Journal of Biomedical Optics

BiomedicalOptics.SPIEDigitalLibrary.org

Reconstruction-classification method for quantitative photoacoustic tomography

Emma Malone
Samuel Powell
Ben T. Cox
Simon Arridge

SPIE.

Reconstruction-classification method for quantitative photoacoustic tomography

Emma Malone,^{a,*} Samuel Powell,^b Ben T. Cox,^a and Simon Arridge^b

^aUniversity College London, Department of Medical Physics and Biomedical Engineering, Gower Street, WC1E 6BT London, United Kingdom

^bUniversity College London, Department of Computer Science, Gower Street, WC1E 6BT London, United Kingdom

Abstract. We propose a combined reconstruction-classification method for simultaneously recovering absorption and scattering in turbid media from images of absorbed optical energy. This method exploits knowledge that optical parameters are determined by a limited number of classes to iteratively improve their estimate. Numerical experiments show that the proposed approach allows for accurate recovery of absorption and scattering in two and three dimensions, and delivers superior image quality with respect to traditional reconstruction-only approaches. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JBO.20.12.126004](https://doi.org/10.1117/1.JBO.20.12.126004)]

Keywords: quantitative photoacoustic tomography; image reconstruction; chromophores.

Paper 150520R received Aug. 4, 2015; accepted for publication Oct. 30, 2015; published online Dec. 14, 2015.

1 Introduction

Photoacoustic tomography (PAT) is an emerging technique for *in vivo* imaging of soft biological tissue.¹ This hybrid modality uses ultrasound to detect optical contrast, combining the high resolution of acoustic methods with the spectroscopic capability of optical imaging. To generate a PA image, a short laser pulse is shined into the object, the ultrasonic waves emitted following the heating of the tissue are measured, and an image of the absorbed optical energy field is recovered. Whereas purely optical methods suffer from poor spatial resolution, acoustic waves propagate with minimal scattering, and PAT can achieve 100 μm resolution at depths of several centimeters. However, PA images provide only qualitative information about the tissue, and are not directly related to tissue morphology and functionality. The principal difficulty is that the PA image is the product of both the optical absorption coefficient (which is directly related to underlying tissue composition) and the light distribution (which is not). This severely restricts the range of applications for which PAT is suitable.

Quantitative photoacoustic tomography (QPAT) aims to provide clinically valuable images of the optical absorption and scattering coefficients, or chromophore (light-absorbing molecules) concentrations from conventional PA images via an image reconstruction method.^{2,3} A model of light propagation is required to relate the absorbed optical energy to the light fluence and tissue parameters. The primary challenge of QPAT is solving the nonlinear imaging problem. In particular, recovering the scattering coefficient is especially difficult due to the weak dependence of the absorbed energy density on scattering.

In this paper, we develop a method for solving the image reconstruction problem for QPAT by alternating reconstruction and segmentation steps in an automated iterative process. We introduce a probabilistic model that describes optical properties in terms of a limited number of optically distinct classes, which may correspond to tissues or chromophores. These are

identified and characterized by a classification, or segmentation, algorithm. This approach allows for the use of information retrieved by the classification in the reconstruction stage and vice versa. The aim of the reconstruction is to choose solutions for which the image parameters take values close to a finite set of discrete points. The aim of the classification algorithm is to progressively improve the parametric optical model and correct for errors in the initial assumptions. Multinomial models have been employed previously in the related fields diffuse optical tomography⁴ and electrical impedance tomography.⁵ For QPAT, the main advantage is that this approach enables accurate recovery of both the absorption and scattering coefficients simultaneously.

2 Numerical Methods

2.1 Quantitative Photoacoustic Imaging

A conventional PAT image is proportional to the absorbed optical energy

$$H(\mathbf{r}) = \hat{\Gamma}(\mathbf{r})\mu_a(\mathbf{r})\phi[\mu_a(\mathbf{r}), \mu_s'(\mathbf{r})] \quad \mathbf{r} \in \Omega, \quad (1)$$

where $\mathbf{r}$ is a position vector within the domain Ω , μ_a and μ_s' are the optical absorption and reduced scattering coefficients, ϕ is the optical fluence, and $\hat{\Gamma}$ is the Grüneisen parameter. The Grüneisen parameter represents the efficiency with which the tissue converts heat into acoustic pressure, and is often taken to be the constant $\hat{\Gamma}(\mathbf{r}) = 1 \quad \forall \mathbf{r} \in \Omega$. The fluence is dependent on the optical parameters and illumination pattern in the whole domain. The problem of recovering the optical parameters (μ_a, μ_s') from a conventional PAT image is known as the “quantitative” problem. The optical absorption μ_a is of particular interest because it is fundamentally related to underlying tissue physiology and functionality, and encodes clinically useful information such as tissue oxygenation levels and chromophore concentrations. Conversely, the absorbed energy density H depends nontrivially on optical absorption and thus is not

*Address all correspondence to: Emma Malone, E-mail: e.malone@ucl.ac.uk

directly related to tissue morphology because it is distorted structurally and spectrally by the nonuniform light fluence.

2.2 Diffusion Model of Light Transport

In order to recover the optical parameters (μ_a, μ_s') , a model of light propagation within the tissue is required. For highly scattering media and those far from boundaries and sources, a low-order spherical harmonic approximation to the “radiative transfer equation” is suitable. The “diffusion approximation” is given by⁶

$$[\mu_a - \nabla \cdot \kappa(\mathbf{r})\nabla]\phi(\mathbf{r}) = q(\mathbf{r}), \quad (2)$$

where $q(\mathbf{r})$ is an isotropic source term and $\kappa = 1/3\mu_s'$ is the diffusion coefficient.

