

Three-dimensional quantitative photoacoustic tomography using an adjoint radiance Monte Carlo model and gradient descent

Jens Buchmann
Bernhard A. Kaplan
Samuel Powell
Steffen Prohaska
Jan Laufer

Three-dimensional quantitative photoacoustic tomography using an adjoint radiance Monte Carlo model and gradient descent

Jens Buchmann,^{a,†} Bernhard A. Kaplan,^{b,†} Samuel Powell,^c Steffen Prohaska,^b and Jan Laufer^{d,*}

^aTechnische Universität Berlin, Institut für Optik und Atomare Physik, Berlin, Germany

^bZuse Institute Berlin, Visual Data Analysis, Berlin, Germany

^cKing's College London, Biomedical Engineering and Imaging Sciences, Becket House, London, United Kingdom

^dMartin-Luther-Universität Halle-Wittenberg, Institut für Physik, Halle (Saale), Germany

Abstract. Quantitative photoacoustic tomography aims to recover maps of the local concentrations of tissue chromophores from multispectral images. While model-based inversion schemes are promising approaches, major challenges to their practical implementation include the unknown fluence distribution and the scale of the inverse problem. We describe an inversion scheme based on a radiance Monte Carlo model and an adjoint-assisted gradient optimization that incorporates fluence-dependent step sizes and adaptive moment estimation. The inversion is shown to recover absolute chromophore concentrations, blood oxygen saturation, and the Grüneisen parameter from *in silico* three-dimensional phantom images for different radiance approximations. The scattering coefficient is assumed to be homogeneous and known *a priori*. © The Authors. Published by SPIE under a Creative Commons Attribution 4.0 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.24.6.066001](https://doi.org/10.1117/1.JBO.24.6.066001)]

Keywords: quantitative photoacoustic imaging; blood oxygen saturation; inverse problem; model based inversion; Monte Carlo; spectral unmixing.

Paper 190010R received Jan. 9, 2019; accepted for publication Apr. 24, 2019; published online Jun. 6, 2019.

1 Introduction

Biomedical photoacoustic (PA) tomography is a hybrid soft-tissue imaging modality that combines the high spatial resolution of ultrasound with the high contrast and specificity of optical imaging techniques.¹⁻³ It relies on the generation of acoustic waves inside the tissue, which result from the absorption of intensity-modulated light, such as laser pulses or frequency chirps, by the tissue chromophores. From time-resolved PA signals recorded at multiple measurement points around the object, three-dimensional (3-D) data sets of the initial pressure distribution (i.e., PA images) can be calculated using acoustic reconstruction algorithms. Quantitative PA tomography (QPAT) aims to exploit the wavelength dependence of the image intensity to recover the local concentrations of endogenous tissue chromophores and exogenous contrast agents from which functional parameters, such as blood oxygen saturation, can be derived. To relate the PA image intensity to local chromophore concentrations, computational models of the physical processes during the image generation in conjunction with inversion schemes represent one approach to QPAT.^{4,5} A major challenge in QPAT is the unknown light fluence in the tissue,⁵⁻⁷ which is a nonlinear function of the concentrations and the scattering coefficient. Its effects on PA images have been described as spectral coloring and structural distortion.⁵ For an accurate quantification of concentrations and their ratios (e.g., blood oxygenation), the wavelength-dependent fluence distribution has to be accounted for.

Commonly used fluence correction methods include the application of empirical correction factors⁸ or simple models under the assumption of homogeneous optical properties.⁹ This is deemed sufficient to recover the absorption coefficient distribution from which maps of the chromophore concentrations can then be calculated using linear matrix inversions. The main limitation of these methods is the reliance on *a priori* knowledge of the distribution and wavelength dependence of the fluence, i.e., $\Phi(\vec{x}, \lambda)$. For *in vivo* images, this assumption is often invalid and can lead to significant quantification errors especially at greater depth. Alternative approaches include data-driven methods. In the work reported by Tzoumas et al.¹⁰ the wavelength-dependent fluence is written on the basis of eigenspectra obtained using a principle component analysis of *in silico* training data, and Kirchner et al.¹¹ calculated the fluence maps by applying deep learning and local context encoding to a large number of training data. While these methods have the potential to offer fast inversions, they require large training sets and may thus lack generality.

Model-based inversions, incorporating light transport models to predict the fluence as a function of the spatial distribution of the absorption and scattering coefficients, remain the most promising approach to QPAT. The initial pressure distribution is obtained by multiplying the fluence with the distribution of the absorption coefficient and the Grüneisen parameter, and PA image data sets can be obtained using acoustic propagation models. The difference between the model output and the measured data, i.e., the objective function, is minimized by iteratively updating the absorption and scattering coefficients during the inversion until convergence is reached. To overcome the nonuniqueness that arises from the use of single-wavelength images,⁵ multi-illumination approaches^{12,13} or multiwavelength image acquisition, in combination with *a priori* knowledge of

*Address all correspondence to Jan Laufer, E-mail: jan.laufer@physik.uni-halle.de

[†]These authors contributed equally to this work.

the wavelength dependence of the absorption and scattering coefficients,^{14,15} are employed.

A major challenge in high-resolution 3-D QPAT is the large number of variables ($>10^6$). While gradient-free methods can be applied to small-scale problems,¹⁶ inversions of a larger scale (tens of variables and more) quickly become computationally unfeasible. Gradient-based methods have the potential to overcome these limitations. They have been implemented using the adjoint formalism,^{17–19} which is applied using a finite element model of the diffusion approximation (DA) to the inversion of measured 3-D phantom images.²⁰ While the DA is valid in the diffusive regime and can be implemented efficiently,^{21,22} high-resolution PAT can cover depths in the ballistic and quasiballistic regimes, where the DA may not be valid.²³ Methods that aim to solve the radiative transfer equation (RTE) directly are computationally expensive and have only been demonstrated in two-dimensional (2-D) images so far.^{24–27} Monte Carlo (MC) models have recently gained attention^{28–31} due to their highly parallelized architecture and advances in graphics-processing units and have already been applied to 2-D QPAT inversions¹⁹ and in initial studies with limited parameter space in 3-D.^{16,32} In this paper, a method for inverting multiwavelength 3-D images based on an adjoint formulation of a radiance MC model is demonstrated *in silico*. The challenges that are addressed in this work are (1) the optimization of the objective function using inherently noisy gradients, (2) accounting for the effect of the concentration-dependent Grüneisen parameter, and (3) the representation of the radiance in terms of spherical harmonics. The capability of this approach to recover absolute chromophore concentrations and their ratios, e.g., blood oxygen saturation (blood SO_2), from high-resolution 3-D image data sets is demonstrated.

2 Methods

The forward model of the generation of the initial pressure shown in tomographic PA images is described in Sec. 2.1. The adjoint formalism²⁶ with which the gradients of the objective function are calculated is described in Sec. 2.2. The approximation of the radiance field as a finite sum of spherical harmonics³¹ within an MC light transport model is described in Secs. 2.3 and 2.4. The numerical phantom and the simulation parameters are outlined in Secs. 2.5 and 2.6, respectively. To reduce the impact of the inherent MC noise on the parameter update and to maximize the convergence speed of the gradient descent, an adaptive moment estimation (Adam) optimization algorithm³³ is employed (Sec. 2.7).

2.1 Forward Model

Assuming that the effects of the limited detection aperture and acoustic propagation can be neglected, the image intensity represents the initial pressure distribution, p_0 , which is given as

$$p_0(\vec{r}, \lambda) = \Gamma(\vec{r})H(\vec{r}), \quad (1)$$

where Γ is the Grüneisen parameter, which describes the PA efficiency; H is the absorbed optical energy density; \vec{r} is the spatial coordinate; and λ is the excitation wavelength. The absorbed energy density is defined as

$$H(\vec{r}) = \mu_a(\vec{r}, \lambda)\Phi(\vec{r}, \lambda), \quad (2)$$

where μ_a is the absorption coefficient and Φ is the light fluence, which is the radiance ϕ integrated over all angles:

$$\Phi(\vec{r}) = \int \phi(\vec{r}, \hat{s}')d\hat{s}'. \quad (3)$$

The absorption coefficient is related to the chromophore concentrations via the specific absorption coefficient, $\alpha_k(\lambda)$, i.e.,

$$\mu_a(\lambda, \vec{r}) = \sum_k^{N_k} c_k(\vec{r})\alpha_k(\lambda), \quad (4)$$

where N_k is the number of chromophores and k indicates the chromophore type. The Grüneisen parameter Γ is assumed to be linearly dependent on chromophore concentrations,^{15,34,35} i.e.,

$$\Gamma(\vec{r}) = \Gamma_{\text{water}} \left[1 + \sum_k \beta_k c_k(\vec{r}) \right], \quad (5)$$

where β_k is an empirical and chromophore-specific coefficient. The MC method is chosen for modeling the light fluence as it provides an accurate approximation of the radiative transport equation for superficial (1 to 2 cm), high-resolution QPAT.⁵ This involves the launch of photons (typically represented as packets of energy³⁶) according to a predefined source distribution. Their propagation within the domain is determined by the optical coefficients $\mu_a(\vec{r})$ and $\mu_s(\vec{r})$, the scattering phase function $\Theta(\hat{s}, \hat{s}', \vec{r})$, and the refractive index n . The deposition of energy when a photon traverses a voxel is determined by the absorption coefficient μ_a . The angular dependence of scattering events is described by the Henyey–Greenstein phase function.³⁷

2.2 Adjoint-Assisted Optimization

Assuming Gaussian noise on the measured data, an estimate of the chromophore distributions is found by minimizing the objective function ε , which is given as

$$\varepsilon = \sum_l^{N_\lambda} \int_{\Omega} \frac{1}{2} [p_0^m(\lambda_l, \vec{r}) - p_0(\lambda_l, \vec{r})]^2 d\Omega, \quad (6)$$

where Ω is the imaged volume domain, $p_0^m(\lambda_l, \vec{r})$ is the measured PA image at wavelength λ_l , p_0 is the PA image obtained from the forward model, and N_λ is the number of excitation wavelengths.

