

Stochastic Proximal Gradient Algorithms for Multi-Source Quantitative Photoacoustic Tomography

Simon Rabanser

Department of Mathematics, University of Innsbruck
Technikerstraße 13, 6020 Innsbruck, Austria
E-mail: simon.rabanser@uibk.ac.at

Lukas Neumann

Institute of Basic Sciences in Engineering Science, University of Innsbruck
Technikerstraße 13, 6020 Innsbruck, Austria
E-mail: lukas.neumann@uibk.ac.at

Markus Haltmeier

Department of Mathematics, University of Innsbruck
Technikerstraße 13, 6020 Innsbruck, Austria
E-mail: markus.haltmeier@uibk.ac.at

Abstract

The development of accurate and efficient image reconstruction algorithms is a central aspect of quantitative photoacoustic tomography (QPAT). In this paper, we address this issues for multi-source QPAT using the radiative transfer equation (RTE) as accurate model for light transport. The tissue parameters are jointly reconstructed from the acoustical data measured for each of the applied sources. We develop stochastic proximal gradient methods for multi-source QPAT, which are more efficient than standard proximal gradient methods in which a single iterative update has complexity proportional to the number applies sources. Additionally, we introduce a completely new formulation of QPAT as multilinear (MULL) inverse problem which avoids explicitly solving the RTE. The MULL formulation of QPAT is again addressed with stochastic proximal gradient methods. Numerical results for both approaches are presented. Besides the introduction of stochastic proximal gradient algorithms to QPAT, we consider the new MULL formulation of QPAT as main contribution of this paper.

Keywords: Photoacoustic tomography; image reconstruction; radiative transfer equation; multilinear inverse problem; limited view; stochastic gradient method; limited data; Dykstra algorithm

1 Introduction

Photoacoustic tomography (PAT) is an emerging imaging modality, which combines the benefits of pure ultrasound imaging (high resolution) with those of pure optical tomography (high contrast); see [9, 54]. The basic principle of PAT is as follows (see Figure 1.1): A semitransparent sample such as a part of a human patient is illuminated with short pulses of optical radiation. A fraction of the optical energy is absorbed inside the sample, which causes thermal heating, expansion, and a subsequent acoustic pressure wave depending on the interior absorbing structure of the sample. The acoustic pressure is measured outside of the sample and used to reconstruct an image of the interior.

One important reconstruction problem in PAT is recovering the initial pressure distribution (see, for example, [1, 12, 29, 34, 35, 36, 45, 48]). The initial pressure distribution only provides qualitative information about the tissue-relevant parameters, as it is the product of the optical absorption coefficient and the spatially varying optical intensity, which again indirectly depends on the tissue parameters. Quantitative photoacoustic tomography (QPAT) addresses this issue and aims at quantitatively estimating the tissue parameters by supplementing the inversion of the acoustic wave equation with an inverse problem for light propagation (see, for example, [3, 6, 7, 13, 16, 17, 18, 33, 32, 41, 43, 44, 47, 49, 52, 56]).




Figure 1.1: Basic principles of PAT. **Left:** the investigated object is illuminated with a short optical pulse; **Middle:** due to the thermoelastic effect, the absorbed light distribution induces an acoustic pressure wave depending on internal tissue properties; **Right:** the acoustic pressure wave is measured outside the object and used to reconstruct an image of the interior.

1.1 Multi-Source QPAT

In this paper, we consider image reconstruction in QPAT using multiple sources. We allow limited view measurements, where, for each illumination, partial data are collected only from a certain angular domain. For modeling the light transport, we use the radiative transfer equation (RTE), which is commonly considered as a very accurate model for light transport in tissue (see, for example, [5, 19, 23, 40]). In particular, opposed to the diffusion approximation, the RTE allows for modeling directed optical radiation, which is required for a reasonable QPAT forward model. Additionally, it allows for including internal voids as regions of low scattering. As proposed in [33], we work with a single-stage reconstruction procedure for QPAT, where the optical parameters are reconstructed directly from the measured acoustical data. The image reconstruction problem of multi-source QPAT using N different sources can be formulated as a system of nonlinear

equations (see, for example, [33, 28])

$$\mathbf{F}_i(\boldsymbol{\mu}) = v_i \quad \text{for } i = 1, \dots, N. \quad (1.1)$$

Here, \mathbf{F}_i is the operator that maps the unknown parameter pair $\boldsymbol{\mu} = (\mu_a, \mu_s)$ consisting of the absorption coefficient $\mu_a: \Omega \rightarrow \mathbb{R}$ and the scattering coefficient $\mu_s: \Omega \rightarrow \mathbb{R}$ to the measured acoustic data v_i corresponding to the i -th source distribution (see Section 2 for precise definitions). There are two main classes of methods for solving the nonlinear inverse problem (1.1), namely, Tikhonov type regularization on the one and iterative regularization methods on the other hand [24, 39, 51]. Both approaches are based on rewriting (1.1) as a single equation $\mathbf{F}(\boldsymbol{\mu}) = v$ with forward operator $\mathbf{F} = (\mathbf{F}_i)_{i=1}^N$ and data $v = (v_i)_{i=1}^N$. In Tikhonov regularization, one defines approximate solutions as minimizers of the penalized least squares functional $\frac{1}{2} \|\mathbf{F}(\boldsymbol{\mu}) - v\|^2 + \lambda R(\boldsymbol{\mu})$. Here, $R(\cdot)$ is an appropriate regularization functional included to stabilize the inversion process and λ a regularization parameter that has to be carefully chosen depending on the data and the noise. In iterative regularization methods, stabilization is achieved via early stopping of iterative schemes. In such a situation, one usually applies iterative optimization techniques designed for minimizing the un-regularized least squares functional $\frac{1}{2} \|\mathbf{F}(\boldsymbol{\mu}) - v\|^2$, and the iteration index plays the role of the regularization parameter.

Tikhonov type as well as iterative regularization methods can both be formulated as finding a solution of the optimization problem

$$\begin{cases} \min & \frac{1}{2} \sum_{i=1}^N \|\mathbf{F}_i(\boldsymbol{\mu}) - v_i\|^2 + G(\boldsymbol{\mu}), \\ \text{with} & \boldsymbol{\mu} \in L^2(\Omega) \times L^2(\Omega). \end{cases} \quad (1.2)$$

In iterative regularization methods, one takes $G = \boldsymbol{\chi}_{\mathbb{D}}$, the characteristic function of the domain of definition \mathbb{D} of the forward operator (taking the value 0 in and the value ∞ outside of \mathbb{D}). In Tikhonov regularization, we take $G = \boldsymbol{\chi}_{\mathbb{D}} + \lambda R$. Well established algorithms for solving Equation (1.2) are proximal gradient algorithms [15, 8], which can be written in the form

$$\boldsymbol{\mu}_{k+1} = \text{prox}_{s_k G} \left(\boldsymbol{\mu}_k - s_k \sum_{i=1}^N \mathbf{F}'_i(\boldsymbol{\mu}_k)^* (\mathbf{F}_i(\boldsymbol{\mu}_k) - v_i) \right). \quad (1.3)$$

Here, $\text{prox}_{s_k G}$ is the proximity operator and s_k the positive step size; $\mathbf{F}'_i(\boldsymbol{\mu}_k)$ denotes the derivative of the i -th forward operator evaluated at $\boldsymbol{\mu}_k$ with $\mathbf{F}'_i(\boldsymbol{\mu}_k)^*$ being its Hilbert space adjoint.

1.2 Stochastic Proximal Gradient Algorithms

Each iteration in the proximal gradient algorithm (1.3) can be numerically quite expensive, since it requires solving the forward and adjoint problems for all N equations in (1.1). In many cases, stochastic (proximal) gradient methods turn out to be more efficient since these methods only consider one of the equations in (1.1) per iteration. The stochastic proximal gradient method (see, for example, [10, 11, 55, 21, 42, 46] and the references therein) for solving (1.2) is defined by

$$\boldsymbol{\mu}_{k+1} = \text{prox}_{s_k G} \left(\boldsymbol{\mu}_k - s_k \mathbf{F}'_{i(k)}(\boldsymbol{\mu}_k)^* (\mathbf{F}_{i(k)}(\boldsymbol{\mu}_k) - v_{i(k)}) \right),$$

where $i(k) \in \{1, \dots, N\}$ corresponds to one of the equations in (1.1) that is selected randomly for the update in the k -th iteration. In opposition to the standard proximal gradient method, this requires solving only one forward and one adjoint problem per iteration. Therefore, one iterative

step is much cheaper for the stochastic gradient method than for the full gradient method. In the case of no regularization, $\lambda = 0$, the stochastic proximal gradient method reduces to the Kaczmarz method for inverse problems studied in [20, 31, 30].

The computationally most expensive task in the above methods is the numerical solution of the RTE. In this paper, we therefore additionally study a reformulation of the inverse problem of QPAT avoiding the computation of a solution of the RTE. For this purpose, the inverse problem is reformulated as multilinear inverse problem (4.1), where the RTE is added as a constraint instead of explicitly including its solution. The new formulation will be again addressed by Tikhonov regularization in combination with proximal stochastic gradient methods as discussed in Section 4.

Note that, in QPAT, it has often been assumed that the initial pressure distribution (corresponding to each illumination) is already recovered from acoustic measurements (see, for example, [2, 4, 7, 17, 50, 52, 47, 53, 56]). Research was focused on inverting the light propagation in tissues either modeled by the RTE or the diffusion approximation. In the case that acoustic measurements are only known on parts of the boundary, reconstruction of the initial pressure distribution is not possible in a stable manner. In order to obtain stable reconstruction results in [33], we propose a single-stage approach for QPAT, where the optical parameters are directly recovered from the acoustic boundary data. Throughout this paper, we will make use of this approach, which delivers stable results especially in the limited view situation. In opposition to [33], in this paper, we introduce (proximal) stochastic gradient methods, which effectively exploit the multi-illumination structure and turn out to be faster than the standard proximal gradient methods.

1.3 Outline

The remainder of this paper is organized as follows. In Section 2, we provide the mathematical model for QPAT (the forward problem) using the RTE. We allow multiple sources and partial acoustic measurements. We also recall known results for QPAT including differentiability of the forward problem. In Section 3, we address the inverse problem of QPAT using Tikhonov regularization and study the proximal stochastic gradient method for its solution. The new reformulation of the inverse problem of QPAT as a multilinear inverse problem is presented in Section 4. For the solution of the proposed formulation, we again develop proximal gradient methods. Numerical results are presented in Section 5. The paper is concluded with a summary and outlook presented in Section 6.

2 The Forward Problem in QPAT

The image reconstruction problem of QPAT can be written as the system (1.1) of nonlinear equations, where the forward operators \mathbf{F}_i map tissue relevant parameters to acoustic data sets recorded in specific regions outside the tissue. Precise formulations will be given in this section.

2.1 Mathematical Notation

We fix some mathematical notation that is used throughout this paper. We denote by $\Omega \subseteq \mathbb{R}^d$ a convex domain with piecewise smooth boundary modeling our domain of interest, with $d \in \{2, 3\}$

denoting the spatial dimension. In order to be able to impose appropriate boundary conditions for the RTE, it is convenient to split the set $\Gamma := \partial\Omega \times \mathbb{S}^{d-1}$ into inflow and outflow boundaries,

$$\begin{aligned}\Gamma_- &:= \{(x, \theta) \in \partial\Omega \times \mathbb{S}^{d-1} \mid \nu(x) \cdot \theta \leq 0\}, \\ \Gamma_+ &:= \{(x, \theta) \in \partial\Omega \times \mathbb{S}^{d-1} \mid \nu(x) \cdot \theta > 0\},\end{aligned}$$

with $\nu(x)$ denoting the outward pointing unit normal at $x \in \partial\Omega$ and $x \cdot y$ the standard inner product in \mathbb{R}^d . We write $B_R = \{x \in \mathbb{R}^d \mid \|x\| < R\}$ for the ball of radius R centered at the origin and suppose $B_R \supseteq \Omega$.

