

Research



Cite this article: Schoeder S, Olefir I, Kronbichler M, Ntziachristos V, Wall WA. 2018 Optoacoustic image reconstruction: the full inverse problem with variable bases. *Proc. R. Soc. A* **474**: 20180369. <http://dx.doi.org/10.1098/rspa.2018.0369>

Received: 6 June 2018

Accepted: 8 October 2018

Subject Areas:

computer modelling and simulation, image processing, medical computing

Keywords:

optoacoustic imaging, image reconstruction, acoustical heterogeneities, ill-conditioned inverse problems, mouse brain imaging

Author for correspondence:

Svenja Schoeder

e-mail: schoeder@nm.mw.tum.de

Electronic supplementary material is available online at <http://dx.doi.org/10.6084/m9.figshare.c.4272632>.

Optoacoustic image reconstruction: the full inverse problem with variable bases

S. Schoeder¹, I. Olefir^{2,3}, M. Kronbichler¹,
V. Ntziachristos^{2,3}, and W. A. Wall¹

¹Institute for Computational Mechanics, and ²School of Bioengineering, Technical University of Munich, Garching, Germany

³Helmholtz Zentrum München, Institute for Biological and Medical Imaging, Neuherberg, Germany

SS, 0000-0002-0157-6760; WAW, 0000-0001-7419-3384

Optoacoustic imaging was for a long time concerned with the reconstruction of energy density or optical properties. In this work, we present the full inverse problem with respect to optical absorption and diffusion as well as speed of sound and mass density. The inverse problem is solved by an iterative gradient-based optimization procedure. Since the ill-conditioning increases with the number of sought parameters, we propose two approaches to improve the conditioning. The first approach is based on the reduction of the size of the basis for the parameter spaces, that evolves according to the particular characteristics of the solution, while maintaining the flexibility of element-wise parameter selection. The second approach is a material identification technique that incorporates prior knowledge of expected material types and uses the acoustical gradients to identify materials uniquely. We present numerical studies to illustrate the properties and functional principle of the proposed methods. Significant convergence speed-ups are gained by the two approaches countering ill-conditioning. Additionally, we show results for the reconstruction of a mouse brain from *in vivo* measurements.

1. Introduction

Optoacoustic imaging (also called photoacoustic imaging) as a comparably young medical imaging technique allows for new insights in several research fields, for example, cancer detection [1], arthritis [2] or atherosclerosis imaging [3]. This imaging technique

gives rise to a complex inverse problem to reconstruct images from measurement data. In optoacoustics, an object of interest is illuminated with a short pulse laser light that distributes within the object according to its optical material properties. The light is absorbed, the energy transforms to heat causing a thermal expansion and hence a respective pressure variation. The pressure propagates through the object and is eventually measured. The purpose of the optoacoustic image reconstruction is to interpret the measurement data such that one can conclude on physical properties of the body. Basically, four different types of image reconstruction are used: backprojection [4], model-based inversion [5,6], time-reversal [7–9] or iterative gradient-based schemes [10,11]. The first three of these four reconstruction types commonly reconstruct the pressure rise or the optical energy while gradient-based and analytical methods also allow to reconstruct absorption as well as diffusion coefficients simultaneously in a quantitative manner [12–17]. As a matter of fact, the measurements additionally include information on the speed of sound and mass density distribution. Sound speed heterogeneities were studied and sometimes even reconstructed in [10,18–22] but those studies are limited by several assumptions on the tomographic set-up or the object. In [23], sound speed variations are modelled as random fluctuations to a constant in the context of coherent interferometry algorithms.

The progress in numerical modelling and implementation as well as the increase in computing power render it possible to model the physical processes unfolding in an optoacoustic scanner during acquisition more accurately. Our approach includes an optoacoustic image reconstruction algorithm that does not only recover the optical properties of an object, namely absorption and diffusion coefficients, but also the mass density and speed of sound without restrictions on the underlying geometry, the tomographic set-up or assumptions on parameter distributions. The algorithm includes an iterative gradient-based optimization scheme and is based on a physical description of the underlying problems taking the primary physical effects into account by means of the optical diffusion approximation and the acoustical wave equation. The physical approach raises the computational effort but this is countered with high-performance optimized finite-element implementations [24] for the acoustical solver. The joint reconstruction of source and speed of sound from boundary measurements, that arises in optoacoustic image reconstruction, has been shown to be unstable for the linearized problem in [21]. The recovery of either the source or the speed of sound, however, is addressed extensively in the literature (e.g. [25,26] and references therein). Herein, the different parameter fields are reconstructed successively avoiding problems due to different scaling and avoiding instabilities due to joint reconstructions, see [25] on stability criteria for recovery of either a source or a speed.

The increase in the number of optimization variables worsens the conditioning of the inverse problem and we propose two approaches to oppose the ill-conditioning. In the first approach, we aim to optimize the basis functions used for the parameter discretizations. Typically, voxel- or pixel-based basis functions are chosen. However, there are several other possibilities, e.g. local radially symmetric basis functions as proposed in [27,28]. The authors of [28] claim an improvement of convergence compared to local polynomial basis functions due to implicit regularization of a low-dimensional solution space if the shape parameters of the bases are chosen appropriately. In [29,30], computationally expensive level-set functions are used to reconstruct inclusion shapes and values. Their application is limited to problems with only a few inclusions of regular shape. In [31,32], group sparse reconstruction algorithms of iterative shrinkage type are used to optimize for the sparsity of a coefficient vector with a morphological component algorithms and Haar wavelets representing piece-wise smooth parameter distributions. Another approach to regularize the problem is presented in [33] where the restriction to small absorbers regularizes the solution. We present a basis approach that starts from a classical voxel-based basis and consolidates voxels to patches, thereby introducing implicit regularization and making use of the biologically reasonable partitioning into several clustered materials. Additionally, material distribution patterns are communicated from the most sensitive parameter to the less sensitive parameters, thereby improving the conditioning of the inverse problem. This approach is more flexible and robust compared to level-set reconstructions and does not require the identification of appropriate shape parameters for the bases. In this way, we are able to reconstruct

distinct inclusions and introduce an implicit regularization while maintaining full flexibility. The approach can be thought of as a segmentation during reconstruction without strong parameter dependence, e.g. as in [34].

The second method we propose to oppose the ill-conditioning is an extension of the previous work [35] on material identification. Material identification is reasonable when the composition of an object and typical values for the material properties are known in advance. This is often the case in medical applications where there is only a limited number of tissue types, e.g. soft tissue, skin, bone, brain, muscle, organs, etc. and typical values are given in the literature [36,37]. In [35], absorption and diffusion coefficients are reconstructed and are used to identify materials from a user-specified material catalogue to update the acoustical properties accordingly. A drawback, however, is that the diffusion coefficient has low sensitivity and its reconstruction is error prone. In this work, we propose to reconstruct only the absorption coefficient and then use the acoustical gradients to find a unique assignment to materials from the catalogue. One benefit is reduced computational time because only one parameter is reconstructed and the acoustical parameters are updated on the fly. The other benefit is that available prior knowledge is brought into the reconstruction scheme without restricting the generality of the approach.

This work is structured as follows. In §2, the physical model of the optoacoustic imaging procedure is presented. The imaging problem with its objective function and objective function gradients is introduced in §3. Preliminary insights for this complex optimization procedure are given in §4. The variable bases approach and the material identification are proposed in §5a,b, respectively. In both of these sections, numerical examples are given to demonstrate the working principle of the approaches. The algorithm is applied to reconstruct images of a mouse brain in §6. Finally, conclusions and an outlook are given in §7.

2. The physical problem

The physical description of the optoacoustic imaging procedure is composed of three main parts: laser light propagates in a strongly scattering medium; optical energy is transformed into a pressure field via the optoacoustic effect; the pressure propagates as a set of sound waves through the potentially heterogeneous medium and is eventually detected. In the following, $\hat{\Omega}$ denotes the domain of the object of interest and Ω the object's domain combined with the domain of the coupling medium, which is used to fill the gap between the object and the detectors.

Light propagation is described by the diffusion approximation

$$\begin{aligned}\mu_a \phi - \nabla \cdot (D \nabla \phi) &= 0 \quad \text{in } \hat{\Omega}, \\ \phi &= \hat{\phi} \quad \text{on } \partial \hat{\Omega}_D, \\ \phi + 2D \nabla \phi \cdot \mathbf{n} &= 0 \quad \text{on } \partial \hat{\Omega}_R,\end{aligned}$$

with the absorption coefficient μ_a , the diffusion coefficient D and the light flux ϕ [38]. The object's boundary $\partial \hat{\Omega}$ consists of a part $\partial \hat{\Omega}_D$ where the light flux $\hat{\phi}$ is prescribed by a Dirichlet condition and of a part $\partial \hat{\Omega}_R$ where a Robin type boundary condition is prescribed with the outward pointing normal vector \mathbf{n} . The diffusion approximation yields sufficiently accurate results in regions far from boundaries and sources where the scattering is stronger than the absorption. In the context of optoacoustic imaging has been used, e.g. in [15,16,39,40].

