Multispectral reconstruction methods for quantitative photoacoustic tomography

Emma Malone[a], Ben Cox[a], and Simon Arridge[b]

[a]University College London, Department of Medical Physics and Biomedical Engineering, Gower Street, London, UK
[b]University College London, Department of Computer Science, Gower Street, London, UK

ABSTRACT

We propose a novel multispectral reconstruction-classification method for simultaneously recovering absorption and scattering coefficients from images of absorbed optical energy. In contrast with pre-existing chromophore reconstruction methods, this approach does not require prior knowledge of the characteristic spectra of the absorbers, which is not always available. Numerical experiments performed on anatomically realistic 3D phantoms show that this approach allows for improved recovery of both the optical absorption and scattering with respect to reconstruction-only methods, and accurate classification of chromophores of clinical interest.

Keywords: quantitative photoacoustic tomography, image reconstruction, chromophores

1. INTRODUCTION

1.1 Quantitative photoacoustic imaging

Quantitative photoacoustic tomography (QPAT) aims to recover the optical absorption and scattering coefficients of biological tissues in order to obtain clinically relevant information about tissue morphology and functionality.[1,2] Conventional photoacoustic (PA) images are obtained from measurements of the acoustic waves resulting from the rapid heating of tissue by a laser pulse, this is known as the *photoacoustic effect*.[3] From these measurements, an image of the initial acoustic pressure distribution p_0 is recovered, which in turn is related to the absorbed optical energy density H by the Grüneisen parameter $\hat{\Gamma}$:

$$H(\boldsymbol{r}) = \hat{\Gamma} p_0(\boldsymbol{r}) \qquad \boldsymbol{r} \in \Omega. \tag{1}$$

The photoacoustic imaging problem is well understood, and several methods exist for recovering PA images, such as FFT,[4] time-reversal,[5,6] model-based,[7–9] and sparsity-based methods.[10,11] However, PA images are not directly representative of the underlying tissue structure because spatial variations in the illumination pattern cause the optical energy to be absorbed non-uniformly throughout the tissue. If $\phi(\boldsymbol{r})$ is the optical fluence generated by the laser within the tissue, then we have

$$H(\boldsymbol{r}) = \hat{\Gamma} \mu_a(\boldsymbol{r}) \phi(\boldsymbol{r}) \qquad \boldsymbol{r} \in \Omega, \tag{2}$$

where μ_a is the absorption coefficient. The *quantitative* photoacoustic problem, is to recover an image of the optical absorption $\mu_a(\boldsymbol{r})$ (which is directly related to tissue morphology) from estimates of the absorbed energy density $H(\boldsymbol{r})$ (which is not). The problem is non-linear because the fluence in turn depends on the optical parameters $\phi(\mu_a(\boldsymbol{r}), \mu_s'(\boldsymbol{r}))$, where $\mu_s'(\boldsymbol{r})$ is the reduced scattering coefficient.

QPAT can provide clinically valuable images of endogenous chromophores, such as oxy- or deoxy-hemoglobin, melanin, lipids and water. Exogenous contrast agents, enzymes, or proteins linked to the expression of a gene of interest, may also be imaged.

Further author information: (Send correspondence to E.M.)
E.M.: E-mail: e.malone@ucl.ac.uk, Telephone: +44 (0)20 7679 0203

Photons Plus Ultrasound: Imaging and Sensing 2016, edited by Alexander A. Oraevsky, Lihong V. Wang
Proc. of SPIE Vol. 9708, 970827 · © 2016 SPIE · CCC code: 1605-7422/16/$18 · doi: 10.1117/12.2212440

1.2 Multispectral reconstruction-classification

We have recently proposed a novel method for reconstructing QPAT images via a reconstruction-classification algorithm.[12] This method exploits knowledge that optical parameters are determined by a limited number of classes to iteratively improve their estimate. Here, we extend the reconstruction-classification method to multi-wavelength measurements, which allows for the identification of individual chromophores. In contrast with pre-existing chromophore reconstruction methods, this approach does not rely on accurate prior knowledge of the characteristic spectra of the absorbers. As such, it is robust to experimental errors committed in measuring chromophore spectra, in addition to unknown local or physiological variations. Instead, a probabilistic model for the tissue properties is initialized and updated automatically on the basis of the data by a classification algorithm. In this paper, we describe the multispectral reconstruction-classification algorithm, present results obtained in numerical experiments using an anatomically realistic 3D phantom, and compare the results to a typical reconstruction-only method.

