Forward and Adjoint Radiance Monte Carlo Models for Quantitative Photoacoustic Imaging

Roman Hochuli [a], Samuel Powell [b], Simon Arridge [b] and Ben Cox [a]

[a]Department of Medical Physics & Biomedical Engineering, University College London, UK.
[b]Department of Computer Science, University College London, UK.

ABSTRACT

In quantitative photoacoustic imaging, the aim is to recover physiologically relevant tissue parameters such as chromophore concentrations or oxygen saturation. Obtaining accurate estimates is challenging due to the non-linear relationship between the concentrations and the photoacoustic images. Nonlinear least squares inversions designed to tackle this problem require a model of light transport, the most accurate of which is the radiative transfer equation. This paper presents a highly scalable Monte Carlo model of light transport that computes the radiance in 2D using a Fourier basis to discretise in angle. The model was validated against a 2D finite element model of the radiative transfer equation, and was used to compute gradients of an error functional with respect to the absorption and scattering coefficient. It was found that adjoint-based gradient calculations were much more robust to inherent Monte Carlo noise than a finite difference approach. Furthermore, the Fourier angular discretisation allowed very efficient gradient calculations as sums of Fourier coefficients. These advantages, along with the high parallelisability of Monte Carlo models, makes this approach an attractive candidate as a light model for quantitative inversion in photoacoustic imaging.

Keywords: quantitative, photoacoustics, Monte Carlo, radiance

1. INTRODUCTION

Photoacoustic imaging (PAI) is a technique that utilises nanosecond, near-infrared pulses of laser light to generate MHz frequency acoustic waves in tissue which can then be detected at the tissue surface, thereby exploiting highly selective optical absorption in biological tissue and the low scattering undergone by ultrasonic waves. There are two inverse problems in PAI: the acoustic inverse problem, which describes the reconstruction of the initial acoustic pressure distribution from boundary acoustic measurements, and the optical inverse problem, which involves the recovery of tissue optical properties or absorber concentrations from the initial acoustic pressure. It is the optical inverse problem, sometimes called *quantitative PAI*, that is of interest here. Solving it - which remains a significant challenge - is of considerable interest as it would allow high resolution, 3D, 'molecular imaging': the quantification of physiologically relevant properties such as oxygen saturation, haemoglobin concentration, or the concentrations of other endogenous or exogenous chromophores. This paper proposes radiance Monte Carlo as an accurate and highly scalable forward model, well suited to tackling this inverse problem.

2. QUANTITATIVE PHOTOACOUSTIC IMAGING

The inversion in *quantitative* PAI is not straightforward as the initial pressure distribution (the photoacoustic image) is not proportional to the absorption coefficient but to the absorbed energy density, H, which is the product of the absorption coefficient and the fluence:

$$H(\mathbf{x}, \lambda) = \mu_a \Phi(\mu_a) = \Phi(\mathbf{x}, \lambda; c_k) \sum_{k=1}^{K} \alpha_k(\lambda) c_k(\mathbf{x}), \tag{1}$$

where Φ is the light fluence, μ_a is the absorption coefficient, α_k is the molar absorption coefficient for the k^{th} chromophore, c_k is the concentration of the k^{th} chromophore, λ is the wavelength of the light fluence, and

Send correspondence to: r.hochuli@ucl.ac.uk

x is position (the fluence will also, in general, depend on how the light is scattered). Thus, although the chromophore concentrations are linearly related to the absorption coefficient, the absorption coefficient is not linearly related to the photoacoustic image, and knowledge of the fluence is required. The internal fluence is unknown in most situations as it depends on the optical properties of the medium, making it both positionally- and wavelength-dependent. A general approach to estimating μ_a or c_k is to iteratively minimise the difference between the measured photoacoustic image and the output of a forward model, for example, by solving the nonlinear least-squares problem

$$\mu_a = \underset{\mu_a}{\arg\min} \, \epsilon(\mu_a), \quad \epsilon = \frac{1}{2} \int_\Omega \left(H_{\text{meas}} - H(\mu_a, \mu_s) \right)^2 d\Omega, \tag{2}$$

where H_{meas} is the measured image, $H(\mu_a, \mu_s)$ is the simulated image at the current estimate of absorption coefficient μ_a and scattering coefficient μ_s, and Ω is the domain. As PA images are high resolution, this is a large scale inversion, and computationally efficient methods are required to solve it. For quasi-Newton inversion or gradient-descent schemes, estimates of the functional gradients, $\partial\epsilon/\partial\mu_a$ and $\partial\epsilon/\partial\mu_s$, are required. Gradient-based schemes, such as quasi-Newton techniques, have been used to invert simulated PAT data using the radiative transfer equation,[1] and the diffusion approximation,[2] as fluence models.