We set Robin boundary conditions

$$\phi(\mathbf{r}) + \frac{1}{2A}\kappa(\mathbf{r})\hat{\mathbf{n}} \cdot \nabla\phi(\mathbf{r}) = 0 \quad \mathbf{r} \in \delta\Omega, \quad (3)$$

where A accounts for the refractive index mismatch at the boundary.

2.3 Minimization-Based Quantitative Photoacoustic Tomography Imaging

In this paper, we adopt a gradient-based minimization approach to image reconstruction. Typically, both μ_a and μ_s' are unknown and need to be recovered simultaneously from the absorbed energy density. An objective function is defined, which measures the distance between the conventional PAT image H^m and the data predicted by the model for the current estimates $H(\mu_a, \mu_s')$

$$\mathcal{E} = \frac{1}{2} \int_{\Omega} [H^m - H(\mu_a, \mu_s')]^2 d\Omega. \quad (4)$$

In order to treat the problem for a generic geometry, the finite element method is employed, whereby a weak formulation of the diffusion approximation [Eq. (2)] is considered. A discretization of the domain is defined, and the fluence and optical parameters are expressed in terms of the same piecewise-linear basis functions $u_i(\mathbf{r})$: $\chi \approx \sum_i \chi_i u_i(\mathbf{r})$ for $\chi \in \{\mu_a, \mu_s', \phi\}$, where χ_i are nodal coefficients and $i = 1, \dots, N$.

We assume that the data $\mathbf{d}^m$ is the absorbed energy density H^m , projected onto a particular basis $\{\Psi_j\}$,

$$\begin{aligned} \mathbf{d}^m &= \{d_j^m, j = 1, \dots, N\}, \\ d_j^m &= \int_{\Omega} H^m(\mathbf{r})\Psi_j(\mathbf{r})d\Omega = \langle \Psi_j, H^m \rangle. \end{aligned} \quad (5)$$

Choices for $\{\Psi_j\}$ include:

1. Point sampling $\Psi_j(\mathbf{r}) = \delta(\mathbf{r} - \mathbf{r}_j)$,
2. Piecewise-linear sampling $\Psi_j = u_j$,
3. Sinc sampling $\Psi_j = \text{sinc}(|\mathbf{r} - \mathbf{r}_j|)$.

Substituting into the objective function [Eq. (4)] leads to the discrete form of the objective function

$$\begin{aligned} \mathcal{E} &= \frac{1}{2} \sum_j [d_j^m - \langle \Psi_j, H(\mu_a, \mu_s') \rangle]^2 \\ &= \frac{1}{2} \sum_j [d_j^m - \langle \Psi_j, \mu_a \phi \rangle]^2. \end{aligned} \quad (6)$$

If a single illumination source is used and both absorption and scattering are undetermined, the problem is ill posed.² In this study, the nonuniqueness of the solution was removed by using multiple illumination patterns,⁷⁻⁹ thus the objective function must be summed over the number of sources. In the following, we have omitted this sum for ease of notation. Prior information regarding the solution can be included by adding a regularization term

$$\mathcal{E} = \frac{1}{2} \sum_j (d_j^m - \langle \Psi_j, \mu_a \phi \rangle)^2 + \mathcal{R}(\mu_a', \mu_s'). \quad (7)$$

In the Bayesian framework, an image is obtained by maximizing the posterior probability of the parameters, given the data

$$p(\mu_a, \mu_s' | \mathbf{d}^m) \propto p(\mathbf{d}^m | \mu_a, \mu_s') p(\mu_a, \mu_s'). \quad (8)$$

Under this interpretation, the regularization term $\mathcal{R}$ is given by the negative log of the prior probability distribution

$$\mathcal{R}(\mu_a, \mu_s') = -\log p(\mu_a, \mu_s'). \quad (9)$$

2.4 Gradient Calculations

Cox et al.¹⁰ have shown that, for the continuous case, the gradient of Eq. (4) with respect to μ_a at position $\mathbf{r}^0$ is given by

$$\left. \frac{\partial \mathcal{E}}{\partial \mu_a} \right|_{\mathbf{r}^0} = -\phi(H^m - H)|_{\mathbf{r}^0} + \phi \cdot \phi^*|_{\mathbf{r}^0}, \quad (10)$$

where the “adjoint” light field ϕ^* is the solution to the equation

$$[\mu_a - \nabla \cdot \kappa(\mathbf{r})\nabla]\phi^*(\mathbf{r}) = \mu_a(H^m - H). \quad (11)$$

In the following, we derive the expression for the gradient in the discrete case. The sampled forward model can be expressed as a vector $\mathbf{H} = \{H_j, j = 1, \dots, N\}$

$$\begin{aligned} H_j &= \int_{\Omega} H(\mathbf{r})\Psi_j(\mathbf{r})d\Omega = \langle \Psi_j, H \rangle \\ &= \sum_{ik} \mu_{ai} \phi_k \int_{\Omega} \Psi_j(\mathbf{r})u_i(\mathbf{r})u_k(\mathbf{r})d\Omega = \boldsymbol{\phi}^T \mathbf{C}^j \boldsymbol{\mu}_a, \end{aligned} \quad (12)$$

where $\mathbf{C}^j$ is a sparse matrix indexed by i, k where the support of the basis functions $\Psi_j(\mathbf{r})$, $u_i(\mathbf{r})$, $u_k(\mathbf{r})$ overlap. Taking the derivative of Eq. (6) with respect to μ_{ai} , we have

$$\frac{\partial \mathcal{E}}{\partial \mu_{ai}} = -\sum_j \left(\frac{\partial H_j}{\partial \mu_{ai}} \right) (d_j^m - H_j). \quad (13)$$