2. METHODS

2.1 Diffusion model

A model is required to describe light transport inside the tissue, and predict the expected optical fluence. For highly scattering media we can use the *diffusion approximation*[13]

$$(\mu_a - \nabla \cdot \kappa(\boldsymbol{r})\nabla)\,\phi(\boldsymbol{r}) = q(\boldsymbol{r}), \tag{3}$$

where $q(\boldsymbol{r})$ is an isotropic source term, and $\kappa = 1/3\mu_s'$ is the diffusion coefficient. We take the Grüneisen parameter to be constant $\hat{\Gamma} = 1$. In order to treat the image reconstruction problem, we project the physical quantities onto a Finite Element basis $\{u_i(\boldsymbol{r}); i = 1, \ldots, I\}$, and consider the discretisation $(\boldsymbol{\mu_a}, \boldsymbol{\mu_s'}) = \{(\mu_a, \mu_s'); i = 1, \ldots, I\}$. We solve by minimizing an error functional

$$\mathcal{E} = \frac{1}{2}\sum_{i}^{I}\left[d_i^m - \left\langle u_i,\ H(\boldsymbol{\mu_a}, \boldsymbol{\mu_s'})\right\rangle\right]^2 + \mathcal{R}(\boldsymbol{\mu_a}, \boldsymbol{\mu_s'}). \tag{4}$$

where $\mathcal{R}(\boldsymbol{\mu_a}, \boldsymbol{\mu_s'})$ is a regularising function, $H(\boldsymbol{\mu_a}, \boldsymbol{\mu_s'})$ is the absorbed energy density predicted by the diffusion model, and

$$d_i^m = \int_\Omega H^m(\boldsymbol{r})u_i(\boldsymbol{r})\,\mathrm{d}\Omega = \langle u_i,\ H^m\rangle \tag{5}$$

is the projection of the estimated energy density H^m onto the basis function u_i.

If measurements are taken at multiple wavelengths $l = 1, \ldots L$ and illumination patterns $q = 1, \ldots, Q$, then these can be considered simultaneously by taking the sum of the data term

$$\mathcal{E} = \frac{1}{2}\sum_{ilq}\frac{1}{w_l}\left[d_{ilq}^m - \left\langle u_i,\ H_q(\boldsymbol{\mu_{a}}_l, \boldsymbol{\mu_s'}_l)\right\rangle\right]^2 + \mathcal{R}(\boldsymbol{\mu_a}_1, \ldots, \boldsymbol{\mu_a}_L, \boldsymbol{\mu_s'}_1 \ldots, \boldsymbol{\mu_s'}_L). \tag{6}$$

The weighting $w_l = \frac{1}{L}\sum_{iq}\left[d_{ilq}^m - \left\langle u_i,\ H_q({\boldsymbol{\mu_a}_l}^0, {\boldsymbol{\mu_s'}_l}^0)\right\rangle\right]^2$, where $({\boldsymbol{\mu_a}}^0, {\boldsymbol{\mu_s'}}^0)$ is the initial guess of the optical parameters, is chosen so that $\mathcal{E}|_{({\boldsymbol{\mu_a}}^0, {\boldsymbol{\mu_s'}}^0)} = 1$. For gradient calculations for the diffusion model see Malone *et al.*[12]

2.2 Spectral model for the optical parameters

In the following, we propose a method to aid the recovery of the absorption and scattering coefficients by assuming a spectral model for the optical parameters . We build a probabilistic model, which links the optical properties of the object of interest to the characteristic spectra of its components. It is assumed that each chromophore (such as oxy- and deoxy-hemoglobin, fat, water etc.) has a unique optical spectrum. For a given chromophore j,

the absorption and scattering are drawn from a multinomial Gaussian distribution $\mathcal{N}(\boldsymbol{m}_j, \boldsymbol{\Sigma}_j)$, where $\boldsymbol{m}_j \in \mathbb{R}^{2L}$, $\boldsymbol{\Sigma}_j \in \mathbb{R}^{2L\times 2L}$, and L is the number of wavelengths. We define the auxiliary labels

$$\zeta_{ij} = \begin{cases} 1 & \text{if the } i\text{th node is assigned to the } j\text{th class;} \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

for nodes $i = 1\ldots, I$ and chromophores $j = 1, \ldots, J$. The probability of a set of labels $\boldsymbol{\zeta}_i = \{\zeta_{ij}; \forall j\}$ being assigned to the ith node is given by a multinomial distribution $\mathrm{p}(\boldsymbol{\zeta}_i|\boldsymbol{\lambda}) = \prod_j \lambda_j^{\zeta_{ij}}$, where λ_j is the overall probability that a node is assigned to the jth class. Assuming independence between the nodes, then the joint probability for the optical parameters and labels is given by the product

$$\mathrm{p}(\boldsymbol{x}, \boldsymbol{\zeta}|\boldsymbol{m}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}) = \prod_i \mathrm{p}(\boldsymbol{x}_i|\boldsymbol{\zeta}_i, \boldsymbol{m}, \boldsymbol{\Sigma})\mathrm{p}(\boldsymbol{\zeta}_i|\boldsymbol{\lambda}) = \prod_{ij} [\lambda_j \mathrm{p}(\boldsymbol{x}_i|\boldsymbol{m}_j, \boldsymbol{\Sigma}_j)]^{\zeta_{ij}}. \tag{8}$$

where $\boldsymbol{x}_i = \{(\mu_{ail}, \mu'_{sil}); \forall l\} \in \mathbb{R}^{2L}$ and $\mathrm{p}(\boldsymbol{x}_i|\boldsymbol{m}_j, \boldsymbol{\Sigma}_j) \propto \mathcal{N}(\boldsymbol{m}_j, \boldsymbol{\Sigma}_j)$. Finally, by marginalizing over all possible values of the labels, a *mixture of Gaussians* model for the optical parameters is obtained

$$\mathrm{p}(\boldsymbol{x}|\boldsymbol{m}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}) = \int_{\boldsymbol{\zeta}} \mathrm{p}(\boldsymbol{x}, \boldsymbol{\zeta}|\boldsymbol{m}, \boldsymbol{\Sigma}, \boldsymbol{\lambda})\mathrm{d}\boldsymbol{\zeta} = \prod_i \sum_j \lambda_j \mathrm{p}(\boldsymbol{x}_i|\boldsymbol{m}_j, \boldsymbol{\Sigma}_j). \tag{9}$$