The physical description of the photoacoustic (optoacoustic) effect in thermal and stress confinement [38] is given by

$$\begin{aligned}p_0 &= -\Gamma \mu_a \phi \quad \text{in } \hat{\Omega}, \\ p_0 &= 0 \quad \text{in } \Omega \setminus \hat{\Omega} \\ \text{and} \quad v_0 &= \mathbf{0} \quad \text{in } \Omega,\end{aligned}$$

with the Grüneisen coefficient Γ , the pressure p and the velocity v . The sign of the initial pressure is negative by definition.

The pressure propagation is described by the acoustic wave equation written as a first-order system

$$\begin{aligned}\rho \frac{\partial \mathbf{v}}{\partial t} + \nabla p &= 0 \quad \text{in } \Omega \times (0, T], \\ \frac{1}{c^2 \rho} \frac{\partial p}{\partial t} + \nabla \cdot \mathbf{v} &= 0 \quad \text{in } \Omega \times (0, T], \\ \mathbf{v} \cdot \mathbf{n} - \frac{1}{c \rho} p &= 0 \quad \text{on } \partial \Omega_A \times (0, T], \\ \mathbf{v} \cdot \mathbf{n} &= 0 \quad \text{on } \partial \Omega_N \times (0, T],\end{aligned}$$

with the speed of sound c , the mass density ρ , the final time T and the two parts of the domain's boundary $\partial \Omega_A$ and $\partial \Omega_N$, where a first-order absorbing boundary condition and a Neumann condition are enforced, respectively. In many cases, the second-order formulation of the acoustic wave equation neglecting spatial density variations is used, see e.g. [19,41–43]. In [10], a formulation additionally considering acoustical absorption is used.

For the numerical scheme, the optical problem is discretized with the standard finite-element method, while the acoustical problem is discretized with the hybridizable discontinuous Galerkin method (HDG) with polynomial degree of the shape functions between one and nine [24]. For a detailed explanation of the physical problems and their discretizations, see [35].

3. The imaging problem

With the equations introduced in the preceding section, pressure values at the detector boundary $\partial \Omega_M$ are simulated and stored in n^t vectors \mathbf{P}_i^s , with n^t being the maximum number of time steps. Each vector \mathbf{P}_i^s contains the pressure values at time t_i for all detectors. Analogously, physical measurements provide measured pressure values \mathbf{P}_i^m . The objective is to minimize the difference between these two quantities by adaption of the absorption coefficient μ_a , the diffusion coefficient D , the speed of sound c and the mass density ρ :

$$\min_{\mu_a, D, c, \rho} J \quad \text{with } J = \frac{1}{2} \sum_{i=1}^{n^t} \|\mathbf{P}_i^s - \mathbf{P}_i^m\|^2, \quad (3.1)$$

with $\|\bullet\|$ indicating the Euclidean vector norm. The vectors μ_a, D, c and ρ hold discrete values, such that the parameter distributions are, for the example of the absorption coefficient, approximated by

$$\mu(x)_a = \mathbf{b}(x) \cdot \mu_a, \quad (3.2)$$

with the spatial coordinate x and a vector of basis functions \mathbf{b} . A convenient choice of basis functions is crucial for the conditioning of the inverse problem, its solvability and its convergence behaviour. Some reconstruction algorithms might require a certain level of smoothness of the parameter representation. Generally speaking, the conditioning of an inverse problem degrades with an increasing number of parameters motivating a reduction of degrees of freedom and hence a reduction of the number of basis functions. These considerations led, e.g. to the approach of radially symmetric basis functions or level set approaches, see [28,29], respectively. A common choice is element-wise basis functions taking the value 1 in one element and 0 in all others. However, this is not the best choice as will be shown in §5a.

(a) The gradients

The gradients of the objective function could be calculated with a finite difference scheme, which is yet prohibitively expensive and cumbersome considering the large number of optimization parameters n^{param} . Thus, the gradients are calculated via the adjoint approach. A detailed

description of the derivation of the optical gradients $dJ/d\mu$ and dJ/dD can be found in [35]. Here, only the resulting equations shall be stated

$$\frac{dJ}{d\mu_a} = \tilde{\Phi}^T \frac{\partial \mathbb{K}_o}{\partial \mu_a} \Phi - \tilde{P}_0^T \frac{\partial \mathbb{K}_m}{\partial \mu_a} \Phi \quad (3.3)$$

and

$$\frac{dJ}{dD} = \tilde{\Phi}^T \frac{\partial \mathbb{K}_o}{\partial D} \Phi. \quad (3.4)$$

Herein, Φ and $\tilde{\Phi}$ are the solution vectors of the forward and adjoint light problem, \tilde{P} is the pressure solution vector of the adjoint acoustical problem, and \mathbb{K}_o and \mathbb{K}_m are the ‘stiffness matrices’ of the optical problem and the optoacoustic mapping, see [35] for details.

The acoustical gradients are derived from the Lagrangian defined in [35] by partial differentiation:

$$\frac{dJ}{dc} = \sum_{i=1}^{n^t} \tilde{P}_i^T \frac{\partial \mathbb{M}_a}{\partial c} (P_i - P_{i-1}) \quad (3.5)$$

and

$$\frac{dJ}{d\rho} = \sum_{i=1}^{n^t} \tilde{V}_i^T \frac{\partial \mathbb{A}_a}{\partial \rho} (V_i - V_{i-1}) + \sum_{i=1}^{n^t} \tilde{P}_i^T \frac{\partial \mathbb{M}_a}{\partial \rho} (P_i - P_{i-1}). \quad (3.6)$$

The acoustical ‘mass matrices’ for pressure and velocity are denoted \mathbb{M}_a and \mathbb{A}_a , respectively, and result from the discretization of the acoustic wave equation with HDG, for details see [24]. The acoustical gradients are calculated from the combination of forward (P_i, V_i) and adjoint (\tilde{P}_i, \tilde{V}_i) solution vectors at each time step. The difficulty is that the adjoint problem has to be solved backwards in time after the solution of the forward problem. A naive implementation would have to store all solution vectors P_i, V_i from the forward problem in order to be able to evaluate the above stated equations, which is impossible with the available computational resources for realistic image reconstructions. For this reason, a checkpointing approach as proposed in [44] is implemented: during the forward run each n^{check} time steps, a restart is written; in the subsequent adjoint run (which is executed backwards in time), the accordant forward restart is read, n^{check} forward time steps are performed, the n^{check} solution vectors are stored, the adjoint and forward solutions are combined to get the gradient contribution, and the adjoint run is continued. Doing so, the gradient calculation requires three times the computational effort of a normal solve and this is $(n^{\text{param}} + 1)/3$ times less expensive than finite differences.

(b) The algorithm

The minimization problem given by (3.1) is solved iteratively using a line search along the steepest descent direction or a low storage BFGS descent direction. The directions are calculated from the gradients of the objective function. The parameters are not optimized simultaneously but one after another because they differ in sensitivity as well as in scaling. Therefore, each parameter needs an individually determined step length to evolve to its optimal distribution. The algorithm starts with a first evaluation of the objective function by solving the entire forward problem. Subsequently, the adjoint problem is solved and the gradients are evaluated. Then, the optimization loop is entered which, in turn, consists of a loop over the four sought parameter fields invoking the line search procedure. The line search is constructed to fulfil the strong Wolfe conditions according to § 3.5 of [45]. Within the line search, the forward and adjoint problem are evaluated several times and the entire parameter field, i.e. its spatial distribution described by a finite but potentially large number of scalar values is optimized. The algorithm terminates as soon as the maximal number of iterations is reached or the objective function value is below the user-defined tolerance. The results depend on the initial guess of the parameters. In common reconstruction algorithms, acoustical homogeneity is assumed. We propose to choose the properties of water (or general soft tissue) for the acoustical parameters. Acoustical parameter updates will improve the objective function over acoustical homogeneous material. For the

absorption coefficient, a small value is proposed as initial guess. This renders the first update very similar to a time reversal and subsequent reconstruction of μ_a from the initial pressure p_0 . For successful reconstruction of the diffusion coefficient, it must be given on the boundary of the domain, see e.g. [39,46], otherwise the solution is not necessarily unique.

4. Preliminary insights

In this section, some preliminary insights concerning basic properties and solution behaviour of the complex optoacoustic inverse problem are presented.

A simple numerical model as shown in figure 1a is used to generate measurement pressure curves P_i^m . It consists of a circle with diameter 16 mm representing the tomograph's interior at whose boundary the sound signals are measured. A circular object of diameter 10 mm with a rectangular inclusion of size 7 mm \times 1.5 mm is located at the tomograph's centre. Here and in all following examples, two-dimensional geometries are used. To each region, a different material $\mathcal{M} = (\mu_a, D, c, \rho)$ is assigned: \mathcal{M}_i in the rectangular region, \mathcal{M}_o in the inner circle and \mathcal{M}_w in the outer ring. Material \mathcal{M}_w is fixed for the rest of the paper to $\mathcal{M}_w = (0.01 \text{ mm}^{-1}, 0.0 \text{ mm}, 1.48 \text{ mm } \mu\text{s}^{-1}, 1.0 \text{ mg mm}^{-3})$ representing water as the coupling medium between the object of interest and the acoustical detectors.