To find the chromophore concentration maps, $c_k(\vec{r})$, the derivative of ε with respect to c_k at any position \vec{r}_i is required, i.e., $\frac{\partial \varepsilon}{\partial c_i} \equiv \frac{\partial \varepsilon}{\partial c_k(\vec{r}_i)}$. For the sake of brevity, only one chromophore will be considered in the remaining description, i.e., $\frac{\partial}{\partial c_i} \equiv \frac{\partial}{\partial c_1(\vec{r}_i)}$. The objective function ε represents the sum of the objective functions, ε_{λ_l} , at each excitation wavelength, which is given as $\varepsilon = \sum_l^{N_\lambda} \varepsilon_{\lambda_l}$. For simplicity, only one wavelength, λ_l , will be considered to describe the derivative:

$$\frac{\partial \varepsilon_{\lambda_l}}{\partial c_i} = - \int_{\Omega} [p_0^m(\vec{r}) - p_0(\vec{r})] \frac{\partial p_0}{\partial c_i} d\Omega. \quad (7)$$

Since $p_0 = \Gamma\mu_a\Phi$, after applying the chain rule, $\frac{\partial p_0}{\partial c_i}$ becomes

$$\frac{\partial p_0}{\partial c_i} = \frac{\partial \Gamma}{\partial c_i} \mu_a \Phi + \Gamma \frac{\partial \mu_a}{\partial c_i} \Phi + \Gamma \mu_a \frac{\partial \Phi}{\partial c_i} \quad (8)$$

$$\begin{aligned}
 &= \beta_k \delta(\vec{r} - \vec{r}_i) H(\vec{r}) + \Gamma(\vec{r}) \alpha_k(\lambda_i) \delta(\vec{r} - \vec{r}_i) \Phi(\vec{r}) \\
 &\quad + \Gamma(\vec{r}) \mu_a(\vec{r}) \frac{\partial \Phi}{\partial c_i}, \tag{9}
 \end{aligned}$$

where $\delta(\vec{r} - \vec{r}_i)$ is the Dirac delta function. The gradient of the fluence with respect to chromophore concentration at a particular position, $\frac{\partial \Phi}{\partial c_i}$, is generally unknown, and one can make use of the adjoint formalism.^{17,26} Briefly, the adjoint approach defines a source term q^* for the adjoint radiance ϕ^*

$$q^*(\vec{r}, \lambda) = \Gamma(\vec{r}) \mu_a(\vec{r}, \lambda) [p_0^m(\vec{r}, \lambda) - p_0(\vec{r}, \lambda)], \tag{10}$$

in which the adjoint radiance is expressed in terms of the difference between the measured images and the modeled images of p_0 . Using this adjoint source definition enables the substitution of the term containing the unknown $\frac{\partial \Phi}{\partial c_i}$ in Eq. (8) with a term containing the radiance ϕ and its adjoint counterpart ϕ^*

$$\begin{aligned}
 &\int_{\Omega} \Gamma(\vec{r}) \mu_a(\vec{r}) (p_0^m - p_0) \frac{\partial \Phi}{\partial c_i} d\Omega \\
 &= -\alpha_k(\lambda_i) \left[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s} \right]_{\vec{r}=\vec{r}_i} V_{\text{vox}}, \tag{11}
 \end{aligned}$$

where \int_{S^2} is the integral over solid angle \vec{s} , and S^2 is the unit sphere. The term V_{vox} is the result of the discretization of the data using a piecewise constant basis to sample (see Sec. 9). The gradient required to update the concentration of one chromophore at one wavelength is therefore given as

$$\begin{aligned}
 \frac{\partial \epsilon_{\lambda_i}}{\partial c_i} &= -[p_0^m(\vec{r}_i) - p_0(\vec{r}_i)] V_{\text{vox}} [\beta_k H(\vec{r}_i) + \Gamma(\vec{r}_i) \alpha_k \Phi(\vec{r}_i)] \\
 &\quad + \alpha_k \left[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s} \right]_{\vec{r}=\vec{r}_i} V_{\text{vox}}. \tag{12}
 \end{aligned}$$

The adjoint model and its derivation are described in general and in more detail in Secs. 6 and 7.

2.3 Radiance Approximation

As the radiance $\phi(\vec{s})$ and the adjoint radiance $\phi^*(\vec{s})$ are functions of solid angle and are defined on the surface of the unit sphere, both quantities can be expressed on the basis of spherical harmonic functions. The expansion of the radiance into spherical harmonics is based on previous work³¹ and is inspired by the P_n approximations, described in Ref. 22 (similar to a Fourier expansion in 1-D or 2-D) and is outlined in detail in Sec. 8. Using a finite expansion of $[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s}]_{\vec{r}=\vec{r}_i}$ in real spherical harmonics, the gradient equation is given as

$$\begin{aligned}
 \frac{\partial \epsilon_{\lambda_i}}{\partial c_i} &= V_{\text{vox}} \left\{ -[p_0^m(\vec{r}_i) - p_0(\vec{r}_i)] [\beta_k H(\vec{r}_i) + \Gamma(\vec{r}_i) \alpha_k \Phi(\vec{r}_i)] \right. \\
 &\quad \left. + \alpha_k \sum_{l=0}^{N_L} \sum_{m=-l}^l \psi_{lm}(\vec{r}_i) \psi_{lm}^*(\vec{r}_i) \right\}, \tag{13}
 \end{aligned}$$

where $\psi_{lm}(\vec{r}_i)$ is the radiance field approximated by the spherical harmonics function of degree l , order m at position \vec{r}_i and ψ_{lm}^* its adjoint counterpart. Equation (13) is implemented in a gradient-based optimization scheme (described in Sec. 2.7), which updates the concentrations iteratively to minimize the mismatch between measured and modeled data [Eq. (6)].

The last term of the gradient in Eq. (13) contains an expansion of the radiance and the adjoint radiance in spherical harmonics

$$\alpha_k \sum_{l=0}^{N_L} \sum_{m=-l}^l \psi_{lm}(\vec{r}_i) \psi_{lm}^*(\vec{r}_i). \tag{14}$$

Here, $N_L = \infty$ would give the most accurate radiance approximation, but due to constraints with respect to computation time and memory, the degree of the spherical harmonics, N_L , needs to be finite. However, it is not clear *a priori* up to which value of N_L the corresponding coefficients needs to be stored to approximate the radiance with sufficient accuracy. This has been investigated by evaluating the inversion scheme for three different configurations: (1) $N_L = 0$, i.e., using only the fluence, (2) for $N_L = 4$, i.e., the most accurate representation of the radiance, and (3) omitting the radiance, i.e., $\psi_{lm} = \psi_{lm}^* = 0$ for all l, m , thus neglecting the gradient term provided by the adjoint radiance. All other parameters remain the same during the inversions.

2.4 Radiance Monte Carlo Simulations

Most MC simulators provide only the light fluence, which is the radiance integrated over all directions and time. To satisfy Eq. (12), a radiance MC algorithm (RMC)^{19,31} is used. To obtain the radiance $\phi(\vec{r})$, the directional information of the photon traversing a voxel at position \vec{r} is stored by depositing the photon weight into the relevant spherical harmonics coefficients $\psi_{lm}(\vec{r})$. The RMC code is implemented in the Julia programming language,³⁸ which controls and dispatches the execution of kernels on both the CPU (written in Julia), and the GPU (written in NVIDIA's CUDA language).³¹ Because the definition of the adjoint model does not change the dynamics of photon propagation, the same RMC simulation code provides the adjoint radiance $\phi^*(\vec{r})$.

2.5 In Silico Phantom

An MC model has been used to calculate the multiwavelength 3-D PA images that represented measured data and are referred to as reference images or reference data throughout this paper. The domain of the model is divided into subvolumes (SVs) that represented simplified anatomical structures, such as a subcutaneous tumor and a number of discrete blood vessels. The depths of the structures are similar to those observed in *in vivo* images acquired using a Fabry-Perot scanner with a planar detection geometry.³⁹⁻⁴² Absorption is assumed to originate from three chromophores, i.e., oxyhemoglobin (HbO₂), deoxyhemoglobin (HHb), and methylene blue (Mb), as an exogenous contrast agent. It should be noted that the method can potentially be applied to any number of chromophores. The computational model of the phantom consists of nine different SVs, each with homogeneous optical properties (Fig. 1). This includes six tube-like structures to mimic blood vessels, a tumor consisting of an ellipsoidal rim and core SV, and the background. The tubes are positioned adjacent to the tumor at depths of 1.5 and 7.5 mm, have a circular cross section with a radius of 0.4 mm, and are filled with HHb and HbO₂. The blood oxygen saturation (sO₂) is defined as the ratio of the concentration of oxyhemoglobin and the total hemoglobin concentration

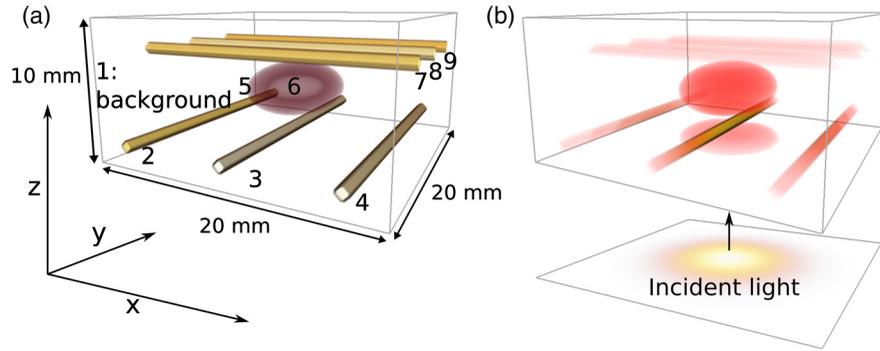


Fig. 1 A 3-D view of the *in silico* phantom and reference p_0^n . (a) Segmented phantom with SV IDs; SV1: homogeneous background material, SV2 to SV4 and SV7 to SV9: tubes representing blood vessels, SV5 and SV6 represent an ellipsoidal tumor consisting of an inner and outer SV. (b) Initial pressure distribution $p_0^n(\vec{r})$ calculated using physiological hemoglobin concentrations and blood sO_2 at an excitation wavelength of $\lambda = 798$ nm. The $x - y$ Gaussian profile of the excitation beam is shown below.

$$sO_2 = \frac{c_{HbO_2}}{c_{HbO_2} + c_{HHb}}. \quad (15)$$

The tube-like structures are assumed to contain blood sO_2 ranging from 75% to 98% to represent the typical values found in veins and arteries.⁴³ The total hemoglobin concentration is 2.3 mM.⁴⁴ Two concentric ellipsoids represent the tumor at a depth ranging from 3.0 mm to 6.0 mm. The outer (inner) ellipsoid's axes are $a = 4.5(2.5)$ mm, $b = 3.5(2.0)$ mm, and $c = 2.0(0.8)$ mm, respectively. The tumor SVs contained 20% blood (0.46 mM).⁴⁵ The tumor shell has an sO_2 of 80%, whereas that of the core is 40% to mimic necrotic tissue. The tumor also contains an exogenous contrast agent, Mb, at a concentration of 10 μ M. The background SV contains a blood volume fraction of 1.5% with an sO_2 of 60%. Other parameters, such as the scattering anisotropy ($g = 0.9$), the refractive index ($n_i = 1.33$ inside the domain, $n_e = 1.5$ outside of the domain), and $\mu_s(\vec{r}, \lambda)$ are held constant and uniform across the domain. The absorption spectra of HHb, HbO₂, and Mb are shown in Fig. 2. The wavelength dependence of the reduced scattering coefficient $\mu'_s(\lambda) = \mu_s(\lambda)(1 - g)$ is approximated using

$\mu_s(\lambda) = 6.65 \cdot 10^3 \cdot \lambda_+^{-1.317} \text{ mm}^{-1}$ with $\lambda_+ = \lambda/\text{nm}$, which resulted in a μ'_s of $\sim 1 \text{ mm}^{-1}$ at $\lambda = 800 \text{ nm}$. The Grüneisen parameter of water is set to 0.124. The coefficient β_k describing the total hemoglobin concentration dependence of the Grüneisen parameter is set to $\beta_{Hb, HbO} = 0.02146 \text{ L/mmol}$.³⁵ It is assumed that the Mb concentration distribution does not change the Grüneisen parameter ($\beta_{Mb} = 0$). The remaining parts of the volume are assumed to be filled with materials, such as water and lipids, whose absorption is considered negligible at the selected excitation wavelengths.