2.1 Models of Light Transport

For light travelling through highly scattering media such as biological tissue, solving Maxwell's electrodynamic equations to model the light distribution is not tractable, and an energy-based approach is usually used, summarised in the time-independent case by the radiative transfer equation (RTE):

$$(\hat{\mathbf{s}} \cdot \nabla + \mu_a(\mathbf{x}) + \mu_s(\mathbf{x}))\phi(\mathbf{x}, \hat{\mathbf{s}}) - \mu_s(\mathbf{x}) \int_{\mathcal{S}^{N-1}} P_\theta(\hat{\mathbf{s}}, \hat{\mathbf{s}}')\phi(\mathbf{x}, \hat{\mathbf{s}}')d\hat{\mathbf{s}}' = q(\mathbf{x}, \hat{\mathbf{s}}), \tag{3}$$

where ϕ is the radiance, μ_a and μ_s are the absorption and scattering coefficients, respectively, $\mathbf{x}$ is position, $\hat{\mathbf{s}}'$ and $\hat{\mathbf{s}}$ are the original and scattered propagation directions, $P_\theta(\hat{\mathbf{s}}, \hat{\mathbf{s}}')$ is the scattering phase function, $q(\mathbf{x}, \hat{\mathbf{s}})$ is a source term and $\mathcal{S}^{N-1}$ is used to indicate integration over angle in $N-1$ dimensions ($N=2$ in 2D). The RTE is often approximated in highly diffuse regions, where the radiance is not strongly dependent on angle, by the diffusion approximation (DA) for the fluence (the angle-integrated radiance).[3] This is straightforward to solve using standard numerical approaches, but is inaccurate at regions close to the source or boundaries, which are often of interest in PAI. On the other hand, the RTE accurately models light in both the ballistic and diffusive regimes, but due to the domain having both spatial and angular discretisation, the memory requirements for commonly used solution approaches, such as finite elements, grow rapidly. For example, a regular cuboid mesh containing 100 nodes along one edge (a typical value for a PAT image) with a 32 angular bins, and $\sim 10^6$ elements, forms a system matrix containing a maximum number of non-zero elements of ~ 60 million which, to single precision, requires ~ 250GB of memory merely to store.

An alternative approach to estimating solutions to the RTE is to use a Monte Carlo (MC) simulation,[4] a stochastic model that involves propagating 'packets' of energy through the tissue, depositing energy along their trajectory. As the number of photons used grows, the fluence estimate tends to a solution of the RTE. As the individual photons do not interact, this method is highly parallelisable, which is very attractive as parallel computing becomes increasingly mainstream. In this paper, we describe a 2D MC code that computes the angle-dependent radiance, rather than the fluence (the essential principles are directly applicable in 3D); spatial discretisation is performed in a piecewise constant basis, and angular discretisation in a Fourier-basis. The advantages of Fourier discretisation are two-fold: first, it allows compact storage of the radiance and improves convergence of radiance estimates due to the sparsity of a diffuse field in this basis. Second, it allows efficient and straightforward computation of gradients of the error function with respect to model parameters. These advantages are discussed in Sections 3 and 4.2.

2.2 Functional Gradient Calculation with the RTE

The adjoint approach for calculating the functional gradients $\partial\epsilon/\partial\mu_a$ and $\partial\epsilon/\partial\mu_s$ when using the RTE as the light model is described by Saratoon et al.[1] The derivatives of ϵ with respect to μ_a and μ_s can be expressed

using an adjoint solution of the RTE, ϕ^*, which solves the equation:

$$(-\hat{\mathbf{s}} \cdot \nabla + \mu_a(\mathbf{x}) + \mu_s(\mathbf{x}))\phi^*(\mathbf{x}, \hat{\mathbf{s}}) - \mu_s(\mathbf{x}) \int_{\mathcal{S}^1} P_\theta(\hat{\mathbf{s}}, \hat{\mathbf{s}}')\phi^*(\mathbf{x}, \hat{\mathbf{s}}')d\hat{\mathbf{s}}' = \mu_a(H_{\text{meas}} - H(\mu_a, \mu_s)), \tag{4}$$

where $\mathcal{S}^1$ is used to indicate that the integral is performed over a circumference in 2D. The source term is isotropic, so in a MC simulation photons must be launched isotropically from all positions where $\mu_a(H_{\text{meas}} - H(\mu_a, \mu_s)) \neq 0$ with their weights scaled such that $\sum_p W_p \propto \mu_a(H_{\text{meas}} - H(\mu_a, \mu_s))$, for p photons per source position. Once the radiance ϕ and adjoint radiance ϕ^* have been calculated, the functional gradients can be found using[1]

$$\frac{d\epsilon}{d\mu_a} = -\Phi(\mu_a, \mu_s)(H_{\text{meas}} - H(\mu_a, \mu_s)) + \int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}; \mu_a, \mu_s)\phi^*(\hat{\mathbf{s}}; \mu_a, \mu_s)d\hat{\mathbf{s}} \tag{5}$$

$$\frac{d\epsilon}{d\mu_s} = \int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}'; \mu_a, \mu_s)\phi^*(\hat{\mathbf{s}}; \mu_a, \mu_s)d\hat{\mathbf{s}} - \int_{\mathcal{S}^1}\int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}; \mu_a, \mu_s)P_\theta(\hat{\mathbf{s}}, \hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}; \mu_a, \mu_s)d\hat{\mathbf{s}}d\hat{\mathbf{s}}'. \tag{6}$$

This approach has the significant benefit that only *two* model runs are required to compute the functional gradients with respect to absorption and scattering; by contrast, the number of model runs required to compute the gradient via a finite difference scheme (discussed in Section 4.3) would be greater than two times the number of pixels in the domain.