The two other material sets are defined as

$$\mathcal{M}_i = (0.11 \text{ mm}^{-1}, 0.5 \text{ mm}, 2.0 \text{ mm } \mu\text{s}^{-1}, 2.0 \text{ mg mm}^{-3}) \quad (4.1)$$

and

$$\mathcal{M}_o = (0.011 \text{ mm}^{-1}, 0.5 \text{ mm}, 1.5 \text{ mm } \mu\text{s}^{-1}, 1.1 \text{ mg mm}^{-3}). \quad (4.2)$$

The values are in the typical range of biological tissues.

Sound signals P_i^m are generated with different combinations of the defined materials on a mesh consisting of 5063 elements and a time step size of 0.01 μs , which gives a maximal Courant number of about 0.1. For reconstruction, an initial geometry as shown in figure 1b is created. The outer ring again represents the coupling medium and material parameters \mathcal{M}_w are set. The initial material values for the inner circle are set to the values of material \mathcal{M}_o . The geometry is meshed with 4024 elements and the time step size is set to 0.0125 μs to get the same Courant number. Additionally, a set-up with a finer discretization is created: the signal generation is performed with a mesh consisting of 20 567 elements and a time step size $\Delta t = 0.002 \mu\text{s}$; the reconstruction is carried out on a mesh consisting of 14 363 elements and a time step size $\Delta t = 0.0025 \mu\text{s}$. By use of coarser spatial and temporal discretizations for the reconstruction, the inverse crime is avoided [47].

In a first step, we evaluated the gradients according to the formulae derived in §3a. Figure 2 shows from left to right the absorption, diffusion, speed of sound and mass density gradient. In the rows, three different set-ups are printed. In the first row, the gradients are calculated for measurement signals generated with the material for the inclusion as combination of \mathcal{M}_o and \mathcal{M}_i . The absorption coefficient gradient is shown for measurement signals generated with only the absorption coefficient from \mathcal{M}_i , the diffusion coefficient is shown for measurement signals generated with only the diffusion coefficient from \mathcal{M}_i and so on. The second row shows the four gradients but for measurement signals generated with only the absorption coefficient from \mathcal{M}_i . The last row shows the gradients for measurement signals where all parameters for the inclusion are set according to \mathcal{M}_i . Obviously, the gradient of the absorption coefficient reproduces the shape of the inclusion best. The sensitivity of the other material parameters is (as expected) lower: the gradients do not reproduce the shape of the inclusion, which can be seen from the last three gradients in the first row. From the second row, one can see how the gradients of the other parameters react to changes in the absorption coefficient. The diffusion gradient reproduces the shape of the inclusion and the acoustical parameters form a region around the inclusion with positive values. The results in the last row of figure 2 give an impression of the dependence of the absorption coefficient on changes in the other parameters, positive values appear next to the

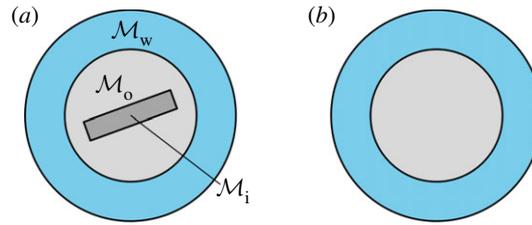


Figure 1. Geometries. (a) Signal generation and (b) reconstruction (Online version in colour.)

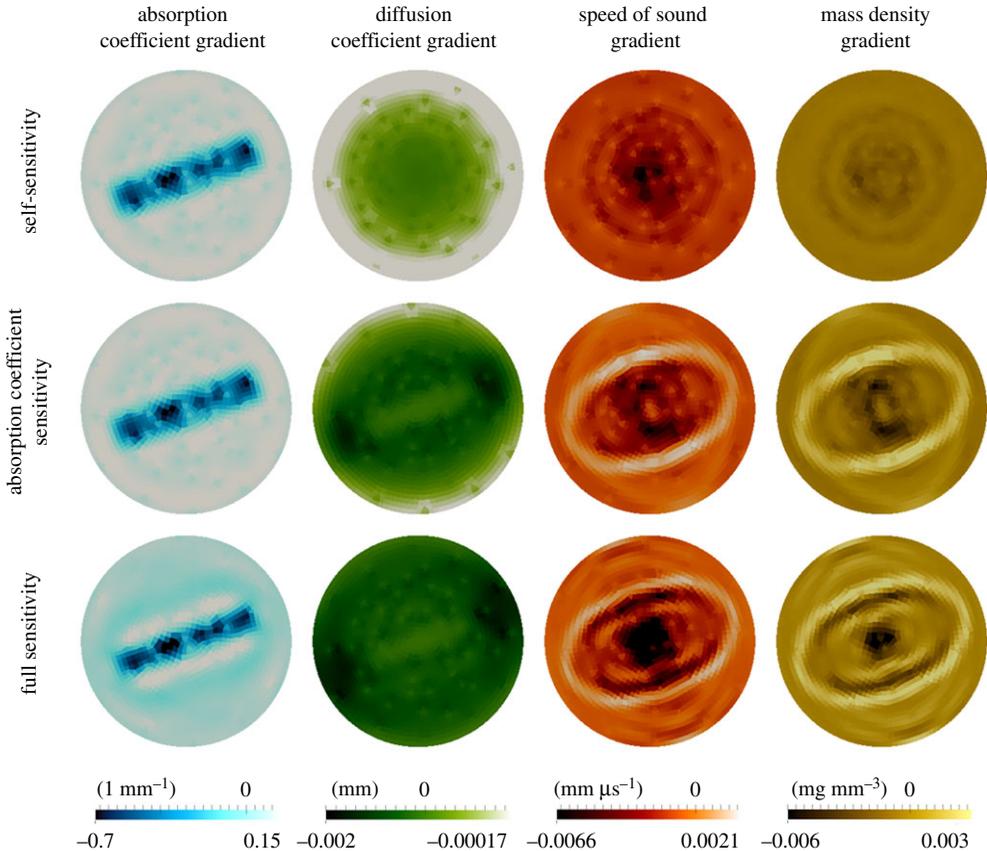


Figure 2. The absorption coefficient, diffusion coefficient, speed of sound and mass density gradients (from left to right) for three different set-ups (from top to bottom): the first row shows the self-sensitivity, i.e. the absorption coefficient only is varied to show the absorption coefficient gradient, the diffusion coefficient only is varied to show the diffusion coefficient gradient and so on. The second row shows the four gradients for variations in the absorption coefficient, i.e. how the gradients react to variations in the absorption coefficient, while the other parameters do not vary. The last row shows the gradients for variations in all parameter fields. (Online version in colour.)

inclusion. The other gradients reproduce similar patterns as in the second row but with slight modifications in the interior.

In a second step, image reconstruction is performed with element-wise discretization of the parameter fields for measurement signals obtained with material parameters \mathcal{M}_i in the rectangular inclusion. The results are given in figure 3 in terms of the relative objective function J/J_0 and in terms of the relative error in the parameter fields for the coarse and fine discretizations

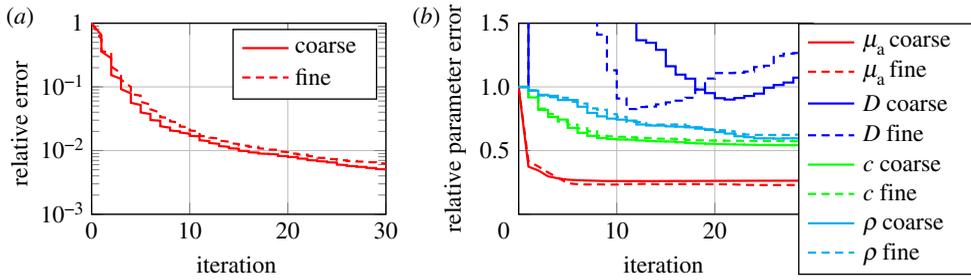


Figure 3. Convergence behaviour of the full optoacoustic problem. (a) Relative objective function and (b) relative error of the parameters (Online version in colour.)

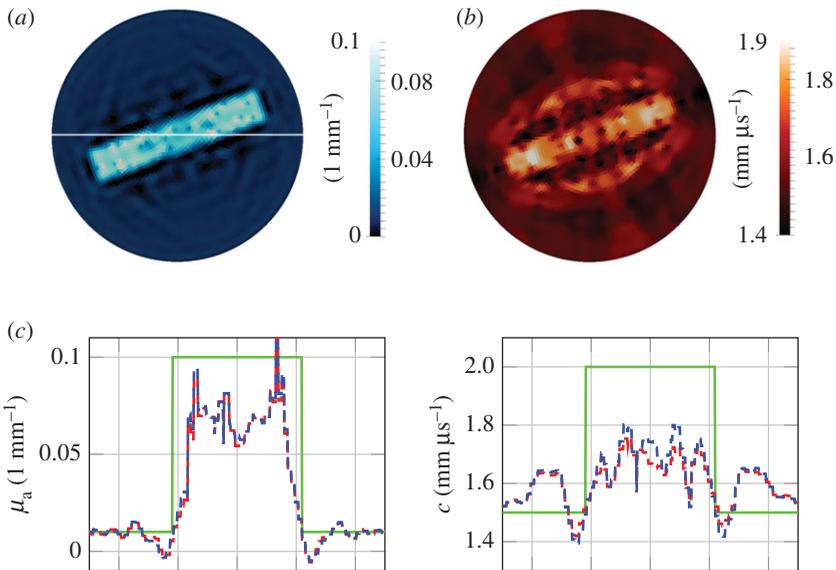


Figure 4. Reconstructions of absorption coefficient and speed of sound. (a) Reconstruction of the absorption coefficient with 8% noise added to the signals, (b) reconstruction of the speed of sound with 8% noise added to the signals and (c) plot over line indicated in (a) of μ_a and c , the sample solution is represented by a green solid line, the reconstruction without noise by a red dashed line and the reconstruction with 8% noise by a blue dashed line. (Online version in colour.)

