UNIVERSITY COLLEGE LONDON

Monte Carlo Methods in Quantitative Photoacoustic Tomography

Roman Hochuli

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the
Photoacoustic Imaging Group
Department of Medical Physics and Biomedical Engineering

September 2016

Declaration of Authorship

I, Roman Hochuli, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed: _____

Date: _____

Abstract

Quantitative photoacoustic tomography (QPAT) is a hybrid biomedical imaging technique that derives its specificity from the wavelength-dependent absorption of near-infrared/visible laser light, and its sensitivity from ultrasonic waves. This promising technique has the potential to reveal more than just structural information, it can also probe tissue function. Specifically, QPAT has the capability to estimate concentrations of endogenous chromophores, such as the concentrations of oxygenated and deoxygenated haemoglobin (from which blood oxygenation can be calculated), as well as the concentrations of exogenous chromophore, e.g. near-infrared dyes or metallic nanoparticles. This process is complicated by the fact that a photoacoustic image is not directly related to the tissue properties via the absorption coefficient, but is proportional to the wavelength-dependent absorption coefficient times the internal light fluence, which is also wavelength-dependent and is in general unknown. This thesis tackles this issue from two angles; firstly, the question of whether certain experimental conditions allow the impact of the fluence to be neglected by assuming it is constant with wavelength, a 'linear inversion', is addressed. It is demonstrated that a linear inversion is appropriate only for certain bands of illumination wavelengths and for limited depth. Where this assumption is not accurate, an alternative approach is proposed, whereby the fluence inside the tissue is modelled using a novel Monte Carlo model of light transport. This model calculates the angle-dependent radiance distribution by storing the field in Fourier harmonics, in 2D, or spherical harmonics, in 3D. This thesis demonstrates that a key advantage of computing the radiance in this way is that it simplifies the computation of functional gradients when the estimation of the absorption and scattering coefficients is cast as a nonlinear least-squares problem. Using this approach, it is demonstrated in 2D that the estimation of the absorption coefficient can be performed to a useful level of accuracy, despite the limited accuracy in reconstruction of the scattering coefficient.

Acknowledgements

I would like to begin by giving special thanks to my primary supervisor, Ben Cox, who provided valuable insight, support and encouragement over the duration of my MRes and PhD. It's hard to imagine having had someone better for the role than Ben.

Thanks goes to my secondary supervisor, Paul Beard, and my collaborators, Simon Arridge and Samuel Powell. Paul and Simon were always able to ask the most difficult questions, challenges without which this thesis would not have reached this standard. Sam was always on hand to make sure I didn't get ahead of myself.

I am grateful to all the members of the Department of Medical Physics and Biomedical Engineering at UCL. Having started my undergraduate in the department in September 2008, there have been many people who helped me over the past 7 years. I have made a number of friends in the group and throughout the rest of the department who have provided many interesting chats over tea or other beverages, and helped me solve many tough challenges. My thanks go to: Thomas Allen, Lu An, Felix Bragman, Sophie Brand, Jo Brunker, Luke Dunne, Rob Ellwood, Nick Everdell, Martina Fonseca, Adam Gibson, Jamie Guggenheim, Jem Hebden, Jan Laufer, Emma Malone, Tom Millard, Eoin O'Finnerty, Dan O'Flynn, Ollie Ogunlade, George Randall, Andy O'Reilly, Gary Royle, Thomas Stahl, Brad Treeby and Edward Zhang.

It is impossible to forget the contribution of my parents, Sandra and Pierre, my brother, Alex, and girlfriend, Róisín. They have stood by me through the difficult times during my degree and my PhD, and I'm grateful for their unwavering support.