Using the expression for the absorbed energy density [Eq. (12)],

$$\frac{\partial H_j}{\partial \mu_{ai}} = \mathbf{e}_i^T \mathbf{C}^j \boldsymbol{\phi} + \boldsymbol{\mu}_a^T \mathbf{C}^j \frac{\partial \boldsymbol{\phi}}{\partial \mu_{ai}}, \quad (14)$$

where $\mathbf{e}_i$ is a vector of zeros with a single 1 in position i . Substituting into Eq. (13) gives

$$\frac{\partial \mathcal{E}}{\partial \mu_{ai}} = -\sum_j \left(\mathbf{e}_i^T \mathbf{C}^j \boldsymbol{\phi} + \boldsymbol{\mu}_a^T \mathbf{C}^j \frac{\partial \boldsymbol{\phi}}{\partial \mu_{ai}} \right) (d_j^m - H_j). \quad (15)$$

The first term in Eq. (15) is

$$\begin{aligned} \sum_j \mathbf{e}_i^T \mathbf{C}^j \boldsymbol{\phi} (d_j^m - H_j) &= \sum_{j,i,k} e_i C_{ik}^j \phi_k (d_j^m - H_j) \\ &= \sum_{j,k} \phi_k (d_j^m - H_j) \int_{\Omega} \Psi_j(\mathbf{r}) u_i(\mathbf{r}) u_k(\mathbf{r}) d\Omega \\ &= \boldsymbol{\phi}^T \mathbf{E}^i (d^m - \mathbf{H}), \end{aligned} \quad (16)$$

where $\mathbf{E}^i$ is given by a reordering of C_{ik}^j

$$E_{kj}^i = \int_{\Omega} \Psi_j(\mathbf{r}) u_i(\mathbf{r}) u_k(\mathbf{r}) d\Omega. \quad (17)$$

Note that while $\mathbf{C}^j$ is symmetric, in general, $\mathbf{E}^i$ is not.

It remains to determine $\partial \boldsymbol{\phi} / \partial \mu_{ai}$. The discrete form of the DA model [Eq. (2)] assumes the form¹¹

$$(\mathbf{M} + \mathbf{K} + \mathbf{F}) \boldsymbol{\phi} = \mathbf{Q}, \quad (18)$$

where

$$M_{jk} = \sum_i \mu_{ai} \int_{\Omega} u_i u_j u_k d\Omega, \quad (19)$$

$$K_{jk} = \sum_i \kappa_i \int_{\Omega} u_i \nabla u_j \cdot \nabla u_k d\Omega, \quad (20)$$

$$F_{jk} = \sum_i \frac{1}{2A} \int_{\partial\Omega} u_j u_k dS, \quad (21)$$

$$Q_j = \sum_i q_i \int_{\Omega} u_i u_j d\Omega. \quad (22)$$

Taking the derivative of Eq. (18) with respect to the i 'th coefficient of $\boldsymbol{\mu}_a$,

$$(\mathbf{M} + \mathbf{K} + \mathbf{F}) \frac{\partial \boldsymbol{\phi}}{\partial \mu_{ai}} = -\mathbf{V}_{\mu_a}^i \boldsymbol{\phi}, \quad (23)$$

where

$$V_{\mu_a, jk}^i = \int_{\Omega} u_i u_j u_k d\Omega \quad (24)$$

is given by the derivative of the system matrix. We define the adjoint field $\boldsymbol{\phi}^*$ as the solution to the equation

$$(\mathbf{M} + \mathbf{K} + \mathbf{F}) \boldsymbol{\phi}^* = \mathbf{Q}^*, \quad (25)$$

where

$$\mathbf{Q}^* = \sum_j \boldsymbol{\mu}_a^T \mathbf{C}^j (d_j^m - H_j) \quad (26)$$

is the adjoint source. Taking $\boldsymbol{\phi}^*$ Equation (23) $-(\partial \boldsymbol{\phi} / \partial \mu_{ai})$. Eq. (25), we obtain

$$\sum_j \boldsymbol{\mu}_a^T \mathbf{C}^j \frac{\partial \boldsymbol{\phi}}{\partial \mu_{ai}} (d_j^m - H_j) = -\boldsymbol{\phi}^T \mathbf{V}_{\mu_a}^i \boldsymbol{\phi}^*. \quad (27)$$

Substituting into Eq. (15) gives the expression for the derivative with respect to μ_{ai}

$$\frac{\partial \mathcal{E}}{\partial \mu_{ai}} = \boldsymbol{\phi}^T [\mathbf{V}_{\mu_a}^i \boldsymbol{\phi}^* - \mathbf{E}^i (d^m - \mathbf{H})]. \quad (28)$$

The derivative with respect to μ_{si}' can be derived analogously

$$\frac{\partial \mathcal{E}}{\partial \mu_{si}'} = -\frac{\partial \kappa_i}{\partial \mu_{si}'} \boldsymbol{\phi}^T \mathbf{V}_{\mu_s'}^i \boldsymbol{\phi}^*, \quad (29)$$

where

$$V_{\mu_s', jk}^i = \int_{\Omega} u_i \nabla u_j \cdot \nabla u_k d\Omega \quad (30)$$

and $(\partial \kappa_i / \partial \mu_{si}') = -1/3 \mu_{si}'^2$. Note that calculation of the gradient requires only two runs of the forward model. The forward problem was solved using the Toast++ software package.¹¹

Choosing point-sampling $\Psi_j(\mathbf{r}) = \delta(\mathbf{r} - \mathbf{r}_j)$ gives simply $\mathbf{C}^j = \mathbf{E}^j = \mathbf{I}$. In this study, we chose piecewise-linear sampling $\Psi_j = u_j$, so we had $\mathbf{C}^j = \mathbf{E}^j = \mathbf{V}_{\mu_a}^j$ and

$$\frac{\partial \mathcal{E}}{\partial \mu_{ai}} = \boldsymbol{\phi}^T \mathbf{V}_{\mu_a}^i (\boldsymbol{\phi}^* - d^m + \mathbf{H}). \quad (31)$$

3 Reconstruction-Classification Method for Quantitative Photoacoustic Tomography

A reconstruction-classification scheme is devised, which enables the recovery μ_a and μ_s' by approaching the image reconstruction and segmentation problems simultaneously. At each reconstruction step, we minimize a regularized objective function, where the regularization term is given by a mixture model. At each classification step, the result of the previous reconstruction step is employed to update the class parameters for the multinomial model. We alternate between reconstruction and classification steps for a fixed number of iterations (Fig. 1).