2.6 Monte Carlo Simulation Parameters

To obtain reference images, i.e., data sets that represent measured multiwavelength PA images, the domain was discretized into $200 \times 200 \times 100$ (i.e., 4×10^6) isotropic voxels of size $V_{\text{vox}} = 10^{-3} \text{ mm}^3$ yielding a total volume of $20 \times 20 \times 10 \text{ mm}^3$. The source profile was a 2-D Gaussian function with a width of $\sigma = 4 \text{ mm}$. About $2 \cdot 10^9$ photon packets were used in the MC simulation of the light fluence. The angle-dependent radiance was not calculated. Reference image data sets were calculated for three different excitation wavelengths that coincided with the absorption peaks of Mb (664 nm), HHb (758 nm), and the isosbestic point of hemoglobin absorption (798 nm). Gaussian noise ($\sigma = 0.1\%$ of the maximum image intensity) was added to the reference data, resulting in negative image intensities in regions of low p_0 .

The domain discretization used during the inversion was identical to that used for the reference data set, which may raise the question of whether this constitutes a so-called inverse crime. In MC models, the discretization is used merely as a basis for sampling physical quantities while photon packets can propagate freely in continuous space. This is in contrast to other methods, such as finite elements, where the discretization has a direct impact on the accuracy of the solution. Taking also into account the stochastic nature of MC models, it can be concluded that using identical discretizations does not constitute an inverse crime.

During an inversion, 1×10^7 photon packets were used for the calculation of the radiance and fluence; 5×10^6 photons were used for the calculation of the corresponding adjoint quantities. The typical running time for one inversion iteration, including the adjoint model with $N_L = 4$, was 84 s on a high-end consumer GPU (NVIDIA GeForce Titan X Pascal). This was reduced to 17 s when the radiance term in Eq. (13) was neglected, i.e.,

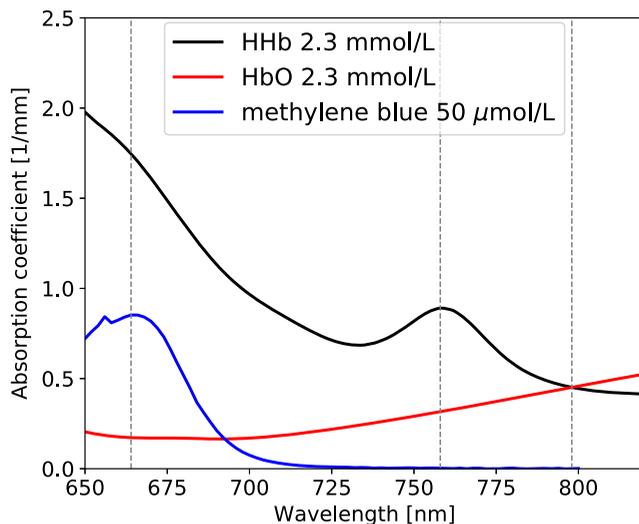


Fig. 2 Spectra of HHb, HbO₂, and Mb used for forward simulation and inversion. The excitation wavelengths are indicated by gray vertical lines. The absorption spectra are taken from Ref. 46.

$N_L = 0$ and no adjoint MC runs. Since $k = 3$ independent chromophore concentrations were associated with each voxel, the model contained a total of 12 million variables.

2.7 Gradient-Based Optimization

The gradient-based optimization was initialized assuming a homogeneous $c_{\text{HHb}} = c_{\text{HbO}_2} = 0.023$ mM, i.e., $s\text{O}_2 = 50\%$, and $c_{\text{Mb}} = 0.0$ mM. Owing to the stochastic nature of MC models, the gradients [Eq. (13)] were subjected to noise. To compensate for this, the Adam optimization algorithm³³ was employed. It was developed for the first-order optimization of noisy objective functions in high-dimensional parameter spaces and can be seen as an extension of the momentum algorithm.^{47,48} The Adam algorithm calculated an exponential moving average of the gradient (first moment) and the squared gradient (second moment) using the decay rates $\beta_{1,\text{Adam}}$ and $\beta_{2,\text{Adam}}$ with an additional bias-correction step. The final update was calculated by multiplying the step-size parameter γ_{Adam} with the first-order moment divided by the square root of the second-order moment. The detailed description can be found in Ref. 33. This algorithm was found to dramatically increase the convergence speed, compared to standard gradient descent. Its efficiency depended on the values of a set of parameters, including the step-size γ_{Adam} , the decay rates $\beta_{1,\text{Adam}}$ and $\beta_{2,\text{Adam}}$, and an additional ϵ_{Adam} , which avoided division by zero. The decay rates and ϵ were set to recommended default values ($\beta_{1,\text{Adam}} = 0.9$, $\beta_{2,\text{Adam}} = 0.999$, and $\epsilon_{\text{Adam}} = 10^{-8}$), and the step size was assigned different values depending on the chromophore according to Eqs. (16) and (17).

To ensure fast convergence for all chromophores, the step size for Mb is set to be significantly smaller than that of HHb/HbO₂ since Mb concentrations are in the range of μM while those of HHb/HbO₂ are in the range of mM. The chromophore-dependent step size is calculated as follows:

$$\gamma_{\text{chrom}} = \gamma_{\text{ref}} \frac{\sum_{\lambda} \alpha_{\text{chrom},\lambda} \cdot c_{\text{max,ref}}}{\sum_{\lambda} \alpha_{\text{ref},\lambda} \cdot c_{\text{max,chrom}}}, \quad (16)$$

where γ_{ref} is the step size of a reference chromophore, the value of which was set *ad hoc* to $\gamma_{\text{ref}} = \gamma_{\text{HbO}_2} = 100$; $\alpha_{\text{chrom}/\text{ref},\lambda}$ is the specific absorption coefficient of the respective chromophore at wavelength λ ; and $c_{\text{max,ref}/\text{chrom}}$ is the anticipated maximum concentration of the respective chromophore, which is set to physiological reasonable values ($c_{\text{max,HbO}_2} = c_{\text{max,HHb}} = 2.3$ mM, $c_{\text{max,Mb}} = 30$ μM).

The gradient is also expressed as a function of fluence in Eq. (13), either directly or via H and Δp . This would result in slow convergence in regions of low fluence. To compensate for this, a spatially dependent step size is used, which increases the step size by normalizing it by the mean fluence over all wavelengths:

$$\gamma(\vec{r})_{\text{chrom,scaled}} = \frac{\gamma_{\text{chrom}}}{\tilde{\Phi}_{\text{norm}}(\vec{r}) + \epsilon_{\Phi}}. \quad (17)$$

Note that $\tilde{\Phi}_{\text{norm}}(\vec{r})$ is the normalized mean fluence:

$$\tilde{\Phi}_{\text{norm}}(\vec{r}) = \frac{\sum_{\lambda} \Phi(\vec{r}, \lambda)}{\sum_{\lambda} \Phi(\vec{r}_{\text{max}}, \lambda)}, \quad (18)$$

where $\vec{r}_{\text{max}} = \arg \max_{\vec{r}} \sum_{\lambda} \Phi(\vec{r}, \lambda)$ represents the location where the total fluence is at a maximum. The parameter ϵ_{Φ}

avoids division by zero and determines the maximum change in the step size in regions of low fluence. In this study, $\epsilon_{\Phi} = 10^{-4}$ leads to a sufficient speed-up in convergence for the deepest tubes. The optimization of the step size with respect to the different chromophores and the local fluence is often referred to as preconditioning.

A single iteration of the gradient-based update consisted of the execution of the MC model and its adjoint counterpart for all three wavelengths. The inversion was run for 1500 iterations to investigate the convergence. After each iteration, the updated concentrations for HHb, HbO₂, and Mb were limited to a reasonable range of values (also known as projected gradient descent) to avoid spurious overshooting or undershooting that could lead to physiologically unrealistic concentrations. To compensate for the effects of noise in low-absorbing regions (i.e., negative p_0 in reference images), negative chromophore concentrations, and hence negative μ_a values, were allowed during the gradient descent. To ensure stability, the range of negative concentration values was limited to a tenth of the maximum positive concentrations (-0.35 to 3.5 mM for c_{HHb} and c_{HbO_2} , -0.02 to 0.2 mM for c_{Mb}). The 3-D blood sO₂ maps were calculated from the recovered c_{HbO_2} and c_{HHb} images. The scattering distribution was assumed to be known *a priori*.

To verify that the inversion scheme is valid over a range of physiologically plausible parameter values, multiwavelength reference images were calculated and inverted for different sO₂ values in the inner tumor region and the background. Two scenarios were chosen, each comprising five different combination of sO₂ values. First, as the tumor core (being completely enclosed by the rim) may be assumed to be most strongly affected by spectral coloring, its sO₂ value was varied from 10% to 90% in 20% increments (SV 6), whereas all other parameters remained fixed. Second, the sO₂ value of the background (SV ID 1) was varied from 10% to 90% in 20% increments.

3 Results

The accuracy of the recovered chromophore concentration and sO₂ maps are reported in Secs. 3.1 and 3.2. In Sec. 3.3, the results obtained using multiple inversions of image data sets, in which the chromophore concentrations and their ratios are varied, are reported.

3.1 Absolute Concentrations

Examples of 3-D volume-rendered images of the absolute concentrations of HbO₂, HHb, and Mb recovered after 1500 iterations are shown in Figs. 3(a)–3(c). The color scales are thresholded to render the background with its comparatively low chromophore concentrations transparent. To reduce the effect of the added Gaussian noise on the rendered images (particularly in regions of low fluence), a 3-D median filter is applied (non-iterative, edge- and face-connected). The error functional as a function of the number of iterations is shown in Fig. 3(d). The value of the error functional cannot reach zero due to the inherent MC noise of the forward model used during the inversion. The dashed orange line indicates the minimum value of the error functional that is reached with 1×10^7 photons, and the green dotted line indicates the total noise level consisting of MC noise and Gaussian noise added to the reference data. The MC noise is obtained from forward calculations with prior knowledge of the correct chromophore distributions.