3. MONTE CARLO MODELING OF RADIANCE

A traditional MC algorithm to simulate light transport involves launching a photon from the source at position $\mathbf{x}$ in direction $\hat{\mathbf{s}}'$ with weight W_0, the photon then travels a distance dl in the current direction; at the new position, $\mathbf{x}' = \mathbf{x} + \hat{\mathbf{s}}'dl$, some weight $\delta W = W_0(1 - \exp(-\mu_a dl))$ is deposited (neglecting the photon packet's angle of incidence), and the photon continues this process until the photon has travelled a total distance $l = \int dl$, at which point it is scattered into the new direction $\hat{\mathbf{s}}$, which in 2D is determined by sampling the phase function:

$$P_\theta(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'; g) = \frac{1}{2\pi} \frac{1 - g^2}{(1 + g^2) - 2g(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}')}, \tag{7}$$

where g is the anisotropy factor. This process is repeated until the photon weight has decayed below some threshold value, for all N_p photons. As $N_p \to \infty$, the quantity being simulated tends to a solution of the RTE. The details of absorption and scattering in MC are discussed by Sassaroli[5] who provides a clear derivation of different, equivalent methods, as well as Wang and Jacques.[4]

In order to simulate the angularly-dependent radiance, the deposition of weight must not only occur at a certain position $\mathbf{x}'$, it must also be deposited in a direction $\hat{\mathbf{s}}'$. Spatial discretisation can be achieved through a voxelised grid or a mesh, while perhaps the most intuitive angular discretisation is to separate the unit circle into segments, the 'discrete ordinate' method. When light first enters the tissue it shows strong directional coherence, but after a scattering mean-free-path into the tissue, it becomes directionally decorrelated, i.e. diffuse. Thus, near the source, in the discrete ordinate method, a large number of angular bins are required to accurately simulate ballistic propagation, and far from the source, due to near-isotropic light propagation, the radiance is dense in this basis. Memory requirements can therefore become significant – GB in size in 3D – for a mesh or grid with both high spatial and high angular discretisation. While this does not exceed what is typically available on CPUs nowadays, given that MC routines are most efficiently run on GPUs, we wish to store this in the GPU memory, which is limited to 12GB on leading devices.

Instead, a Fourier description may be more appropriate as only few orders (1-2) would be required for diffuse propagation, while ballistic propagation would require higher orders, compared with an order of magnitude more terms being required using traditional angular discretisation. In this approach, the radiance is expressed as a Fourier series,

$$\phi(\mathbf{x}, \theta') = \frac{1}{2}a_0(\mathbf{x}) + \sum_{n=1}^{N=\infty} a_n(\mathbf{x})\cos(n\theta') + \sum_{n=1}^{N=\infty} b_n(\mathbf{x})\sin(n\theta'), \tag{8}$$

where a_0 is the weight deposited in the zeroth harmonic, representing isotropic light propagation, and the higher orders, a_n and b_n, are the Fourier coefficients, scaled by the weight deposited:

$$a_0 = \frac{\delta W}{\pi} \int_{-\pi}^{\pi} \delta(\theta - \theta') d\theta = \frac{\delta W}{\pi} \tag{9}$$

$$a_n = \frac{\delta W}{\pi} \int_{-\pi}^{\pi} \delta(\theta - \theta') \cos(n\theta) d\theta = \frac{\delta W}{\pi} \cos(n\theta') \tag{10}$$

$$b_n = \frac{\delta W}{\pi} \int_{-\pi}^{\pi} \delta(\theta - \theta') \sin(n\theta) d\theta = \frac{\delta W}{\pi} \sin(n\theta'), \tag{11}$$