Contents

List of Figures

Chapter 1

Introduction and background

Medical imaging techniques find application in experimental or clinical settings depending on the type of information, structural or functional, they are able to provide [1]. Magnetic resonance imaging provides strong contrast between soft tissues making it a highly specific functional and structural imaging technique [2] for anatomies such as the brain, lungs and muscle tissues. However, the limited ($\sim$1mm) resolution of MRI means X-ray based techniques, planar or CT [3], may be more appropriate when accurate sub-millimetre localisation is required, despite the fact that endogenous soft tissue contrast is typically lower than that of MRI. Ultrasound imaging has acoustic impedance mismatch as its source of contrast, which does not provide a high degree of specificity, though resolution is typically sub-millimetre due to low scattering of ultrasonic waves [4]. When highly specific chemical or functional information is required, optical methods, such as functional near-infrared spectroscopy, are a powerful tool; fNIRS offers the ability to perform time-resolved measurements to obtain concentrations of endogenous or exogenous absorbers, but low penetration depth and limited resolution due to high optical scattering in tissue limit its application in adults. Photoacoustic (PA) imaging is a technique that combines the high specificity of purely optical techniques with ultrasound resolution. This is highly attractive as it offers the capability of imaging micron-scale subcutaneous structures containing endogenous optical absorbers, such as blood. PA imaging is achieved by illuminating the tissue using short pulses of light in

the near-infrared (NIR)/ visible range to excite an acoustic signal through the deposition of optical energy as heat. The spectral and positional information carried by the acoustic wave, detected at the tissue surface using an ultrasound detection system, can in principle be used to determine physiologically relevant parameters such as blood oxygen saturation or molecular concentrations in high-resolution 3D datasets. In order for multiwavelength photoacoustic measurements to reveal such information, the origin of the signal must be located and the individual contributions of each absorber to each PA signal must be ascertained, thus presenting two inverse problems. The acoustic inverse problem, which describes the recovery of the PA image from acoustic measurements made over time, has been studied extensively and there are many methods by which it can be solved. The optical inverse problem, which involves the estimation of physiologically important parameters given one or multiple PA images, is large-scale, nonlinear and potentially ill-posed, and has a range of challenges associated with finding suitably tractable and accurate solution methods. This thesis addresses a number of these challenges and proposes a method for solving the optical inverse problem using a novel and highly scalable model of light transport.

1.1 Biomedical photoacoustic imaging

The photoacoustic effect was discovered by Alexander Graham Bell in 1880 while he was investigating a means for long distance sound transmission. He found that sunlight intermittently blocked off by a rotating wheel, incident upon a metal surface generated acoustic waves by deposition of light energy as heat [5]. The advent of the laser in the early 1960s, and high peak power pulsed lasers shortly after, led rapidly to the photoacoustic effect being revisited through laser-generation of ultrasound in a wide range of applications [6], including 'gas phase' photoacoustic spectroscopy [7]. In biomedical photoacoustic imaging, a field developed over a century after Bell's 'photophone' experiments [8], the photoacoustic effect is exploited to generate ultrasonic waves in tissue. This is achieved by illuminating the tissue using short pulses of light in the near-infrared (NIR)-visible range, whose energy is then thermalised, resulting in a pressure increase which travels to the tissue surface as an acoustic wave where it is detected.

The information carried by the acoustic wave has the potential to reveal structural as well as functional features of the anatomy and physiology such as blood oxygen saturation or molecular concentrations in high-resolution 3D datasets. There are two inverse problems that must be solved in order to reveal such information; the acoustic inverse problem involves reconstructing an image from acoustic measurements made at the surface, while the optical inverse problem tackles the retrieval of functional or chemical information from photoacoustic images. Solving the optical inverse problem is often termed quantitative photoacoustic imaging (QPAT). There exist a vast number of solution methods to the acoustic inversion in the literature, but the optical inverse problem has not yet been solved in a sufficiently general way such that useful and accurate quantitative measures can be readily obtained from photoacoustic images. In fact, there are many challenges standing in a way of a practical solution. Despite these numerous and significant challenges, of the $\sim$1000 papers published on the subject of photoacoustics annually, only a small percentage of these papers tackle the optical inversion. If a suitable solution method were found, QPAT would be a hugely valuable clinical and preclinical tool, given that it would have the ability to provide high-resolution 3D images of functional information (e.g. blood oxygenation).

This thesis considers the optical problem in photoacoustic imaging in detail in later chapters; in this chapter, the formation of the photoacoustic image is discussed first, the process illustrated by moving from the top to the bottom of Fig. 1.1, followed by the reconstruction of the image from acoustic time-series, illustrated by the acoustic inversion step in Fig. 1.1. Brief consideration is given to how physiological parameters might be estimated from multiwavelength photoacoustic images at the end of this chapter.