as mentioned earlier. The error in the parameter fields is calculated as the square of the difference between the actual and the expected parameter values. The elements intersected by the inclusion interface are skipped for the error evaluation. From figure 3a, it can be seen that the objective function value decreases monotonically. Every fourth update, the decrease in the objective function is comparably high, which corresponds to updates of the most sensitive parameter—the absorption coefficient. Figure 3b shows that the relative error in the absorption coefficient distribution converges best, followed by the speed of sound and the mass density. After 30 iterations, the reconstruction is not yet fully converged. The diffusion coefficient starts with a disadvantageous update and slowly recovers but the algorithm fails to improve the diffusion coefficient distribution significantly. For the given set-up and comparable applications, the diffusion coefficient's sensitivity is much lower than for the other parameters as can be seen from figures 2 and 3, and results presented in [35]. Therefore, the reconstruction of the diffusion coefficient is not considered in the remainder of this work. The convergence behaviour for the coarse and fine discretizations is similar. Figure 4 shows the reconstructions of the absorption

coefficient and the speed of sound for a reconstruction with noise of 8% amplitude compared to the noiseless data added to the measurements. Figure 4c plots the reconstruction over the indicated line and compares results obtained without and with superimposed noise. It appears that the speed of sound is more sensitive to noise because its reconstructions vary more strongly compared to the absorption coefficient reconstructions. As already indicated in figure 3b, the reconstruction of the speed of sound is not yet completely converged.

5. Overcoming the ill-conditioning

Reconstructing not only the absorption coefficient but two additional material parameters with lower sensitivity deteriorates the conditioning of the optoacoustic inverse problem further. The deterioration is countered by two approaches presented in the following. One approach addresses the parameter discretization and is presented in §5a. The other approach relies on the idea of material identification in objects of known composition. It is an extension of [35] and is presented in §5b.

(a) Variable bases for parameter spaces

Up to this point, basis functions $b(x)$ for the parameters with value 1 in one element and 0 in all others were used. This is an intuitive approach when using finite elements since parameters are commonly stored element-wise. However, this is not necessarily the best approach. Usually, the solver for the physical problem and the solver for the inverse problem have different demands on the spatial discretization. If the gradient of the objective function is calculated using finite differences, one favours to keep the number of model parameters to a minimum, i.e. use large elements, since every additional parameter requires an additional evaluation of the forward problem thereby increasing computational expenses. On the other hand, element-wise parameter discretizations in combination with a physical solver are restricted to certain accuracy criteria implying the usage of small discretization elements and hence a high number of model parameters. Also, inverse problems commonly suffer from ill-conditioning. This can be counteracted by basis functions that are tailored to the respective inverse problem, e.g. by incorporating prior knowledge concerning the parameter distribution properties.

In this section, a new type of parameter basis consisting of Patched Basis Functions (PBF) is introduced. Biological bodies as well as engineering components consist of different tissue/material types with distinct differences in their properties. Within one tissue type, slight variations of the material properties are present, while discontinuities may appear from one type to the other. Also, one tissue type is spatially connected, i.e. clustered and generally not scattered. The proposed basis imitates these features by clustering several neighbouring elements to a patch. Patches are built from an element-wise quantity by collecting connected elements of similar value. The patch summarizes all element-wise basis functions to one patched basis function. After setting up all patches, one does no longer have n^{ele} but n^{Patch} basis functions that are 1 inside the respective patch and 0 outside the patch. Thereby, implicit regularization is introduced, while full flexibility is maintained, because the patches can be rebuilt in every iteration. Additionally, the approach allows to transport information of a body's composition from the most sensitive parameter to less sensitive parameters by creating PBFs on the most sensitive parameter and reusing the basis for the others. Thus the ill-conditioning is strongly improved.

In the following, the symbol ϖ is used as a placeholder for parameter element values, e.g. exemplary for the absorption coefficient $\varpi = \mu_a$, such that the parameter field is described by the product of basis functions and parameter values $\varpi(x) = b_e(x)\varpi$ as in equation (3.2). An element-wise parameter gradient is denoted by g_e with $g_e = dJ/d\varpi$, e.g. exemplary for the absorption coefficient $g_e = dJ/d\mu_a$.

The starting point to build patches is an element-wise quantity q_e , either an element-wise gradient $q_e = g_e$ or element-wise parameter values $q_e = \varpi$. In the first iteration, the initial parameters are commonly assumed spatially constant, which is why they cannot be used to set

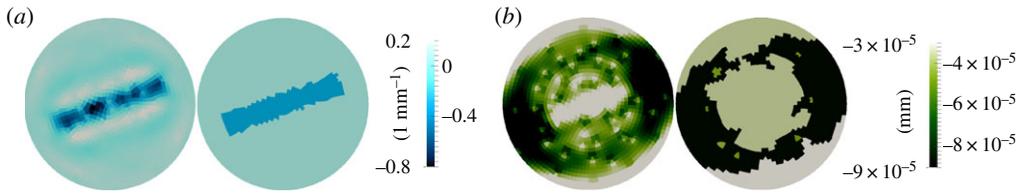


Figure 5. Two exemplary element-wise gradients and their patched counterparts. (a) Gradient 1 (absorption coefficient) and (b) gradient 2 (diffusion coefficient). (Online version in colour.)

up patches. The gradient, however, already contains information of the measurements in the first iteration. For all following iterations, it depends on the parameter properties and the inverse problem itself, which quantity is more beneficial (in §5a(i), several possibilities to choose q_e and results are presented). The procedure is as follows:

- The range of the quantity is calculated as $r = |\max q_e - \min q_e|$ and a subrange is defined according to $r_{\text{sub}} = \alpha r$. The parameter α must be in the interval $(0, 1)$ and is reasonably chosen between $[0.1, 0.9]$ depending on the problem set-up.
- The element with the maximum value of the quantity q_e is identified and the element ID e_{check} as well as the gradient value g_{check} are stored. To this element, the patch ID $P_{\text{ID}} = 1$ is assigned.
- It is checked if the neighbours e_n of element e_{check} should belong to the same patch. The criterion is that their value q_{e_n} of the quantity q_e is in the range of $[q_{e_{\text{check}}} - r_{\text{sub}}, q_{e_{\text{check}}}]$. The patch ID $P_{\text{ID}} = 1$ is set and their corresponding neighbours are checked next, and so on.
- After that, the largest unassigned entry in the vector q_e defines the element with patch ID $P_{\text{ID}} = 2$ and the recursive check of neighbours and neighbours of neighbours is repeated.
- The overall routine is executed until every element is assigned to a patch.
- The next step is to calculate the average of the gradient values in each patch and to replace the original value g_e by the patch averaged value to eventually get the gradient g_p .
- The line search procedure can be entered as in the original algorithm. The calculation of the transformation $g_e \rightarrow g_p$ is performed in every evaluation of the gradient, i.e. in each iteration a different number of patches and different shapes of patches can be present.

To clarify the purpose of the PBF, figure 5 compares element-wise gradients to patched gradients for two different examples: in (a) the gradient is transformed to two patches; in (b) the gradient is transformed to three large patches and seven smaller ones. The algorithm does not require the number of expected patches as input. The single parameter α controls the allowed differences within one patch and consequently the size of the patches. In the limit $\alpha \rightarrow 0$, the element-wise parameter discretization is recovered. For $\alpha \rightarrow 1$, all elements are clustered into one patch.

The described approach is favourable for parameter gradients that directly reproduce inclusion shapes like the absorption gradient in figure 2. For the other parameters, we additionally propose an advanced approach, which is not only suited for optoacoustics but also for other image reconstruction schemes or inverse problems involving more than one parameter with different sensitivities. The approach makes use of the fact that the absorption coefficient has the highest sensitivity in optoacoustic imaging. In a set-up phase, the absorption coefficient is optimized in $n_{\text{setup,iter}}$ iterations with patches built according to its own gradient. After the set-up phase, the other parameters are optimized as well but they do not build patches according to their own gradients but according to the absorption coefficient distribution. Doing so, the PBF method allows to transport information via the basis from the most sensitive parameter field to the others. Several other types of PBF in the different parameters fields are possible and the most relevant ones are examined in the following section.