In Fig. 4, cross-sectional xz -images of the true and recovered concentration of the three chromophores and the absolute error

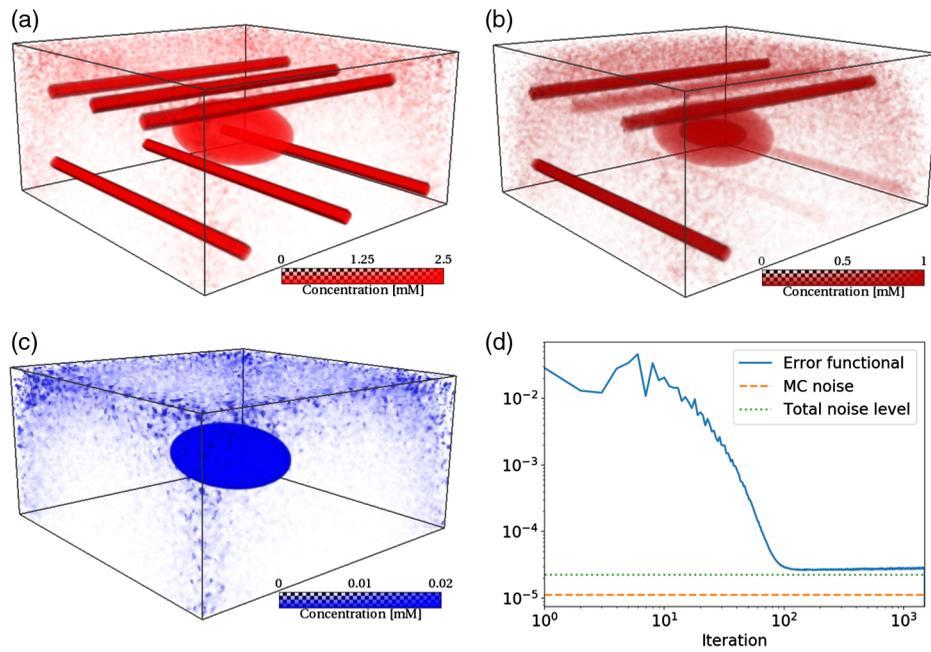


Fig. 3 Absolute chromophore concentration maps recovered using the inversion scheme and its convergence. (a)–(c) The 3-D volume-rendered images of the concentrations of HbO_2 , (b) HHb , and (c) Mb . (d) Values of the error functional (solid blue line), the baseline of the MC noise (dashed orange line), and added Gaussian noise (dotted green line).

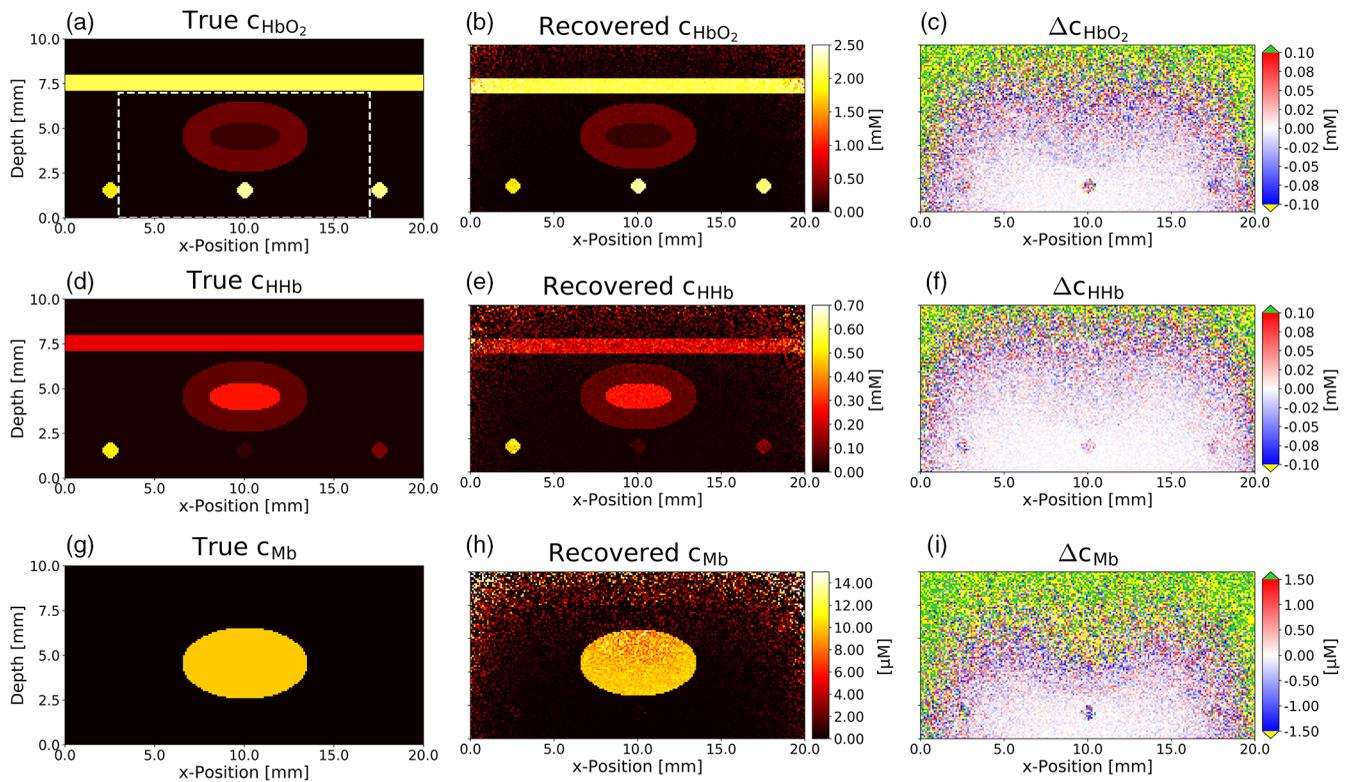


Fig. 4 The 2-D cross-sectional xz -images of the true and recovered chromophore concentrations, together with the absolute error at $y = 10$ mm. Left column: true chromophore concentrations. Center column: recovered concentrations. Left and center column share the same color bar. Right column: absolute concentration error. Voxels where the error exceeds the limits of the color scale are rendered in green and yellow. In (a), the reduced background SV 1^* is illustrated as a white dashed rectangle. Because background voxels far away from the source exhibit large errors due to low SNR, this reduced SV is used for calculating average concentrations.

at $y = 10$ mm (center plane) are shown. Excellent agreement is found in regions of high fluence [e.g., corresponding to white, light blue, and light red pixels in Figs. 4(c), 4(f), and 4(i)]. By contrast, significant quantification errors [corresponding to yellow and green pixels in Figs. 4(c), 4(f), and 4(i)] are observed in regions of low fluence. The recovered c_{Mb} appears to exhibit larger quantification errors in the background compared to c_{HbO_2} and c_{HHb} .

Table 1 contains the true and recovered concentrations averaged over all voxels of each SV defined in Sec. 2.5. The brackets indicate the standard deviation (concise notation). The recovered c_{HHb} and c_{HbO_2} in SV 2 to 9 are in excellent agreement with the true values. While the recovered c_{Mb} are generally in good agreement with the true values, SV 2 to 4 and SV 7 to 9 exhibit negative and small concentrations with large standard deviations. The background (SV 1) also exhibits large errors and standard deviations due to the low signal-to-noise ratio (SNR). In the background region closer to the light source [$140 \times 140 \times 70$ central voxels inside the $200 \times 200 \times 100$ image domain, illustrated in Fig. 4(a)] where the SNR is greater, the recovered concentrations are in good agreement with the true values (SV 1* in Table 1).

As described in Sec. 2.4, the inversion is implemented using an approximation of the radiance based on spherical harmonics of varying degree N_L , including the omission of the adjoint term. The inversions are found to converge to the same final values while propagating along different routes. Convergence is reached irrespective of the radiance approximation (see Fig. 5).

3.2 Blood Oxygen Saturation

Figure 6 shows the cross-sectional xz -images ($y = 10$ mm) of the known and recovered blood sO_2 together with the absolute error. The accuracy is clearly affected by noise as shown in the difference image in Fig. 6(c). While most voxels in SV 2 to 9

exhibit an error within $\pm 5\%$ sO_2 , it is noticeably larger for objects at greater depth. Concentrations in the background SV 1 are also affected by noise, particularly in regions further away from the source. However, the accuracy improves near the source where SNR is increased. The average values of blood sO_2 for each SV are also calculated and are summarized in Table 1. Blood sO_2 in SV 2 to 9, i.e., corresponding to blood-filled tubes and the tumor SVs, are found to lie within 0.3% of the true values (i.e., while sO_2 errors of individual voxels can be quite large due to noise, averaging over lots of voxels greatly improves accuracy). The average blood sO_2 for SV 1 (i.e., the entire background SV) is found to show significant errors (18.3% sO_2) and is attributed mainly to the adverse effects of noise in low fluence regions. By contrast, the inversion results are accurate for the reduced background SV 1* due to higher SNR.

3.3 Validation over a Range of Blood Oxygen Saturation

The inversion scheme is validated on image data sets where sO_2 is varied over a range of physiologically plausible parameter values (Sec. 2.7). The inversions are computed without including the gradient term of the radiance and $N_L = 0$ in order to minimize the computation time. To obtain the final sO_2 value for each image data set, the inversion is run for 1500 iterations after which the average sO_2 is obtained from the SVs. Figure 7(a) shows the true and recovered sO_2 values for all SVs and all image data sets together with the line of unity (dashed line) and a $\pm 5\%$ error interval (dotted lines). All recovered sO_2 values are in good agreement with the known values and exhibit an average error below 0.3% sO_2 . Figure 7(b) shows the difference between true and recovered sO_2 for all SVs and image data sets sorted by SV. Only the results corresponding to the reduced background SV 1* are shown as this region exhibits sufficient SNR.

Table 1 True and recovered (inversion) absolute chromophore concentrations and blood sO_2 in the SVs of the phantom. The concentrations represent the average value over all voxels within each SV; the values in brackets indicate the standard deviations (concise notation). SV 1—background, SV 1*—background close to the source, SV 5 and SV 6—outer and inner tumor SVs, respectively, SV 2 to 4 and SV 7 to 9—blood-filled tubes.