2.3 Reconstruction-classification method

In the Bayesian framework, a spectral regularization term for the image reconstruction problem is obtained by taking

$$\mathcal{R}(\boldsymbol{\mu_a}, \boldsymbol{\mu'_s}) = -\log \mathrm{p}(\boldsymbol{x}|\boldsymbol{m}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}) = -\log \mathcal{N}(\bar{\boldsymbol{x}}, \Sigma_{\bar{x}}) = \frac{\tau}{2}\|\mathbf{L}_{\bar{x}}(\boldsymbol{x} - \bar{\boldsymbol{x}})\|^2, \tag{10}$$

where τ is the regularization parameter; $\bar{\boldsymbol{x}} \in \mathbb{R}^{2LN}$ indicates for each node the mean of the class assigned by a certain labelling $\bar{\boldsymbol{\zeta}}$; and $\mathbf{L}_{\bar{x}} \in \mathbb{R}^{2LN\times 2LN}$ contains a $2L\times 2L$ block for each node, indicating the assigned covariance matrix. Note that it is necessary in this case to consider multi-wavelength measurements simultaneously because the covariance in the regularisation term may have non-zero off-diagonal values.

We propose to perform alternating reconstruction and classification steps for a fixed number of outer iterations $t = 1, \ldots, T$. In reconstruction step t, the objective function obtained by inserting equation (10) into (6) is minimised using l-BFGS,[14] and an updated set of images $\boldsymbol{x}^t$ is obtained. In classification step t, the labelling $\bar{\boldsymbol{\zeta}}^t$ and class parameters $(\boldsymbol{m}^t, \boldsymbol{\Sigma}^t, \boldsymbol{\lambda}^t)$ are updated on the basis of the new $\boldsymbol{x}^t$, and a new regularization term $\mathcal{R}^{t+1}(\boldsymbol{\mu_a}, \boldsymbol{\mu'_s})$ (to be used in the following reconstruction step) is computed.

An expectation-maximization algorithm is employed to perform *fuzzy k-means*[15] classification and update the class parameters. First, the labelling $\bar{\boldsymbol{\zeta}}^t$ is determined by computing the expected value of the labels

$$\begin{aligned} \int \zeta_{ij}\mathrm{p}(\zeta_{ij} = 1|\boldsymbol{x}_i^t, \boldsymbol{m}^t, \boldsymbol{\Sigma}^t, \boldsymbol{\lambda}^t)\,\mathrm{d}\zeta_{ij} &= 0 \times \mathrm{p}(\zeta_{ij} = 0|\boldsymbol{x}_i^t, \boldsymbol{m}^t, \boldsymbol{\Sigma}^t, \boldsymbol{\lambda}^t) + 1 \times \mathrm{p}(\zeta_{ij} = 1|\boldsymbol{x}_i^t, \boldsymbol{m}^t, \boldsymbol{\Sigma}^t, \boldsymbol{\lambda}^t) \\ &= \frac{\mathrm{p}(\boldsymbol{x}_i|\zeta_{ij} = 1, \boldsymbol{m}^t, \boldsymbol{\Sigma}^t)\mathrm{p}(\zeta_{ij} = 1)}{\mathrm{p}(\boldsymbol{x}_i|\boldsymbol{\theta}, \boldsymbol{\lambda})} \\ &= \frac{\lambda_j^t \mathrm{p}(\boldsymbol{x}_i^t|\boldsymbol{m}_j^t, \boldsymbol{\Sigma}_j^t)}{\sum_j \lambda_j^t \mathrm{p}(\boldsymbol{x}_i^t|\boldsymbol{m}_j^t, \boldsymbol{\Sigma}_j^t)} = r_{ij}^t, \end{aligned} \tag{11}$$

and setting

$$\bar{\zeta}_{ij}^{t+1} = \begin{cases} 1 & \text{if } r_{ij}^t \text{ is maximum } \forall j, \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