(i) Numerical study of patched basis functions

On the example from §4, we assess the convergence and solution quality for different parameter basis functions. In order to put the results of this examination into perspective, we also compare with the standard approach of element-wise parameters and also with two approaches using global basis functions:

- element-wise: element-wise basis functions
- PBF self: all parameters build patches according to their own gradient ($q_e^{\mu_a, c, \rho} = dJ/d\boldsymbol{w}$)
- PBF abs grad: the absorption coefficient builds patches according to its own gradient, the other parameters use these patches ($q_e^{\mu_a, c, \rho} = dJ/d\boldsymbol{\mu}_a$)
- PBF abs vals: the absorption coefficient builds patches according to its own gradient ($q_e^{\mu_a} = dJ/d\boldsymbol{\mu}_a$), the other parameters build patches according to the absorption coefficient ($q_e^{c, \rho} = \boldsymbol{\mu}_a$)
- PBF mixed: the absorption coefficient does not build patches, speed of sound and mass density build patches according to the absorption coefficient ($q_e^{c, \rho} = \boldsymbol{\mu}_a$)
- harmonic: global harmonic functions
- Lagrange: global Lagrange polynomials

For the Lagrange basis, the i th Lagrange polynomial of degree k is denoted l_i^k and is mapped to the global size L of the underlying geometry. For the given two-dimensional set-up the Lagrange basis is of size $(k + 1)^2$ with equidistant points

$$\mathbf{b}_l(\mathbf{x}) = [l_1^k(x_1)l_1^k(x_2) \quad \dots \quad l_{k+1}^k(x_1)l_{k+1}^k(x_2)].$$

For the harmonic basis, the following basis functions are used

$$\begin{aligned} & \sin\left(\frac{n\pi x_1}{L}\right) \cdot \sin\left(\frac{m\pi x_2}{L}\right) \\ & \cos\left(\frac{n\pi x_1}{L}\right) \cdot \sin\left(\frac{m\pi x_2}{L}\right) \\ & \cos\left(\frac{n\pi x_1}{L}\right) \cdot \cos\left(\frac{m\pi x_2}{L}\right) \end{aligned}$$

with $n, m = 0, 1, 2, \dots, k - 1$, such that the basis is of size $3k^2$ for a given k .

Measurement signals are generated with the numerical model described in §4. Reconstruction is performed with respect to three parameters, absorption coefficient, speed of sound and mass density, using the described procedure. The parameter α is set to $\alpha = 2/3$ and a maximum of 20 iterations are performed. In the scenarios ‘PBF abs vals’ and ‘PBF mixed’, the absorption coefficient is optimized 10 times before the standard optimization iterations, to have a good parameter set-up for the formation of the basis according to the absorption coefficient distribution. For the ‘harmonic’ and ‘Lagrange’ basis, k is set to $n^{\text{iter}} + 2$. Different sizes for these bases have been tested, but the chosen results present the best solution. Since the results for the fine discretization do not differ significantly, only the results for the coarse discretization are presented here.

Figure 6 shows the development of the parameter error and the objective function over the iteration count, respectively. In every outer iteration of the optimizer, three optimization steps with respect to the three parameter fields are performed and for each step, the iteration counter is incremented.

The red lines ‘element-wise’ are similar to those shown in figure 3. Figure 7 shows the final speed of sound distribution for the element-wise and ‘PBF mixed’ reconstruction.

(ii) Discussion of numerical results

All performed reconstructions (except for the one with the Lagrange basis) were successful in the sense of convergence of the objective function (figure 6). The set-ups reduced the objective

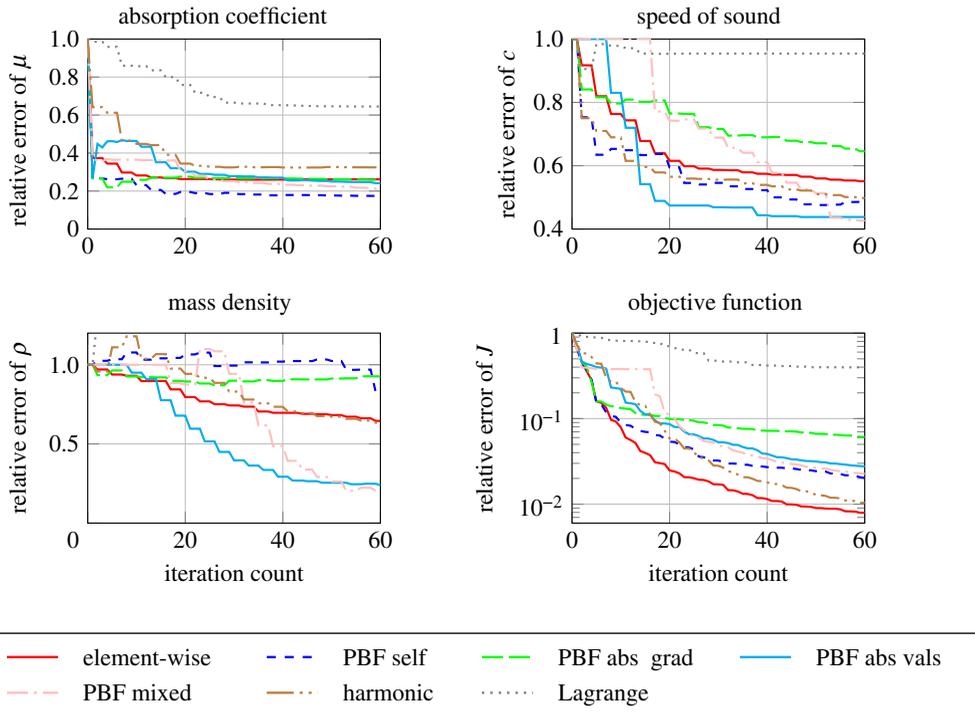


Figure 6. Convergence behaviour of the relative objective function for different patch types. (Online version in colour.)

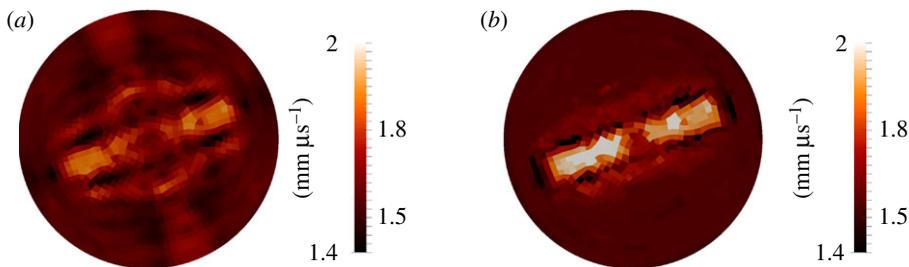


Figure 7. Speed of sound reconstruction on coarse discretization for different patch types after 20 iterations. (a) Element-wise parameter discretization and (b) 'PBF mixed' parameter discretization. (Online version in colour.)

function by 16–120 times. The element-wise parameter optimization gave the most significant decrease of the objective function and in practice, this is the only available measure.

When comparing the three PBF set-ups 'PBF self', 'PBF abs grad' and 'PBF abs vals' that all build patches according to the absorption coefficient gradient, a distinct difference in convergence is observed. This is due to the difference in the interaction of the parameters. Even the influence of the less sensitive parameters towards the most sensitive parameter is not negligible.

The two acoustical material parameters show slightly different behaviour: sensitivity is higher for the speed of sound than for the mass density which changes the signal speed in addition to the reflections also present for mass density. That is why the error in the mass density converges not as well as the error in the speed of sound in the standard set-up. For both, the PBF approach improves the final result: the final errors in the speed of sound and mass density distribution decrease to 43 and 24%, respectively. 'PBF abs vals' and 'PBF mixed' seem most suitable for the acoustical parameters. As can be seen from figure 7 displaying the speed of sound reconstruction

for element-wise and ‘PBF mixed’ parameter discretization, the patch approach improves the reconstruction quality significantly. The inclusion is more pronounced and fewer artefacts appear in the remaining domain.

The Lagrange basis performed poorly in all metrics, while the harmonic basis performed good in terms of the objective function value. The objective function decreased most for the ‘element-wise’ optimization but note that the axis of ordinates in figure 6 is logarithmic and the final differences in the objective function are small.

Figure 6 clearly shows how the objective function stagnates after some iterations in the initial set-up phase of ‘PBF abs vals’ and ‘PBF mixed’. After the set-up phase, the additional optimization with respect to the acoustical parameters enables a fast decrease in the objective function value. ‘PBF mixed’ outperforms ‘PBF abs vals’ in all four measures.

(b) Material identification

In [35], a material identification procedure for optoacoustic image reconstruction was presented. It relies on the definition of expected materials in a material catalogue. Generally, a reconstruction is run on objects of known composition and the utilization of the user’s prior knowledge is added to the reconstruction algorithm in terms of a material catalogue. This is beneficial to counter the ill-conditioning of the imaging problem. A drawback of [35] however was that it required the reconstruction of the absorption as well as the diffusion coefficient to find a unique mapping between optical values and listed materials in case similar absorption coefficients are catalogued. The diffusion coefficient however has a very low sensitivity and its reconstruction is error prone. Here, we propose a new material identification procedure that does not require the diffusion coefficient but instead uses the acoustical gradients to identify materials correctly. The absorption coefficient is reconstructed with the iterative reconstruction procedure on a pixel-based or PBF discretization and the available acoustical gradients are then used for the identification process.