By $L^2(\Omega)$ and $L^2(\Omega \times \mathbb{S}^{d-1})$, we denote the Hilbert spaces of square integrable functions on Ω and $\Omega \times \mathbb{S}^{d-1}$, respectively. By $L^2(\Gamma_-, |\nu \cdot \theta|)$, we denote the space of all $q_o: \Gamma_- \rightarrow \mathbb{R}$ for which $\|q_o\|_{L^2(\Gamma_-, |\nu \cdot \theta|)}^2 := \int_{\Gamma_-} |q_o(x, \theta)|^2 |\nu \cdot \theta| d(x, \theta)$ is finite. We further write

$$\begin{aligned}\|\Phi\|_{\mathbb{W}}^2 &:= \|\Phi\|_{L^2(\Omega \times \mathbb{S}^{d-1})}^2 + \|\theta \cdot \nabla_x \Phi\|_{L^2(\Omega \times \mathbb{S}^{d-1})}^2 + \|\Phi|_{\Gamma_-}\|_{L^2(\Gamma_-, |\nu \cdot \theta|)}^2, \\ \|v\|_{\mathbb{Y}}^2 &:= \int_0^\infty \int_{\partial B_R} |v(x, t)|^2 t dx dt,\end{aligned}$$

and define

$$\begin{aligned}\mathbb{Q} &:= L^2(\Omega \times \mathbb{S}^{d-1}) \times L^2(\Gamma_-, |\nu \cdot \theta|), \\ \mathbb{X} &:= L^2(\Omega) \times L^2(\Omega), \\ \mathbb{W} &:= \{\Phi: \Omega \times \mathbb{S}^{d-1} \rightarrow \mathbb{R} \mid \|\Phi\|_{\mathbb{W}} < \infty\}, \\ \mathbb{Y} &:= \{v: \partial B_R \times (0, \infty) \rightarrow \mathbb{R} \mid \|v\|_{\mathbb{Y}} < \infty\}.\end{aligned}$$

The inner products in \mathbb{Q} , \mathbb{X} , \mathbb{W} , \mathbb{Y} will be denoted by $\langle \cdot, \cdot \rangle_{\mathbb{Q}}$, $\langle \cdot, \cdot \rangle_{\mathbb{X}}$, $\langle \cdot, \cdot \rangle_{\mathbb{W}}$, $\langle \cdot, \cdot \rangle_{\mathbb{Y}}$, respectively. The subspace of all $\Phi \in \mathbb{W}$ with $\Phi|_{\Gamma_-} = 0$ will be denoted by \mathbb{W}_0 .

Elements in \mathbb{X} will be written in the form $\mu = (\mu_a, \mu_s)$ and are the parameters we aim to determine. They are actually required to be contained in the convex subset

$$\mathbb{D}(\mathbf{T}) := \{\mu \in \mathbb{X} \mid 0 \leq \mu_a \leq \bar{\mu}_a, 0 \leq \mu_s \leq \bar{\mu}_s\}, \quad (2.1)$$

where $\bar{\mu}_a, \bar{\mu}_s > 0$. Elements in \mathbb{Q} will be written in the form $q = (q_o, q_i)$ and model the optical sources. Elements in \mathbb{W} describe the optical radiation, and elements in \mathbb{Y} the measured acoustic data.

2.2 The Radiative Transfer Equation

To specify the forward operators, we require mathematical models for the light propagation, the conversion of optical into acoustic energy, and the propagation of the acoustic waves. These models will be presented in the rest of this section.

We model the optical radiation by a function $\Phi: \Omega \times \mathbb{S}^{d-1} \rightarrow \mathbb{R}$, where $\Phi(x, \theta)$ is the density of photons at position $x \in \Omega$ and propagating into direction $\theta \in \mathbb{S}^{d-1}$. The interaction of the photons with the background are described by absorption coefficient $\mu_a: \Omega \rightarrow \mathbb{R}$, the scattering coefficient $\mu_s: \Omega \rightarrow \mathbb{R}$, and the scattering operator $\mathbf{K}: \Phi \mapsto \mathbf{K}\Phi$, taking the form (see [5, 40])

$$\forall (x, \theta) \in \Omega \times \mathbb{S}^{d-1}: \quad \mathbf{K}\Phi(x, \theta) = \int_{\mathbb{S}^{d-1}} k(\theta, \theta') \Phi(x, \theta') d\theta', \quad (2.2)$$

with scattering kernel $k: \mathbb{S}^{d-1} \times \mathbb{S}^{d-1} \rightarrow \mathbb{R}$. The absorption coefficient describes the ability of the background to absorb photons and the scattering coefficient describes the amount of photon

scattering. The scattering kernel $k(\theta, \theta')$ describes the redistribution of velocity directions due to interaction of the photons with the background. From physical considerations, it is natural to assume k to be measurable, symmetric, nonnegative, and to satisfy $\int_{\mathbb{S}^{d-1}} k(\cdot, \theta') d\theta' = 1$. In this article, we are concerned with the situation when the kernel is known a priori. Additionally, we assume k to be essentially bounded.

The photon density $\Phi(x, \theta)$ is supposed to satisfy the stationary radiative transfer equation (RTE),

$$(\theta \cdot \nabla_x + \mu_a + \mu_s(\mathbf{I} - \mathbf{K})) \Phi(x, \theta) = q_i(x, \theta) \quad \text{for } (x, \theta) \in \Omega \times \mathbb{S}^{d-1} \quad (2.3)$$

with boundary conditions

$$\Phi|_{\Gamma_-}(x, \theta) = q_o(x, \theta) \quad \text{for } (x, \theta) \in \Gamma_- . \quad (2.4)$$

Here, $q_i: \Omega \times \mathbb{S}^{d-1} \rightarrow \mathbb{R}$ denotes an internal photon source and $q_o: \Gamma_- \rightarrow \mathbb{R}$ a prescribed boundary source pattern. Note that PAT uses very short light pulses (below microseconds) and that light propagation happens on time scales much shorter than the scale of acoustic wave propagation. This justifies the use of the stationary case for the RTE; see [6] for a more complete discussion.

Theorem 2.1 (Well-posedness of the RTE). *For every $\mu \in \mathbb{D}(\mathbf{T})$ and $q \in \mathbb{Q}$, the stationary RTE (2.3) admits a unique solution $\Phi \in \mathbb{W}$. Moreover, there exists a constant C only depending on the parameters $\bar{\mu}_a, \bar{\mu}_s > 0$ (defining the domain $\mathbb{D}(\mathbf{T})$), such that*

$$\|\Phi\|_{\mathbb{W}} \leq C (\|q_i\|_{L^2} + \|q_o\|_{L^2}) . \quad (2.5)$$

Proof. See [22]. □

Definition 2.2 (Solution operator for the RTE). *The solution operator for the RTE is defined by*

$$\mathbf{T}: \mathbb{Q} \times \mathbb{D}(\mathbf{T}) \rightarrow \mathbb{W}: (q, \mu) \mapsto \mathbf{T}(q, \mu) := \Phi ,$$

where Φ denotes the unique solution of (2.3).

Theorem 2.1 guarantees that the operator \mathbf{T} , mapping $(q, \mu) \in \mathbb{Q} \times \mathbb{D}(\mathbf{T})$ to the solution of the RTE, is well defined. Note that in the actual application $q = (q_i, q_o) \in \mathbb{Q}$ are prescribed sources, and $\mu = (\mu_a, \mu_s) \in \mathbb{D}(\mathbf{T})$ are the unknown parameters to be recovered.

2.3 Heating Operator

Due to the spatially varying absorption of photons, the tissue is locally heated and emits an acoustic pressure wave. The acoustic source is proportional to the amount of absorbed photons, the light intensity and the so-called Grüneisen parameter γ describing the efficiency of conversion of optical to acoustical energy. We assume γ to be constant and after appropriate re-scaling we take $\gamma = 1$; for more details about the Grüneisen parameter, we refer to [18]. Therefore, the conversion of the optical energy into acoustic pressure wave is described by the heating operator defined as follows.

Definition 2.3 (Heating operator). *The heating operator is defined by*

$$\mathbf{H}: \mathbb{Q} \times \mathbb{D}(\mathbf{T}) \rightarrow L^2(\Omega) \quad (2.6)$$

$$(q, \mu) \mapsto \mu_a \int_{\mathbb{S}^{d-1}} \mathbf{T}(q, \mu)(\cdot, \theta) d\theta .$$

If one introduces the averaging operator $\mathbf{A}: \mathbb{W} \mapsto L^2(\Omega)$ defined by $\mathbf{A}\Phi = \int_{\mathbb{S}^{d-1}} \Phi(\cdot, \theta) d\theta$ one may write the heating operator in the form

$$\mathbf{H}(q, \mu) = \mu_a \mathbf{A} \circ \mathbf{T}(q, \mu) \quad \text{for } (q, \mu) \in \mathbb{Q} \times \mathbb{D}(\mathbf{T}).$$

Because $\mathbf{T}(q, \mu)$ models the photon density, $\mathbf{A} \circ \mathbf{T}(q, \mu)$ actually models the total light intensity. The heating operator is therefore given by the product of the absorption coefficient and the light intensity. The averaging operator \mathbf{A} is well defined and bounded and therefore the heating operator is well defined as a mapping between $\mathbb{Q} \times \mathbb{D}(\mathbf{T})$ and $L^2(\Omega)$.

2.4 The Wave Equation

The local heating causes an acoustic pressure wave, where the initial pressure distribution p_0 is proportional to a fraction of the absorbed energy. Assuming constant speed of sound and after rescaling, the induced acoustic pressure $p: \mathbb{R}^d \times (0, \infty) \rightarrow \mathbb{R}$ satisfies the free-space wave equation:

$$\begin{cases} (\partial_t^2 - \Delta)p(x, t) = 0 & \text{for } (x, t) \in \mathbb{R}^d \times (0, \infty), \\ p(x, 0) = p_0(x) & \text{for } x \in \mathbb{R}^d, \\ \partial_t p(x, 0) = 0 & \text{for } x \in \mathbb{R}^d. \end{cases} \quad (2.7)$$

Here, the function p_0 vanishes outside B_R , the ball of radius R , and acoustic data are collected on a subset of $\partial B_R \times (0, \infty)$ that we denote by $\Lambda \times (0, \infty)$. Recall that coupling of the RTE and the wave equation happens in such a way that the result of the heating operator $\mathbf{H}(q, \mu)$ acts as initial sound source p_0 depending on tissue parameters; see Definition 2.6. Standard existence and uniqueness theory for hyperbolic equations guarantees that, for any $p_0 \in H^1$, (2.7) has a unique solution $p \in H^1$, which continuously depends on p_0 . Taking the trace results in loss of regularity by degree 1/2. Therefore, $p_0 \mapsto p|_{\partial B_R \times (0, \infty)}$ is continuous between H^1 and $H^{1/2}$. The following Lemma implies the much stronger result that $p_0 \mapsto p|_{\partial B_R \times (0, \infty)}$ is actually an L^2 -isometry.

Lemma 2.4. *Let $p_0 \in C^\infty(\mathbb{R}^d)$ have support in B_R and let p denote the solution of (2.7). Then,*

$$\int_{B_R} \int_0^\infty |p(x, t)|^2 dt dx = \frac{R}{2} \int_{B_R} |p_0(x)|^2 dx. \quad (2.8)$$

Proof. See [27] for d odd and [26] for d even. □

Definition 2.5. *We define the solution operator with full boundary data for the wave Equation (2.7) by*

$$\mathbf{U}: C^\infty(B_R) \subseteq L^2(B_R) \rightarrow \mathbb{Y}: p_0 \mapsto p|_{\partial B_R \times (0, \infty)}, \quad (2.9)$$

where p denotes the solution of (2.7).

According to Lemma 2.4, the operator \mathbf{U} can be uniquely extended to a bounded linear operator defined on $L^2(B)$, denoted again by $\mathbf{U}: L^2(B_R) \rightarrow \mathbb{Y}$. The partial acoustic measurements made on $\Lambda \subseteq \partial B_R$ are then modeled by $\chi_{\Lambda \times (0, \infty)} \mathbf{U} p_0$.

2.5 Analysis of the Forward Problem in Multi-Source QPAT

We assume that we perform N individual experiments, where each experiment consists of separate optical sources and separate acoustic measurements. For the i -th experiment, we denote the

source term by $q_i \in \mathbb{Q}$ and assume the acoustic measurements are made on $\Lambda_i \times (0, T_i) \subseteq \partial B_R \times (0, \infty)$.