This gives the mean of the regularization term $\mathcal{R}^{t+1}(\boldsymbol{\mu_a}, \boldsymbol{\mu'_s})$ (10)

$$\bar{\boldsymbol{x}}_i^{t+1} = \sum_j \bar{\zeta}_{ij}^{t+1} \cdot \boldsymbol{m}_j, \tag{13}$$

and determines the corresponding covariance $\mathbf{L}_{\bar{x}}^{t+1}$. Then, the class parameters are updated by maximising the log posterior

$$(\boldsymbol{m}^{t+1}, \boldsymbol{\Sigma}^{t+1}, \boldsymbol{\lambda}^{t+1}) = \arg \max_{(\boldsymbol{m},\boldsymbol{\Sigma},\boldsymbol{\lambda})} \log \mathrm{p}(\boldsymbol{x}^t|\boldsymbol{m}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}) + \log \mathrm{p}(\boldsymbol{m}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}), \tag{14}$$

which, for non-informative priors, gives the update rules:

$$\lambda_j^{t+1} = \frac{\sum_i r_{ij}^t}{I}, \tag{15}$$

$$\boldsymbol{m}_j^{t+1} = \frac{\sum_i r_{ij}^t \boldsymbol{x}_i^t}{\sum_i r_{ij}^t}, \tag{16}$$

$$\Sigma_j^{t+1} = \frac{\sum_i r_{ij}^t (\boldsymbol{x}_i^t - \boldsymbol{m}_j^{t+1})(\boldsymbol{x}_i^t - \boldsymbol{m}_j^{t+1})^T + \Gamma_j}{\sum_i r_{ij}^t + \nu_j + d + 1}, \tag{17}$$

where d is the dimension of the domain, and ν_j and Γ_j are the parameters of the inverse Wishardt distribution, i.e. the prior of the covariance. For more details on the reconstruction-classification method, see Malone *et al.*[12,16] and Hiltunen *et al.*[17]

2.4 Initialising the class parameters

To avoid reliance on prior knowledge of the chromophores, we propose an automated process by which to initialize the classification algorithm. The mean optical parameters for each class are initialized by segmenting the first reconstructed image for the absorption coefficient at each wavelength. Given the difficulty in recovering the scattering, we ignore the images of μ_s at this stage. First, we inspect a histogram of the values of $\boldsymbol{\mu_a}$, and find the most common occurrence $\boldsymbol{\mu_a}_h$. Then we select a covariance matrix $\boldsymbol{\Sigma}_{\mu_a}$ and compute a Gaussian probability density function pdf $= \mathcal{N}(\boldsymbol{\mu_a}_h, \boldsymbol{\Sigma}_{\mu_a})$ for all values. We set the tolerance level to the value of the pdf at a distance of $3\sqrt{\Sigma_{\mu_a}^{1,1}}$ from the mean along the first dimension, and put any node with pdf greater than the tolerance into the first class. We repeat this process on the remaining nodes until 99.99% have been classified. We set the number of classes to the number of iterations performed. Once the whole image has been segmented, the mean optical parameters of the classes $\boldsymbol{m}_j$ are initialized to the mean value of the corresponding image across each area found, and the covariances are set to

$$\boldsymbol{\Sigma}_j = \begin{pmatrix} \boldsymbol{\Sigma}_{\mu_a} & 0 \\ 0 & \boldsymbol{\Sigma}_{\mu_s} \end{pmatrix} \tag{18}$$

for all classes. The diagonal entries of the scattering covariance $\boldsymbol{\Sigma}_{\mu_s}$ are chosen to be significantly larger than those of $\boldsymbol{\Sigma}_{\mu_a}$ to reflect the higher level of uncertainty in its recovery.

3. RESULTS

3.1 Numerical phantom and data simulation

We created an anatomically realistic numerical phantom from a vascular cast of a rat brain, obtained from a μCT scan. The image was segmented to identify the background and the main vasculature. We subsequently added a tumour inclusion, surrounded by a vascularised region (figure 2). The phantom was mapped to a $87 \times 77 \times 106$ regular grid with resolution of a 200 μm. The total size of the domain was $17.4 \times 15.4 \times 21.2$ mm. We assigned optical properties to the segmented regions based on the values reported by Jacques *et al.*[18] for the absorption and scattering coefficients of brain, oxy- and deoxy-haemoglobin, and water (figure 1). The vessels and tumour were set to 45% full blood and 55% water. Blood oxygenation was set to 90% $\mathrm{Hb0_2}$ to Hb for the vessels, and 60% for the tumour. The average mass concentration of haemoglobin in full blood was taken to be 150 L · g. The background was set to the optical properties of brain tissue. The average mass concentration of haemoglobin in the brain was 3.75% that of full blood, with 60% $\mathrm{Hb0_2}$ and 40% Hb. The scattering was set to that of blood in the vessels and tumour, and brain in the background.

Two illumination sources with Gaussian profile and radius 6 mm were placed in the centre of the yz-plane (East), and the xz-plane (North). Data was simulated for each source and two wavelengths, 650 nm and 800 nm, and 1% white Gaussian noise was added (figure 2).