3.1 Mixture Model for μ_a and μ_s'

In this section, we introduce a probability model for μ_a and μ_s' , which encodes prior knowledge about the optical parameters and allows us to bias the solution of the imaging problem accordingly. We assume that an array of labels ζ_i can be determined for each node, such that

$$\zeta_{ij} = \begin{cases} 1 & \text{if the } i\text{'th node is assigned to the } j\text{'th class,} \\ 0 & \text{otherwise.} \end{cases} \quad (32)$$

The labels constitute ‘‘hidden variables’’ on which the image parameters are dependent. For each class $j = 1, \dots, J$, a mean vector $\mathbf{m}_j = (\bar{\mu}_{aj}, \bar{\mu}_{sj}') \in \mathbb{R}^2$ is defined, and the covariance of each class is described by matrix $\Sigma_j \in \mathbb{R}^{2 \times 2}$.

We assume that if $\zeta_{ij} = 1$, the probability distribution for $\mathbf{x}_i = (\mu_{ai}, \mu'_{si})$ is given by a multivariate Gaussian distribution

$$p(\mathbf{x}_i|\boldsymbol{\theta}_j) = \mathcal{N}(\mathbf{m}_j, \Sigma_j), \quad (33)$$

where $\boldsymbol{\theta}_j$ indicates the set of class parameters $(\mathbf{m}_j, \Sigma_j)$.

The prior probability distribution of the class properties $\boldsymbol{\theta}_j$ is given by the conjugate prior to the Gaussian distribution. Prior information about the distribution of the class means or covariances can be encoded by choosing the parameters of the conjugate prior accordingly. Using a noninformative prior for the class means we have $p(\mathbf{m}_j) \propto 1$. The conjugate prior distribution for the covariance of a normal distribution is given by the normal inverse Wishart distribution (NIW):

$$\text{NIW}(\nu_j, \Gamma_j) = |\Sigma_j|^{-(\nu+d+1)/2} \exp\left[-\frac{1}{2} \text{Tr}(\Gamma_j \Sigma_j^{-1})\right], \quad (34)$$

where d is the dimension of the domain, ν_j indicates the number of degrees of freedom, and Γ_j is a scaling matrix. If the prior is noninformative, then $\nu_j = 0$ and $\Gamma_j = 0$, and the probability distribution of the class parameters becomes

$$p(\boldsymbol{\theta}_j) \propto |\Sigma_j|^{-(d+1)/2}, \quad (35)$$

which is known as Jeffreys prior.

The probability that the set of labels $\boldsymbol{\zeta}_i = \{\zeta_{i1}, \dots, \zeta_{ij}, \dots, \zeta_{iJ}\}$ is assigned to the i 'th node is given by a multinomial distribution

$$p(\boldsymbol{\zeta}_i|\boldsymbol{\lambda}) = \prod_j \lambda_j^{\zeta_{ij}}, \quad (36)$$

where λ_j is the overall probability that a node is assigned to the j 'th class. Therefore, the joint probability for $(\mathbf{x}_i, \boldsymbol{\zeta}_i)$ is given by the product

$$p(\mathbf{x}_i, \boldsymbol{\zeta}_i|\boldsymbol{\theta}, \boldsymbol{\lambda}) = p(\mathbf{x}_i|\boldsymbol{\zeta}_i, \boldsymbol{\theta})p(\boldsymbol{\zeta}_i|\boldsymbol{\lambda}) = \prod_j [\lambda_j p(\mathbf{x}_i|\boldsymbol{\theta}_j)]^{\zeta_{ij}}. \quad (37)$$

By marginalizing over all possible values of the indicator variables ζ_{ij} , a “mixture of Gaussians” model for the optical parameters is obtained

$$p(\mathbf{x}_i|\boldsymbol{\theta}, \boldsymbol{\lambda}) = \int_{\boldsymbol{\zeta}_i} p(\mathbf{x}_i, \boldsymbol{\zeta}_i|\boldsymbol{\theta}, \boldsymbol{\lambda}) d\boldsymbol{\zeta}_i = \sum_j \lambda_j p(\mathbf{x}_i|\boldsymbol{\theta}_j). \quad (38)$$

Finally, for independent nodes, the prior of the image is given by

$$p(\mathbf{x}|\boldsymbol{\theta}, \boldsymbol{\lambda}) = \prod_i \sum_j \lambda_j p(\mathbf{x}_i|\boldsymbol{\theta}_j). \quad (39)$$