Definition 2.6. For any $i \in \{1, \dots, N\}$, we denote

$$\begin{aligned} \mathbf{T}_i &: \mathbb{D}(\mathbf{T}) \rightarrow \mathbb{W}: \mu \mapsto \mathbf{T}(q_i, \mu), \\ \mathbf{H}_i &: \mathbb{D}(\mathbf{T}) \rightarrow L^2(\Omega): \mu \mapsto \mu_a(\mathbf{A} \circ \mathbf{T}_i)(\mu), \\ \mathbf{U}_i &: L^2(\Omega) \rightarrow \mathbb{Y}: p_0 \mapsto \chi_{\Lambda_i \times (0, T_i)} \mathbf{U} p_0, \\ \mathbf{F}_i &: \mathbb{D}(\mathbf{T}) \rightarrow \mathbb{Y}: \mu \mapsto (\mathbf{U}_i \circ \mathbf{H}_i)(\mu). \end{aligned}$$

Here, \mathbf{T}_i denotes the i -th solution operator for the RTE, \mathbf{H}_i the i -th heating operator, \mathbf{U}_i the i -th partial solution operator for the wave equation, and \mathbf{F}_i the i -th forward operator.

Recall that \mathbf{T} stands for the solution operator for the RTE (2.3) given in Definition 2.2, $\mathbf{A}\Phi = \int_{\mathbb{S}^{d-1}} \Phi(\cdot, \theta) d\theta$ is the averaging operator, and \mathbf{U} the solution operator for the wave Equation (2.7); see Definition 2.5. The operator \mathbf{T}_i models the photon transport and its solution (via the heating operator) acts as input for the solution of the wave equation and thereby couples the optical with the acoustical part.

Next, we recall continuity and differentiability of the forward operators. For that purpose, we call $h \in \mathbb{X}$ a feasible direction at $\mu \in \mathbb{D}(\mathbf{T})$, if there exists some $\epsilon > 0$ with $\mu + \epsilon h \in \mathbb{D}(\mathbf{T})$.

Theorem 2.7 (Continuity and Differentiability).

- (a) The operators \mathbf{T}_i , \mathbf{F}_i and \mathbf{H}_i are sequentially continuous and Lipschitz-continuous.
(b) For every $\mu \in \mathbb{D}(\mathbf{T})$, the one-sided directional derivatives $\mathbf{T}'_i(\mu)(h)$, $\mathbf{F}'_i(\mu)(h)$ of \mathbf{T}_i , \mathbf{F}_i at μ in any feasible direction h exist, and are given by

$$\mathbf{T}'_i(\mu)(h) = \mathbf{T}(0, -(h_a + h_s - h_s \mathbf{K})\mathbf{T}(\mu), \mu), \quad (2.10)$$

$$\mathbf{F}'_i(\mu)(h) = \mathbf{U}_{\Lambda_i, T_i}(h_a \mathbf{A} \mathbf{T}_i(\mu) + \mu_a \mathbf{A}(\mathbf{T}'_i(\mu)(h))). \quad (2.11)$$

Proof. See [33]. □

Equations (2.10) and (2.11) define a bounded linear operator $\mathbf{F}'_i(\mu): \mathbb{X} \rightarrow \mathbb{Y}$, which we call the derivative of \mathbf{F}_i at $\mu \in \mathbb{D}(\mathbf{T})$. Numerical minimization schemes actually require the adjoint of $\mathbf{F}'_i(\mu)$, which we compute next.

Theorem 2.8 (Adjoint of $\mathbf{F}'_i(\mu)$). Let $i \in \{1, \dots, N\}$ and $\mu \in \mathbb{D}(\mathbf{T})$. Furthermore, set $\Phi_i := \mathbf{T}_i(\mu)$ and let Φ_i^* denote the solution of the adjoint problem

$$(-\theta \cdot \nabla_x + (\mu_a + \mu_s - \mu_s \mathbf{K})) \Phi_i^* = -\mathbf{A}^*(\mu_a(\mathbf{U}_i^* v)) \quad (2.12)$$

with $\Phi_i^*|_{\Gamma_+} = 0$. Then, $\mathbf{F}'_i(\mu)^*: \mathbb{Y} \rightarrow \mathbb{X}$ is given by

$$\mathbf{F}'_i(\mu)^* v = \begin{bmatrix} \mathbf{A}(\Phi_i^* \Phi_i) + (\mathbf{A} \Phi_i)(\mathbf{U}_i^* v) \\ \mathbf{A}([(I - \mathbf{K})(\Phi_i^*)] \Phi_i) \end{bmatrix}. \quad (2.13)$$

Proof. See [33]. □

Given data $v_1, \dots, v_N \in \mathbb{Y}$, most numerical schemes for QPAT use gradients of the partial data-fidelity terms $F_l: \mathbb{D}(\mathbf{T}) \rightarrow \mathbb{R}$ for $l \subseteq \{1, \dots, N\}$, where

$$F_l(\mu) = \sum_{i \in l} F_i(\mu) \quad \text{with} \quad F_i(\mu) := \frac{1}{2} \|\mathbf{F}_i(\mu) - v_i\|_{\mathbb{Y}}^2. \quad (2.14)$$

By the chain rule, the gradient of F_l is given by $\nabla F_l(\mu) = \sum_{i \in l} \nabla F_i(\mu)$ with $\nabla F_i(\mu) = \mathbf{F}'_i(\mu)^* (\mathbf{F}_i(\mu) - v_i)$, where $\mathbf{F}'_i(\mu)^*$ can be computed by Theorem 2.8. Convergence of schemes such as the (stochastic) proximal gradient method considered in the following section require the Lipschitz continuity of ∇F_l , which will be shown in the following theorem.

Theorem 2.9 (Lipschitz continuity of ∇F_l). *For any data $v_1, \dots, v_N \in \mathbb{Y}$ and any subset $l \subseteq \{1, \dots, N\}$, the map $\mu \mapsto \nabla F_l(\mu)$ is Lipschitz-continuous.*

Proof. Without loss of generality, we assume $N = 1$, $l = \{1\}$ and write $v = v_1$, $F = F_{\{1\}}$, $\mathbf{T} = \mathbf{T}_1$, $\mathbf{U} = \mathbf{U}_1$, and $v(\mu) = \mathbf{F}(\mu) - v$. For any $\mu \in \mathbb{D}(\mathbf{T})$, let $\mathbf{T}^*(\mu)$ denote the solution of (2.12) with $v(\mu)$ in place of v . Then, for any $\mu, \tilde{\mu} \in \mathbb{D}(\mathbf{T})$,

$$\begin{aligned} & \|\nabla F(\mu) - \nabla F(\tilde{\mu})\|_{\mathbb{X}}^2 \\ &= \|\mathbf{A}(\mathbf{T}^*(\mu)\mathbf{T}(\mu)) + (\mathbf{A}\mathbf{T}(\mu))(\mathbf{U}^*v(\mu)) - \mathbf{A}(\mathbf{T}^*(\tilde{\mu})\mathbf{T}(\tilde{\mu})) - (\mathbf{A}\mathbf{T}(\tilde{\mu}))(\mathbf{U}^*v(\tilde{\mu}))\|_{L^2(\Omega)}^2 \\ & \quad + \|\mathbf{A}((\mathbf{I} - \mathbf{K})(\mathbf{T}^*(\mu))\mathbf{T}(\mu)) - \mathbf{A}((\mathbf{I} - \mathbf{K})(\mathbf{T}^*(\tilde{\mu}))\mathbf{T}(\tilde{\mu}))\|_{L^2(\Omega)}^2. \end{aligned} \quad (2.15)$$

The second term in (2.15) can be bounded by

$$\begin{aligned} & 2 \|\mathbf{A}((\mathbf{I} - \mathbf{K})(\mathbf{T}^*(\mu))\mathbf{T}(\mu)) - \mathbf{A}((\mathbf{I} - \mathbf{K})(\mathbf{T}^*(\mu))\mathbf{T}(\tilde{\mu}))\|_{L^2(\Omega)}^2 \\ & \quad + 2 \|\mathbf{A}((\mathbf{I} - \mathbf{K})(\mathbf{T}^*(\mu))\mathbf{T}(\tilde{\mu})) + \mathbf{A}((\mathbf{I} - \mathbf{K})(\mathbf{T}^*(\tilde{\mu}))\mathbf{T}(\tilde{\mu}))\|_{L^2(\Omega)}^2 \\ & \leq 4(\|k\|_{\infty} + 1)^2 |\mathbb{S}^{d-1}| \|\mathbf{T}^*(\mu)\|_{L^{\infty}(\Omega \times \mathbb{S}^{d-1})}^2 \|\mathbf{T}(\mu) - \mathbf{T}(\tilde{\mu})\|_{L^2(\Omega)}^2. \end{aligned}$$

The squared norm of the difference $\mathbf{A}(\mathbf{T}^*(\mu)\mathbf{T}(\mu)) - \mathbf{A}(\mathbf{T}^*(\tilde{\mu})\mathbf{T}(\tilde{\mu}))$ in the first term in (2.15) is estimated in a similar manner. Furthermore, we have

$$\begin{aligned} & \|(\mathbf{A}\mathbf{T}(\mu))(\mathbf{U}^*v(\mu)) - (\mathbf{A}\mathbf{T}(\tilde{\mu}))(\mathbf{U}^*v(\mu)) + (\mathbf{A}\mathbf{T}(\tilde{\mu}))(\mathbf{U}^*v(\mu)) - (\mathbf{A}\mathbf{T}(\tilde{\mu}))(\mathbf{U}^*v(\tilde{\mu}))\|_{L^2(\Omega)}^2 \\ & \leq 2 \|(\mathbf{U}^*v(\mu))\mathbf{A}[\mathbf{T}(\mu) - \mathbf{T}(\tilde{\mu})]\|_{L^2(\Omega)}^2 + 2 \|(\mathbf{A}\mathbf{T}(\tilde{\mu}))(\mathbf{U}^*(v(\mu) - v(\tilde{\mu})))\|_{L^2(\Omega)}^2. \end{aligned}$$

Noting that \mathbf{A} , \mathbf{U} and \mathbf{U}^* are linear and bounded, Theorem 2.7, an analog of Theorem 2.1 for the supremums norm, and the computations above yield the Lipschitz continuity of ∇F . \square

3 The Stochastic Proximal Gradient Method for QPAT

3.1 Formulation of the Inverse Problem

The inverse problem of multi-source QPAT consists in finding $\mu^* \in \mathbb{X}$ from measured data

$$v_i = \mathbf{F}_i(\mu^*) + z_i \quad \text{for } i = 1, \dots, N. \quad (3.1)$$

Here, $\mu^* = (\mu_a^*, \mu_s^*)$ are the unknowns to be estimated, z_i are the unknown error vectors, and v_1, \dots, v_N are the given noisy data. Using the notation

$$v := (v_1, \dots, v_N) \in \mathbb{Y}^N,$$

$$\mathbf{F} := (\mathbf{F}_1, \dots, \mathbf{F}_N): \mathbb{D}(\mathbf{T}) \rightarrow \mathbb{Y}^N,$$

we can write (3.1) in the alternative form

$$\text{Estimate } \mu^* \in \mathbb{X} \text{ from } v = \mathbf{F}(\mu^*) + z. \quad (3.2)$$

Here, $z \in \mathbb{Y}^N$ denotes the error vector.

There are, at least, two different strategies to address such an inverse problem: Tikhonov type regularization on the one and iterative methods on the other hand. In this section, we give an overview of such methods. In particular, we describe proximal stochastic gradient methods (for minimizing the Tikhonov functional), which seem particularly well suited for multi-source QPAT but have not been investigated yet for that purpose.