where δW is the weight deposited by the photon and θ' is the direction of propagation of the photon. A flowchart of the code is shown in Figure 1.
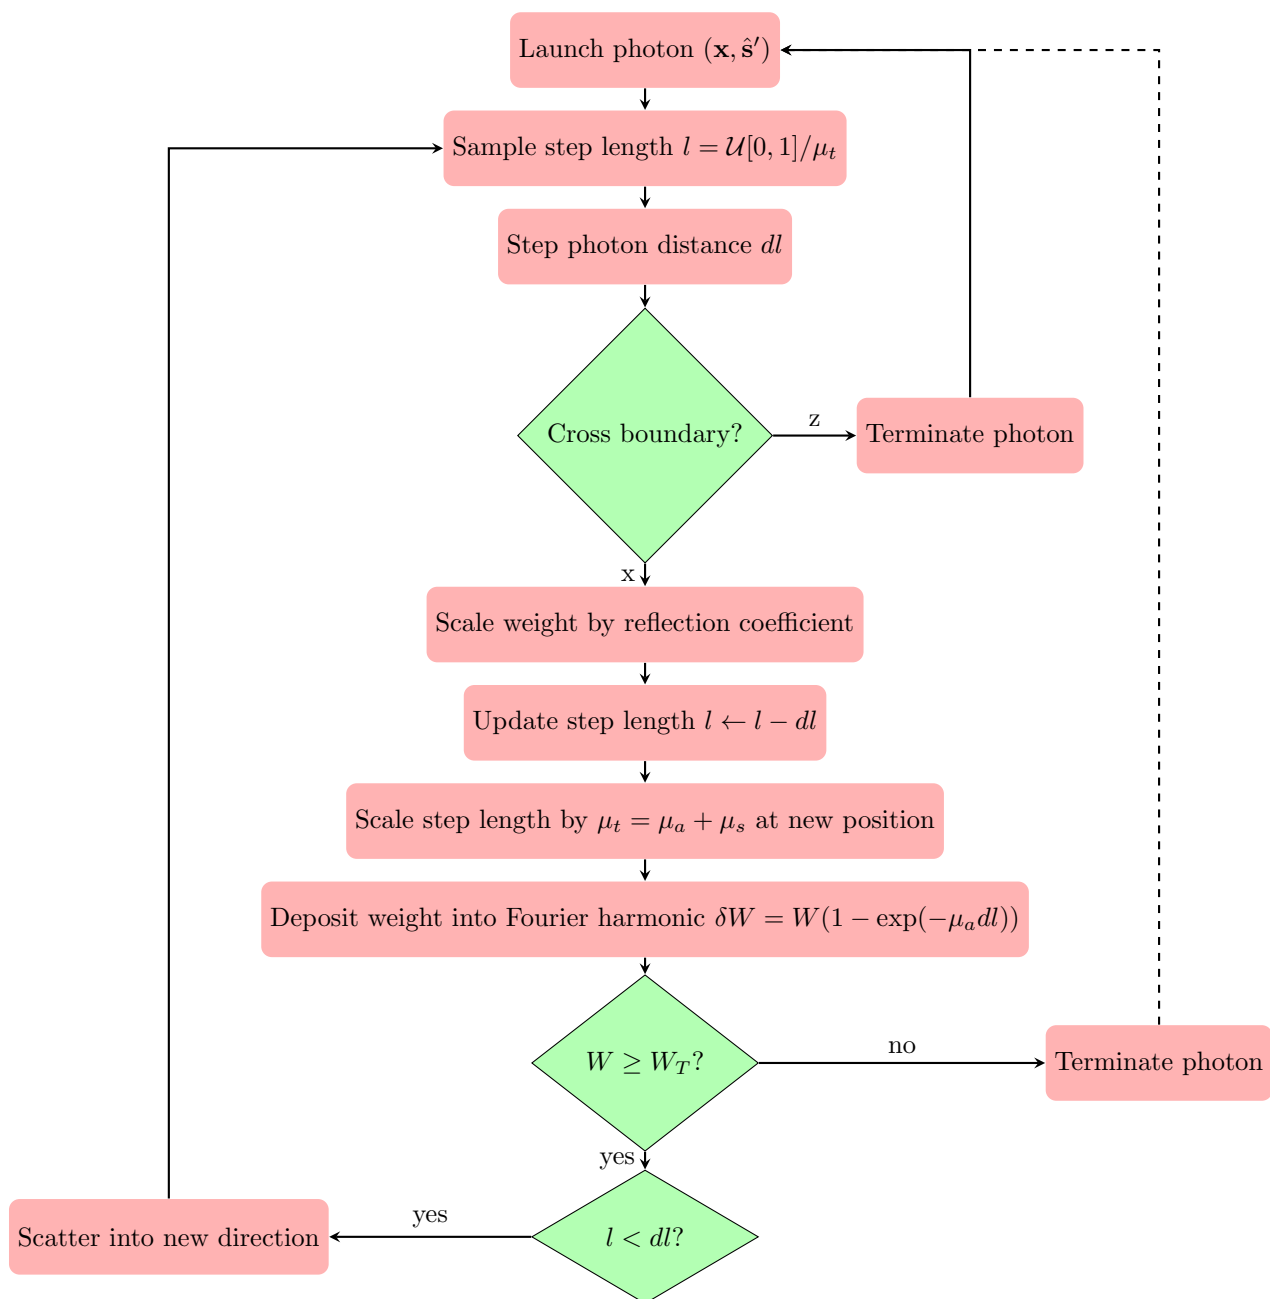

Figure 1. Flowchart of RMC algorithm with index mismatch on x-boundaries and index-matched z-boundaries.