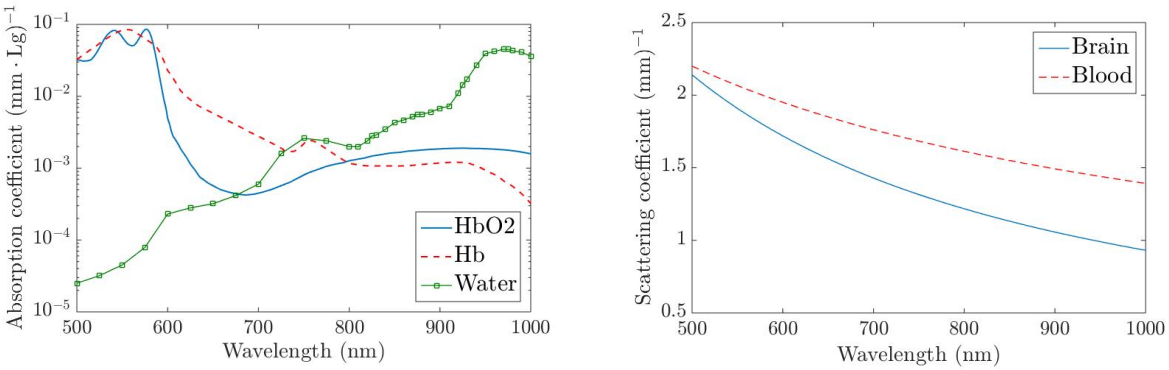

Figure 1. Absorption and scattering spectra of tissues used in numerical phantom.[18]

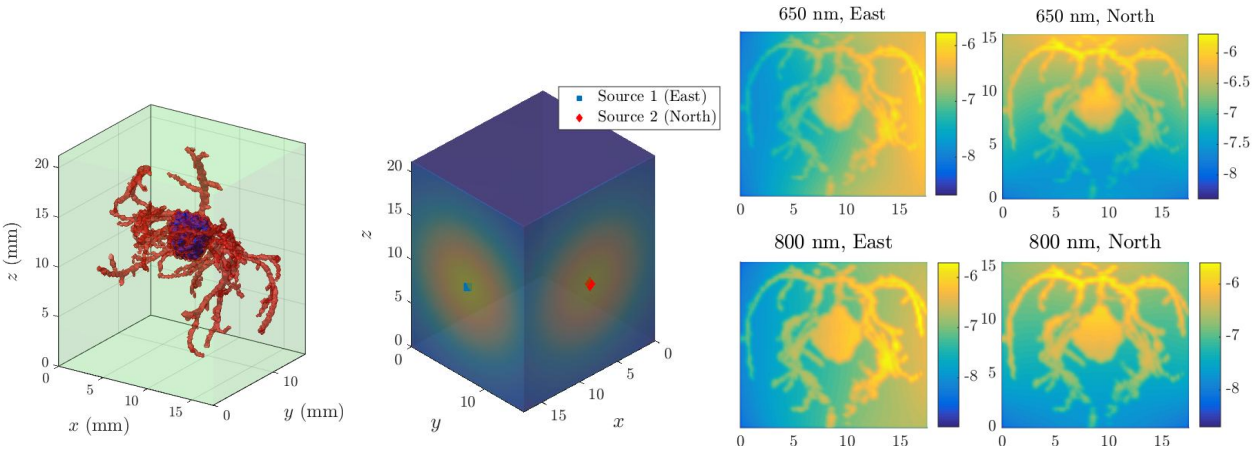

Figure 2. Left: 3D representation of the numerical phantom (blood vessels are marked in red, the tumour is marked in blue). Centre: Boundary illumination profiles for each source position (East and North). Right: Logarithm of the data (absorbed energy density H) simulated for each wavelength and source; maximum intensity projection images along the z-direction.

3.2 Algorithm parameters and class initialization

We applied the L-BFGS-B algorithm to minimize the objective function in the reconstruction step.[19,20] Convergence tolerance was set to $10^{11}\times$ machine precision. We first performed one reconstruction step using no regularization (figure 5, row 1), and then initialized the mean optical parameters for each class as described in section 2.4 (figure 3). Three classes were found, and the covariances were initialised to

$$\mathbf{\Sigma}_j = \begin{pmatrix} 10^{-4} & 0 & 0 & 0 \\ 0 & 5\cdot 10^{-4} & 0 & 0 \\ 0 & 0 & 10^{-1} & 0 \\ 0 & 0 & 0 & 10^{-1} \end{pmatrix} \begin{matrix} 650\text{ nm} \\ 800\text{ nm} \\ 650\text{ nm} \\ 800\text{ nm} \end{matrix} \begin{matrix} \} \mu_a \\ \\ \} \mu_s' \\ \\ \end{matrix} \tag{19}$$

for all classes $j = \{1, 2, 3\}$. The parameters of the inverse Wishardt distribution (equation (17)) were set to $\nu_1 = 1$ for the backround, $\nu_{2,3} = 10^3$ for the vessels and tumour, and $\Gamma_j = 10^3\mathbf{\Sigma}_j\ \forall j$. The regularitazion parameter (equation (10)) was chosen by inspection of logarithmically spaced values (powers of 10), and set to $\tau = 10^{-11}$.