The detailed procedure for material identification is displayed in algorithm 1. For input, the user has to specify a list of expected materials containing values of absorption coefficient, speed of sound and mass density. Only the absorption coefficient is optimized in a line search procedure and the material identification is carried out in every line search iteration. First, step lengths α_c, α_ρ for c, ρ are calculated as the smallest step for which $c - \alpha_c \nabla c$ covers the range $[\min c, \max c]$ from the materials in the catalogue and analogously for the mass density. Then, materials are identified that match the current element absorption coefficient. If no material is applicable, the default values for soft tissue are set. If exactly one material is applicable, the corresponding acoustical values are set. If several materials are applicable, the material is chosen for which the acoustic gradients indicate the correct trend, i.e. decrease or increase. If the acoustic gradients cannot identify any material or cannot identify one material uniquely, the one with the smallest distance to the values of the range covering distribution $c^{\text{ran}} = c - \alpha_c \nabla c$ and $\rho^{\text{ran}} = \rho - \alpha_\rho \nabla \rho$ is set.

(i) Numerical study of material identification

We again re-use the example from §4 to analyse the gains stemming from the material identification approach. Two different reconstructions are carried out on the measurement data with different material catalogues. In the most simple set-up, the catalogue matches the expected materials \mathcal{M}_o and \mathcal{M}_i as in the definitions (4.1) and (4.2). In the second set-up, a third material appears in the catalogue that is not present in the numerical phantom used for the measurement data:

$$\mathcal{M}_{\text{artificial}} = (0.11 \text{ mm}^{-1}, 0.5 \text{ mm}, 1.0 \text{ mm } \mu\text{s}^{-1}, 1.0 \text{ mg mm}^{-3}).$$

Figure 8 compares the absorption coefficient and speed of sound reconstructions after 10 iterations for the standard approach without patches and without material identification to the results when material identification is applied with the matching material catalogue and the catalogue

Algorithm 1 Material identification.

```

determine step length for  $c, \rho$ 
for all elements  $e$  do
  find all materials  $\mathcal{M}_j$  that contain  $\mu_a$ 
  if no material found then
    set default acoustical materials
  else if one material is found then
    set acoustical materials to the found material
  else
    check if the gradient directions indicate materials
    if no material direction applies then
      determine material with smallest distance and set  $c, \rho$ 
    else if one material supplies indicated direction then
      set  $c, \rho$ 
    else
      determine material with smallest distance and set  $c, \rho$ 
    end if
  end if
end for
update speed of sound and density distribution

```

with one artificial material. The standard reconstruction yields a good absorption coefficient reconstruction but the speed of sound distribution gives too low values in the inclusion. The material identification results in improved images. For the matching material catalogue, the speed of sound is recovered perfectly which in turn causes an improved absorption coefficient. For the material catalogue with an artificial material, the inclusion appears smaller in the speed of sound image and four pixels are identified with the wrong material. In terms of the image quality, the reconstruction is still significantly better than the standard reconstruction. In terms of the relative objective function, the standard approach yields a reduction to 1.5% after 10 iterations while the material identification yields 1.9 and 8.4% for material catalogue of two and three materials, respectively. The results using the material identification are obtained in a third of the computational time, because only one field is optimized with a line search procedure.

6. Mouse brain

To demonstrate the applicability of the proposed method, an *in vivo* image of a mouse brain was reconstructed by imaging a female Hsd:ATHymic Nude-Foxn1nu/nu mouse. The animal was anesthetized with a mixture of 1.8% isoflurane in 100% O₂ at 0.8 ml min⁻¹ flow rate and placed in supine position. The measurement data are generated with the multispectral optoacoustic tomograph MSOT inVision256-TF (iThera Medical GmbH, München, Germany) with a transducer ring of 80 mm diameter covering 270° with 256 array elements. The centre frequency of the transducers is 5 MHz and the signal is band-pass filtered between 100 kHz and 6 MHz. The acoustical signal is recorded with a sampling frequency of 40 MHz corresponding to a time step size of $\Delta t = 0.025 \mu\text{s}$. To decrease the influence of noise, the final signal is an average of 20 measurements. The light source is provided by a laser system with adjustable wavelength. Ten fibres provide the light from different directions to obtain a uniform illumination. A wavelength of 700 nm is chosen.

For the reconstruction, a computational model is built by a circular computational domain of 80 mm diameter meshed with 2×10^6 quadrilateral elements with quadratic shape functions. The number of elements in the reconstruction area (a circle of diameter 16 mm) is 67 960, which is the number of pixels in the image. Hence, the pixel size is about 0.06 mm. The maximal number of

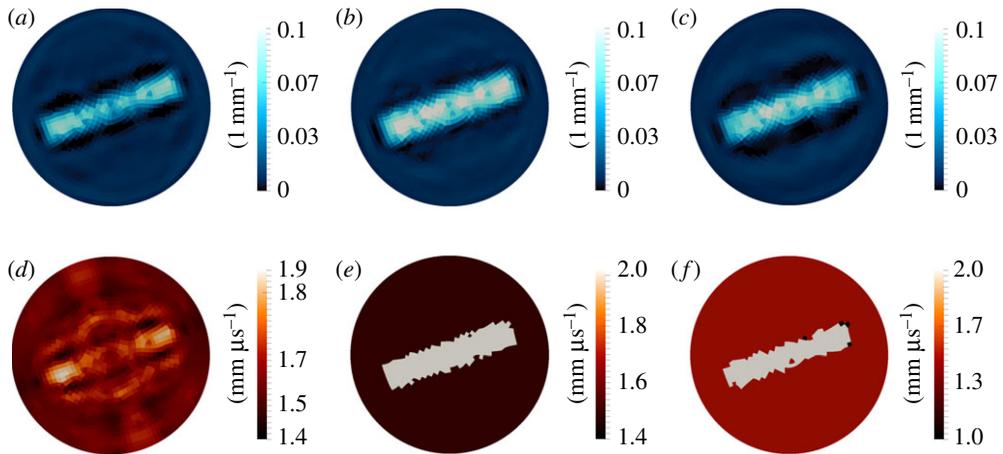


Figure 8. Comparison of reconstructed images with material identification (b),(c),(e) and (f) compared to the standard approach in terms of the absorption coefficient (a–c) and the speed of sound (d–f). (a) Standard reconstruction, μ_a , (b) two materials in catalogue, μ_a , (c) three materials in catalogue, μ_a , (d) standard reconstruction, c , (e) two materials in catalogue, c and (f) three materials in catalogue, c . (Online version in colour.)

iterations is set to 3. The numerical time step size is $\Delta t = 5 \times 10^{-4} \mu\text{s}$. In contrast to the numerical results presented in the previous sections, the impulse response of the detectors was incorporated into the reconstruction algorithm in order to represent the imaging procedure as realistically as possible. No regularization like Tikhonov or total variation was applied. For the simulations with material identification, a material catalogue consisting of water, general soft tissue, skin and bone was supplied.

Figure 9 shows a cryoslice through the mouse brain for reference, identifying the basic anatomic features, and the reconstruction of the absorbed energy map is shown using model-based inversion as described in [6]. Figure 10 shows the reconstructions of the absorption coefficient, the speed of sound and the mass density for reconstructions with element-wise parameter discretization, ‘PBF mixed’ with $\alpha = 0.1$, and material identification. Since this is the first time a general reconstruction of acoustical parameters in an optoacoustic tomograph from an *in vivo* experiment are presented, not only the PBF but also the element-wise reconstructions are presented. In figure 9, a plot along the indicated lines compares the result from the model-based inversion and the element-wise reconstruction with the new approach and detected acoustical heterogeneities. Note that all results are marked with physical units but the measurement data is in arbitrary units and the measurement data is calibrated in the process of reconstruction. The mass density contributes only over relative differences and not over its absolute value, which renders the scaling dependent on initial values.

The new method successfully highlights all the features as in the model-based image like the cortex or the temporal artery. More oscillations and artefacts at the exterior of the mouse brain occur due to the wave propagation over long distances with numerical errors and because no regularization is applied. For ‘PBF mixed’, the acoustical images are segmented during optimization. The number of spatial variations is decreased while the contrast is increased. The element-wise mass density and speed of sound reconstructions show several fluctuations of high amplitude. The fluctuations are highest in the regions of the skull. In the centre of the interior, low values are found. The ‘PBF mixed’ approach alters the mass density image significantly compared to the element-wise reconstruction: the basic features of the absorption coefficient distribution are retrieved, making it more representative for the tissue types. The material identification identifies 0.3% of the image pixels as bone which are located at the actual position of the skull. However, not the entire skull is identified correctly. Note that the scales in

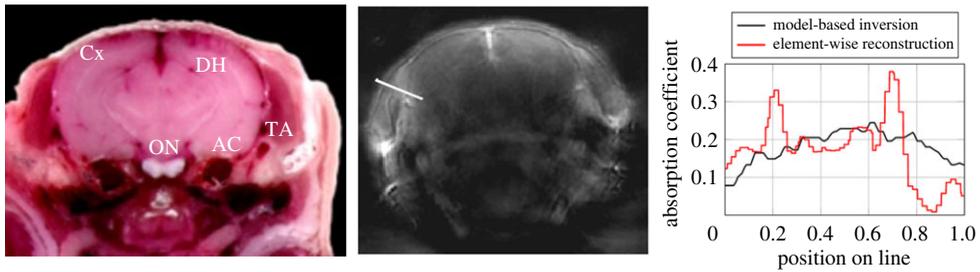


Figure 9. Mouse brain imaging, from left to right: cryoslice (Cx, cortex; DH, dorsal hippocampus; TA, temporal artery; AC, amygdala complex; ON, optic nerve), the absorbed energy map reconstructed using model-based inversion, and a plot of the absorption coefficient along line indicated in the model-based inversion and the element-wise reconstruction (figure 10). (Online version in colour.)