3.2 Tikhonov Regularization in QPAT

In this section, we consider a quadratic Tikhonov regularization term for solving (3.1). Let

$$\mathbf{L}: \mathbb{D}(\mathbf{L}) \subseteq \mathbb{X} \rightarrow \mathbb{Z}: \mu \mapsto \mathbf{L}\mu$$

be a linear, densely defined, and possibly unbounded operator between \mathbb{X} and another Hilbert space $(\mathbb{Z}, \langle \cdot, \cdot \rangle_{\mathbb{Z}})$ and set $\mathbb{D} := \mathbb{D} \cap \mathbb{D}(\mathbf{L})$. In this context, any element $\mu^+ \in \mathbb{D}$ with $\|\mathbf{L}\mu^+\| = \min\{\|\mathbf{L}\mu\| \mid \mathbf{F}(\mu) = v\}$ is called an $\|\mathbf{L}(\cdot)\|$ -minimizing solution of $\mathbf{F}\mu = v$. Tikhonov regularization with regularization term $\frac{1}{2} \|\mathbf{L}\mu\|_{\mathbb{Z}}^2$ consists in computing a minimizer of the generalized Tikhonov functional $T_{v,\lambda}: \mathbb{X} \rightarrow \mathbb{R} \cup \{\infty\}$, defined by

$$T_{v,\lambda}(\mu) := \begin{cases} \frac{1}{2} \|\mathbf{F}(\mu) - v\|^2 + \frac{\lambda}{2} \|\mathbf{L}\mu\|^2, & \text{if } \mu \in \mathbb{D}, \\ \infty, & \text{otherwise.} \end{cases} \quad (3.3)$$

Here, $\lambda > 0$ denotes the regularization parameter that acts as a trade-off between the data fitting term and stability.

Theorem 3.1 (Well-posedness and convergence).

- (a) For any $v \in \mathbb{Y}$ and any $\lambda > 0$, the Tikhonov functional $T_{\lambda,v}$ has at least one minimizer.
- (b) Let $v \in \text{ran}(\mathbf{F})$, $(\delta_m)_{m \in \mathbb{N}} \in (0, \infty)^{\mathbb{N}}$, $(v^m)_{m \in \mathbb{N}} \in \mathbb{Y}^{\mathbb{N}}$ with $\|v - v^m\| \leq \delta_m$. Suppose further that $(\lambda_m)_{m \in \mathbb{N}} \in (0, \infty)^{\mathbb{N}}$ satisfies $\lambda_m \rightarrow 0$ and $\delta_m^2/\lambda_m \rightarrow 0$ as $m \rightarrow \infty$. Then:
 - Every sequence $(\mu^m)_{m \in \mathbb{N}}$ with $\mu^m \in \arg \min T_{v^m, \lambda_m}$ has a weakly converging subsequence.
 - The limit of every weakly convergent subsequence of $(\mu^m)_{m \in \mathbb{N}}$ is an $\|\mathbf{L}(\cdot)\|$ -minimizing solution of $\mathbf{F}\mu = v$.
 - If the $\|\mathbf{L}(\cdot)\|$ -minimizing solution of $\mathbf{F}\mu = v$ is unique and denoted by μ^+ , then $(\mu^m) \rightharpoonup \mu^+$.

Proof. See [33]. □

3.3 The Proximal Stochastic Gradient Algorithm for QPAT

Depending on the particular choice of \mathbf{L} , the Tikhonov functional (3.3) may be ill-conditioned. To address this issue in [33], we proposed the proximal gradient algorithm for minimizing (3.3), which is a very flexible algorithm for minimizing functionals of the form $F + G$, where F is smooth and G is convex (see, for example, [15, 8]). Here, we extend the approach to the proximal stochastic gradient algorithm. Additionally, we propose computing the proximal step using Dykstra's projection algorithm.

■ **Proximal gradient algorithm:** The proximal gradient algorithm is a splitting method that iteratively computes explicit gradient steps for F and implicit proximal steps for G . In our context, we take F as the data fidelity term and

$$G(\mu) = G_\lambda(\mu) := g_\lambda(\mu) + \chi_{\mathbb{D}}(\mu) := \frac{\lambda}{2} \|\mathbf{L}\mu\|^2 + \chi_{\mathbb{D}}(\mu), \quad (3.4)$$

where $\chi_{\mathbb{D}}$ is the characteristic function taking the value zero inside \mathbb{D} and ∞ outside. The proximal gradient algorithm for minimizing the QPAT-Tikhonov functional (3.3) reads

$$\mu^{k+1} = \text{prox}_{s_k G_\lambda} \left(\mu^k - s_k \sum_{i=1}^N \nabla F_i(\mu^k) \right). \quad (3.5)$$

Here, $\text{prox}_{s_k G_\lambda} : \mathbb{X} \rightarrow \mathbb{D}$ denotes the proximal mapping corresponding to the functional $s_k G_\lambda$,

$$\text{prox}_{s_k G_\lambda}(x) = \arg \min \left\{ \frac{1}{2} \|x - (\cdot)\|^2 + s_k G_\lambda \right\}. \quad (3.6)$$

Furthermore, $\nabla F_i(\mu^k)$ is the gradient of the i -th data fidelity term computed in Theorem 2.8.

■ **Dykstra's projection algorithm:** The constraint quadratic optimization problem (3.6) can efficiently be solved by a proximal variant of Dykstra's projection algorithm [14, 15, 8]. For that purpose, we write $s_k G_\lambda = \chi_{\mathbb{D}} + g$ with $g(x) := \frac{s_k \lambda}{2} \|\mathbf{L}x\|^2$. Setting $x_0 = \mu$, $p_0 = 0$ and $q_0 = 0$, Dykstra's projection algorithm for (3.6) reads, for $m \in \mathbb{N}$,

$$y_m = \text{prox}_g(x_m + p_m), \quad (3.7)$$

$$x_{m+1} = \mathbf{P}_{\mathbb{D}}(y_m + q_m), \quad (3.8)$$

$$p_{m+1} = x_m + p_m - y_m, \quad (3.9)$$

$$q_{m+1} = y_m + q_m - x_{m+1}. \quad (3.10)$$

Both proximal mapping in (3.7) and the projection in (3.8) can be computed explicitly. In fact, one readily verifies that

$$\text{prox}_g(x) = (\mathbf{I}_{\mathbb{X}} + s_k \lambda \mathbf{L}^* \mathbf{L})^{-1} x, \quad (3.11)$$

$$\mathbf{P}_{\mathbb{D}}(\mu) = \min \{ \bar{\mu}, \max \{ 0, \mu \} \}. \quad (3.12)$$

Here, $\mathbf{I}_{\mathbb{X}}$ is the identity operator on \mathbb{X} and $\mathbf{P}_{\mathbb{D}}$ the projection onto \mathbb{D} .

■ **Proximal stochastic gradient algorithm:** The methods described so far require in any iterative step the computation of the full gradient

$$\nabla F(\mu) = \sum_{i=1}^N \nabla F_i(\mu) \quad \text{with} \quad \nabla F_i(\mu) = \mathbf{F}'_i(\mu)^* (\mathbf{F}_i(\mu) - v_i).$$

The evaluation of each $\nabla F_i(\mu)$ requires the solution of the RTE and an adjoint problem and therefore is quite time-consuming. For multi-source QPAT, where $N > 1$, a significant acceleration may be obtained by a Kaczmarz strategy, where in each iterative step only one of the summands $\nabla F_i(\mu)$ is used. The resulting proximal stochastic gradient method for minimizing the Tikhonov functional (3.3) in QPAT reads

$$\mu^{k+1} = \text{prox}_{s_k G_\lambda} \left(\mu^k - s_k \nabla F_{i(k)}(\mu^k) \right), \quad (3.13)$$

where $i(k) \in \{1, \dots, N\}$ is selected randomly for the update in the k -th iteration. Furthermore, $\text{prox}_{s_k G_\lambda}$ is the proximal mapping of $s_k G_\lambda$ that can be computed by Dykstra algorithm (3.7)–(3.10) and $\nabla F_i(\mu^k)$ is the gradient of the i -th data fidelity term that can be computed by Theorem 2.8.

One can also incorporate a block-iterative (or mini-batch) strategy in the stochastic gradient method, meaning that a small subset of $\{1, \dots, N\}$ of equations is used per iteration instead of a single one. Such a variant could be especially useful in the case of a large number of different illumination patterns. For more details about stochastic gradient methods, see [10, 11, 55, 21, 42, 46] and the references therein. Note that, in general, convergence of stochastic gradient methods requires asymptotically vanishing step size [46].

3.4 Iterative Regularization Methods

An alternative class of algorithms to address nonlinear inverse problems are iterative techniques. The most basic iterative method for solving the nonlinear inverse problem $v = \mathbf{F}(\mu)$ is the Landweber iteration. In the case that the domain of definition \mathbb{D} is a proper subset, we have to combine the Landweber iteration with a projection step onto \mathbb{D} as presented in this subsection. The projected Landweber iteration applied to multi-source QPAT reads

$$\mu^{k+1} = \mathbf{P}_{\mathbb{D}}\left(\mu^k - s_k \sum_{i=1}^N \nabla F_i(\mu^k)\right), \quad (3.14)$$

where ∇F_i is the gradient of F_i (see Equation (2.14)), and $\mathbf{P}_{\mathbb{D}}(\mu) = \min\{\bar{\mu}, \max\{0, \mu\}\}$ denotes the projection onto \mathbb{D} . In Tikhonov regularization, the regularity of solutions is enforced by an explicitly included penalty. In opposition to that, in iterative regularization methods, a stabilization effect is enforced by early stopping of the iteration. A common stopping rule is the discrepancy principle, where iteration is stopped at the smallest index $k \in \mathbb{N}$ satisfying $\|v - \mathbf{F}(\mu^k)\| \leq \tau\delta$, where δ is an estimate for the noise and $\tau \geq 1$. Formally, the projected Landweber iteration (3.14) arises as a special case of the proximal gradient iteration (3.5) for minimizing the Tikhonov functional, where the regularization parameter is taken as $\lambda = 0$ and where the proximal mapping (3.6) reduces to the orthogonal projection onto \mathbb{D} .

In a similar manner, one can also use a stochastic version of the projected Landweber iteration. Using the loping strategy of [20, 31, 30] in order to stabilize the iterative process, the resulting projected loping Landweber–Kaczmarz iteration reads

$$\mu^{k+1} = \mathbf{P}_{\mathbb{D}}\left(\mu^k - s_k \omega_k \nabla F_{i(k)}(\mu^k)\right), \quad (3.15)$$

$$\omega_k := \begin{cases} 1, & \|\mathbf{F}_{i(k)}(\mu^k) - v_{i(k)}\|_{\mathbb{X}} > \tau\delta_{i(k)}, \\ 0, & \text{otherwise.} \end{cases} \quad (3.16)$$

Here, $i(k) \in \{1, \dots, N\}$ for any $k \in \mathbb{N}$ may be randomly selected, $\tau > 1$ is an appropriately chosen positive constant, $\nabla F_i(\cdot)$ is the gradient of the i -th data fidelity term computed in Theorem 2.8 and $\mathbf{P}_{\mathbb{D}}(\mu) = \min\{\bar{\mu}, \max\{0, \mu\}\}$ denotes the projection onto \mathbb{D} . The iteration (3.15), (3.16) terminates if $\|\mathbf{F}_i(\mu^k) - v_i\| \leq \tau\delta$ for all $i \in \{1, \dots, N\}$. It is worth mentioning that, for noise free data, we have $\omega_k = 1$ for all k and, therefore, in this special situation, the iteration becomes $\mu^{k+1} = \mathbf{P}_{\mathbb{D}}(\mu^k - s_k \nabla F_{i(k)}(\mu^k))$, which formally arises from the proximal stochastic gradient method (3.13) with $\lambda = 0$. A convergence analysis of the loping Landweber–Kaczmarz method can be found in [20, 31].

4 QPAT as Multilinear Inverse Problem

Since the RTE is time-consuming to solve, we are looking for a suitable reformulation of the inverse problem in multi-source QPAT avoiding computation of a solution of the RTE in each iterative step. In this paper, we propose to write (3.1) as a multilinear inverse problem, where we add the RTE as a constraint instead of explicitly including its solution. The new formulation will again be addressed by Tikhonov regularization and proximal stochastic gradient methods.