SV	HHb (mM)		HbO ₂ (mM)		Mb (μ M)		sO ₂ (%)	
	True	Inversion	True	Inversion	True	Inversion	True	Inversion
1	0.0138	0.0087(1169)	0.0207	0.0320(212)	0	0.2(58)	60	78.3
1*	0.0138	0.0138(219)	0.0207	0.0205(384)	0	-0.002(922)	60	59.8
2	0.506	0.503(70)	1.79	1.78(11)	0	-0.08(266)	78	77.9
3	0.115	0.117(63)	2.18	2.16(11)	0	-0.04(244)	95	94.9
4	0.046	0.047(31)	2.25	2.25(6)	0	-0.05(106)	98	97.9
5	0.092	0.092(18)	0.368	0.367(31)	10	9.76(99)	80	79.9
6	0.276	0.275(18)	0.184	0.183(29)	10	9.51(100)	40	39.8
7	0.575	0.560(128)	1.73	1.71(20)	0	-0.35(510)	75	75.3
8	0.230	0.229(90)	2.07	2.04(14)	0	-0.41(374)	90	89.9
9	0.460	0.450(125)	1.84	1.82(20)	0	-0.37(509)	80	80.0

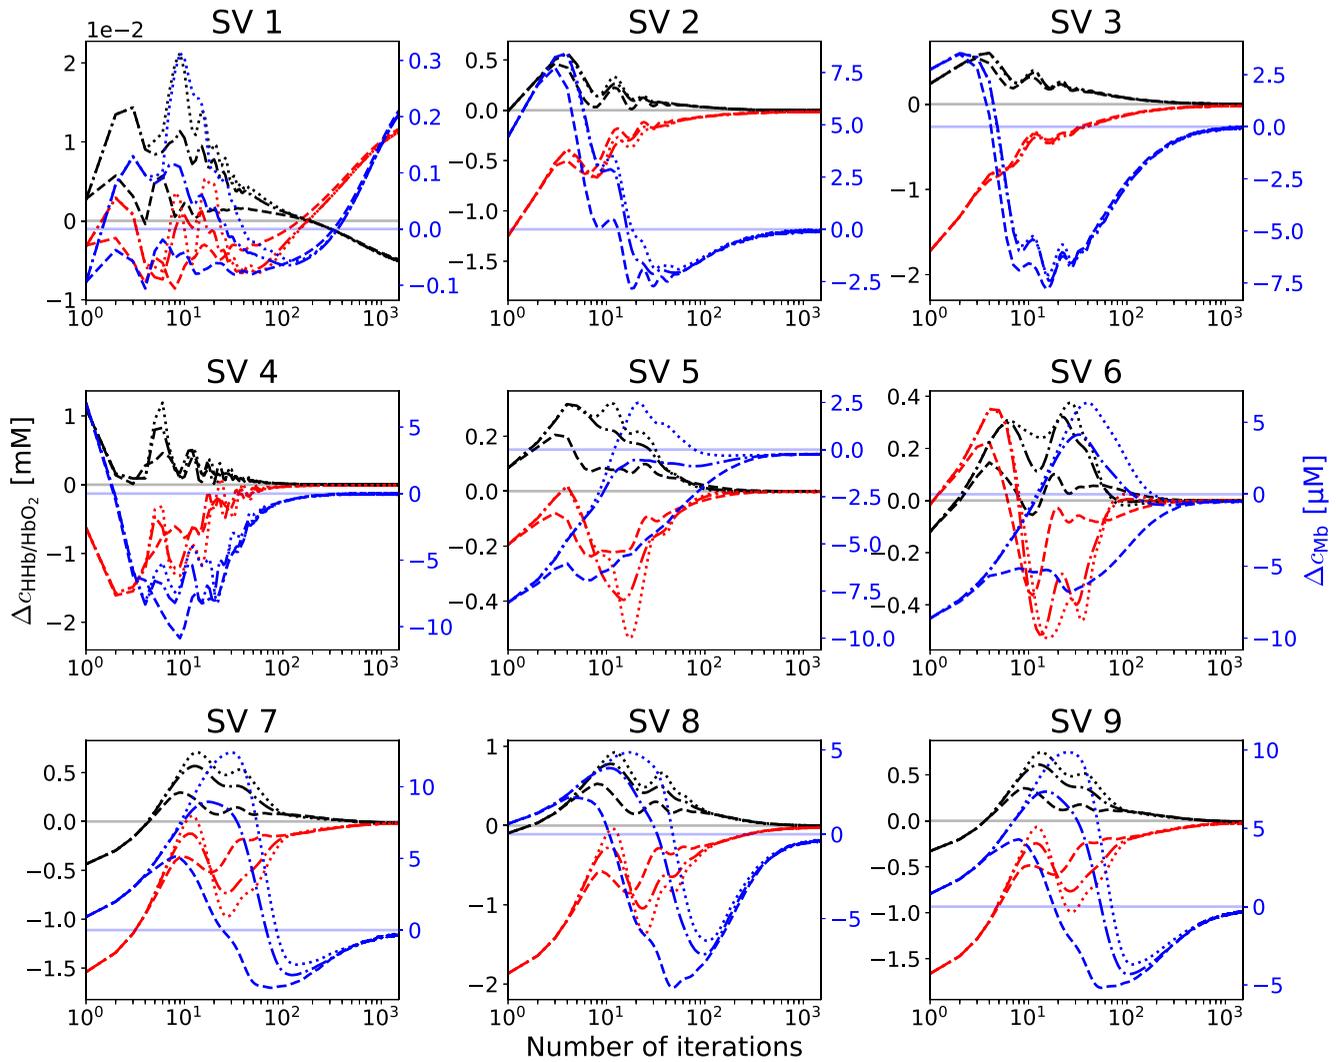


Fig. 5 Convergence of the true and recovered concentrations for inversions incorporating the gradient term of the adjoint radiance for $N_L = 4$ (dotted line, \cdots), $N_L = 0$ (dashed-dotted line, $-\cdot-$), and without the adjoint term (dashed line, $--$). The chromophores are shown in blue (Mb), red (HbO₂), and black (HHb). Horizontal lines in light color depict perfect convergence, i.e., $\Delta c = 0$.

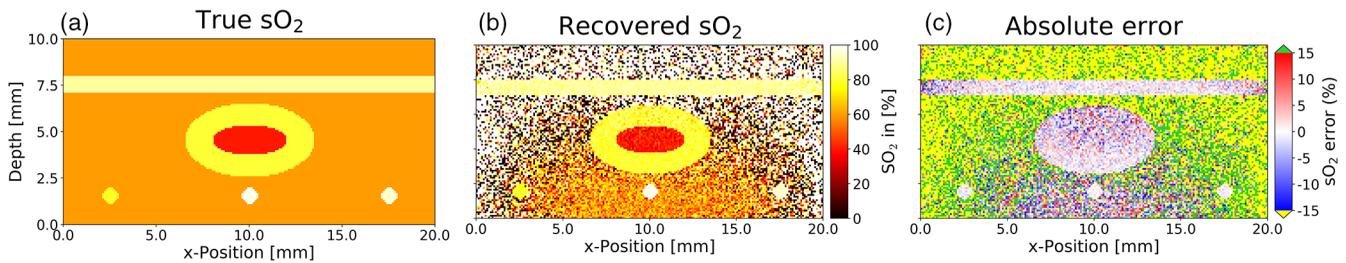


Fig. 6 The 2-D cross-sectional yz -images of the true (a) and recovered sO₂ (b), together with the absolute error (c) at the center of the phantom ($y = 10$ mm). The images on the left and in the center share the same color bar. In the right color bar, voxels exceeding (dropping below) an absolute difference between true and recovered sO₂ of 15% are shown in green (yellow).

4 Discussion

The 3-D maps of absolute concentrations of HbO₂, HHb and Mb, and the resulting blood sO₂ recovered using a gradient-based MC inversion scheme showed excellent agreement with the true values. To achieve the best possible match of

noise-affected PA images and the model, the inversion scheme was implemented without a non-negativity constraint for the chromophore concentrations. Even though negative concentrations were physiologically implausible, it was found that incorporating a non-negativity constraint greatly affected the

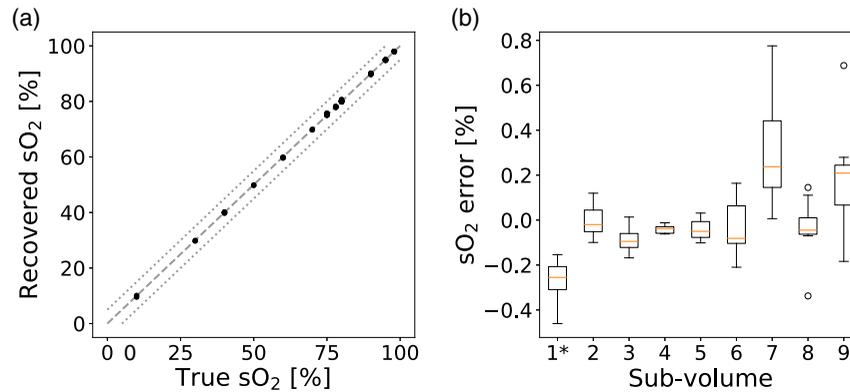


Fig. 7 Validation of the inversion scheme over a range of blood sO_2 . (a) Average recovered sO_2 as a function of the true sO_2 of all SVs and image data sets together with the line of unity (dashed line) and a $\pm 5\%$ error interval (dotted lines). In the five image data sets, the sO_2 in the inner tumor material varies from 10% to 90% in 20% increments. In another five image data sets, the background sO_2 ranges from 10% to 90% (20% increments). (b) Box-and-whisker-plot of the absolute difference between true and recovered sO_2 for each SV.

recovered average concentrations and blood sO_2 values. However, negative concentrations can lead to negative absorption coefficients. In the MC model, it led to the photon packets gaining weight as they traversed a voxel with negative absorption. If too many voxels exhibited negative μ_a , unstable inversions can be observed as the photon weight diverges. While this was occasionally observed in this study, it was found that a reduction of the step size and an increase in the number of iterations remedied this problem. The inversion scheme described in this paper included an expression of the radiance and its adjoint in a basis of spherical harmonics. The influence of the adjoint formalism and the spherical harmonics approximation on the accuracy and convergence of the inversion was evaluated under the assumption that the scattering coefficient was known *a priori*. It was found that neither accuracy nor convergence speed were affected by the radiance term and its adjoint, i.e., the last term in Eq. (13). This was also observed when the radiance term was omitted (Fig. 5) and was confirmed by the relative magnitudes of the individual gradient terms. The radiance term (irrespective of the spherical harmonic approximations) was always significantly smaller than the remaining terms of Eq. (13). Omitting the computation of the adjoint radiance resulted in a major increase in computational speed. However, from the limited investigation presented here, it can only be concluded that the adjoint term may be neglected if the scattering coefficient was known. If the recovery of the scattering coefficient was of interest, a radiance approximation to a minimum of $N_L = 1$ deg may be necessary.³¹

The gradient-based inversion was found to benefit greatly from optimization algorithms in which parameters, such as the step sizes and the exponential decay rates of the Adam algorithm, were predefined to enable a fast and accurate convergence. A potential drawback of such methods was the need to test several sets of these parameters prior to an inversion to assess whether they have a positive impact on the convergence speed. Within the scope of this study, only minor and *ad hoc* parameter tuning was conducted. A more thorough investigation, including the development of automated parameter selection algorithms, may yield significantly faster convergence.

The chromophore-dependent step size and the fluence-dependent spatial step size scaling [Eqs. (16) and (17)] proved to be vital

for achieving convergence. Without chromophore-dependent step sizes, Mb concentrations diverged to the upper and lower fit limits. Similarly, the fluence-dependent spatial step size scaling was crucial for achieving fast convergence in the regions of low fluence.

The selection of the excitation wavelengths could also be optimized further to improve inversion accuracy and convergence speed.⁴⁹ However, such a study would exceed the scope of this paper. Despite potentially suboptimal excitation wavelengths, the inversion is shown to recover blood sO_2 over a wide range [Fig. 7(a)] with high accuracy ($<1\%$ error in sO_2) across the domain.