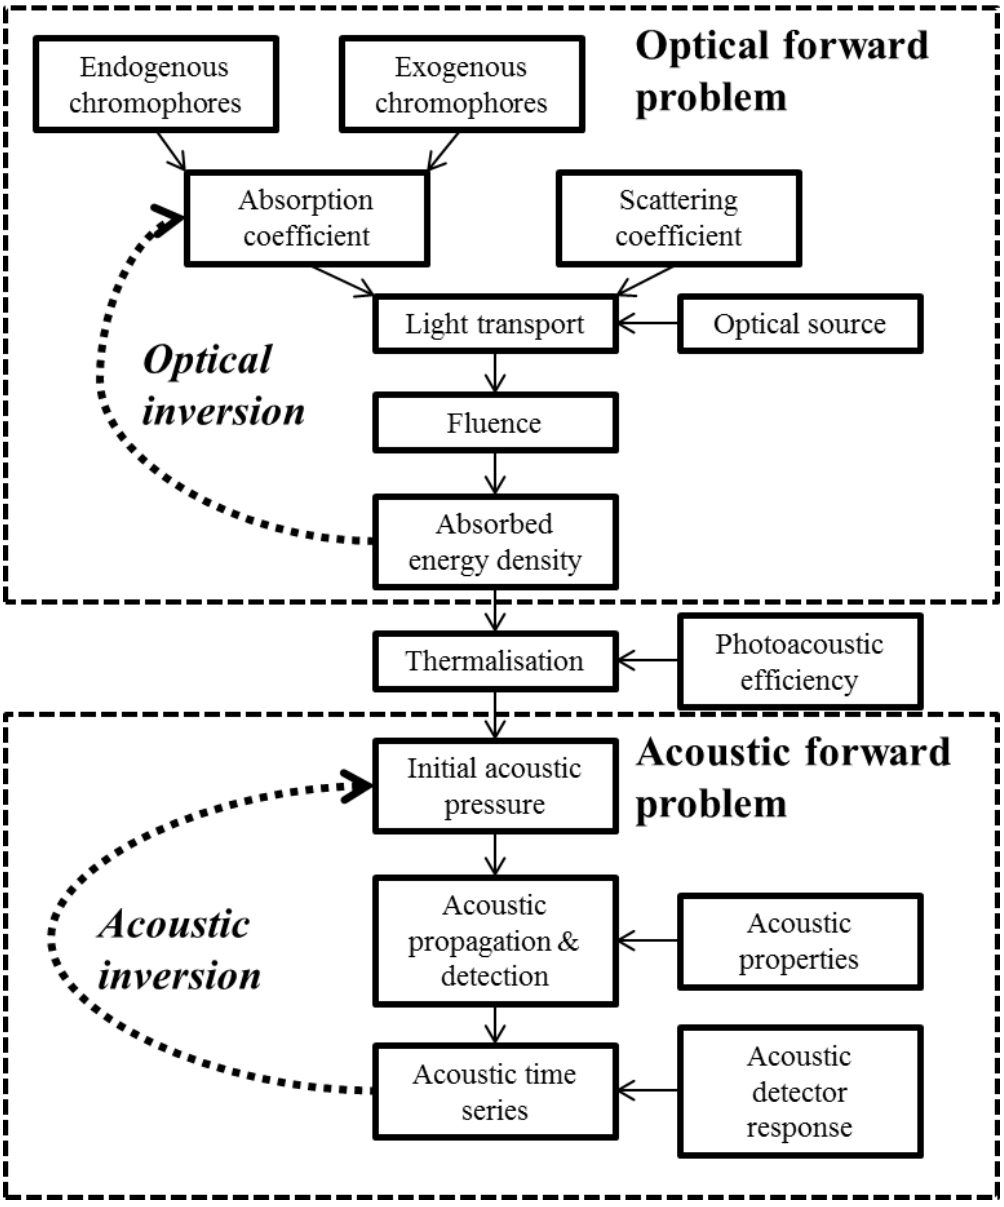

FIGURE 1.1: Diagram of optical and acoustic forward and inverse problems.