The photon weight deposition scheme described above relating to Eq. (9) can be visualised by placing an isotropic point source at the centre of a 4mm×4mm domain with 0.1mm pixel size and up to second order Fourier harmonics. Figure 2 demonstrates a_n and b_n for $n \in [0, 2]$. The zeroth-order term, a_0, is proportional to the fluence and represents isotropic light propagation, while the higher order terms in the series represent the directional propagation of light, as is evident from Figure 2.
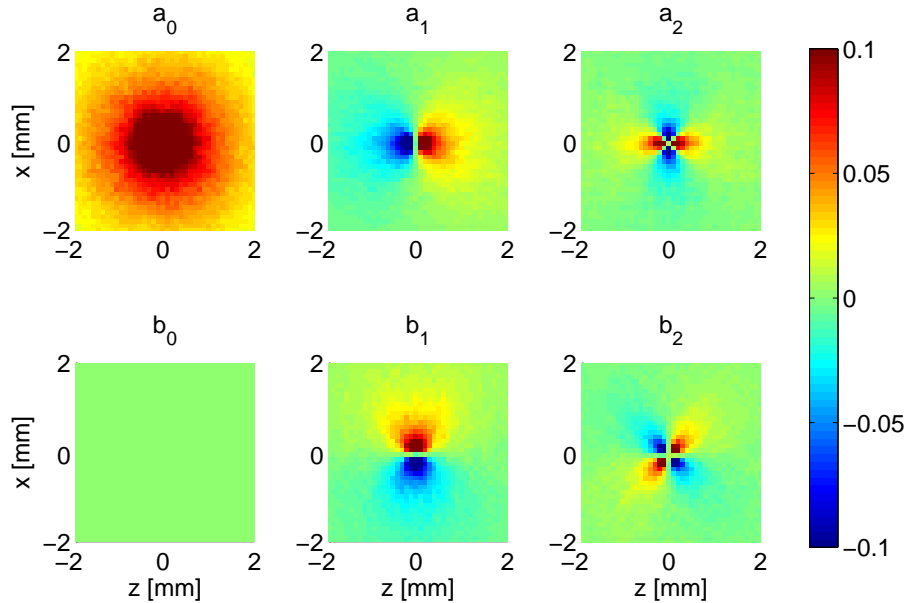

Figure 2. Fourier harmonics a_n (left column) and b_n (right column), for $n \in [0, 2]$, simulated using MC model with 10^6 photons in 4mm×4mm domain and isotropic point source at centre. Images normalised by maximum value of a_0.

3.1 Validation of 2D RMC

Light transport in 2D is fundamentally different to that in 3D due to a difference in phase function; a solution to the RTE in 2D cannot be obtained by integrating a 3D solution along one spatial axis. There are many available 3D models of light transport in turbid media such as diffusion models[6] and MC models,[4,7] but there is a lack of 2D models of light transport against which 2D MC models can be compared. In this section, we compare the accuracy of the 2D RMC algorithm against a 2D finite element (FE) model of the RTE.[8]

The domain simulated was 5mm×5mm in size and 0.1mm isotropic voxels and was illuminated using an isotropic source at the centre of the x-direction at z=0mm. The medium properties were homogeneously distributed with μ_a=0.01mm^{-1}, μ_s=10mm^{-1} and g=0.5. The MC model used 10mm^7 photons and 12 Fourier orders, while the RTE simulation consisted of 64 angular bins. Figure 3 contains plots of the radiance as a function of angle at two positions in the domain: x=0mm (along beam axis), z=0.6mm and x=-1mm, z=0.5mm.
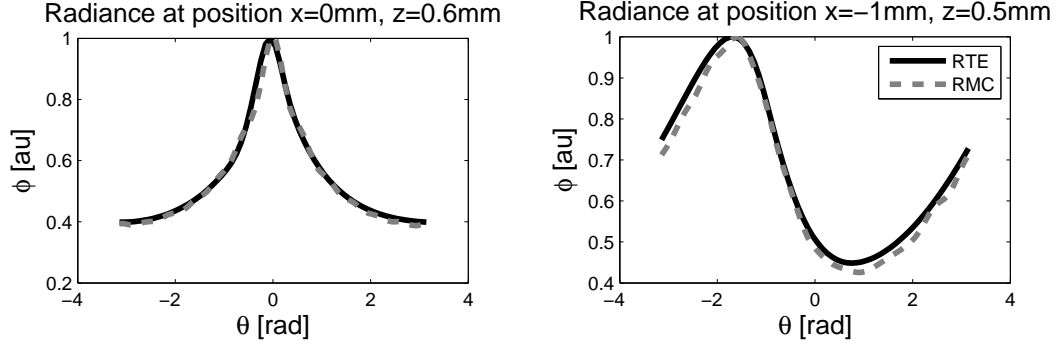
Figure 3. Comparison of radiance as a function of angle in 5mm×5mm domain with 0.1mm isotropic voxel size, illuminated by an isotropic point source on the boundary at x=0mm (centre of edge), z=0mm. RTE simulation has 64 angles and RMC had 12 Fourier orders used to compute radiance in 64 angles. Two positions considered: x=0mm (along beam axis), z=0.6mm (left) and x=-1mm, z=0.5mm (right). Radiance values normalised by maximum value.

It can be observed that the light maintains some of its forward-directionality at a distance of 0.6mm from the source along the beam axis, while off-axis, at x=-1mm, z=0.5mm, the radiance is orientated predominantly at approximately -1.5 radians. Very good agreement is obtained between the MC code and the RTE simulations. It is anticipated that remaining discrepancies would be eliminated with finer spatial and angular discretisation.