Gradient-based methods do not guarantee convergence to a global minimum, especially when the inversion is adversely affected by a noisy gradient. While the Adam optimization algorithm (compared to, for example, standard or momentum gradient descent) has been shown to greatly reduce the likelihood of finishing the inversion in a local minimum or on a saddle point, such a result cannot be ruled out entirely. It should also be noted that the application of this method to measured PA images does not require their segmentation into subregions. While this makes the method generally valid, some form of image segmentation may still be advantageous as it would reduce the number of variables and the risk of convergence to a local minimum and increase convergence speed. Moreover, explicit regularization of the objective function could further improve the accuracy and speed of the convergence.

While the general methodology of a gradient-based inversion using an adjoint formulation of an MC model has been demonstrated *in silico*, the application of this approach to experimental 3-D PA images, especially those acquired *in vivo*, requires further investigation. One of the perhaps most critical points is the recovery of the scattering coefficient as it is likely to have an impact on the importance of the radiance approximation using spherical harmonics. Other issues, such as the choice of inversion parameters and the selection of optimal excitation wavelengths, are also important in the translation of QPAT methods toward applications in the medical and life sciences.

5 Conclusions

An inversion scheme for recovering absolute chromophore concentrations and their ratios, such as blood sO_2 from 3-D

multiwavelength PA images was developed and validated *in silico*. The scheme was based on an adjoint formulation of an MC light-transport model and allowed an approximation of the radiance using spherical harmonics. It was found that the adjoint radiance was not required to obtain accurate inversion results, provided the scattering coefficient was constant. The speed of convergence was increased by incorporating the Adam optimization algorithm, chromophore-dependent step sizes, and fluence-dependent step-size scaling. This work represented an important step in the development of robust and generally applicable methods for quantitative functional and molecular PA imaging.

6 Appendix A: Definition of the Adjoint Model

The idea of the adjoint formalism is to define nonphysical quantities, an adjoint source $q^*(\vec{r}, \lambda)$, adjoint radiance $\phi^*(\vec{r}, \vec{s}, \lambda)$, and an adjoint fluence $\Phi^*(\vec{r}, \lambda) = \int_{S^2} \phi^*(\vec{r}, \vec{s}, \lambda) d\vec{s}$ that help in replacing the integral term containing the unknown $\frac{\partial \Phi}{\partial c_i}$ in the definition of the gradient. In our case the gradient equation is

$$\begin{aligned} \frac{\partial \epsilon_{\lambda_i}}{\partial c_i} = & -[p_0^m(\vec{r}_i) - p_0(\vec{r}_i)] V_{\text{vox}} [\beta_k H(\vec{r}_i) + \Gamma(\vec{r}_i) \alpha_k \Phi(\vec{r}_i)] \\ & + \int_{\Omega} (p_0^m - p_0) \Gamma(\vec{r}) \mu_a(\vec{r}) \frac{\partial \Phi}{\partial c_i} d\Omega. \end{aligned} \quad (19)$$

The adjoint approach has been used in the context of PAT earlier, see e.g., Refs. 17–19, 26, and 27.

As a first step, the adjoint source term is defined as

$$q^*(\vec{r}, \lambda) = [p_0^m(\vec{r}, \lambda) - p_0(\vec{r}, \lambda)] \Gamma(\vec{r}) \mu_a(\vec{r}, \lambda). \quad (20)$$

Since the approach is targeted for multispectral QPAT, each wavelength requires its own definition of an adjoint source based on the difference between modeled $p_0(\vec{r}, \lambda)$ and measured data $p_0^m(\vec{r}, \lambda)$. We only denote one wavelength here and omit the dependence on λ for the sake of brevity.

The adjoint source is usually defined as the “prefactor” of the unknown term that is to be replaced and contains the error between modeled and measured data. By defining the behavior of the adjoint radiance ϕ^* and adjoint fluence Φ^* , a relationship between these and the desired $\frac{\partial \Phi}{\partial c_i}$ can be established, which is outlined in the following derivation. One important aspect of the definition of the adjoint quantity is to leave the equations underlying the development similar to the ones of their physical counterpart, which is in this context the time-independent RTE.

The RTE is given as

$$\begin{aligned} [\vec{s} \cdot \nabla + \mu_a(\vec{r}) + \mu_s(\vec{r})] \phi(\vec{r}, \vec{s}) - \mu_s \int_{S^2} \Theta(\vec{s}, \vec{s}') \phi(\vec{r}, \vec{s}') d\vec{s}' \\ = q(\vec{r}, \vec{s}). \end{aligned} \quad (21)$$

Similarly, we define the adjoint RTE (ARTE) as

$$\begin{aligned} [-\vec{s} \cdot \nabla + \mu_a(\vec{r}) + \mu_s(\vec{r})] \phi^*(\vec{r}, \vec{s}) \\ - \mu_s \int_{S^2} \Theta(\vec{s}, \vec{s}') \phi^*(\vec{r}, \vec{s}') d\vec{s}' = q^*(\vec{r}, \vec{s}). \end{aligned} \quad (22)$$

One advantage of defining the adjoint radiance in that way is that the propagation dynamics are identical to that of the normal radiance defined by the RTE since the left-hand side of both the

RTE and the ARTE are practically identical, the only difference being the negative sign in the ARTE indicating a change of direction in light propagation. One way to interpret the negative sign is to follow the propagation of photons in the opposite direction, which does not affect the photon’s movement and the final results in terms of energy deposit. Thus, the mechanisms for absorption and scattering remain the same as in the RTE. Because the light transport is dominated by scattering and absorption and not whether photons move in the forward or backward direction, light propagation can be seen as reciprocal and hence the numerical framework implementing the ARTE is unaffected by the additional negative sign. Thus, the same simulation code as for the RTE can be used for the ARTE. Only the difference in the source distributions needs to be taken into account, with the adjoint source being 3-D, whereas normal source distributions are usually 2-D.

It is important to note that the adjoint fluence and adjoint radiance have different physical units as their physical counterparts. The adjoint radiance’s unit is $J/(m^3 \text{ sr})$ and the adjoint fluence’s unit is $J/(m^3)$.

7 Appendix B: Adjoint-Assisted Derivation of the Gradient

Using the above definition of the adjoint source in Eq. (20), the substitution of the unknown term $\frac{\partial \Phi}{\partial c_i}$

$$\begin{aligned} \int_{\Omega} (p_0^m - p_0) \Gamma(\vec{r}) \mu_a(\vec{r}) \frac{\partial \Phi}{\partial c_i} d\Omega \\ = -\alpha_k \left[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s} \right]_{\vec{r}=\vec{r}_i} V_{\text{vox}}, \end{aligned} \quad (23)$$

can be derived.

The derivation follows the ideas presented in previous works, in particular Ref. 27. In Ref. 27, $\frac{\partial \text{RTE}}{\partial c_i}$ is combined with the ARTE

$$\phi^* \cdot \frac{\partial \text{RTE}}{\partial c_i} - \frac{\partial \phi}{\partial c_i} \cdot \text{ARTE}. \quad (24)$$

The term $\frac{\partial \text{RTE}}{\partial c_i}$ is

$$\begin{aligned} [\vec{s} \cdot \nabla + \mu_a(\vec{r}) + \mu_s(\vec{r})] \frac{\partial \phi(\vec{r}, \vec{s})}{\partial c_i} + \frac{\partial \mu_a}{\partial c_i} \phi(\vec{r}, \vec{s}) \\ - \mu_s \int_{S^2} \Theta(\vec{s}, \vec{s}') \frac{\partial \phi(\vec{r}, \vec{s}')}{\partial c_i} d\vec{s}' = 0, \end{aligned} \quad (25)$$

as $\frac{\partial q}{\partial c_i} = 0$ since the external light source does not depend on chromophore concentrations. The basic idea underlying the following steps is to rearrange all the terms in Eq. (24) so that all terms on the left-hand side can be set to zero after integrating over space and angles. First, we insert $\frac{\partial \mu_a}{\partial c_i} = \alpha \delta(\vec{r} - \vec{r}_i)$ and rearrange Eq. (25)

$$\begin{aligned} (\vec{s} \cdot \nabla + \mu_a(\vec{r}) + \mu_s(\vec{r})) \frac{\partial \phi(\vec{r}, \vec{s})}{\partial c_i} - \mu_s \int_{S^2} \Theta(\vec{s}, \vec{s}') \frac{\partial \phi(\vec{r}, \vec{s}')}{\partial c_i} d\vec{s}' \\ = -\alpha(\vec{r}) \phi(\vec{r}, \vec{s}) \delta(\vec{r} - \vec{r}_i). \end{aligned} \quad (26)$$

The combination of the ARTE and the RTE from Eq. (24) is (for brevity we omit the dependency on \vec{r} and \vec{s} for the moment)

$$\begin{aligned}
 & \phi^*(\vec{s} \cdot \nabla + \mu_a + \mu_s) \frac{\partial \phi}{\partial c_i} - \phi^* \mu_s \int_{S^2} \Theta(\vec{s}, \vec{s}') \frac{\partial \phi}{\partial c_i} d\vec{s}' \\
 & - \frac{\partial \phi}{\partial c_i} (-\vec{s} \cdot \nabla + \mu_a + \mu_s) \phi + \frac{\partial \phi}{\partial c_i} \mu_s \int_{S^2} \Theta(\vec{s}, \vec{s}') \phi^* d\vec{s}' \\
 & = -\alpha \phi^* \phi \delta(\vec{r} - \vec{r}_i) - \frac{\partial \phi}{\partial c_i} q^*. \quad (27)
 \end{aligned}$$

The left-hand side can be simplified and can be given as

$$\begin{aligned}
 & \phi^*(\vec{s} \cdot \nabla) \frac{\partial \phi}{\partial c_i} + \frac{\partial \phi}{\partial c_i} (\vec{s} \cdot \nabla) \phi^* - \phi^* \mu_s \int_{S^2} \Theta(\vec{s}, \vec{s}') \frac{\partial \phi}{\partial c_i} d\vec{s}' \\
 & + \frac{\partial \phi}{\partial c_i} \mu_s \int_{S^2} \Theta(\vec{s}, \vec{s}') \phi^* d\vec{s}' \\
 & = -\alpha \phi^* \phi \delta(\vec{r} - \vec{r}_i) - \frac{\partial \phi}{\partial c_i} q^*. \quad (28)
 \end{aligned}$$

The next steps of the derivation are from now on identical to Eqs. (4.11 ff) in Ref. 27. The left-hand side of Eq. (28) equates to zero, which can be seen after integrating first over all angles $s \in S^{n-1}$ and over the volume Ω with surface $\partial\Omega$:

$$\begin{aligned}
 & \int_{\Omega} \int_{S^2} \phi^*(\vec{s} \cdot \nabla) \frac{\partial \phi}{\partial c_i} d\vec{s} d\Omega + \int_{\Omega} \int_{S^2} \frac{\partial \phi}{\partial c_i} (\vec{s} \cdot \nabla) \phi^* d\vec{s} d\Omega \\
 & - \int_{\Omega} \mu_s \int_{S^2} \phi^*(\vec{s}) \int_{S^2} \Theta(\vec{s}, \vec{s}') \frac{\partial \phi(\vec{s}')}{\partial c_i} d\vec{s}' d\vec{s} d\Omega \\
 & + \int_{\Omega} \mu_s \int_{S^2} \frac{\partial \phi(\vec{s})}{\partial c_i} \int_{S^2} \Theta(\vec{s}, \vec{s}') \phi^*(\vec{s}') d\vec{s}' d\vec{s} d\Omega \\
 & = - \int_{\Omega} \alpha \delta(\vec{r} - \vec{r}_i) \int_{S^2} \phi^* \phi d\vec{s} d\Omega - \int_{\Omega} q^* \int_{S^2} \frac{\partial \phi}{\partial c_i} d\vec{s} d\Omega. \quad (29)
 \end{aligned}$$