3.1.1 Reconstruction step

The objective function takes the form of Eq. (7), where at iteration t of the reconstruction-classification algorithm, the regularization is given by Eqs. (9) and (39)

$$\begin{aligned} \mathcal{R}^t(\boldsymbol{\mu}_a, \boldsymbol{\mu}'_s) &= -\log p(\mathbf{x}|\boldsymbol{\theta}^t, \boldsymbol{\lambda}^t) = -\log \mathcal{N}(\bar{\mathbf{x}}, \Sigma_{\bar{\mathbf{x}}}) \\ &= \frac{\tau}{2} \|\mathbf{L}_{\bar{\mathbf{x}}}(\mathbf{x} - \bar{\mathbf{x}})\|^2, \end{aligned} \quad (40)$$

where τ is a regularization parameter and

$$\bar{\mathbf{x}}_i = \sum_j \zeta_{ij} \cdot \mathbf{m}_j \Big|_{\text{MAP}(\boldsymbol{\zeta})} = \mathbf{m}_{j'} \in \mathbb{R}^2 \quad (41)$$

is obtained by fixing the labels to the “maximum a posteriori” estimate, given the results of the previous iteration

$$\text{MAP}(\boldsymbol{\zeta}) = \arg \max_{\boldsymbol{\zeta}} p(\boldsymbol{\zeta}|\mathbf{x}^{t-1}, \boldsymbol{\theta}^{t-1}, \boldsymbol{\lambda}^{t-1}), \quad (42)$$

which is calculated in the classification step (see Sec. 3.1.2). The weighting matrix $\mathbf{L}_{\bar{\mathbf{x}}}$ is the Cholesky decomposition of $\Sigma_{\bar{\mathbf{x}}}^{-1}$, where $\Sigma_{\bar{\mathbf{x}}} \in \mathbb{R}^{2N \times 2N}$ is a sparse matrix of which the i 'th 2×2 block along the diagonal is $\Sigma_{j'}$ if the i 'th element belongs to the j 'th class.

In order to sphere the solution space, that is, to render the space dimensionless, we performed a change of variables $\boldsymbol{\mu}_a \rightarrow \boldsymbol{\mu}_a/\mu_{a0}$ and $\boldsymbol{\mu}'_s \rightarrow \boldsymbol{\mu}'_s/\mu'_{s0}$, where (μ_{a0}, μ'_{s0}) is the initial guess for the optical parameters (in this study, we initialized to the homogeneous background). Given the size of the problem, we chose a gradient-based optimization method in order to reduce memory use and computational expense.¹² The minimization was performed using the limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) method,¹³ with a storage memory of six iterations.

3.1.2 Classification

The purpose of the classification step is to update the multinomial model using the result of the previous reconstruction step. First, the expected values of the labels $\boldsymbol{\zeta}^{t+1}$ are computed for the current class parameters $(\boldsymbol{\theta}^t, \boldsymbol{\lambda}^t)$ and image $\mathbf{x}^t = (\boldsymbol{\mu}_a^t, \boldsymbol{\mu}'_s{}^t)$ (E-step). Then the model parameters are updated by maximizing the posterior probability (M-step)

$$p(\boldsymbol{\theta}, \boldsymbol{\lambda}|\mathbf{x}^t) \propto p(\mathbf{x}^t|\boldsymbol{\theta}, \boldsymbol{\lambda})p(\boldsymbol{\theta}, \boldsymbol{\lambda}). \quad (43)$$

E-step: The “responsibility” r_{ij}^t is a measure of the probability that the i 'th node is assigned to the j 'th class. Using Bayes' theorem and the Gaussian mixture model [Eq. (38)], we have

$$\begin{aligned} p(\zeta_{ij} = 1|\mathbf{x}_i^t, \boldsymbol{\theta}^t, \boldsymbol{\lambda}^t) &= \frac{p(\mathbf{x}_i|\zeta_{ij} = 1, \boldsymbol{\theta}^t)p(\zeta_{ij} = 1)}{p(\mathbf{x}_i|\boldsymbol{\theta}, \boldsymbol{\lambda})} \\ &= \frac{\lambda_j^t p(\mathbf{x}_i^t|\boldsymbol{\theta}_j^t)}{\sum_j \lambda_j^t p(\mathbf{x}_i^t|\boldsymbol{\theta}_j^t)} = r_{nj}^t. \end{aligned} \quad (44)$$

The expectation for the indicator values is

$$\begin{aligned} E(\zeta_{ij}|\mathbf{x}_i^t, \boldsymbol{\theta}^t, \boldsymbol{\lambda}^t) &= \int \zeta_{ij} p(\zeta_{ij} = 1|\mathbf{x}_i^t, \boldsymbol{\theta}^t, \boldsymbol{\lambda}^t) d\zeta_{ij} \\ &= 0 \times p(\zeta_{ij} = 0|\mathbf{x}_i^t, \boldsymbol{\theta}^t, \boldsymbol{\lambda}^t) + 1 \times p(\zeta_{ij} = 1|\mathbf{x}_i^t, \boldsymbol{\theta}^t, \boldsymbol{\lambda}^t) \\ &= r_{ij}^t. \end{aligned} \quad (45)$$

Therefore, the MAP estimate for the labels is

$$\zeta_{ij}^{t+1} = \begin{cases} 1 & \text{if } r_{ij}^t \text{ is maximum } \forall j, \\ 0 & \text{otherwise,} \end{cases} \quad (46)$$

which can be used in Eq. (42).

M-step: The parameters (θ, λ) are chosen in order to maximize the log posterior

$$(\theta^{t+1}, \lambda^{t+1}) = \arg \max_{(\theta, \lambda)} \log p(\mathbf{x}^t | \theta, \lambda) + \log p(\theta, \lambda). \quad (47)$$

Averaging over all possible values of ζ gives

$$\log p(\mathbf{x}^t | \theta, \lambda) + \log p(\theta, \lambda) = \int_{\zeta} \log p(\mathbf{x}^t, \zeta | \theta, \lambda) d\zeta + \log p(\theta, \lambda). \quad (48)$$

Using “Jensen’s inequality”¹⁴ and ignoring terms that do not depend on (θ, λ) , we obtain a lower bound for the log prior

$$\begin{aligned} \mathcal{B}(\theta, \lambda) &= \sum_i \sum_j r_{ij}^t \log[\lambda_j p(\sigma_n | \theta_j)] + \log p(\lambda) + \log p(\theta) \\ &= \sum_i \sum_j r_{ij}^t \left[\log(\lambda_j) + \log(|\Sigma_j|) \right. \\ &\quad \left. - \frac{1}{2} (\mathbf{x}_{i(n)} - \mathbf{m}_j)' \Sigma_j^{-1} (\mathbf{x}_{i(n)} - \mathbf{m}_j) \right] \\ &\quad + \sum_j \left[(\alpha_j - 1) \log(\lambda_j) - \frac{\nu_j + d + 1}{2} \log|\Sigma_j| \right]. \end{aligned} \quad (49)$$