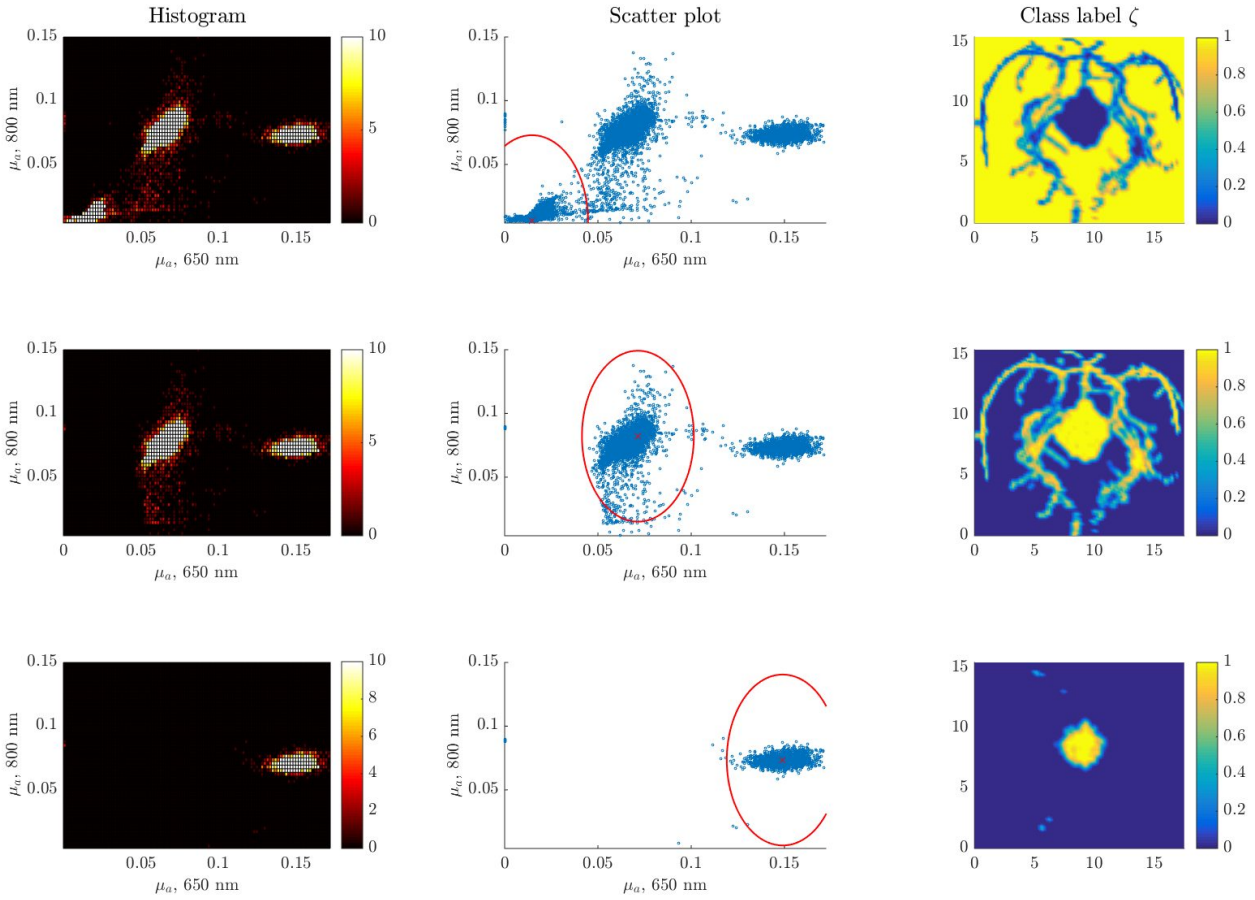

Figure 3. Initialization of the class means for the classification algorithm. Each row corresponds to an iteration of the initialization process 2.4. Left: 2D histogram of values of μ_a recovered for each wavelength after one reconstruction iteration. Centre: scatter plot showing nodes included in the class inside red line, and mean of the class indicated by a cross. Right: result of the image segmentation, node i assigned to the class j is given $\zeta_{ij} = 1$, and 0 otherwise; maximum intensity projection images along the z-direction.

3.3 Reconstruction-classification results

The image error was calculated by taking the norm of the difference between the reconstructed and simulated (figure 4) optical properties

$$\text{ReconErr}^t_{\mu_a} = \frac{\sum_l \left\| \boldsymbol{\mu_a}_l^t - \boldsymbol{\mu_a}_l^{\text{model}} \right\|}{\sum_l \left\| \boldsymbol{\mu_a}_l^{\text{model}} \right\|}, \tag{20}$$

$$\text{ReconErr}^t_{\mu_s'} = \frac{\sum_l \left\| \boldsymbol{\mu_s'}_l^t - \boldsymbol{\mu_s'}_l^{\text{model}} \right\|}{\sum_l \left\| \boldsymbol{\mu_s'}_l^{\text{model}} \right\|}. \tag{21}$$

We also considered a classification error

$$\text{ClassErr}^t = \frac{\sum_i^I \bar{\boldsymbol{\zeta}}_i^t \cdot \boldsymbol{\zeta}_i^{model}}{I}, \tag{22}$$

that is the ratio of misclassified nodes.

We performed 10 iterations of reconstruction-classification steps The result obtained without using a spectral model is displayed in the first row (iteration 1) in figure 5. The final result (iteration 10) is displayed in the third row. The results of the classification step are displayed in the fifth column. All the measures of image error decreased over the iterations (figure 6).