To see that the left-hand side equates to zero the volume integral with the terms involving ∇ can be transformed into a surface integral using this form of the divergence theorem:

$$\int_{\Omega} ab \cdot \nabla cd\Omega + \int_{\Omega} cb \cdot \nabla ad\Omega = \int_{\partial\Omega} b \cdot \hat{n} acd\Omega, \quad (30)$$

along with the following substitutions:

$$a = \phi^*(\vec{s}), \quad b = \vec{s} \quad \text{and} \quad c = \frac{\partial \phi(\vec{s})}{\partial c_i}. \quad (31)$$

Hence, using the divergence theorem, the first two terms in Eq. (28) are replaced by a single term

$$\begin{aligned}
 & \int_{\Omega} \int_{S^2} \phi^*(\vec{s} \cdot \nabla) \frac{\partial \phi}{\partial c_i} d\vec{s} d\Omega + \int_{\Omega} \int_{S^2} \frac{\partial \phi}{\partial c_i} (\vec{s} \cdot \nabla) \phi^* d\vec{s} d\Omega \\
 & = \int_{\partial\Omega} \int_{S^2} (\vec{s} \cdot \hat{n}) \phi^*(\vec{s}) \frac{\partial \phi(\vec{s})}{\partial c_i} d\vec{s} d\Omega.
 \end{aligned}$$

By definition, both $\phi^* \rightarrow 0$ and $\partial\Phi/\partial c_i \rightarrow 0$ on the boundary of the volume $\partial\Omega$. Thus, the integrand on the right-hand side and hence the integral equate to zero.

This reduces Eq. (28) to

$$\begin{aligned}
 & - \int_{\Omega} \mu_s \int_{S^2} \phi^*(\vec{s}) \int_{S^2} \Theta(\vec{s}, \vec{s}') \frac{\partial \phi(\vec{s}')}{\partial c_i} d\vec{s}' d\vec{s} d\Omega \\
 & + \int_{\Omega} \mu_s \int_{S^2} \frac{\partial \phi(\vec{s})}{\partial c_i} \int_{S^2} \Theta(\vec{s}, \vec{s}') \phi^*(\vec{s}') d\vec{s}' d\vec{s} d\Omega \\
 & = - \int_{\Omega} \alpha \delta(\vec{r} - \vec{r}_i) \int_{S^2} \phi^* \phi d\vec{s} d\Omega - \int_{\Omega} q^* \int_{S^2} \frac{\partial \phi}{\partial c_i} d\vec{s} d\Omega. \quad (32)
 \end{aligned}$$

Because we can assume that all functions are integrable, the terms on the left-hand side of Eq. (32) can be rearranged after changing the order of integration, yielding

$$\begin{aligned}
 & \int_{\Omega} \mu_s \int_{S^2} \int_{S^2} \Theta(\vec{s}, \vec{s}') \phi^*(\vec{s}) \frac{\partial \phi(\vec{s}')}{\partial c_i} d\vec{s}' d\vec{s} d\Omega \\
 & - \int_{\Omega} \mu_s \int_{S^2} \int_{S^2} \Theta(\vec{s}, \vec{s}') \phi^*(\vec{s}') \frac{\partial \phi(\vec{s})}{\partial c_i} d\vec{s} d\vec{s}' d\Omega.
 \end{aligned}$$

Hence, the left-hand side of Eq. (32) equates to zero, which leaves us with

$$\int_{\Omega} q^* \int_{S^2} \frac{\partial \phi}{\partial c_i} d\vec{s} d\Omega = - \int_{\Omega} \alpha \delta(\vec{r} - \vec{r}_i) \int_{S^2} \phi^* \phi d\vec{s} d\Omega. \quad (33)$$

The fluence Φ is by definition the integral of the time-integrated radiance $\phi(\vec{s})$ over all directions \vec{s} , that is,

$$\Phi(\vec{r}) = \int_{S^2} \phi(\vec{r}, \vec{s}) d\vec{s}. \quad (34)$$

Applying the derivative with respect to c_i yields

$$\frac{\partial \Phi(\vec{r})}{\partial c_i} = \int_{S^2} \frac{\partial \phi(\vec{r}, \vec{s})}{\partial c_i} d\vec{s}, \quad (35)$$

which lets us write Eq. (33) as

$$\int_{\Omega} q^* \frac{\partial \Phi}{\partial c_i} d\Omega = -\alpha(\vec{r}_i) \left[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s} \right]_{\vec{r}=\vec{r}_i} V_{\text{vox}}. \quad (36)$$

Because we have defined the adjoint source term as $q^* = (p_0^m - p_0) \Gamma \mu_a$ the left-hand side of Eq. (36) is exactly the last term in Eq. (19)

$$\int_{\Omega} (p_0^m - p_0) \Gamma \mu_a \frac{\partial \Phi}{\partial c_i} d\Omega = -\alpha(\vec{r}_i) \left[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s} \right]_{\vec{r}=\vec{r}_i} V_{\text{vox}}. \quad (37)$$

Inserting this result into the error-gradient Eq. (19) provides us with the subgradient term obtained using one wavelength

$$\begin{aligned}
 \frac{\partial \epsilon_{\lambda_l}}{\partial c_i} & = -[p_0^m(\vec{r}_i) - p_0(\vec{r}_i)] \alpha(\lambda_l) \Phi(\vec{r}_i) V_{\text{vox}} \\
 & + \alpha(\vec{r}_i) \left[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s} \right]_{\vec{r}=\vec{r}_i} V_{\text{vox}}. \quad (38)
 \end{aligned}$$

In summary, the adjoint formalism enables the update of the concentration distribution using only terms obtained from

running the forward model implemented by the RMC algorithm. The following section will focus on the question as to how the term $\phi^* \phi$ can be computed and approximated.

8 Appendix C: Radiance Approximations

To simulate the radiance, some discretization over angle is required. One option is to use a set of piecewise constant basis functions. However, due to the importance of ballistic and quasi-ballistic light propagation in PAT, a high number of discretization orders would be required to capture the directionality of the radiance in regions near the source. This would result in very large memory demands in finite-element implementations (see Refs. 29, 50, and 51, for more details on different approaches for estimating the radiance, their limitations, and suggested solutions. For a summary thereof, see Ref. 52). Inspired by the P_n approximation²² and continuing the work presented by Refs. 19 and 52 we approximate the radiance in 3-D using spherical harmonics as basic functions, as in Ref. 31. Instead of discretizing the angular domain into segments, the radiance field ϕ at any position \vec{r} can be expanded using a series of spherical harmonics^{53,54}

$$\phi(\vec{r}, \vec{s}) = \sum_{l=0}^{N_L} \sum_{m=-l}^l \psi_{lm}(\vec{r}) Y_{lm}(\vec{r}, \vec{s}), \quad (39)$$

where $\psi_{lm}(\vec{r})$ are the coefficients corresponding to the real spherical harmonics $Y_{lm}(\vec{r}, \vec{s})$, expressed as

$$Y_{lm} = \begin{cases} \sqrt{2} \sqrt{\frac{2l+1}{4\pi}} \sqrt{\frac{(l-|m|)!}{(l+|m|)!}} P_l^{|m|} \cos(\theta) \sin(|m|\phi) & \text{if } m < 0, \\ \sqrt{\frac{2l+1}{4\pi}} P_l^m [\cos(\theta)] & \text{if } m = 0, \\ \sqrt{2} \sqrt{\frac{2l+1}{4\pi}} \sqrt{\frac{(l-m)!}{(l+m)!}} P_l^m \cos(\theta) \cos(m\phi) & \text{if } m > 0, \end{cases} \quad (40)$$

where l is the degree of the spherical harmonic, m is the order, and P_l^m is the associated Legendre polynomials. The coefficient $\psi_{lm}(\vec{r})$ scales the total weight deposited by all simulated photons at position (voxel) \vec{r} for the associated spherical harmonic.

The advantage of expressing $\phi^* \phi$ in a spherical harmonics expansion lies in the fact that the Y_{lm} forms an orthonormal basis, i.e.,

$$\int_{S^2} Y_{lm} \cdot Y_{l'm'} d\vec{s} = \delta_{ll'} \delta_{mm'}. \quad (41)$$

Using this orthonormality condition greatly simplifies the term $[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s}]_{\vec{r}=\vec{r}_i}$ from Eq. (38), when the radiance is expressed in spherical harmonics

$$\begin{aligned} \left[\int_{S^2} \phi^*(\vec{s}) \phi(\vec{s}) d\vec{s} \right]_{\vec{r}=\vec{r}_i} &= \int_{S^2} \left[\sum_{l=0}^{N_L} \sum_{m=-l}^l \psi_{lm}^*(\vec{r}_i) Y_{lm}^*(\vec{r}_i, \vec{s}) \right] \\ &\quad \times \left[\sum_{l'=0}^{N_L} \sum_{m'=-l'}^{l'} \psi_{l'm'}(\vec{r}_i) Y_{l'm'}(\vec{r}_i, \vec{s}) \right] d\vec{s}, \end{aligned} \quad (42)$$

$$= \left[\sum_{l'=0}^{N_L} \sum_{m'=-l'}^{l'} \psi_{l'm'}(\vec{r}_i) \right] \left[\underbrace{\int_{S^2} Y_{lm}^*(\vec{r}_i, \vec{s}) Y_{l'm'}(\vec{r}_i, \vec{s}) d\vec{s}}_{\delta_{ll'} \delta_{mm'}} \right] \quad (43)$$

$$= \sum_{l=0}^{N_L} \sum_{m=-l}^l \psi_{lm}(\vec{r}_i) \psi_{lm}^*(\vec{r}_i). \quad (44)$$

Hence, with this approximation the gradient to update the distribution of chromophore k finally becomes

$$\begin{aligned} \frac{\partial \varepsilon_{\lambda_i}}{\partial c_i} &= V_{\text{vox}} \left\{ -[p_0^m(\vec{r}_i) - p_0(\vec{r}_i)] [\beta_k H(\vec{r}_i) + \Gamma(\vec{r}_i) \alpha_k \Phi(\vec{r}_i)] \right. \\ &\quad \left. + \alpha_k \sum_{l=0}^{N_L} \sum_{m=-l}^l \psi_{lm}(\vec{r}_i) \psi_{lm}^*(\vec{r}_i) \right\}. \end{aligned} \quad (45)$$

9 Appendix D: Discretization

To solve the given equations numerically, the bases in which the data and model output are represented must be defined. Assuming a sampling of continuous fields in a point-wise basis $\Psi_j(\vec{r}) = \delta(\vec{r} - \vec{r}_j)$, as shown in Ref. 18. Hence, the data projected onto this basis becomes a vector of coefficients \vec{p}_0^m

$$h_j = \langle \Psi_j, p_0^m(\vec{r}) \rangle = p_0^m(\vec{r}_j). \quad (46)$$

Equally, all other continuous fields are discretized. Transforming an integral of any integrable function $f(\vec{r})$ over the continuous domain Ω into discretized space introduces a volume element $d\Omega = V_{\text{vox}}$

$$\int_{\Omega} f(\vec{r}) d\Omega = \sum_j^{N_{\text{vox}}} \langle \Psi_j, f(\vec{r}) \rangle V_{\text{vox}} = \sum_j^{N_{\text{vox}}} f(\vec{r}_j) V_{\text{vox}}. \quad (47)$$

Disclosures

The authors declare that there are no conflicts of interest related to this article.