Maximizing $\mathcal{B}(\theta, \lambda)$ for $\sum_j \lambda_j = 1$ and using noninformative priors, we obtain the update rules for the model parameters

$$\lambda_j^{t+1} = \frac{\sum_i r_{ij}^t}{N}, \quad (50)$$

$$\mathbf{m}_j^{t+1} = \frac{\sum_i r_{ij}^t \mathbf{x}_i}{\sum_i r_{ij}^t}, \quad (51)$$

$$\Sigma_j^{t+1} = \frac{\sum_i r_{ij}^t (\mathbf{x}_i - \mathbf{m}_j)(\mathbf{x}_i - \mathbf{m}_j)^T + \Gamma_j}{\sum_i r_{ij}^t + \nu_j + d + 1}. \quad (52)$$

3.2 Class Means Initialization

The number of classes J and the class means $\mathbf{m}_j$ were initialized by automatically segmenting the result of the first reconstruction step and averaging over the segmented areas. To segment the image [e.g., see Fig. 2(a)], we looked at a binned histogram of the image of μ_a and chose the value μ_{a_h} for which the number of occurrences was highest [Fig. 2(c), column 1]. We found the first node index h for which the value μ_{a_h} occurs, and identified the corresponding scattering value μ_{s_h}' . Having chosen a covariance matrix Σ_h , we computed a map of the multivariate normal probability of the (μ_a, μ_s') images, with mean (μ_{a_h}, μ_{s_h}') [Fig. 2(c), column 2]. A suitable choice for Σ_h is the initial covariance of the classes. Then we selected a tolerance level

tol_h at which to truncate the probability map, and selected all nodes with probability higher than the tolerance as belonging to the same class as node h [Fig. 2(c), column 3]. We repeated this process on the remaining nodes until all nodes were classified. Thus, the number of classes was set to the number of iterations, and the average of the optical parameters over each class was used to initialize the class means [Fig. 2(b)].

3.3 Visualization of the Results

Results obtained using the reconstruction-classification method are displayed alongside scatter plots of the nodal values recovered in the two-dimensional (2-D) feature space (μ_a, μ_s') [e.g., see Fig. 2(c), final column in 4]. The positions of the class means $\mathbf{m}_j = (\bar{\mu}_{a_j}, \bar{\mu}_{s_j}')$ are identified by a cross, and the class covariances Σ_j are represented by ellipses. These are color-coded by class, and are indicative of the clustering of image nodal values around the class means.

4 Results

4.1 Two-Dimensional Validation and Reconstruction

We chose a numerical phantom defined on a 2-D circular mesh with 1331 nodes and radius 25 mm. Four illumination sources were placed on the boundary at angles 0, $\pi/2$, π , and $3\pi/2$ rad. In all cases, the illumination profile was a normalized Gaussian with radius (distance from the center at which the profile drops to $1/e$) 6 mm. The background optical parameters were set to $\mu_a = 0.01 \text{ mm}^{-1}$ and $\mu_s' = 1 \text{ mm}^{-1}$. Two circular perturbations of radius 6 mm were added in positions (6 mm, 10 mm) and (−6 mm, −10 mm) [Fig. 3(a)]. The values of the perturbations were $\mu_a = 0.02 \text{ mm}^{-1}$, $\mu_s' = 1.5 \text{ mm}^{-1}$ and $\mu_a = 0.03 \text{ mm}^{-1}$, $\mu_s' = 1.25 \text{ mm}^{-1}$, respectively. The absorbed energy field was simulated for each illumination, and 1% white Gaussian noise was added [Fig. 3(b)]. The class covariances were initialized to

$$\Sigma_j = \begin{pmatrix} 10^{-6} & 0 \\ 0 & 10^{-1} \end{pmatrix} \quad \forall j = 1, \dots, 3, \quad (53)$$

where the first variable was the absorption and the second was the reduced scattering. The parameters of the Jeffreys prior were set to $\Gamma_j = \Sigma_j \forall j$, $\nu(1) = 1$ for the background class and $\nu(2,3) = 10$ for the perturbation classes. The number of classes and optical parameters was initialized using the class means initialization method (Sec. 3.2) with $\text{tol}_h = 10^{-5}$ and $\Sigma_h = \Sigma_j$ [Eq. (53)], and the labels were initialized to 1 for the background class and zero for all other classes. The tolerance of the L-BFGS algorithm was set to $\text{tol} = 10^{-11}$, and the total number of reconstruction-classification iterations was set to $\text{MaxIt} = 10$ (Fig. 4). The regularization parameter $\tau = 10^{-10}$ was chosen by inspection. For comparison, images were reconstructed without introducing a prior (Fig. 5); the images were reconstructed by minimizing Eq. (6) using the L-BFGS method with $\text{tol} = 10^{-12}$.