1.2 The forward problem

The forward problem in PA imaging is composed of three parts: first, the optical forward problem describes the propagation and absorption of the light energy, which is then thermalised, and produces acoustic waves that propagate outwards to the tissue

surface where they are detected. The wavelengths used in PA imaging are in the near-infrared/visible range where biological tissue has a window of absorption and scattering, allowing the light to penetrate far deeper than wavelengths outside this range. Nanosecond pulsed lasers are used to illuminate the tissue surface which, when the pulses have propagated into the tissue and been absorbed, result in the excitation of chromophores. The amount of light absorbed to give the absorbed energy density, H, depends on the local fluence (in Jmm^{-2}) and the absorption coefficient (in mm^{-1}):

$$H(\mathbf{x}, t) = \mu_a \Phi(\mathbf{x}, t), \tag{1.1}$$

where $\mathbf{x}$ is position and t is time.

The mechanisms that give rise to the absorption of light differ depending on the chromophore's electronic structure or polarisability. In the case of haemoglobin, for instance, absorption of light occurs due to electronic excitation of the heme molecule, whereas water, which can be polarised, undergoes vibrational or rotational excitation. The different interactions by which chromophores absorb light give rise to distinct absorption spectra. The spectra of the dominant chromophores in mammalian tissues (blood, water, fat and melanin) are plotted in Fig. 1.2.

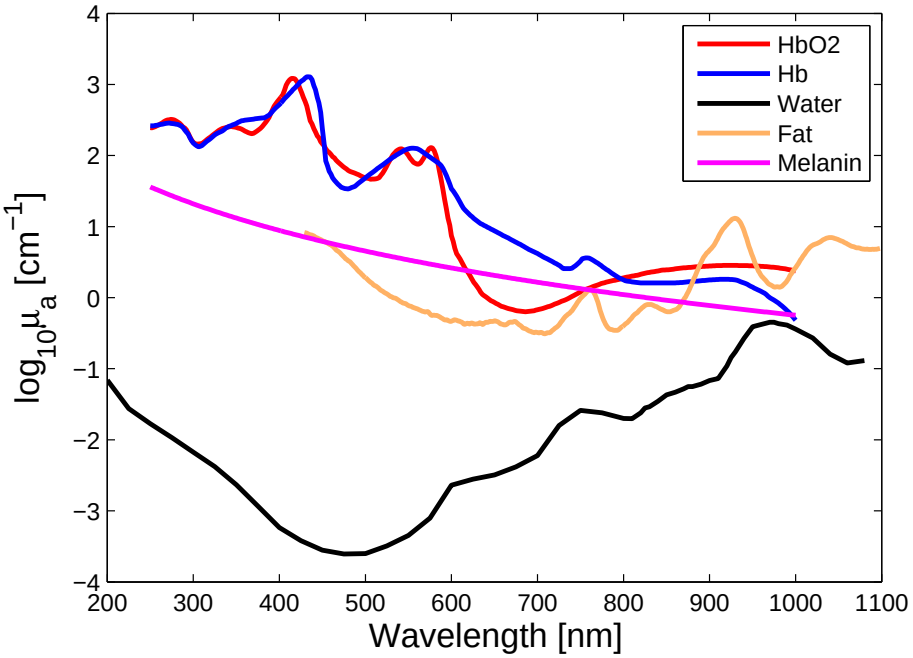

FIGURE 1.2: Absorption spectra of oxygenated haemglobin (HbO2), deoxygenated haemoglobin (Hb), water, fat, melanin in 200-1100nm range [http://omlc.org/spectra/index.html]. Haemoglobin concentrations assumed to be 150gl^{-1} and melanin modelled as skin [9].