4.1 Reformulation as Multilinear Inverse Problem

Recall the forward problem of QPAT governed by the RTE (2.3). With the abbreviation $\mathbf{M}(\boldsymbol{\mu}) := \theta \cdot \nabla_x + \mu_a + \mu_s(\mathbf{I} - \mathbf{K})$, the RTE can be written in compact form $\mathbf{M}(\boldsymbol{\mu})\Phi = q$, where $\boldsymbol{\mu} = (\mu_a, \mu_s)$ is the unknown parameter pair. In the case of exact data, the multi-source problem in QPAT (3.1) then can be reformulated as the problem of finding the tuple $\mathbf{z} := (\boldsymbol{\mu}, (\Phi_i, H_i)_{i=1}^N) \in \mathbb{D} \times (\mathbb{W} \times L^2(\Omega))^N$ such that

$$\begin{aligned} \mathbf{M}(\boldsymbol{\mu})\Phi_i &= q_i & \text{for } i = 1, \dots, N, \\ H_i &= \mu_a \mathbf{A}\Phi_i & \text{for } i = 1, \dots, N, \\ v_i &= \mathbf{U}(H_i) & \text{for } i = 1, \dots, N. \end{aligned} \quad (4.1)$$

Here, the index i indicates the i -th illumination, and $q_i \in \mathbb{Q}$, $\Phi_i \in \mathbb{W}$, $H_i \in L^2(\Omega)$, $v_i \in \mathbb{Y}$ are the corresponding source, photon density, heating and acoustical data, respectively, and $\mathbf{A}\Phi_i = \int_{\mathbb{S}^{d-1}} \Phi_i(\cdot, \theta) d\theta$ is the averaging operator. We call (4.1) and resulting formulations below the multilinear (MULL) formulation of QPAT.

4.2 Application of Tikhonov Regularization

In the case that the data v_i are only known approximately, we use Tikhonov regularization for the stable solution of (4.1). For that purpose, we approximate (4.1) by the constrained optimization problem

$$\begin{aligned} \min_{(\boldsymbol{\mu}, \Phi_i, H_i)_{i=1}^N} & \left(\frac{1}{2} \sum_{i=1}^N \|v_i - \mathbf{U}(H_i)\|^2 + \frac{\lambda}{2} \|\mathbf{L}(\boldsymbol{\mu})\|^2 + \boldsymbol{\chi}_{\mathbb{D}}(\boldsymbol{\mu}) \right), \\ \text{s.t.} & \begin{cases} \mathbf{M}(\boldsymbol{\mu})\Phi_i = q_i, \\ H_i = \mu_a \mathbf{A}\Phi_i \text{ for } i = 1, \dots, N. \end{cases} \end{aligned} \quad (4.2)$$

Here, the operator $\mathbf{L}\boldsymbol{\mu} = (\mathbf{L}_a\mu_a, \mathbf{L}_s\mu_s)$ is possibly unbounded, $\frac{\lambda}{2}\|\mathbf{L}(\boldsymbol{\mu})\|^2$ is the regularization term and $\lambda > 0$ the regularization parameter. Note that (4.2) is equivalent to (3.3) and therefore the well-posedness and convergence results of Theorem 3.1 apply to (4.2) as well.

The constrained optimization problem (4.2) proposed in this paper can be addressed by various solution methods, for example using penalty methods or augmented Lagrangian techniques [37]. In this paper, we use a penalty approach for solving (4.2) where the constraints are included as penalty term. To simplify notation, we introduce the unconstrained functionals

$$\begin{aligned} \mathbf{J}^{(i)}(\mathbf{z}) &:= \frac{a_1}{2} \|\mathbf{M}(\boldsymbol{\mu})\Phi_i - q_i\|^2 + \frac{a_2}{2} \|\mu_a \mathbf{A}\Phi_i - H_i\|^2 \\ &\quad + \frac{a_3}{2} \|v_i - \mathbf{U}(H_i)\|^2 + \frac{\lambda}{2} \|\mathbf{L}(\boldsymbol{\mu})\|^2 + \boldsymbol{\chi}_{\mathbb{D}}(\boldsymbol{\mu}), \end{aligned} \quad (4.3)$$

for certain parameters $a_1, a_2, a_3 > 0$ and $\mathbf{z}_i := (\mu, \Phi_i, H_i) \in \mathbb{Q} \times \mathbb{W} \times L^2(\Omega)$. The sum of the unconstrained functionals (4.3) over all illuminations will actually be minimized in our numerical implementations. For that purpose, we define

$$\begin{aligned}\mathbf{J}_1^{(i)}(\mathbf{z}) &= \frac{1}{2} \|\mathbf{M}(\mu)\Phi_i - q_i\|^2, \\ \mathbf{J}_2^{(i)}(\mathbf{z}) &= \frac{1}{2} \|\mu_a \mathbf{A}\Phi_i - H_i\|^2, \\ \mathbf{J}_3^{(i)}(\mathbf{z}) &= \frac{1}{2} \|v_i - \mathbf{U}(H_i)\|^2, \\ \mathbf{J}_4^{(i)}(\mathbf{z}) &= \frac{1}{2} \|\mathbf{L}(\mu)\|^2.\end{aligned}$$

Then, we have $\mathbf{J}^{(i)}(\mathbf{z}) = \sum_{\ell=1}^4 a_\ell \mathbf{J}_\ell^{(i)}(\mathbf{z}) + \chi_{\mathbb{D}}(\mu)$. For the approximate solution of (4.2), we minimize the unconstrained functional $\mathbf{J}(\mathbf{z}) = \sum_{i=1}^4 \mathbf{J}^{(i)}(\mathbf{z})$, which can be written in the forms

$$\mathbf{J}(\mathbf{z}) = \sum_{i=1}^N \sum_{\ell=1}^4 a_\ell \mathbf{J}_\ell^{(i)}(\mathbf{z}) + \chi_{\mathbb{D}}(\mu), \quad (4.4)$$

$$\mathbf{J}(\mathbf{z}) = \sum_{i=1}^N \sum_{\ell=1}^3 a_\ell \mathbf{J}_\ell^{(i)}(\mathbf{z}) + \frac{\lambda}{2} \|\mathbf{L}(\mu)\|^2 + \chi_{\mathbb{D}}(\mu). \quad (4.5)$$

(Here and below, we also write $a_4 = \lambda$, if it simplifies notation.) The formulations (4.4) as well as (4.5) can be solved by various optimization techniques. In particular, as the functionals are given as the sum of simpler terms, the stochastic (proximal) gradient method is particularly appealing.

4.3 Solution of the MULL Formulation of QPAT Using Stochastic Gradient Methods

For solving QPAT in the novel MULL formulation (4.1), we use stochastic gradient methods similar to previous sections. For that purpose, we require the gradients (determining the steepest descent directions) of the individual functionals $\mathbf{J}_\ell^{(i)}(\mathbf{z}_i)$ with respect to $\mathbf{z}_i = (\mu, \Phi_i, H_i)$, which are given as

$$\nabla_{\mu_a} \mathbf{J}_1^{(i)}(\mathbf{z}) = \Phi_i (\mathbf{M}(\mu)\Phi_i - q_i), \quad (4.6)$$

$$\nabla_{\mu_s} \mathbf{J}_1^{(i)}(\mathbf{z}) = (\mathbf{I} - \mathbf{K})\Phi_i (\mathbf{M}(\mu)\Phi_i - q_i), \quad (4.7)$$

$$\nabla_{\Phi_i} \mathbf{J}_1^{(i)}(\mathbf{z}) = \mathbf{M}(\mu) (\mathbf{M}(\mu)\Phi_i - q_i), \quad (4.8)$$

$$\nabla_{\mu_a} \mathbf{J}_2^{(i)}(\mathbf{z}) = (\mathbf{A}\Phi_i) (\mu_a \mathbf{A}\Phi_i - H_i), \quad (4.9)$$

$$\nabla_{H_i} \mathbf{J}_2^{(i)}(\mathbf{z}) = -(\mu_a \mathbf{A}\Phi_i - H_i), \quad (4.10)$$

$$\nabla_{\Phi_i} \mathbf{J}_2^{(i)}(\mathbf{z}) = \mathbf{A}^* [\mu_a (\mu_a \mathbf{A}\Phi_i - H_i)], \quad (4.11)$$

$$\nabla_{H_i} \mathbf{J}_3^{(i)}(\mathbf{z}) = -\mathbf{U}^T (v_i - \mathbf{U}(H_i)), \quad (4.12)$$

$$\nabla_{\mu_a} \mathbf{J}_4^{(i)}(\mathbf{z}) = \mathbf{L}_a^* \mathbf{L}_a \mu_a, \quad (4.13)$$

$$\nabla_{\mu_s} \mathbf{J}_4^{(i)}(\mathbf{z}) = \mathbf{L}_s^* \mathbf{L}_s \mu_s. \quad (4.14)$$

(All other partial gradients are vanishing.) In the following, let N be the number of illuminations, write $\mathbf{z} = (\mu_a, \mu_s, (\Phi_i, H_i)_{i=1}^N)$ and let $(s_k)_{k \in \mathbb{N}}$ be a sequence of step sizes. In this paper, we propose the following instances of the stochastic proximal gradient method for QPAT based on the multilinear formulation (4.1).

■ **MULL-projected stochastic gradient algorithm:** Here, we consider the form (4.4). For any iteration index $k \in \mathbb{N}$ choose $i(k) \in \{1, \dots, N\}$ and $\ell(k) \in \{1, \dots, 4\}$ and define the sequence of iterates $(\mathbf{z}^k)_{k \in \mathbb{N}}$ by

$$\mathbf{z}^{k+1} = (\mathbf{P}_{\mathbb{D}} \times \mathbf{I}) \left(\mathbf{z}^k - s_k \nabla \mathbf{J}_{\ell(k)}^{(i(k))}(\mathbf{z}^k) \right). \quad (4.15)$$

Here, the mapping $\mathbf{P}_{\mathbb{D}} \times \mathbf{I}$ is the proximal mapping corresponding to $\mathbf{z} \mapsto \boldsymbol{\chi}_{\mathbb{D}}(\mu)$, which equals the projection $\mathbf{P}_{\mathbb{D}}$ in the μ component and equals the identity \mathbf{I} in the other components.

■ **MULL-proximal stochastic gradient algorithm:** Here, we consider the form (4.5). For any iteration index $k \in \mathbb{N}$ choose $i(k) \in \{1, \dots, N\}$ and $\ell(k) \in \{1, \dots, 3\}$ and define sequence of iterates $(\mathbf{z}^k)_{k \in \mathbb{N}}$ by

$$\mathbf{z}^{k+1} = \text{prox}_{s_k G_\lambda} \left(\mathbf{z}^k - s_k \nabla \mathbf{J}_{\ell(k)}^{(i(k))}(\mathbf{z}^k) \right). \quad (4.16)$$

The second step implements the proximal mapping of $\mathbf{z} \mapsto s_k G_\lambda(\mu)$ with $G_\lambda(\mu) = \frac{\lambda}{2} \|\mathbf{L}(\mu)\|^2 + \boldsymbol{\chi}_{\mathbb{D}}(\mu)$. As in the previous section, this can be computed with Dykstra's projection algorithm (3.7)–(3.10).

For better scaling, in our actual numerical implementation, we replace the scalar step sizes s_k by the adaptive step size rule

$$s_k^{i,\ell} := \arg \min \{ \mathbf{z}^k - t \nabla \mathbf{J}_\ell^{(i)}(\mathbf{z}^k) \mid t \in \mathbb{R} \}. \quad (4.17)$$

Note that computing such step sizes does barely increase the computational time of the stochastic gradient method, since all involved calculations are anyhow necessary for computing the gradient for the iterative update. In opposition to that, calculating a similar adaptive step size for the algorithms proposed in Section 3 would require evaluation of the forward operators \mathbf{F}_i and therefore would significantly increase the computation time. This might be seen as an additional advantage of the novel MULL formulation (4.1) and its regularized version (4.2).

5 Numerical Simulations

For the Tikhonov approach to multi-source QPAT, the radiative transfer equation is numerically solved by a streamline diffusion finite element method. Solving the RTE is required to evaluate the forward operator \mathbf{F} and the gradient ∇F of the data fidelity term in every iterative step. For the alternative multilinear approach, these calculations are not necessary. However, the application of the transport operator to Φ has to be calculated for every update of \mathbf{J}_1 . The simulations are performed on the square domain $\Omega = [-1 \text{ cm}, 1 \text{ cm}]^2$, where the absorption and the scattering coefficient are supported.