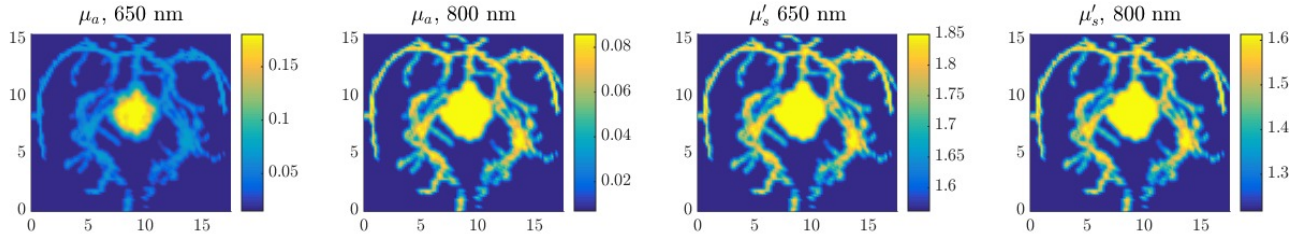

Figure 4. Model of optical properties at simulated wavelengths; maximum intensity projection images along z-direction.

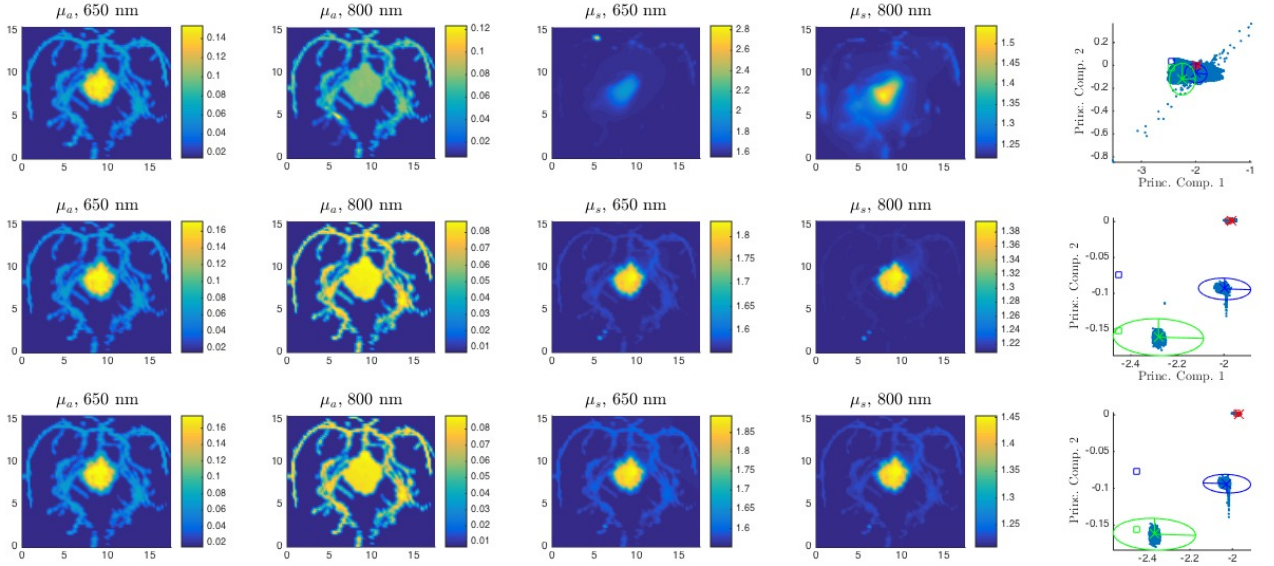

Figure 5. Reconstruction-classification results after 1 (row 1), 5 (row 2), and 10 iterations (row 3). Columns 1-4: Rreconstructed values of μ_a and μ_s' for each wavelength (columns 1-4). Column 5: scatter plot of projections onto first and second principle components of $\boldsymbol{x}$ (column 5). The crosses indicates the position of the mean of each class $\boldsymbol{m}_j$, and the ellipses represent the class covariances $\boldsymbol{\Sigma}_j$ $\forall j = \{1, 2, 3\}$ (9). The squares indicate the position of the simulated values in the same coordinates.

4. CONCLUSIONS

We have proposed a novel method to identify and localize chromophores from photoacoustic images acquired at multiple wavelengths. The method employs a probabilistic spectral model, which is updated iteratively as the recovery of the optical parameters improves. The advantages of this technique are twofold. First, this method does not rely on prior knowledge of the absorption spectra of the chromophores. Second, it allows for the simultaneous recovery of the optical scattering together with the absorption. We have demonstrated the application of our reconstruction-classification method using a realistic numerical phantom, and compared the results to those obtained without employing a spectral model in terms of measurable image error criteria. We have shown that the proposed reconstruction classification method can deliver improved recovery of both the optical absorption and scattering, and the identification of distinct chromophore classes.