Acknowledgments

The authors would like to thank Simon Arridge and Ben Cox for valuable discussions. This project was funded by the Deutsche Forschungsgemeinschaft (DFG project grants LA3273/1-1, PR1226/5-1), the Royal Academy of Engineering (RAEng Fellowship RF1516\15\33), and the Engineering and Physical Sciences Research Council (EPSRC grant EP/N032055/1). We acknowledge financial support within the funding programme Open Access Publishing by the German Research Foundation (DFG).

References

1. P. Beard, "Biomedical photoacoustic imaging," *Interface Focus* **1**(4), 602–631 (2011).
2. J. Xia, J. Yao, and L. Wang, "Photoacoustic tomography: principles and advances," *Electromagn. Waves* **147**, 1–22 (2014).
3. L. Wang and J. Yao, "A practical guide to photoacoustic tomography in the life sciences," *Nat. Methods* **13**, 627–638 (2016).

4. B. Cox, J. Laufer, and P. Beard, "The challenges for quantitative photoacoustic imaging," *Proc. SPIE* **7177**, 717713 (2009).
5. B. Cox et al., "Quantitative spectroscopic photoacoustic imaging: a review," *J. Biomed. Opt.* **17**(6), 061202 (2012).
6. C. Lutzweiler and D. Razansky, "Optoacoustic imaging and tomography: reconstruction approaches and outstanding challenges in image performance and quantification," *Sensors* **13**(6), 7345–7384 (2013).
7. M. Li, Y. Tang, and J. Yao, "Photoacoustic tomography of blood oxygenation: a mini review," *Photoacoustics* **10**, 65–73 (2018).
8. X. Wang et al., "Noninvasive imaging of hemoglobin concentration and oxygenation in the rat brain using high-resolution photoacoustic tomography," *J. Biomed. Opt.* **11**(2), 024015 (2006).
9. D. Razansky et al., "Multispectral opto-acoustic tomography of deep-seated fluorescent proteins in vivo," *Nat. Photonics* **3**, 412–417 (2009).
10. S. Tzoumas et al., "Eigenspectra optoacoustic tomography achieves quantitative blood oxygenation imaging deep in tissues," *Nat. Commun.* **7**, 12121 (2016).
11. T. Kirchner, J. Gröhl, and L. Maier-Hein, "Context encoding enables machine learning-based quantitative photoacoustics," *J. Biomed. Opt.* **23**(5), 056008 (2018).
12. G. Bal and K. Ren, "On multi-spectral quantitative photoacoustic tomography in diffusive regime," *Inverse Prob.* **28**(2), 025010 (2012).
13. R. J. Zemp, "Quantitative photoacoustic tomography with multiple optical sources," *Appl. Opt.* **49**, 3566–3572 (2010).
14. B. T. Cox, S. R. Arridge, and P. C. Beard, "Estimating chromophore distributions from multiwavelength photoacoustic images," *J. Opt. Soc. Am. A* **26**, 443–455 (2009).
15. J. Laufer et al., "Quantitative determination of chromophore concentrations from 2D photoacoustic images using a nonlinear model-based inversion scheme," *Appl. Opt.* **49**(8), 1219–1233 (2010).
16. B. A. Kaplan et al., "Monte-Carlo-based inversion scheme for 3D quantitative photoacoustic tomography," *Proc. SPIE* **10064**, 100645J (2017).
17. B. Cox, S. Arridge, and P. Beard, "Gradient-based quantitative photoacoustic image reconstruction for molecular imaging," *Proc. SPIE* **6437**, 64371T (2007).
18. E. Malone et al., "Reconstruction-classification method for quantitative photoacoustic tomography," *J. Biomed. Opt.* **20**(12), 126004 (2015).
19. R. Hochuli et al., "Quantitative photoacoustic tomography using forward and adjoint Monte Carlo models of radiance," *J. Biomed. Opt.* **21**(12), 126004 (2016).
20. M. Fonseca et al., "Three-dimensional photoacoustic imaging and inversion for accurate quantification of chromophore distributions," *Proc. SPIE* **10064**, 1006415 (2017).
21. G. Bal and G. Uhlmann, "Inverse diffusion theory of photoacoustics," *Inverse Prob.* **26**(8), 085010 (2010).
22. S. R. Arridge, "Optical tomography in medical imaging," *Inverse Prob.* **15**(2), R41 (1999).
23. T. Tarvainen et al., "Reconstructing absorption and scattering distributions in quantitative photoacoustic tomography," *Inverse Prob.* **28**(8), 084009 (2012).
24. A. V. Mamonov and K. Ren, "Quantitative photoacoustic imaging in the radiative transport regime," *Commun. Math. Sci.* **12**(2), 201–234 (2014).
25. S. Rabanser, L. Neumann, and M. Haltmeier, "Stochastic proximal gradient algorithms for multi-source quantitative photoacoustic tomography," *Entropy* **20**(2), 121 (2018).
26. T. Saratoon et al., "A gradient-based method for quantitative photoacoustic tomography using the radiative transfer equation," *Inverse Prob.* **29**(7), 075006 (2013).
27. T. Soonthornsaratoon, "Gradient-based methods for quantitative photoacoustic tomography," PhD Thesis, University College London (UCL) (2014).
28. S. T. Flock et al., "Monte Carlo modeling of light propagation in highly scattering tissues. I. Model predictions and comparison with diffusion theory," *IEEE Trans. Biomed. Eng.* **36**(12), 1162–1168 (1989).
29. L. Wang and S. L. Jacques, "Hybrid model of Monte Carlo simulation and diffusion theory for light reflectance by turbid media," *J. Opt. Soc. Am. A* **10**, 1746–1752 (1993).
30. Q. Fang and D. A. Boas, "Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units," *Opt. Express* **17**(22), 20178–20190 (2009).
31. S. Powell, R. Hochuli, and S. R. Arridge, "Radiance Monte-Carlo for application of the radiative transport equation in the inverse problem of diffuse optical tomography," *Proc. SPIE* **10059**, 100590W (2017).
32. J. Buchmann et al., "Experimental validation of a Monte-Carlo-based inversion scheme for 3D quantitative photoacoustic tomography," *Proc. SPIE* **10064**, 1006416 (2017).
33. D. Kingma and J. Ba, "Adam: a method for stochastic optimization," arXiv:1412.6980 (2014).
34. M. B. Fonseca, L. An, and B. T. Cox, "Sulfates as chromophores for multiwavelength photoacoustic imaging phantoms," *J. Biomed. Opt.* **22**(12), 125007 (2017).
35. D.-K. Yao et al., "Photoacoustic measurement of the Grüneisen parameter of tissue," *J. Biomed. Opt.* **19**(1), 017007 (2014).
36. S. A. Prahl, "A monte carlo model of light propagation in tissue," *Proc. SPIE* **10305**, 1030509 (1989).
37. L. G. Henyey and J. L. Greenstein, "Diffuse radiation in the galaxy," *Astrophys. J.* **93**, 70–83 (1941).
38. J. Beazans et al., "Julia: a fresh approach to numerical computing," *SIAM Rev.* **59**(1), 65–98 (2017).
39. J. Laufer et al., "In vivo preclinical photoacoustic imaging of tumor vasculature development and therapy," *J. Biomed. Opt.* **17**(5), 056016 (2012).
40. J. Märk et al., "Dual-wavelength 3D photoacoustic imaging of mammalian cells using a photoswitchable phytochrome reporter protein," *Commun. Phys.* **1**(1), 3 (2018).
41. E. Zhang, J. Laufer, and P. Beard, "Backward-mode multiwavelength photoacoustic scanner using a planar Fabry-Pérot polymer film ultrasound sensor for high-resolution three-dimensional imaging of biological tissues," *Appl. Opt.* **47**(4), 561–577 (2008).
42. J. Buchmann et al., "Characterization and modeling of Fabry-Pérot ultrasound sensors with hard dielectric mirrors for photoacoustic imaging," *Appl. Opt.* **56**(17), 5039–5046 (2017).
43. B. G. Barratt-Boyes and E. H. Wood, "The oxygen saturation of blood in the venae cavae, right-heart chambers, and pulmonary vessels of healthy subjects," *J. Lab. Clin. Med.* **50**(1), 93–106 (1957).
44. E. Beutler and J. Waalen, "The definition of anemia: what is the lower limit of normal of the blood hemoglobin concentration?" *Blood* **107**(5), 1747–1750 (2006).
45. G. Serša et al., "Contrast enhanced MRI assessment of tumor blood volume after application of electric pulses," *Electro- Magnetobiol.* **17**(2), 299–306 (1998).
46. S. Prahl, "Assorted Spectra," <https://omlc.org/spectra/index.html> (2017).
47. B. T. Polyak, "Some methods of speeding up the convergence of iteration methods," *USSR Comput. Math. Math. Phys.* **4**(5), 1–17 (1964).
48. G. Goh, "Why momentum really works," *Distill* (2017).
49. G. P. Luke, S. Y. Nam, and S. Y. Emelianov, "Optical wavelength selection for improved spectroscopic photoacoustic imaging," *Photoacoustics* **1**(2), 36–42 (2013).
50. T. Tarvainen et al., "Hybrid radiative-transfer-diffusion model for optical tomography," *Appl. Opt.* **44**(6), 876–886 (2005).
51. T. Tarvainen, "Computational methods for light transport in optical tomography," PhD Thesis, Faculty of Natural and Environmental Sciences, Department of Physics, University of Kuopio (2006).
52. R. Hochuli, "Monte Carlo methods in quantitative photoacoustic tomography," PhD Thesis, University College London (UCL) (2016).
53. F. Inanc and A. Rohach, "A nodal solution of the multigroup neutron transport equation using spherical harmonics," *Ann. Nucl. Energy* **15**(10–11), 501–509 (1988).
54. A. G. Buchan et al., "A POD reduced order model for resolving angular direction in neutron/photon transport problems," *J. Comput. Phys.* **296**, 138–157 (2015).

Biographies of the authors are not available.