5.1 Numerical Solution of the RTE

Employing a finite element scheme, we derive the weak formulation of Equation (2.3) by integrating against a test function $w: \Omega \times \mathbb{S}^1 \rightarrow \mathbb{R}$ and replacing the exact solution Φ by a linear combination in the finite element space $\Phi^{(h)} = \sum_{i=1}^{N_h} c_i^{(h)} \psi_i^{(h)}(x, \theta)$ as in [33]. Here, the basis function $\psi_i^{(h)}(x, \theta)$ is the product of a basis function in space and a basis function in velocity. The spatial domain is triangulated uniformly with mesh size h and P_1 -Lagrangian element function for the spatial and velocity domain. By choosing the test function $w(x, \theta) = \sum_{j=1}^{N_h} w_j(\psi_j(x, \theta) + D(x, \theta) \theta \cdot \nabla_x \psi_j(x, \theta))$ with streamline diffusion coefficient $D(x, \theta)$, we obtain

$$\int_{\Omega} \int_{\mathbb{S}^1} (D\theta \cdot \psi_i - \psi_i) \theta \cdot \nabla_x \psi_j d\theta dx + \int_{\Gamma_+} |\theta \cdot \nu| d\sigma + \int_{\Omega} \int_{\mathbb{S}^1} (\mu_a + \mu_s - \mu_s \mathbf{K})(\psi_j + D\theta \cdot \nabla_x \psi_j) \psi_i d\theta dx = \int_{\Gamma_-} |\theta \cdot \nu| \psi_i \psi_j d\sigma. \quad (5.1)$$

Equation (5.1) yields a system of linear equations $M^{(h)} c^{(h)} = b^{(h)}$, where evaluating the left-hand side of (5.1) provides the entries of $M^{(h)}$, the right-hand side gives the components of vector $b^{(h)}$. Note that the sparsity of matrix $M^{(h)}$ is low and solving the linear system for the Tikhonov approach is very time-consuming. On the other hand, the solution via the MULL formulation requires only a matrix vector multiplication, since in this case $\Phi^{(h)}$ is an independent variable. Thus, only the application to Φ has to be calculated and the transport equation does not need to be solved.

5.2 Test Scenario for Multiple Illumination

The sample is illuminated in orthogonal direction at the boundaries of $\Omega = [-1 \text{ cm}, 1 \text{ cm}]^2$. In our simulations, we use $N = 4$ homogenous illuminations and no internal sources. The illuminations are applied separately from each side (left, right, top and bottom) with acoustic data measured on a half circle on the same side as the illumination (see Figure 5.1). For the scattering kernel, we use the two-dimensional Henyey–Greenstein kernel,

$$k(\theta, \theta') := \frac{1}{2\pi} \frac{1 - g^2}{1 + g^2 - 2g \cos(\theta \cdot \theta')} \quad \text{for } \theta, \theta' \in \mathbb{S}^1,$$

where the anisotropy factor is chosen as $g = 0.5$ in all our experiments.

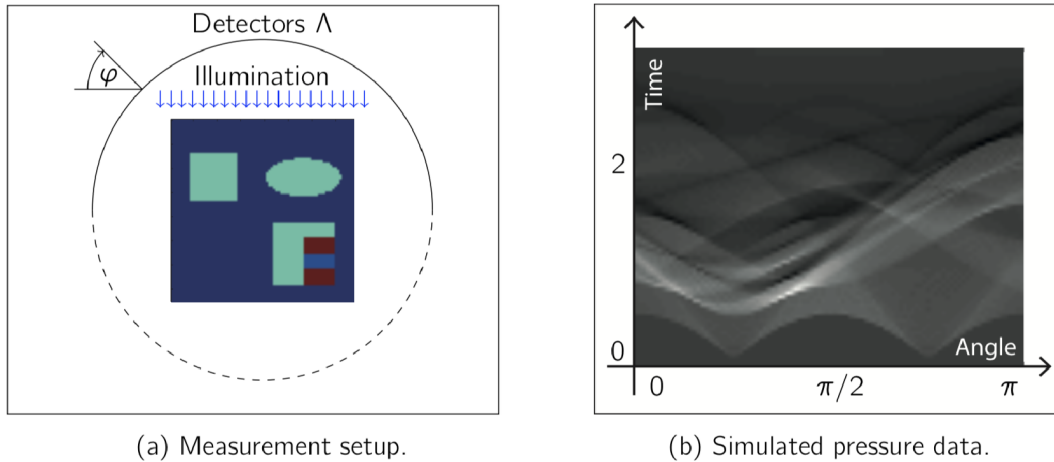


Figure 5.1: (a) The phantom is defined on the square $\Omega = [-1 \text{ cm}, 1 \text{ cm}]^2$ and the acoustic pressure is measured on a semi-circle on the side of the illumination. (b) The simulated pressure correspond to the phantom and the illumination on the right hand side and are represented as gray scale density.

For the simulated data, we choose a spatial mesh size $2/100$, in order to discretize the velocity direction the unit circle is divided in 64 subintervals. In order to avoid inverse crime, for the reconstruction, we use a different spatial mesh size $h = 2/80$ and use $N_\theta = 48$ velocity directions. Calculating the simulated data corresponds to evaluating the forward operators \mathbf{F}_i with perpendicular boundary illumination constant along one side of the boundary square,

$q_o(x, \theta) = \delta(\theta - \theta_i)\chi_i(x)1 \text{ mJ cm}^{-1}$, where δ is an approximation of the Dirac delta function and χ_i the indicator function of side i of Ω . In this way, we simulate data

$$v_i = \mathbf{U} \circ \mathbf{H}_i(\mu) + z_i^{\text{noise}} \quad \text{for } i = 1, \dots, 4.$$

Thereby, the heating operator is computed numerically by solving the RTE as described in Section 5.1. The wave operator \mathbf{U} is evaluated by straightforward discretization of the well-known explicit formulas for (2.7) that can be found, for example, in [38, 25]. In the following, we present results for exact data (where $z_i^{\text{noise}} = 0$) as well as for noisy data. For the noisy data case, we add 0.5% random noise to the simulated data, i.e., we take the maximum value of the simulated pressure and add white noise z_i^{noise} with a standard deviation of 0.5% of that maximal value. The phantom, the setup and the simulated data for one of the four illuminations (top) are shown in Figures 5.1 and 5.2.

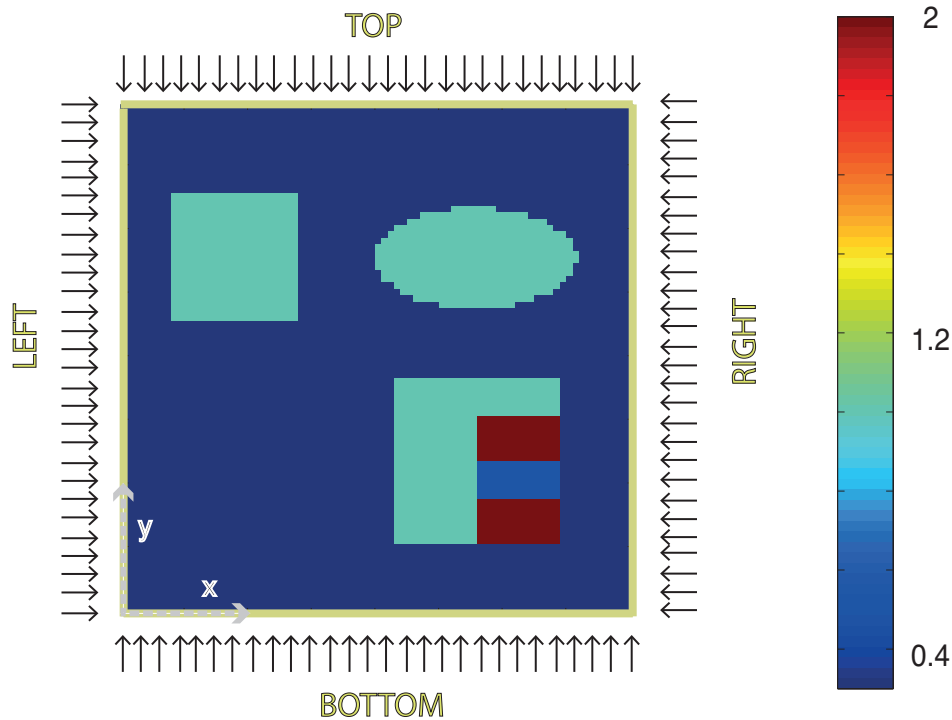


Figure 5.2: Absorption coefficient distribution of the tissue sample used for the numerical examples. Background absorption of the tissue is taken as $\mu_a = 0.3 \text{ cm}^{-1}$, the blue obstacles have $\mu_a = 1 \text{ cm}^{-1}$ and the red stripes $\mu_a = 2 \text{ cm}^{-1}$. The area between the red stripes has absorption coefficient $\mu_a = 0.5 \text{ cm}^{-1}$. The scattering coefficient is constant in the whole sample and chosen to be $\mu_s = 3 \text{ cm}^{-1}$. Illuminations are applied consecutively from top, right, bottom and left. The corresponding boundary sources are given by $q_o(x, \theta) = \delta(\theta - \theta_i)\chi_i(x)1 \text{ mJ cm}^{-1}$. The x - and y -axis cover $[-1 \text{ cm}, 1 \text{ cm}]$.

5.3 Numerical Results

For regularizing the absorption and scattering coefficient, we make use of Laplace regularization and choose $\mathbf{L}_a = \Delta$ and $\mathbf{L}_s = 100\Delta$, respectively. We assume that the coefficient μ is known at the boundary of Ω and is therefore used as the starting value of our iterative schemes. Furthermore, we use the boundary value of μ for regularization; that is, we implement it in the Dykstra projection procedure (3.8) by iteratively projecting on the known boundary value. In the following, we discuss the methods that we have outlined in the previous section.

■ **Standard formulation of QPAT (3.1):** We assume that the scattering coefficient is known and we restrict ourselves to reconstructing the absorption coefficient. Then, the proximal gradient and proximal stochastic gradient algorithm, respectively, read

$$\mu_a^{k+1} = \text{prox}_{s_k G_\lambda} \left(\mu_a^k - \frac{s_k}{4} \sum_{i=1}^4 \nabla F_i(\mu_a^k, \mu_s) \right), \quad (5.2)$$

$$\mu_a^{k+1} = \text{prox}_{s_k G_\lambda} \left(\mu_a^k - s_k \nabla F_{i(k)}(\mu_a^k, \mu_s) \right). \quad (5.3)$$

In contrast, to the full proximal gradient algorithm, the proximal stochastic gradient algorithm avoids evaluating the full gradient ∇F , but selects randomly an illumination number $i \in \{1, \dots, 4\}$ for each iterative step. Because of formula (2.13), each iteration of the above procedures requires the calculation of the solution of the radiative transfer equation Φ as well of its adjoint Φ^* .

The top row in Figure 5.3 shows reconstruction result for the absorption coefficient using the original formulation with the proximal gradient method with $\lambda = 2 \times 10^{-8}$ and 10 iterative steps. The left picture shows the relative error $\|\mu_a - \mu_a^k\|/\|\mu_a\|$. Note that, in this case, solutions of the RTE and its adjoint have to be computed for four illuminations per iterative step. The reconstruction results in the bottom row in Figure 5.3 are obtained by the proximal stochastic gradient method with $\lambda = 2 \times 10^{-7}$. The regularization parameters λ have been selected empirically as a trade-off between stability and accuracy. The total number of iterations is taken as 30. In each iteration, a illumination pattern is chosen randomly and the computation of RTE and its adjoint is executed only for this single illumination. Therefore, the computational effort for the proximal stochastic gradient method is approximately 3/4 of the proximal gradient algorithm using full gradients. For the algorithms based on the standard formulation (3.1), calculating adaptive step sizes similar to (4.17) is time-consuming as this requires another evaluation of the forward operator \mathbf{F}_i and therefore another solution of the RTE. Therefore, we simply use a constant step size rule; in our numerical experiments, it turned out that $s_k = 0.5$ is a suitable choice.