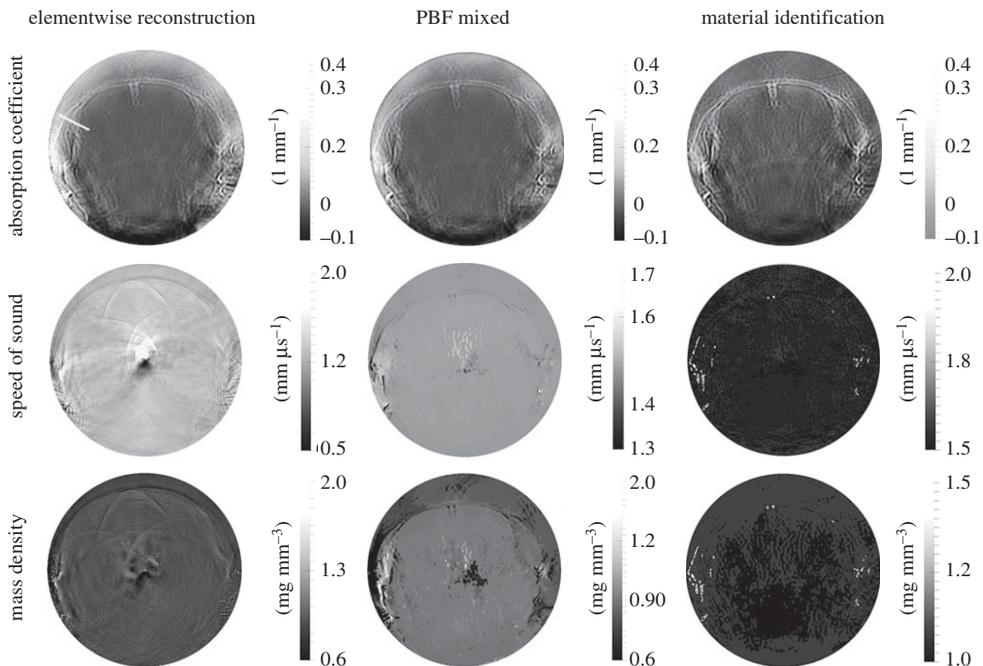


Figure 10. Mouse brain imaging results: absorption coefficient, speed of sound and mass density from top to bottom and element-wise, 'PBF mixed' and material identification results from left to right.

figure 10 do not show the maximal value: for the material identification approach, a maximal value of $c = 4.1 \text{ mm } \mu\text{s}^{-1}$ for the sound speed is reached, while the other reconstructions only show lower maximal values ($c = 3.0 \text{ mm } \mu\text{s}^{-1}$ and $c = 1.8 \text{ mm } \mu\text{s}^{-1}$ for element-wise and 'PBF mixed', respectively). The comparison between the model-based inversion and the element-wise reconstruction in figure 9 shows that the new approach displays distinct layers along the skull, as expected. The peaks are located where the speed of sound and mass density reconstructions reveal acoustical heterogeneities. The model-based inversion fails to reconstruct these features. The objective function after three iterations was lowest for the 'element-wise' set-up, just as for the numerical example in the preceding sections.

7. Conclusion

We have presented a new reconstruction method for the optoacoustic imaging inverse problem. It relies on the diffusion approximation for the optical transport and the acoustic wave equation for sound propagation. Discretization is carried out by continuous finite elements and a hybridizable discontinuous Galerkin method, respectively. Optimization is performed with an iterative gradient-based scheme with the gradient resulting from the adjoint equations.

The presented method allows not only for the reconstruction of optical parameters, namely absorption coefficient and diffusion coefficient but also for the acoustical parameters speed of sound and mass density. The gradient calculation for the acoustical parameters is slightly more elaborate than for the optical parameters since the combination of forward and adjoint quantities in each time step is required for their evaluation. This is realized by a checkpointing approach and a sequential repetition of forward evaluations during the adjoint run. Numerical results demonstrate the convergence of the proposed method.

Optoacoustic imaging is inherently ill-conditioned, especially in relevant applications of e.g. small animal imaging. In order to limit the ill-conditioning introduced by the additional reconstruction parameters, we have proposed to adapt the bases of the parameter spaces during the iterative procedure using patches. Doing so, the parameter space is reduced in each iteration (which improves conditioning) but the overall scheme retains the full flexibility because the bases potentially change in each iteration. We have presented several approaches to build the PBFs and compared their performance to the traditional pixel-based or element-wise approach using a numerical example. Also, global Lagrange polynomials as well as harmonic functions have been evaluated as basis functions for the parameter space but their performance turned out unsatisfactory. The PBFs are beneficial considering the error in the parameter fields. In the objective function value, the convergence is slightly slower which is due to the fact that the adaptation of the basis decelerates the optimization from the point of view of pixels, which are located at the boundary of patches. Especially, beneficial for optoacoustic imaging are the approaches 'PBF abs vals' and 'PBF mixed' because they transport the information of inclusions from the most sensitive parameter field (absorption coefficient) to the less sensitive fields. The approach of PBF can also be applied to other inverse problems, including inverse problems with only one parameter field (as described for 'PBF self') or inverse problems with several parameter fields. Owing to its simplicity, it is easily introduced into existing reconstruction codes and its improvement of the conditioning allows to incorporate more unknowns. Also, PBF inherently provides regularization, reducing or even avoiding the need for suited regularization schemes and weights.

As a second possibility to render the optoacoustic imaging problem less ill-conditioned, we have proposed a material identification where only the absorption coefficient is reconstructed and the acoustical parameters are updated based on the acoustic gradients and a user-specified catalogue of expected materials. Thereby, the computational effort is lower and the ill-conditioning amounts to that of the pure absorption coefficient reconstruction. We have reported significant gains in numerical experiments in terms of computational time as well as image quality. The material identification can easily be applied to other inverse problems where the tissues of an object are known and one is interested in several parameter fields.

We have successfully applied the presented reconstruction algorithm to measurement data from an *in vivo* mouse brain experiment. Several anatomical features were clearly identified. Oscillations occur in the absorption coefficient reconstruction but additional features along the skull were revealed. The material identification approach helps to identify bone structures surrounding the brain. The robustness of the material identification in the presence of slight material variations within one tissue type, as apparent in an experimental set-up, are subject to future research.

Future work addresses a performance comparison between PBF and level-set based optimizations. Applicability and quantitative image quality improvements due to the

reconstruction of acoustical heterogeneities will be studied extensively based on phantoms and for typical optoacoustic imaging scenarios.

Ethics. All procedures involving animal experimentation were conducted according to the guidelines of Helmholtz Zentrum München and the government of Upper Bavaria and complied with German Federal and EU law. Efforts were made to reduce animal usage and suffering.

Data accessibility. The source code for the image reconstruction including the PBF approach and the material identification and measurement data is provided in the electronic supplementary material. The wave equation solver underlying the experiments is open-source software, licenced under the LGPL, and available at <https://github.com/kronbichler/exwave>. For a description of the algorithms, see [24,35,48].

Authors' contributions. S.S., M.K. and W.A.W. developed the numerical methods and performed the numerical experiments, I.O., and V.N. set up the experimental methodology, I.O. and S.S. performed the *in vivo* measurements. S.S. prepared the article.

Competing interests. We declare that we have no competing interests.

Funding. No funding has been received for this article.

Acknowledgements. We thank our colleagues for fruitful discussions.