4.2 Three-Dimensional Validation and Reconstruction

We chose a three-dimensional (3-D) phantom analogous to the 2-D case, defined on a cylinder with 27,084 nodes, radius 25 mm, and height 25 mm. Two spherical inclusions of radius 6 mm were placed in (6, 10, and 0 mm) and (−6, −10, and

```

Set MaxIt, tol
Initialize optical and class parameters  $\mathbf{x}, \zeta, \theta, \lambda$ 
Initialize iteration count  $t = 1$  and regularization term  $\mathcal{R}^0 = 0$  (no regularization)
repeat
  Reconstruction:
    Update  $\mathbf{x}$ ; minimize (7) using L-BFGS until  $\mathcal{E} < \text{tol}$ 
    if  $t = 1$  then
      Initialize class means  $\mathbf{m}_j$  (section 3.2)
    end if
  Classification:
    E-step; compute expected labels  $\zeta$  (44)
    M-step; update class parameters  $(\theta, \lambda)$  (50)(51)(52)
    Update regularization term  $\mathcal{R}^t$  (40)
    Re-initialize  $\mathbf{x} = \bar{\mathbf{x}}$  (41)
     $t \leftarrow t + 1$ 
until  $t \geq \text{MaxIt}$ 

```

Fig. 1 Reconstruction-classification algorithm outline.

0 mm) [Fig. 6(a)]. Illumination sources were Gaussian in the xy -plane constant in the z -axis, with radius 6 mm and length 25 mm [Figs. 6(b) and 6(c)]. PAT images were simulated for four illuminations at the cardinal points, and 1% noise was added to the absorbed energy [Fig. 6(d)]. The optical, covariance, and reconstruction parameters were set to the same values used in the 2-D case. The class initialization parameters were set to $\text{tol}_h = 10^{-7}$ and $\Sigma_h = \Sigma_j$. Images were reconstructed by performing 10 iterations of the reconstruction-classification method (Fig. 7).

5 Discussion

5.1 Summary of Findings

We applied the proposed reconstruction-classification algorithm to a 2-D numerical phantom with three tissues, a background, and two perturbations (Fig. 3). The optical absorption was recovered reliably within a small number of iterations, and the scattering was recovered with sufficient accuracy after approximately 10 iterations (Fig. 4). We compared the optical

Fig. 2 Class initialization example: (a) original image of μ_a to which we apply the segmentation; (b) result of taking average image values over the segmented areas; (c) first column, histogram of occurrences of values of μ_a in the portion of the image requiring segmentation—value with highest number of occurrences is μ_{ah} (indicated by a red cross); second column, probability density function with mean (μ_{ah}, μ'_{sh}) and covariance Σ_h ; third column, labels identifying nodes with probability density higher than tolerance value tol_h ; each row corresponds to an iteration and a distinct class, so in this case, $J = 3$.

Fig. 3 Two-dimensional (2-D) model: (a) circular mesh and (b) absorbed energy for each illumination pattern.

model with images obtained by the reconstruction-classification method and by a traditional reconstruction-only (no regularization) method (Fig. 5). We found that the reconstruction-classification method delivered superior image quality, particularly with regards to the scattering parameter. We applied the reconstruction-classification algorithm to a much larger 3-D problem (Fig. 6) and observed similar results (Fig. 7) as in the 2-D case.

5.2 Choice of Parameters

The parametric optical model and classification algorithm introduce a number of parameters that require tuning by the user. In

addition to the regularization parameter, the parameters of the Jeffreys prior Γ and ν and the initial guess of the class variances Σ_j must be set before performing the classification. However, their significance is fairly intuitive, and with experience of a certain type of problem, the choice of parameters becomes natural. Visualizing the class covariance matrix Σ_j as an ellipse, changing the value of Γ varies its eccentricity, and changing ν varies the length of its axes. Further, given that in the first iteration the optical absorption is recovered with superior accuracy to the scattering, it is preferable to initialize the variance of the former to a smaller value than the latter, indicating greater confidence in the imaging solution.

Fig. 4 2-D reconstruction-classification results at iteration 1 (first row), 5 (second row), and 10 (third row). Reconstructed values of μ_a' and μ_s' (first and second columns), labels recovered for perturbation classes (third and fourth columns), and scatter plot (fifth column).

Fig. 5 2-D model and reconstruction: first column, model of μ_s and μ'_s ; second column, reconstructed values of μ_a and μ'_s without multinomial prior; third column, reconstructed values of μ_a and μ'_s with multinomial prior.

5.3 Initialization of the Class Means

The purpose of the means initialization scheme is to increase automation of the method so that minimum user intervention and no prior knowledge of the number of tissues or their optical properties is required. The algorithm simply performs a segmentation of the image, then takes averages over the segmented areas to initialize the class properties (Fig. 1). Alternative segmentation techniques could have been employed; however, the advantage of the proposed approach is that it directly exploits

the mixture of Gaussians model to identify the tissues. Our choice to investigate a node h with μ_a belonging to the bin with a maximum number of occurrences leads to the background tissue being identified first, followed by the perturbation tissues. The choice of the node index h could have been randomized so that tissues would be identified in random order. This approach is equally valid; however, we found that in cases where tissue values were close together (such as after a single reconstruction-classification iteration), it was preferable to identify the largest classes first because the mean was estimated with

Fig. 6 Three-dimensional (3-D) model: (a) numerical phantom and perturbation locations, (b) all illumination sources, (c) cross section of optical parameters used to simulate the data for $z = 0$, (d) cross section of absorbed energy for each illumination pattern.

Fig. 7 3-D reconstruction-classification results at iteration 1 (first row), 5 (second row), and 10 (third row). Reconstructed values of μ_a and μ'_s (first and second columns), labels recovered for perturbation classes (third column), and scatter plot (fourth column).

greater accuracy for the classes with a larger number of samples. Further, for a given image and tolerance level, our choice renders the result of the segmentation process unique and reproducible.