■ **Novel MULL formulation of QPAT (4.1):** The multilinear approach overcomes the problem of solving the RTE by minimizing (4.4) or (4.5). In both cases, one selects an arbitrary functional and performs a steepest descent step, resulting in an iterative scheme for the variables μ_a, μ_s, Φ and H . Recall that none of the partial gradients (4.6)–(4.14) requires solving the RTE (which is the most time-consuming part for the standard formulation of QPAT). In each iterative step, we take a random illumination number $i \in \{1, \dots, 4\}$ and a random functional number $\ell \in \{1, 2, 3\}$. The gradient step then consists of the update rule

$$\begin{pmatrix} \mu \\ \Phi_i \\ H_i \end{pmatrix}^{k+1} = \begin{pmatrix} \mu \\ \Phi_i \\ H_i \end{pmatrix}^k + s_k \cdot \nabla \mathbf{J}_\ell((\mu, \Phi_i, H_i)^k). \quad (5.4)$$

Dykstra's algorithm for smoothing the μ component is applied after each iterative step when $\ell \in \{1, 2\}$. Iteration (5.4) contains a gradient step for the RTE. Since one gradient step is not enough to obtain an appropriate approximation to the solution of the transport equation, we apply iteration (5.4) 40 times whenever $\ell = 1$ is chosen. In this situation, we apply the Dykstra iteration in the μ component after these 40 iteration steps, whereas the positivity projection is done in every step. Flowcharts of the stochastic gradient algorithms (standard and MULL formulations) are shown in Figure 5.4. For the projected stochastic gradient method, regularization of μ is done by incorporating the regularization functional \mathbf{J}_4 in the random choice of functionals; see (4.4). The positivity restriction is realized with the cut projection $\mathbf{P}_{\mathbb{D}}(\mu) = \max\{0, \mu\}$ applied after every iterative step.




Figure 5.3: Reconstruction results based on standard formulation (3.1). **Top**: proximal gradient method; **Bottom**: proximal stochastic gradient method. The left images show the relative reconstruction errors of the reconstructed absorption coefficient as a function of the number of iterations, whereas the right pictures show the result after the final iteration. (The phantom is as described in Figure 5.2.)

Figure 5.5 shows reconstructions with the stochastic gradient methods for the novel MULL formulation of QPAT (4.1). The results in the top row are for the MULL-proximal stochastic gradient algorithm with $\lambda = 2 \times 10^{-8}$ and in the the bottom row results for the MULL-projected stochastic gradient method with $\lambda = 2 \times 10^{-8}$ are shown. In both cases, we used 1000 iterations.

Remark 5.1. *Note that in the stochastic gradient methods for the novel MULL formulation of QPAT calculating the matrix vector product $M^{(h)} \cdot \Phi$ is the most costly part. In contrast, the standard formulation (3.1) requires the solution of the system $M^{(h)}c(h) = b^{(h)}$. Since the matrix $M^{(h)}$ is sparse only in its spatial domain, this is very time-consuming. On the other hand, the matrix $M^{(h)}$ (which is a discretization of $\theta \cdot \nabla_x + \mu_a + \mu_s(\mathbf{I} - \mathbf{K})$) has a simple dependence on the variables μ_a, μ_s . We therefore can compute the velocity entries of $M^{(h)}$ at the beginning of the iterative process to save computation time.*

The reconstruction times for the final reconstructions using all methods described above are shown in Table 1. For the standard formulation of QPAT, the reconstruction times seem to be in accordance with reported results using gradient or Newton-type methods for QPAT (see, for example, [50].) The methods based on the new MULL formulation (4.1) (after 1000 iter-

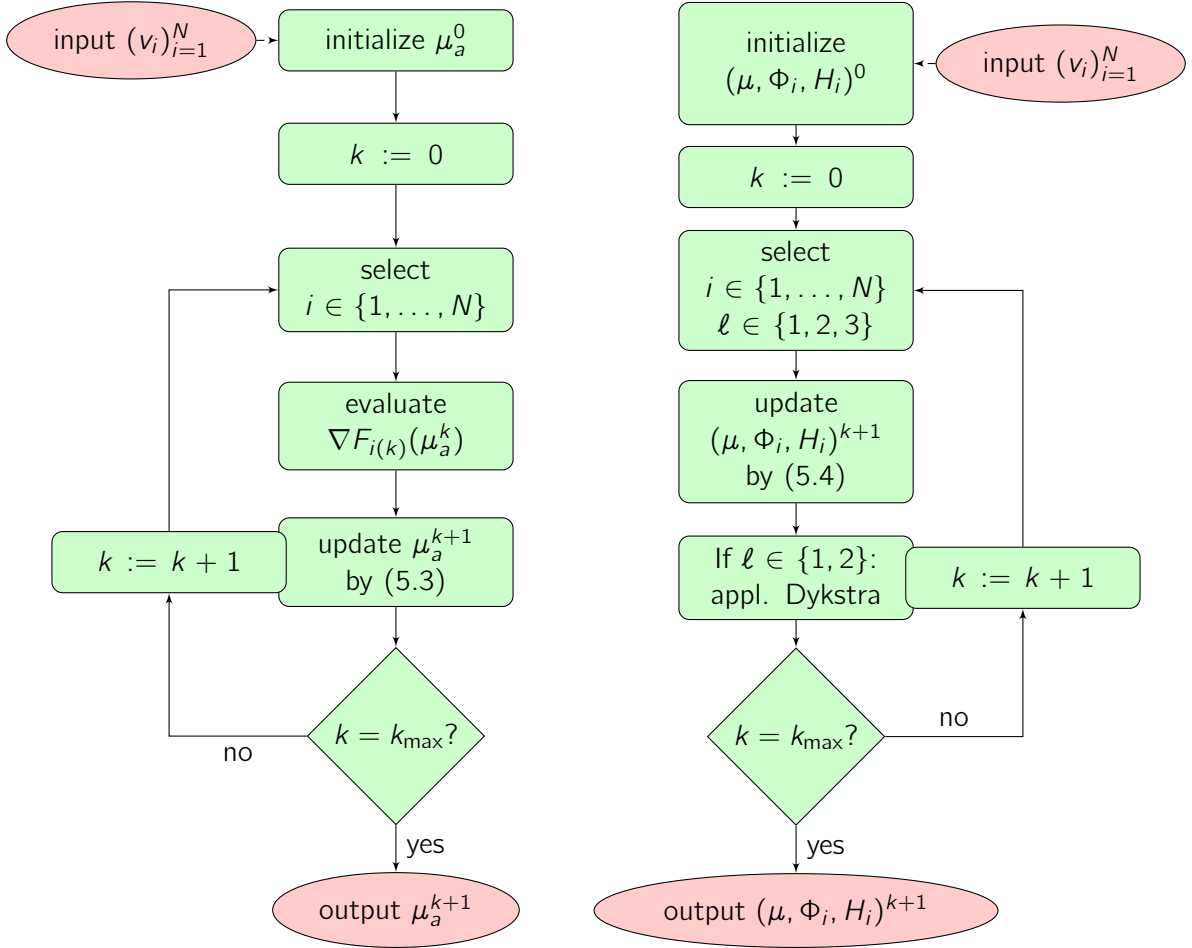


Figure 5.4: Flowcharts of stochastic gradient algorithms for QPAT proposed in this paper. **Left:** algorithm based on the standard formulation (3.1). **Right:** algorithm based on the novel MULL formulation (4.1). The update (5.3) using the standard formulation requires solving the forward RTE and the adjoint RTE, which is not required by (5.4) with the MULL formulation. Simulations are performed with $N = 4$.

ations) are faster than the methods based on the standard formulation (3.1) of QPAT (after 10, respectively, 40 iterations). From the relative reconstruction errors shown in Figures 5.3 and 5.5, one notices that, opposed to the methods based on (3.1), the methods using the MULL formulation could even be stopped much earlier while still obtaining a comparable reconstruction quality. We roughly estimate a speedup of a factor 10 using the novel MULL formulation instead of the standard formulation of QPAT.

In Figure 5.6, we show results for noisy data using the proximal gradient method based on the standard formulation (3.1) (top) and the proximal stochastic gradient method using the MULL formulation for QPAT (bottom). The regularization parameter is chosen as in the exact data case. Finally, in Figure 5.7, we show reconstruction results using only two consecutive illuminations applied from the top and from the left with noisy data. We use 10 iterations for the proximal gradient algorithm based on (3.1) (shown in left image in Figure 5.7) and 500 iterations for the stochastic gradient algorithms based on the MULL formulation (4.1) (shown in the right image in Figure 5.7).

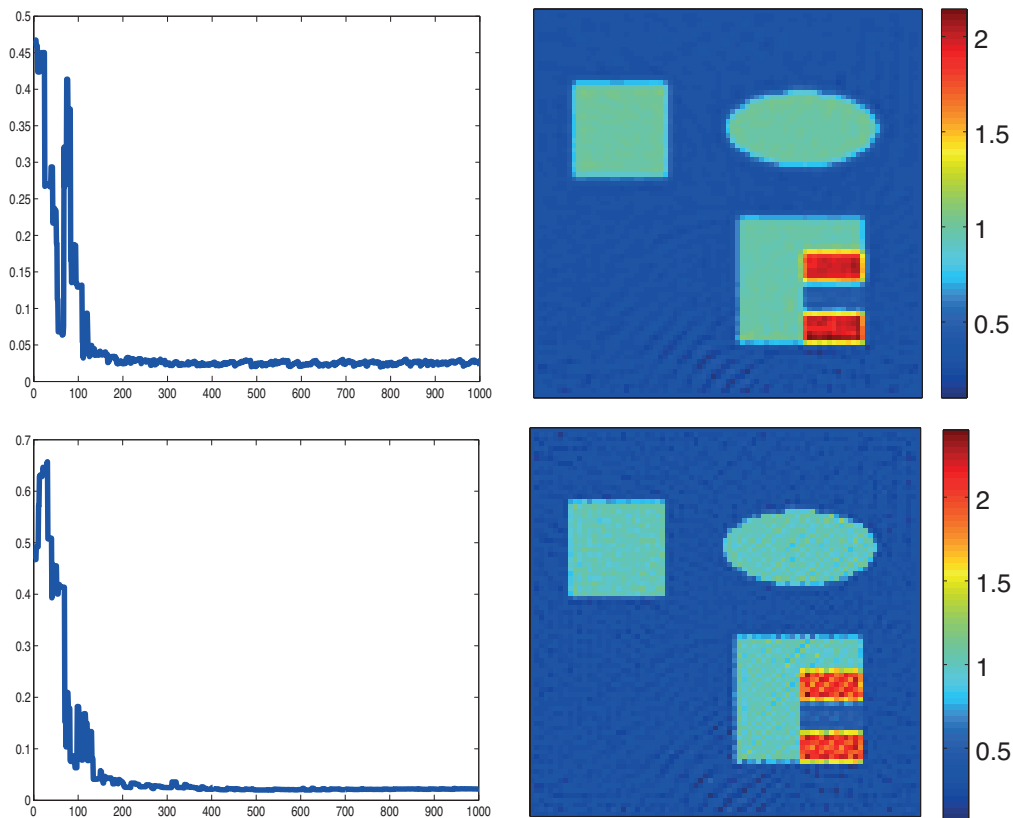


Figure 5.5: Reconstruction results based on the novel MULL formulation (4.1). **Top:** MULL-proximal stochastic gradient method based on the decomposition (4.4). **Bottom:** MULL-projected stochastic gradient method based on the decomposition (4.5). The left images show the relative reconstruction errors of the reconstructed absorption coefficient as a function of the number of iterations, whereas the right pictures show the results after the last iterations. (The phantom is as described in Figure 5.2.)

6 Conclusions

In this paper, we developed efficient proximal stochastic gradient methods for image reconstruction in multi-source QPAT. We used the RTE as an accurate model for light transport and employed the single stage approach for QPAT introduced in [33]. One class of the proximal stochastic gradient methods has been developed based on the standard formulation for QPAT given in (3.1). Additionally, we developed another class using proximal stochastic gradient methods for the new MULL formulations (4.1) and (4.2) for QPAT. Besides proposing proximal stochastic gradient methods for QPAT, we also consider the formulations (4.1) and (4.2) as the main contributions of the present article. These new formulations avoid the time-consuming evaluation of the RTE at each iteration and allow for treating the QPAT problem as a constrained optimization problem, which enables the use of a variety of numerical algorithms. Here, we used a penalty approach in combination with stochastic gradient methods for the solution. Future work will be done in the direction of developing new algorithms based on (4.1) and (4.2). Additionally, we will investigate the use of different regularization terms in (4.2). Finally, the theoretical convergence analysis of proximal gradient algorithms and other iterative algorithms for solving (4.1) will be the subject of future research.