References

1. Taruttis A, van Dam G, Ntziachristos V. 2015 Mesoscopic and macroscopic optoacoustic imaging of cancer. *Cancer Res.* **75**, 1548–1559. (doi:10.1158/0008-5472.CAN-14-2522)
2. Beziere N *et al.* 2014 Optoacoustic imaging and staging of inflammation in a murine model of arthritis. *Arthritis Rheumatol.* **66**, 2071–2078. (doi:10.1002/art.38642)
3. Jansen K, Wu M, van der Steen AF, van Soest G. 2014 Photoacoustic imaging of human coronary atherosclerosis in two spectral bands. *Photoacoustics* **2**, 12–20. (doi:10.1016/j.pacs.2013.11.003)
4. Xu M, Wang L. 2005 Universal back-projection algorithm for photoacoustic computed tomography. *Phys. Rev. E* **71**, 016706. (doi:10.1103/PhysRevE.71.016706)
5. Mohajerani P, Tzoumas S, Rosenthal A, Ntziachristos V. 2015 Optical and optoacoustic model-based tomography: theory and current challenges for deep tissue imaging of optical contrast. *IEEE Signal Proc. Mag.* **32**, 88–100. (doi:10.1109/msp.2014.2352672)
6. Rosenthal A, Ntziachristos V, Razansky D. 2011 Model-based optoacoustic inversion with arbitrary-shape detectors. *Med. Phys.* **38**, 4285–4295. (doi:10.1118/1.3589141)
7. Burgholzer P, Matt GJ, Haltmeier M, Paltauf G. 2007 Exact and approximative imaging methods for photoacoustic tomography using an arbitrary detection surface. *Phys. Rev. E* **75**, 046706. (doi:10.1103/PhysRevE.75.046706)
8. Hristova Y, Kuchment P, Nguyen L. 2008 Reconstruction and time reversal in thermoacoustic tomography in acoustically homogeneous and inhomogeneous media. *Inverse Probl.* **24**, 055006. (doi:10.1088/0266-5611/24/5/055006)
9. Treeby B, Zhang E, Cox B. 2010 Photoacoustic tomography in absorbing acoustic media using time reversal. *Inverse Probl.* **26**, 115003. (doi:10.1088/0266-5611/26/11/115003)
10. Huang C, Wang K, Nie L, Wang L, Anastasio M. 2013 Full-wave iterative image reconstruction in photoacoustic tomography with acoustically inhomogeneous media. *IEEE T. Med. Imaging* **32**, 1097–1110. (doi:10.1109/TMI.2013.2254496)
11. Haltmeier M, Nguyen L. 2017 Analysis of iterative methods in photoacoustic tomography with variable sound speed. *SIAM J. Imaging Sci.* **10**, 751–781. (doi:10.1137/16M1104822)
12. Cox BT, Arridge SR, Beard PC. 2007 Gradient-based quantitative photoacoustic image reconstruction for molecular imaging. In *Proc. SPIE* vol. 6437, p. 64371T.
13. Elbau P, Scherzer O. 2015 Modelling the effect of focusing detectors in photoacoustic sectional imaging. *SIAM J. Imaging Sci.* **8**, 1–18. (doi:10.1137/120882962)
14. Hannukainen A, Hyvönen N, Majander H, Tarvainen T. 2016 Efficient inclusion of total variation type priors in quantitative photoacoustic tomography. *SIAM J. Imaging Sci.* **9**, 1132–1153. (doi:10.1137/15M1051737)
15. Naetar W, Scherzer O. 2014 Quantitative photoacoustic tomography with piecewise constant material parameters. *SIAM J. Imaging Sci.* **7**, 1755–1774. (doi:10.1137/140959705)
16. Tarvainen T, Cox B, Kaipio J, Arridge S. 2012 Reconstructing absorption and scattering distributions in quantitative photoacoustic tomography. *Inverse Probl.* **28**, 084009. (doi:10.1088/0266-5611/28/8/084009)

17. Saratoon T, Tarvainen T, Cox B, Arridge S. 2013 A gradient-based method for quantitative photoacoustic tomography using the radiative transfer equation. *Inverse Probl.* **29**, 075006. (doi:10.1088/0266-5611/29/7/075006)
18. Belhachmi Z, Glatz T, Scherzer O. 2016 A direct method for photoacoustic tomography with inhomogeneous sound speed. *Inverse Probl.* **32**, 045005. (doi:10.1088/0266-5611/32/4/045005)
19. Dean-Ben X, Rui M, Razansky D, Ntziachristos V. 2011 Statistical approach for optoacoustic image reconstruction in the presence of strong acoustic heterogeneities. *IEEE T. Med. Imag.* **30**, 401–408. (doi:10.1109/TMI.2010.2081683)
20. Huang C, Wang K, Schoonover RW, Wang LV, Anastasio MA. 2016 Joint reconstruction of absorbed optical energy density and sound speed distributions in photoacoustic computed tomography: a numerical investigation. *IEEE Trans. Comput. Imaging* **2**, 136–149. (doi:10.1109/TCI.2016.2523427)
21. Stefanov P, Uhlmann G. 2013 Instability of the linearized problem in multiwave tomography of recovery both the source and the speed. *Inverse Probl. Imaging* **7**, 1367–1377. (doi:10.3934/ipi.2013.7.1367)
22. Yuan X, Wang L. 2003 Effects of acoustic heterogeneity in breast thermoacoustic tomography. *IEEE Trans. Ultrason. Ferr.* **50**, 1134–1146. (doi:10.1109/TUFFC.2003.1235325)
23. Ammari H, Bretin E, Garnier J, Jugnon V. 2012 Coherent interferometry algorithms for photoacoustic imaging. *SIAM J. Numer. Anal.* **50**, 2259–2280. (doi:10.1137/100814275)
24. Kronbichler M, Schoeder S, Müller C, Wall W. 2016 Comparison of implicit and explicit hybridizable discontinuous Galerkin methods for the acoustic wave equation. *Int. J. Numer. Eng.* **106**, 712–739. (doi:10.1002/nme.5137)
25. Stefanov P, Uhlmann G. 2013 Recovery of a source term or a speed with one measurement and applications. *Trans. Am. Math. Soc.* **365**, 5737–5758. (doi:10.1090/S0002-9947-2013-05703-0)
26. Wang L (ed.) 2009 *Photoacoustic imaging and spectroscopy*. Boca Raton, FL: CRC Press.
27. Wang K, Schoonover RW, Su R, Oraevsky A, Anastasio MA. 2014 Discrete imaging models for three-dimensional optoacoustic tomography using radially symmetric expansion functions. *IEEE Trans. Med. Imaging* **33**, 1180–1193. (doi:10.1109/tmi.2014.2308478)
28. Schweiger M, Arridge S. 2003 Image reconstruction in optical tomography using local basis functions. *J. Electron. Imaging* **12**, 583–593. (doi:10.1117/1.1586919)
29. Cezaro AD, Leitao A, Tai XC. 2013 On piecewise constant level-set (PCLS) methods for the identification of discontinuous parameters in ill-posed problems. *Inverse Probl.* **29**, 015003. (doi:10.1088/0266-5611/29/1/015003)
30. Schweiger M, Arridge SR, Dorn O, Zacharopoulos A, Kolehmainen V. 2006 Reconstructing absorption and diffusion shape profiles in optical tomography by a level set technique. *Opt. Lett.* **31**, 471–473. (doi:10.1364/OL.31.000471)
31. Alberti G, Ammari H, Jin B, Seo J, Zhang W. 2016 The linearized inverse problem in multifrequency electrical impedance tomography. *SIAM J. Imaging Sci.* **9**, 1525–1551. (doi:10.1137/16M1061564)
32. Alberti GS, Ammari H. 2017 Disjoint sparsity for signal separation and applications to hybrid inverse problems in medical imaging. *Appl. Comput. Harmonic Anal.* **42**, 319–349. (doi:10.1016/j.acha.2015.08.013)
33. Ammari H, Bossy E, Jugnon V, Kang H. 2011 Reconstruction of the optical absorption coefficient of a small absorber from the absorbed energy density. *SIAM J. Appl. Math.* **71**, 676–693. (doi:10.1137/09077905X)
34. Mandal S, Deán-Ben XL, Razansky D. 2016 Visual quality enhancement in optoacoustic tomography using active contour segmentation priors. *IEEE Trans. Med. Imaging* **35**, 2209–2217. (doi:10.1109/TMI.2016.2553156)
35. Schoeder S, Kronbichler M, Wall W. 2017 Photoacoustic image reconstruction: material detection and acoustical heterogeneities. *Inverse Probl.* **33**, 055010. (doi:10.1088/1361-6420/aa635b)
36. Cheong WF, Prahl S, Welch A. 1990 A review of the optical properties of biological tissues. *IEEE J. Quantum Electron.* **26**, 2166–2185. (doi:10.1109/3.64354)
37. Feng D (ed.) 2008 *Biomedical information technology*. Elsevier Inc.
38. Wang L, Wu HI. 2007 *Biomedical optics: principles and imaging*. Hoboken, NJ: John Wiley and Sons.
39. Bal G, Ren K. 2011 Multi-source quantitative photoacoustic tomography in a diffusive regime. *Inverse Probl.* **27**, 075003. (doi:10.1088/0266-5611/27/7/075003)

40. Yuan Z, Jiang H. 2009 Quantitative photoacoustic tomography. *Phil. Trans. R. Soc. A* **367**, 3043–3054. (doi:10.1098/rsta.2009.0083)
41. Haltmeier M, Neumann L, Rabanser S. 2015 Single-stage reconstruction algorithm for quantitative photoacoustic tomography. *Inverse Probl.* **31**, 065005. (doi:10.1088/0266-5611/31/6/065005)
42. Kirsch A, Scherzer O. 2012 Simultaneous reconstructions of absorption density and wave speed with photoacoustic measurements. *SIAM J. Appl. Math.* **72**, 1508–1523. (doi:10.1137/110849055)
43. Jin X, Wang LV. 2006 Thermoacoustic tomography with correction for acoustic speed variations. *Phys. Med. Biol.* **51**, 6437–6448. (doi:10.1088/0031-9155/51/24/010)
44. Griewank A. 1992 Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation. *Optim. Method. Softw.* **1**, 35–54. (doi:10.1080/10556789208805505)
45. Nocedal J, Wright S. 2006 *Numerical optimization*, 2 edn. Berlin, Germany: Springer.
46. Ren K, Gao H, Zhao H. 2013 A hybrid reconstruction method for quantitative PAT. *SIAM J. Imaging Sci.* **6**, 32–55. (doi:10.1137/120866130)
47. Wirgin A. 2008 The inverse crime. Laboratoire de Mécanique et d' Acoustique, France. (<http://arxiv.org/abs/math-ph/0401050v1>).
48. Schoeder S, Kormann K, Wall W, Kronbichler M. 2018 Efficient explicit time stepping of high order discontinuous Galerkin schemes for waves. (<http://arxiv.org/abs/1805.03981>).