4. GRADIENT CALCULATION USING RMC

4.1 Simulating the adjoint radiance

As the source term in the adjoint model is $\mu_a(H_{\mathrm{meas}} - H(\mu_a, \mu_s))$, and this is expressed on a piece-wise constant (voxelised) grid, a brief discussion of how this is implemented is necessary. This source term is isotropic, meaning the source must be distributed homogeneously over each pixel/voxel as well as emit photons in an isotropic manner. To reduce the computational load, sources whose strength was $10^{-5}\times$ less than the maximum value of $\mu_a(H_{\mathrm{meas}} - H)$ were neglected. The spatial distribution was implemented by spreading source positions uniformly randomly over a source voxel; given the corner of a voxel $\mathbf{x_c}$, a uniform random variable, scaled by the voxel size, dx, was added to this: $\mathbf{x_c} + \mathcal{U}[0,1]dx$. The angular distribution was achieved by sampling a uniform random variable over the unit circle: $\mathcal{U}[0,2\pi]$.

4.2 Gradient computation in Fourier domain

Writing the radiance as in Eq. (8) and the adjoint radiance in a similar fashion, but with Fourier coefficients a_0^*, a_m^* and b_m^*, and computing the integral term in Eq. (5), we obtain

$$
\int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}; \mu_a, \mu_s) \phi^*(\hat{\mathbf{s}}; \mu_a, \mu_s) d\hat{\mathbf{s}} = \int_{2\pi} \left[\frac{1}{4} a_0 a_0^* + \frac{1}{2} a_0 \sum_m a_m^* \cos(m\theta') + \frac{1}{2} a_0 \sum_m a_m^* \sin(m\theta') \right.
$$
$$
+ \frac{1}{2} a_0^* \sum_n a_n \cos(n\theta') + \sum_n \sum_m a_n a_m^* \cos(n\theta') \cos(m\theta') + \sum_n \sum_m a_n b_m^* cos(n\theta') \sin(m\theta')
$$
$$
\left. + \frac{1}{2} a_0^* \sum_n b_m \cos(m\theta') + \sum_n \sum_m a_m^* b_n \sin(n\theta') cos(m\theta') + \sum_n \sum_m b_n b_m^* \sin(n\theta') \sin(m\theta') \right] d\theta'. \qquad (12)
$$

Using orthogonality, Eq. (12) reduces to

$$
\int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}; \mu_a, \mu_s) \phi^*(\hat{\mathbf{s}}; \mu_a, \mu_s) d\hat{\mathbf{s}} = \int_{2\pi} \left[\frac{1}{4} a_0 a_0^* + \sum_n a_n a_n^* \cos^2(n\theta') + \sum_n b_n b_n^* \sin^2(n\theta') \right] d\theta', \qquad (13)
$$

which is simply

$$\int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}; \mu_a, \mu_s)\phi^*(\hat{\mathbf{s}}; \mu_a, \mu_s)d\hat{\mathbf{s}} = \frac{\pi}{2}a_0 a_0^* + \pi \sum_n a_n a_n^* + \pi \sum_n b_n b_n^*. \tag{14}$$

This closed form expression of the second term in Eq. (5) is computationally straightforward due to the fact that we are simply summing over products of Fourier coefficients already loaded in memory, meaning numerical integration is not necessary.

A similar approach can be applied to the scattering gradient of Eq. (6). We already have an expression for the convolution of the forward and adjoint radiance fields, given in Eq. (14), which is the first term in Eq. (6). The second term includes the phase function which, if we choose the 2D analogue to the Henyey-Greenstein phase function,[9] has the useful property that it can be expanded as a Fourier series in powers of g, shown in Eq. (15).[10]

$$P_\theta(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'; g) = \frac{1}{2\pi} + \frac{1}{\pi}\sum_{l=1}^{\infty} g^l \cos(l\Delta\theta), \tag{15}$$

where $\Delta\theta = \arccos(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}')$. Thus we can write,

$$\int_{\mathcal{S}^1}\int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}'; \mu_a, \mu_s)P_\theta(\hat{\mathbf{s}}, \hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}; \mu_a, \mu_s)d\hat{\mathbf{s}}d\hat{\mathbf{s}}' = \int_{\mathcal{S}^1}\int_{\mathcal{S}^1} \left[\frac{1}{2}a_0 + \sum_n a_n \cos(n\theta') + \sum_n b_n \sin(n\theta')\right]$$
$$\left[\frac{1}{2\pi} + \frac{1}{\pi}\sum_l g^l \cos(l(\theta - \theta'))\right]$$
$$\left[\frac{1}{2}a_0^* + \sum_m a_m^* \cos(m\theta) + \sum_m b_m^* \sin(m\theta)\right]d\theta d\theta', \tag{16}$$

where we let θ and θ' be the angles between the z-axis and $\hat{\mathbf{s}}$ and $\hat{\mathbf{s}}'$, respectively; as such, the scattering angle between the previous direction $\hat{\mathbf{s}}'$ into the new direction $\hat{\mathbf{s}}$ is given by $(\theta - \theta')$.

