015), 065004.
- [18] H. GAO, S. OSHER, AND H. ZHAO, *Quantitative Photoacoustic Tomography*, in Mathematical Modeling in Biomedical Imaging II, H. Ammari, ed., Springer, 2012, pp. 131–158.
- [19] H. GAO, H. ZHAO, AND S. OSHER, *Bregman Methods in Quantitative Photoacoustic Tomography*, UCLA CAM Report, 10-42, UCLA, Los Angeles, CA, 2010.
- [20] P. GRISVARD, *Elliptic Problems in Nonsmooth Domains*, Pitman, Boston, 1985.
- [21] M. HALTMEIER, L. NEUMANN, AND S. RABANSER, *Single-stage reconstruction algorithm for quantitative photoacoustic tomography*, Inverse Problems, 31 (2015), 065005.
- [22] A. HANNUKAINEN, L. HARHANEN, N. HYVÖNEN, AND H. MAJANDER, *Edge-promoting reconstruction of absorption and diffusivity in optical tomography*, Inverse Problems, 32 (2016), 015008.
- [23] P. CH. HANSEN, *Rank-Deficient and Discrete Ill-Posed Problems*, SIAM, Philadelphia, 1998.
- [24] L. HARHANEN, N. HYVÖNEN, H. MAJANDER, AND S. STABOULIS, *Edge-enhancing reconstruction algorithm for three-dimensional electrical impedance tomography*, SIAM J. Sci. Comput., 37 (2015), pp. B60–B78, doi:10.1137/140971750.
- [25] J. W. HILGERS, *On the equivalence of regularization and certain reproducing kernel Hilbert space approaches for solving first kind problems*, SIAM J. Numer. Anal., 13 (1976), pp. 172–184, doi:10.1137/0713018.
- [26] J. P. KAIPIO AND E. SOMERSALO, *Statistical and Computational Inverse Problems*, Springer-Verlag, 2005.
- [27] P. KUCHMENT AND L. KUNYANSKY, *Mathematics of thermoacoustic tomography*, Eur. J. Appl. Math., 19 (2008), pp. 191–224.
- [28] J. LAUFER, B. COX, E. ZHANG, AND P. BEARD, *Quantitative determination of chromophore concentrations from 2D photoacoustic images using a nonlinear model-based inversion scheme*, Appl. Opt., 49 (2010), pp. 1219–1233.
- [29] C. LI AND L. V. WANG, *Photoacoustic tomography and sensing in biomedicine*, Phys. Med. Biol., 54 (2009), pp. R59–R97.
- [30] A. V. MAMONOV AND K. REN, *Quantitative photoacoustic imaging in radiative transport regime*, Comm. Math. Sci., 12 (2014), pp. 201–234.
- [31] W. NAETAR AND O. SCHERZER, *Quantitative photoacoustic tomography with piecewise constant material parameters*, SIAM J. Imaging Sci., 7 (2014), pp. 1755–1774.
- [32] C. CH. PAIGE AND M. A. SAUNDERS, *Algorithm 583: LSQR: Sparse linear equations and least squares problems*, ACM Trans. Math. Softw., 8 (1982), pp. 195–209.
- [33] C. CH. PAIGE AND M. A. SAUNDERS, *LSQR: An algorithm for sparse linear equations and sparse least squares*, ACM Trans. Math. Softw., 8 (1982), pp. 43–71.
- [34] P. PERONA AND J. MALIK, *Scale-space and edge detection using anisotropic diffusion*, IEEE Trans. Pattern Anal. Mach. Intell., 12 (1990), pp. 629–639.

- [35] A. PULKKINEN, B. T. COX, S. R. ARRIDGE, J. P. KAIPIO, AND T. TARVAINEN, *Direct estimation of optical parameters from photoacoustic time series in quantitative photoacoustic tomography*, IEEE Trans. Med. Imag., to appear.