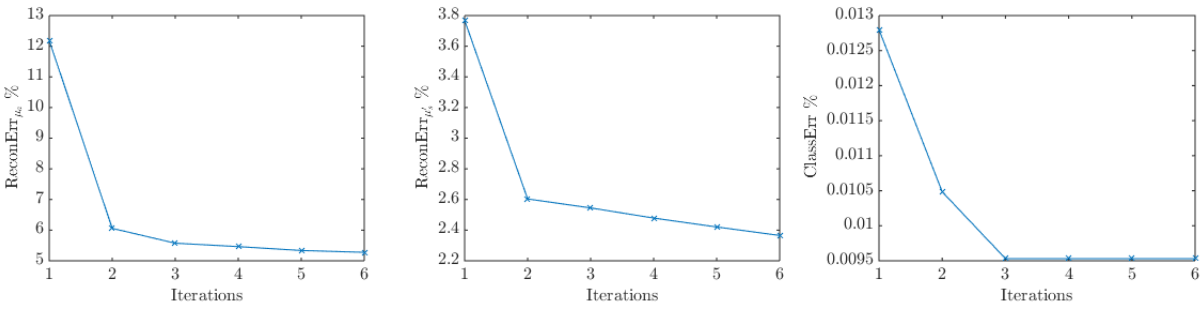

Figure 6. Reconstruction (20) (21) and classification (22) errors displayed as a function of iteration number.

ACKNOWLEDGMENTS

We thank Roman Hochuli and Felix Lucka for providing the numerical phantom used in this study. This work was funded by the EPSRC Doctoral Prize Fellowship EP/M506448/1.

REFERENCES

[1] Cox, B., Laufer, J. G., Arridge, S. R., and Beard, P. C., "Quantitative spectroscopic photoacoustic imaging: a review.," *Journal of Biomedical Optics* **17**(6), 061202 (2012).

[2] Gao, H., Osher, S., and Zhao, H., "Quantitative Photoacoustic Tomography," in [*Mathematical Modeling in Biomedical Imaging II*], ch. 5, 131–158, Springer Berlin Heidelberg (2012).

[3] Beard, P., "Biomedical photoacoustic imaging.," *Interface Focus* **1**(4), 602–31 (2011).

[4] Kostli, K. P., Frenz, M., Bebie, H., Weber, and P, H., "Temporal backward projection of optoacoustic pressure transients using Fourier transform methods," *Physics in Medicine and Biology* **46**, 1863–1872 (2001).

[5] Treeby, B. E., Zhang, E. Z., and Cox, B. T., "Photoacoustic tomography in absorbing acoustic media using time reversal," *Inverse Problems* **26**(11), 115003 (2010).

[6] Treeby, B. E. and Cox, B. T., "k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields," *Journal of Biomedical Optics* **15**(2), 021314 (2010).

[7] Saratoon, T., Tarvainen, T., Cox, B. T., and Arridge, S. R., "A gradient-based method for quantitative photoacoustic tomography using the radiative transfer equation," *Inverse Problems* **29**, 075006 (2013).

[8] Shao, P., Cox, B., and Zemp, R. J., "Estimating optical absorption, scattering, and Grueneisen distributions with multiple-illumination photoacoustic tomography.," *Applied Optics* **50**(19), 3145–54 (2011).

[9] Bal, G. and Ren, K., "Multiple-source quantitative photoacoustic tomography," *Inverse Problems* **27**(7), 1–22 (2011).

[10] Rosenthal, A., Razansky, D., and Ntziachristos, V., "Quantitative optoacoustic signal extraction using sparse signal representation," *IEEE Transactions on Medical Imaging* **28**(12), 1997–2006 (2009).

[11] Alberti, G. S. and Ammari, H., "Disjoint sparsity for signal separation and hybrid inverse problems," *Applied and Computational Harmonic Analysis* (August), 1–25 (2015).

[12] Malone, E., Cox, B. T., Arridge, S., Malone, E., Powell, S., Cox, B. T., and Arridge, S., "Reconstruction-classification method for quantitative photoacoustic tomography," *Journal of Biomedical Optics* **20**(12), 126004 (2015).

[13] Arridge, S., "Optical tomography in medical imaging," *Inverse Problems* **15**, R41 (1999).

[14] Nocedal, J. and Wright, S., [*Numerical optimization*], Springer Series in Operations Research and Financial Engineering , Springer-Verlag, New York (1999).

[15] Prince, S., [*Computer vision: models, learning, and inference*], Cambridge University Press (2012).

[16] Malone, E., Sato dos Santos, G., Holder, D., and Arridge, S., "A reconstruction-classification method for Multifrequency Electrical Impedance Tomography," *IEEE Transactions on Medical Imaging* **34**(7), 1486 – 1497 (2015).

[17] Hiltunen, P., Prince, S. J. D., and Arridge, S., "A combined reconstruction-classification method for diffuse optical tomography.," *Physics in Medicine and Biology* **54**(21), 6457–76 (2009).

[18] Jacques, S. L., "Optical properties of biological tissues: a review," *Physics in Medicine and Biology* **58**(14), 5007–5008 (2013).

[19] Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C., "A limited memory algorithm for bound constrained optimization," *SIAM Journal on Scientific Computing* **16**(5), 1190–1208 (1995).

[20] Zhu, C., Byrd, R. H., Lu, P., and Nocedal, J., "Fortran subroutines for large-scale bound constrained optimization," *ACM Trans Math Softw* **23**(4), 550–60 (1994).