It is possible to expand $\cos(l(\theta - \theta'))$ as $\cos(l\theta)\cos(l\theta') + \sin(l\theta)\sin(l\theta')$ which in turn allows us to employ orthogonality relationships to simplify the above integrals and write

$$\int_{\mathcal{S}^1}\int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}'; \mu_a, \mu_s)P_\theta(\hat{\mathbf{s}}, \hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}; \mu_a, \mu_s)d\hat{\mathbf{s}}d\hat{\mathbf{s}}' = \frac{\pi}{2}a_0 a_0^* + \pi \sum_n a_n a_n^* g^n + \pi \sum_n b_n b_n^* g^n, \tag{17}$$

Substituting this expression into Eq. (6), we can write the full expressions for the functional gradients with respect to the absorption and scattering coefficients:

$$\frac{d\epsilon}{d\mu_a} = -\Phi(\mu_a, \mu_s)(H_{\text{meas}} - H(\mu_a, \mu_s)) + \frac{\pi}{2}a_0 a_0^* + \pi \sum_n a_n a_n^* + \pi \sum_n b_n b_n^* \tag{18}$$

$$\frac{d\epsilon}{d\mu_s} = \pi \sum_{n=1} [a_n a_n^* + b_n b_n^*](1 - g^n). \tag{19}$$

4.3 Comparison with finite difference gradient calculation

In order to determine the accuracy of the adjoint-assisted gradient calculation, we can compare the gradients against those computed using a finite difference approach. We can define a forward difference for the absorption and scattering gradients at a pixel position i as

$$\frac{d\epsilon}{d\mu_a^i} = \frac{\epsilon(\mu_a^i + \delta\mu_a, \mu_s) - \epsilon(\mu_a^i, \mu_s)}{\delta\mu_a} \tag{20}$$

$$\frac{d\epsilon}{d\mu_s^i} = \frac{\epsilon(\mu_a, \mu_s^i + \delta\mu_s) - \epsilon(\mu_a, \mu_s^i)}{\delta\mu_s}, \tag{21}$$

where $\delta\mu_a$ or $\delta\mu_s$ is a perturbation of the relevant parameter. A typical approach to determine a suitable perturbation magnitude is to calculate the functional gradient and iteratively reduce the perturbation size until the gradient converges. However, quantities computed with MC simulations will always exhibit some variance or uncertainty and low-noise estimates of functional gradients using MC simulations requires an impractical number of photons to be simulated. Thus, the functional gradient will not converge beyond some value due to noise in the forward model.

'Measured' data was first simulated by running the forward model using 10^7 photons in a domain 4mm×4mm with 0.1mm isotropic voxels using the true absorption and scattering, shown in Fig. 4 and 5, with $g=0.9$ homogeneously distributed. The illumination was collimated in the +z direction and positioned at x=0mm, z=0mm. Using this data, functional gradients with respect to absorption and scattering were computed; for the gradient with respect to absorption, the current parameter estimates were the background value for absorption, μ_a^{bg}, and the true value for scattering, μ_s^{true}. For scattering it was the opposite: the current parameter estimates were the background value for scattering, μ_s^{bg}, and the true value for absorption, μ_a^{true}.

The finite difference calculations, shown in Eq. (20) and (21), were performed using MC simulations using 10^7 photons, with a perturbation sizes $\delta\mu_a = \mu_a^{bg}/2$ and $\delta\mu_s = \mu_s^{bg}/2$. The finite difference calculation was run for a line along x=-1.1mm (indicated by the dashed white line in Fig. 4 and 5). The adjoint simulations were also run with 10^7 photons and used three Fourier orders for angular discretisation and gradients were computed according to Eq. (18) and (19).
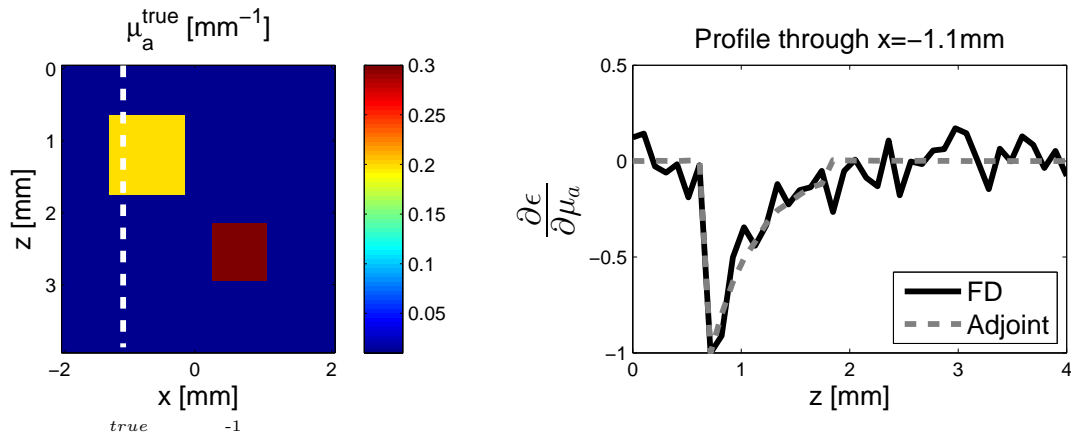

Figure 4. Image of μ_a^{true} in mm^{-1} with white dashed line used to indicate pixels i at which gradients were computed (left); Plots of functional gradient with respect to absorption computed using adjoint-assisted (Adjoint) and finite difference (FD) methods along line parallel to z-axis at x=-1.1mm. Gradients normalised to unity.