- [36] A. PULKKINEN, B. T. COX, S. R. ARRIDGE, J. P. KAIPIO, AND T. TARVAINEN, *A Bayesian approach to spectral quantitative photoacoustic tomography*, Inverse Problems, 30 (2014), 065012.
- [37] A. PULKKINEN, B. T. COX, S. R. ARRIDGE, J. P. KAIPIO, AND T. TARVAINEN, *Quantitative photoacoustic tomography using illuminations from a single direction*, J. Biomed. Opt., 20 (2015), 036015.
- [38] A. PULKKINEN, V. KOLEHMAINEN, J. P. KAIPIO, B. T. COX, S. R. ARRIDGE, AND T. TARVAINEN, *Approximate marginalization of unknown scattering in quantitative photoacoustic tomography*, Inverse Probl. Imaging, 8 (2014), pp. 811–829.
- [39] D. RAZANSKY, J. BAETEN, AND V. NTZIACHRISTOS, *Sensitivity of molecular target detection by multi-spectral optoacoustic tomography (MSOT)*, Med. Phys., 36 (2009), pp. 939–945.
- [40] K. REN, H. GAO, AND H. ZHAO, *A hybrid reconstruction method for quantitative PAT*, SIAM J. Imaging Sci., 6 (2013), pp. 32–55, doi:10.1137/120866130.
- [41] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, Physica D, 60 (1992), pp. 259–268.
- [42] T. SARATOON, T. TARVAINEN, S. R. ARRIDGE, AND B. T. COX, *3D quantitative photoacoustic tomography using the δ -Eddington approximation*, in Proceedings of SPIE - The International Society for Optical Engineering 8581, A. A. Oraevsky and L. V. Wang, eds., Bellingham, WA, 2013, 85810V.
- [43] T. SARATOON, T. TARVAINEN, B. T. COX, AND S. R. ARRIDGE, *A gradient-based method for quantitative photoacoustic tomography using the radiative transfer equation*, Inverse Problems, 29 (2013), 075006.
- [44] P. SHAO, B. COX, AND R. J. ZEMP, *Estimating optical absorption, scattering and Grueneisen distributions with multiple-illumination photoacoustic tomography*, Appl. Opt., 50 (2011), pp. 3145–3154.
- [45] P. SHAO, T. HARRISON, AND R. J. ZEMP, *Iterative algorithm for multiple illumination photoacoustic tomography (MIPAT) using ultrasound channel data*, Biomed. Opt. Express, 3 (2012), pp. 3240–3248.
- [46] N. SONG, C. DEUMIÉ, AND A. D. SILVA, *Considering sources and detectors distributions for quantitative photoacoustic tomography*, Biomed. Opt. Express, 5 (2014), pp. 3960–3974.
- [47] T. TARVAINEN, B. T. COX, J. P. KAIPIO, AND S. R. ARRIDGE, *Reconstructing absorption and scattering distributions in quantitative photoacoustic tomography*, Inverse Problems, 28 (2012), 084009.
- [48] T. TARVAINEN, A. PULKKINEN, B. T. COX, J. P. KAIPIO, AND S. R. ARRIDGE, *Bayesian image reconstruction in quantitative photoacoustic tomography*, IEEE Trans. Med. Imag., 32 (2013), pp. 2287–2298.
- [49] C. R. VOGEL AND M. E. OMAN, *Iterative methods for total variation denoising*, SIAM J. Sci. Comput., 17 (1996), pp. 227–238, doi:10.1137/0917016.
- [50] L. V. WANG, ed., *Photoacoustic Imaging and Spectroscopy*, CRC Press, Boca Raton, FL, 2009.
- [51] M. XU AND L. V. WANG, *Photoacoustic imaging in biomedicine*, Rev. Sci. Instrum., 77 (2006), 041101.
- [52] R. J. ZEMP, *Quantitative photoacoustic tomography with multiple optical sources*, Appl. Opt., 49 (2010), pp. 3566–3572.