Algorithm	Model Update	No. Iterations	Reconstruction Time
Proximal gradient	(3.1) (3.5)	10	27.2 h
Proximal stochastic gradient	(3.1) (3.13)	30	24.4 h
MULL-proximal stochastic gradient	(4.1) (4.15)	1000	14.7 h
MULL-projected stochastic gradient	(4.1) (4.16)	1000	11.9 h

Table 1: Reconstruction times for all methods. Recall that one iteration of the proximal stochastic gradient method is approximately four times cheaper than one iteration of the full proximal gradient method (both based on (3.1)). Further recall that one step in the methods based on the MULL formulation (4.1) is much less time consuming than for the methods based on (3.1); see Remark 5.1.

Acknowledgments

M.H. acknowledges support of the Austrian Science Fund (FWF), project P 30747. S.R. is a recipient of a DOC Fellowship of the Austrian Academy of Sciences (OEAW) and acknowledges support of the OEAW for his PhD project. We thank the referees for helpful comments that increased the readability of the this manuscript, as well as pointing out some interesting references to us.

References

- [1] M. Agranovsky, P. Kuchment, and L. Kunyansky. On reconstruction formulas and algorithms for the thermoacoustic tomography. In L. V. Wang, editor, *Photoacoustic imaging and spectroscopy*, chapter 8, pages 89–101. CRC Press, 2009.
- [2] G.S. Alberti and H. Ammari. Disjoint sparsity for signal separation and applications to hybrid inverse problems in medical imaging. *Appl. Comput. Harmon. Anal.*, 42(2):319–349, 2017.
- [3] H. Ammari, E. Bossy, V. Jugnon, and H. Kang. Reconstruction of the optical absorption coefficient of a small absorber from the absorbed energy density. *SIAM J. Appl. Math.*, 71(3):676–693, 2011.
- [4] H. Ammari, J. Garnier, H. Kang, L.H. Nguyen, and L. Seppecher. *Multi-Wave Medical Imaging, Mathematical Modelling & Imaging Reconstruction*. World Science, 2017.
- [5] S. R. Arridge. Optical tomography in medical imaging. *Inverse Probl.*, 15(2):R41–R93, 1999.
- [6] G. Bal, A. Jollivet, and V. Jugnon. Inverse transport theory of photoacoustics. *Inverse Probl.*, 26:025011, 2010.
- [7] G. Bal and K. Ren. Multi-source quantitative photoacoustic tomography in a diffusive regime. *Inverse Probl.*, 27(7):075003, 20, 2011.
- [8] Heinz H Bauschke and Patrick L Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.
- [9] P. Beard. Biomedical photoacoustic imaging. *Interface focus*, 1(4):602–631, 2011.




Figure 5.6: Reconstruction results from noisy data. **Top:** Proximal gradient method based on (3.1). **Bottom:** MULL-proximal stochastic gradient method. The left images show the relative reconstruction errors of the reconstructed absorption coefficient as a function of the number of iterations, whereas the right pictures show the results after the last iterations. (The phantom is as described in Figure 5.2.)

- [10] D. P. Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. *Optimization for Machine Learning*, 2010(1-38):3, 2011.
- [11] D. P. Bertsekas. Incremental proximal methods for large scale convex optimization. *Mathematical programming*, 129(2):163, 2011.
- [12] P. Burgholzer, J. Bauer-Marschallinger, H. Grün, M. Haltmeier, and G. Paltauf. Temporal back-projection algorithms for photoacoustic tomography with integrating line detectors. *Inverse Probl.*, 23(6):S65–S80, 2007.
- [13] J. Chen and Y. Yang. Quantitative photo-acoustic tomography with partial data. *Inverse Probl.*, 28(11):115014, 2012.
- [14] P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer, 2011.
- [15] Patrick L Combettes and Valérie R Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling & Simulation*, 4(4):1168–1200, 2005.
- [16] B. T. Cox, S. A. Arridge, and P. C. Beard. Gradient-based quantitative photoacoustic image reconstruction for molecular imaging. *Proc. SPIE*, 6437:64371T, 2007.




Figure 5.7: Reconstruction results from noisy data for two illuminations. **Left:** proximal gradient method based on (3.1) using 10 iterations. **Right:** MULL-proximal stochastic gradient method using only 500 iterations. The phantom is as described in Figure 5.2 and, for the reconstruction methods, we use two consecutive illuminations (from the top and from the left). The reconstruction time has been about 14 h for the method based on the standard formulation (3.1) and 3 h for the proposed MULL-proximal stochastic gradient method.

- [17] B. T. Cox, S. R. Arridge, P. Köstli, and P. C. Beard. Two-dimensional quantitative photoacoustic image reconstruction of absorption distributions in scattering media by use of a simple iterative method. *Appl. Opt.*, 45(8):1866–1875, 2006.
- [18] B. T. Cox, J. G. Laufer, S. R. Arridge, and Paul C. Beard. Quantitative spectroscopic photoacoustic imaging: a review. *J. Biomed. Opt.*, 17(6):0612021, 2012.
- [19] R. Dautray and J. Lions. *Mathematical analysis and numerical methods for science and technology. Vol. 6.* Springer-Verlag, Berlin, 1993.
- [20] A. De Cezaro, M. Haltmeier, A. Leitão, and O. Scherzer. On steepest-descent-kaczmarz methods for regularizing systems of nonlinear ill-posed equations. *Applied Mathematics and Computation*, 202(2):596–607, 2008.
- [21] J. Duchi and Y. Singer. Efficient online and batch learning using forward backward splitting. *Journal of Machine Learning Research*, 10:2899–2934, 2009.
- [22] H. Egger and M. Schlottbom. Stationary radiative transfer with vanishing absorption. *Math. Models Methods Appl. Sci.*, 24(5):973–990, 2014.
- [23] H. Egger and M. Schlottbom. Numerical methods for parameter identification in stationary radiative transfer. *Comput. Optim. Appl.*, 62(1):67–83, 2015.
- [24] Heinz Werner Engl, Martin Hanke, and Andreas Neubauer. *Regularization of inverse problems*, volume 375. Springer Science & Business Media, 1996.
- [25] L. C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. Amer. Math. Soc., Providence, RI, 1998.
- [26] D. Finch, M. Haltmeier, and Rakesh. Inversion of spherical means and the wave equation in even dimensions. *SIAM J. Appl. Math.*, 68(2):392–412, 2007.

- [27] D. Finch, S. K. Patch, and Rakesh. Determining a function from its mean values over a family of spheres. *SIAM J. Math. Anal.*, 35(5):1213–1240, 2004.
- [28] H. Gao, J. Feng, and L. Song. Limited-view multi-source quantitative photoacoustic tomography. *Inverse Probl.*, 31(6):065004, 2015.
- [29] M. Haltmeier. Universal inversion formulas for recovering a function from spherical means. *SIAM J. Math. Anal.*, 46(1):214–232, 2014.
- [30] M. Haltmeier, R. Kowar, A. Leitao, and O. Scherzer. Kaczmarz methods for regularizing nonlinear ill-posed equations ii: Applications. *Inverse Probl. Imaging*, 1(3):507, 2007.
- [31] M. Haltmeier, A. Leitao, and O. Scherzer. Kaczmarz methods for regularizing nonlinear ill-posed equations i: convergence analysis. *Inverse Probl. Imaging*, 1(2):289, 2007.
- [32] M. Haltmeier, L. Neumann, L. V. Nguyen, and S. Rabanser. Analysis of the linearized problem of quantitative photoacoustic tomography, 2017. arXiv:1702.04560.
- [33] M. Haltmeier, L. Neumann, and S. Rabanser. Single-stage reconstruction algorithm for quantitative photoacoustic tomography. *Inverse Probl.*, 31(6):065005, 2015.
- [34] M. Haltmeier and L. V. Nguyen. Analysis of iterative methods in photoacoustic tomography with variable sound speed. *SIAM J. Imaging Sci.*, 10(2):751–781, 2017.
- [35] M. Haltmeier, T. Schuster, and O. Scherzer. Filtered backprojection for thermoacoustic computed tomography in spherical geometry. *Math. Method. Appl. Sci.*, 28(16):1919–1937, 2005.
- [36] C. Huang, K. Wang, L. Nie, and M. A. Wang, L. V. and Anastasio. Full-wave iterative image reconstruction in photoacoustic tomography with acoustically inhomogeneous media. *IEEE Trans. Med. Imag.*, 32(6):1097–1110, 2013.
- [37] Kazufumi Ito and Karl Kunisch. *Lagrange multiplier approach to variational problems and applications*. SIAM, 2008.
- [38] F. John. *Partial Differential Equations*, volume 1 of *Applied Mathematical Sciences*. Springer Verlag, New York, fourth edition, 1982.
- [39] B. Kaltenbacher, A. Neubauer, and O. Scherzer. *Iterative Regularization Methods for Nonlinear Ill-Posed Problems*. de Gruyter, Berlin, 1 2008.
- [40] G. Kanschat. Solution of radiative transfer problems with finite elements. In *Numerical methods in multidimensional radiative transfer*, pages 49–98. Springer, Berlin, 2009.
- [41] R. A Kruger, P. Lui, Y. R. Fang, and R. C. Appledorn. Photoacoustic ultrasound (PAUS) – reconstruction tomography. *Med. Phys.*, 22(10):1605–1609, 1995.
- [42] H. Li and M. Haltmeier. The averaged kaczmarz iteration for solving inverse problems, 2017. arXiv:1709.00742.
- [43] A. V. Mamonov and K. Ren. Quantitative photoacoustic imaging in radiative transport regime. *Comm. Math. Sci.*, 12(2):201–234, 2014.
- [44] W. Naetar and O. Scherzer. Quantitative photoacoustic tomography with piecewise constant material parameters. *SIAM J. Imaging Sci.*, 7(3):1755–1774, 2014.

- [45] L. V. Nguyen and L. A. Kunyansky. A dissipative time reversal technique for photoacoustic tomography in a cavity. *SIAM J. Imaging Sci.*, 9(2):748–769, 2016.
- [46] M. Pereyra, P. Schniter, E. Chouzenoux, J.-Y. Pesquet, J.-C. and Tourneret, A. O. Hero, and S. McLaughlin. A survey of stochastic simulation and optimization methods in signal processing. *IEEE J. Sel. Topics Signal Process.*, 10(2):224–241, 2016.
- [47] K. Ren, H. Gao, and H. Zhao. A hybrid reconstruction method for quantitative PAT. *SIAM J. Imaging Sci.*, 6(1):32–55, 2013.
- [48] A. Rosenthal, V. Ntziachristos, and D. Razansky. Acoustic inversion in optoacoustic tomography: A review. *Curr. Med. Imaging Rev.*, 9(4):318–336, 2013.
- [49] A. Rosenthal, D. Razansky, and V. Ntziachristos. Fast semi-analytical model-based acoustic inversion for quantitative optoacoustic tomography. *IEEE Trans. Med. Imag.*, 29(6):1275–1285, 2010.
- [50] 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 Probl.*, 29(7):075006, 2013.
- [51] O. Scherzer, M. Grasmair, H. Grossauer, M. Haltmeier, and F. Lenzen. *Variational methods in imaging*, volume 167 of *Applied Mathematical Sciences*. Springer, New York, 2009.
- [52] T. Tarvainen, B. T. Cox, J. P. Kaipio, and S. A. Arridge. Reconstructing absorption and scattering distributions in quantitative photoacoustic tomography. *Inverse Probl.*, 28(8):084009, 2012.
- [53] C. Wang and T. Zhou. On iterative algorithms for quantitative photoacoustic tomography in the radiative transport regime. *Inverse Problems*, 33(11):115006, 2017.
- [54] L. V. Wang. Multiscale photoacoustic microscopy and computed tomography. *Nature Phot.*, 3(9):503–509, 2009.
- [55] L. Xiao and T. Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24(4):2057–2075, 2014.
- [56] L. Yao, Y. Sun, and J. Huabei. Transport-based quantitative photoacoustic tomography: simulations and experiments. *Phys. Med. Biol.*, 55(7):1917–1934, 2010.