5.4 Recovery of the Scattering

From the comparison with the reconstruction-only case with no regularization (Fig. 5), it is evident that the introduction of the parametric prior enables better recovery of the scattering. The inconsistency between the quality of the recovered absorption and scattering parameters in the nonregularized case is due to the weaker dependence of the latter on the absorbed energy density with respect to the former. This results in the scattering gradient being approximately an order of magnitude smaller than the absorption gradient. Although the problem can be mitigated by sphering the solution space, variations in the data due to the scattering often fall below the noise floor. In the reconstruction-classification case, typically the absorption is recovered with good accuracy within a small number of iterations. Thus, the absorption takes values very close to the class means (resulting in small clusters), and the variance along the μ_a direction converges to a small value. Given that the regularization term is weighted by the inverse of the covariance matrix, the dependence of the absorption gradient on the data becomes weaker at each iteration, until its magnitude is comparable or smaller to

that of the scattering. In the iterations that follow, the descent of the data term of the objective function is primarily due to updates to the scattering, which converges to the correct values.

5.5 Computational Demands

Computational performance was found to be strongly dependent on the problem size. In the 2-D case with 1331 nodes (Fig. 4), the total reconstruction time (10 outer reconstruction-classification iterations) using MATLAB on a 16-processor PC with 128 GB RAM was only 77 s. In the 3-D case with 27,084 nodes (Fig. 7), the total reconstruction time increased linearly with the number of nodes and was approximately 3.7 h on the same workstation. The increase in computation time was mostly due to much longer processing times for the L-BFGS algorithm in the reconstruction step.

5.6 Experimental Application

In experimental situations, prior information on tissue properties may be held, such as knowledge of the characteristic optical absorption and scattering spectra of chromophores of interest. These may be obtained from the literature¹⁵ or gained through tissue sample measurements. This information could be used in one of two ways. First, a library of typical chromophores could be used to initialize the class parameters instead of the proposed

class means initialization method. The classification process could then perform the function of correcting for uncertainty, errors, or local variations in the real optical properties with respect to the prior information. Alternatively, it could be used to label the chromophores found by the segmentation process and identify these as certain tissues such as, e.g., “oxygenated blood” or “fat,” on the basis of the closeness of the recovered means to the characteristic properties.

5.7 Additional Priors

In this study, we assumed independence between nodal values; however, the mixture of Gaussian models could be used in conjunction with a spatial prior. Knowledge of smoothness or sparsity properties of the solution could be employed to introduce a homogeneous spatial regularizer such as first-order Tikhonov¹⁶ or total variation.^{7,17} Knowledge of structural information, such as that provided by an alternative imaging method or anatomical library, could be exploited by introducing a spatially varying probability map for the optical properties.

6 Conclusions

In this paper, we proposed a method for performing image reconstruction in QPAT. We introduced a parametric class model for the optical parameters and implemented a minimization-based reconstruction algorithm. We suggested an automated method by which to initialize the parameters of the class model and proposed a classification algorithm by which to progressively update and improve those parameters after each reconstruction step. We demonstrated through 2-D and 3-D numerical examples that the reconstruction-classification method allows for the simultaneous recovery of optical absorption and scattering. In particular, we found that this approach delivered superior accuracy in the recovery of the scattering with respect to traditional gradient-based reconstruction.

Acknowledgments

This work was funded by the EPSRC Doctoral Prize Fellowship EP/M506448/1.

References

1. P. Beard, “Biomedical photoacoustic imaging,” *Interface Focus* **1**, 602–631 (2011).
2. B. Cox et al., “Quantitative spectroscopic photoacoustic imaging: a review,” *J. Biomed. Opt.* **17**, 061202 (2012).
3. H. Gao, S. Osher, and H. Zhao, “Quantitative photoacoustic tomography,” *Lect. Notes Math.* **2035**, 131–158 (2012).
4. P. Hiltunen, S. J. D. Prince, and S. Arridge, “A combined reconstruction-classification method for diffuse optical tomography,” *Phys. Med. Biol.* **54**, 6457–6476 (2009).
5. E. Malone et al., “A reconstruction-classification method for multifrequency electrical impedance tomography,” *IEEE Trans. Med. Imaging* **34**(7), 1486–1497 (2015).
6. S. Arridge, “Optical tomography in medical imaging,” *Inverse Prob.* **15**, R41 (1999).
7. G. Bal and K. Ren, “Multiple-source quantitative photoacoustic tomography in a diffusive regime,” *Inverse Prob.* **27**(7), (2011).
8. P. Shao, B. Cox, and R. J. Zemp, “Estimating optical absorption, scattering, and Grueneisen distributions with multiple-illumination photoacoustic tomography,” *Appl. Opt.* **50**(19), 3145–3154 (2011).
9. H. Gao, J. Feng, and L. Song, “Limited-view multi-source quantitative photoacoustic tomography,” *Inverse Prob.* **31**(6), 065004 (2015).
10. B. T. Cox, S. R. Arridge, and P. C. Beard, “Gradient-based quantitative photoacoustic image reconstruction for molecular imaging,” *Proc. SPIE* **6437**, 64371T (2007).
11. M. Schweiger and S. Arridge, “The Toast++ software suite for forward and inverse modeling in optical tomography,” *J. Biomed. Opt.* **19**, 040801 (2014).
12. T. Saratoon et al., “3D quantitative photoacoustic tomography using the δ -Eddington approximation,” *Proc. SPIE* **8581**, 85810V (2013).
13. J. Nocedal and S. Wright, “Numerical optimization,” in *Springer Series in Operations Research and Financial Engineering*, Springer-Verlag, New York (1999).
14. S. Prince, *Computer Vision: Models, Learning, and Inference*, Cambridge University Press, Cambridge, United Kingdom (2012).
15. S. L. Jacques, “Optical properties of biological tissues: a review,” *Phys. Med. Biol.* **58**, 5007–5008 (2013).
16. T. Saratoon et al., “A gradient-based method for quantitative photoacoustic tomography using the radiative transfer equation,” *Inverse Prob.* **29**, 075006 (2013).
17. T. Tarvainen et al., “Reconstructing absorption and scattering distributions in quantitative photoacoustic tomography,” *Inverse Prob.* **28**, 084009 (2012).

Biographies for the authors are not available.