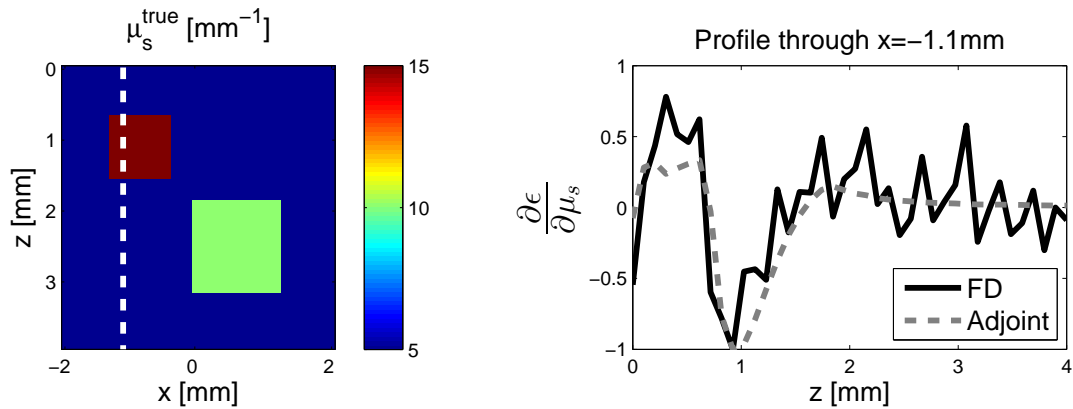Figure 5. Image of μ_s^{true} in mm^{-1} with white dashed line used to indicate pixels i at which gradients were computed (left); Plots of functional gradient with respect to scattering computed using adjoint-assisted (Adjoint) and finite difference (FD) methods along line parallel to z-axis at x=-1.1mm. Gradients normalised to unity.

Despite the variance in the finite difference estimate, it can be seen from Fig. 4 and 5 that there is good agreement between the functional gradients computed using the two different methods. It is clear that variance in the forward simulation is propagated and amplified using the finite difference calculation, whereas it is much less apparent in the adjoint-assisted gradients.

5. CONCLUSIONS

A 2D Monte Carlo (MC) model for computing angle-dependent radiance in scattering media was proposed. It used a Fourier basis for angular discretisation which facilitated compact storage of diffuse radiance fields. This, along with the inherent parallelisability of MC models, makes it an attractive candidate for diffuse optical inverse problems. In particular, it was shown that through the use of an adjoint model, the functional gradients required for quantitative photoacoustic imaging could be calculated much more accurately and efficiently than with a finite-difference approach. This radiance MC approach can be straightforwardly extended to 3D, where the discretisation over angle can be performed using a spherical harmonic basis. It therefore holds great potential as a forward model in gradient-based approaches to the quantitiative photoacoustic imaging problem of quantifying chromophore concentrations from 3D photoacoustic images.

REFERENCES

[1] 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 (July 2013).

[2] Cox, B. T., Arridge, S. R., and Beard, P. C., "Quantitative photoacoustic image reconstruction for molecular imaging," *Proceedings of SPIE* **6086**, 60861M–60861M–9 (2006).

[3] Arridge, S., Schweiger, M., and M, "A finite element approach for modeling photon transport in tissue," *Medical Physics* **20**(2), 299–309 (1993).

[4] Wang, L., Jacques, S., and Zheng, L., "MCML–Monte Carlo modeling of light transport in multi-layered tissues," *Computer methods and programs in biomedicine* **47**, 131–146 (1995).

[5] Sassaroli, A. and Martelli, F., "Equivalence of four Monte Carlo methods for photon migration in turbid media.," *Journal of the Optical Society of America. A, Optics, Image Science, and Vision* **29**, 2110–7 (Oct. 2012).

[6] Schweiger, M., "Comparison of a finite-element forward model with experimental phantom results: application to image reconstruction," *OE/LASE'93* **1888**, 179–190 (Sept. 1993).

[7] Fang, Q. and Boas, D. A., "Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units," *Optics Express* **17**, 20178–90 (Oct. 2009).

[8] Tarvainen, T., Vauhkonen, M., Kolehmainen, V., Kaipio, J. P., and Arridge, S. R., "Utilizing the Radiative Transfer Equation in Optical Tomography," *PIERS Online* **4**(6), 655–660 (2008).

[9] Henyey, L. and Greenstein, J., "Diffuse radiation in the galaxy," in [*Annales d'Astrophysique*], **3**, 70–84 (1940).

[10] Heino, J., Arridge, S., Sikora, J., and Somersalo, E., "Anisotropic effects in highly scattering media," *Physical Review E* **68**, 1–8 (Sept. 2003).