Once in an excited state, chromophores undergo de-excitation, either radiatively by fluorescing or phosphorescing, or non-radiatively by producing vibrations that are dissipated as heat through collisions with neighbouring molecules. Non-radiative relaxation occurs on picosecond time-scales. Thus, the absorbed power density, $\mathcal{H}$, (in Wmm^{-3}) can be written as

$$\mathcal{H}(\mathbf{x}, t) = \mu_a \Phi(\mathbf{x}) f(t), \tag{1.2}$$

where $f(t)$ describes the temporal shape of the pulse. The illumination being nanosecond duration pulses ensures that all the optical energy is thermalised before the light intensity falls to zero. This means that Eq. (1.2) can be re-written as

$$H(\mathbf{x}) = \mu_a \Phi(\mathbf{x}) \int_{t=0}^{T} f(t) dt. \tag{1.3}$$

The absorbed energy density acts as a source term to heat diffusion. However, thermal diffusion occurs on time-scales 10^3-10^6 longer than light absorption and relaxation of the absorbers. The pressure increase, resulting from a localised temperature increase due to thermal relaxation, also only propagates on microsecond time-scales, after the illumination pulse has fallen to zero. This is referred to as stress confinement. The pressure resulting from the absorbed energy density is described by the Grüneisen parameter, Γ, and the acoustic pressure distribution is given by

$$p_0(\mathbf{x}) = (1 - \upsilon(\mathbf{x}))\Gamma(\mathbf{x})H(\mathbf{x}), \tag{1.4}$$

where υ is the quantum efficiency in the medium, describing how much energy is dissipated via radiative relaxation. In most biological tissues the the Grüneisen parameter accurately describes the relationship between acoustic pressure and absorbed energy, although in some cases, e.g. homogeneously distributed nanoscale absorbers where absorption of the light and heating happen in the two different materials, the photoacoustic efficiency is used. The Grüneisen parameter is equal to $\frac{\beta c^2}{C_p}$, where β is the volume expansion coefficient and C_p is the specific heat capacity at constant pressure. The complete optical forward model is given by

$$p_0(\mathbf{x}, \lambda) = (1 - \upsilon)\frac{\beta(\mathbf{x})c(\mathbf{x})^2}{C_p(\mathbf{x})}\Phi(\mathbf{x}, \lambda)\mu_{\mathrm{a}}(\mathbf{x}, \lambda), \tag{1.5}$$

where λ is the wavelength of the incident light pulse and it is assumed the quantum efficiency does not vary spatially inside the tissue.

Following the formation of the initial pressure, the broadband (MHz frequency) acoustic waves begin to propagate outwards from the tissue. The simplest description of this process is given by the wave equation

$$\left(\frac{\partial^2}{\partial t^2} - c_0^2\nabla^2\right)p(\mathbf{x}, t) = 0, \tag{1.6}$$

where p is acoustic pressure. The acoustic forward model derives from two first-order PDEs, momentum conservation and mass conservation, and a pressure-density relation. This model is subject to two initial conditions, $p(\mathbf{x}, t = 0) = p_0$ and $\frac{\partial p(\mathbf{x},t)}{\partial t}|_{t=0} = 0$, where the latter is based on the assumption that no acoustic wave is propagating initially. The acoustic field is then recorded at the tissue surface using an acoustic detector (or an array of detectors) and measures acoustic time-series. A wide range of detection schemes exist from focused, resonant PZT detectors [10], to planar interferometric detector arrays [11] and hydrophones [12]. Acoustic forward propagation can be modelled in many ways; many analytic schemes depend on known source/detector geometries [13–16] while others such as time-domain [17] or pseudospectral-domain algorithms [18] offer more flexibility in terms of heterogeneous media and arbitrary geometries. These algorithms are described in more detail in Section 1.3.2.

1.3 Formation of the photoacoustic image

The characteristics of photoacoustic images are determined by the detection scheme and the method by which images are reconstructed from acoustic time-series. This section contrasts photoacoustic tomography from other photoacoustic imaging modes and considers the reconstruction algorithms necessary to obtain tomographic images. Some of the limitations of tomographic detection and reconstruction are demonstrated at the end of the section.

1.3.1 Photoacoustic imaging configurations

The method for the formation of the PA image, as well as image properties such as scale and resolution, are determined by the imaging configuration. PA imaging has been performed in three basic forms: tomography (PAT), acoustic-resolution photoacoustic microscopy (AR-PAM) and optical-resolution photoacoustic microscopy (OR-PAM). Tomography uses wide-field illumination, flooding the tissue with light, meaning that PA signals originate throughout the illuminated volume. These signals are acquired over a surface surrounding a portion of the imaging volume. There are a variety of detection

configurations used in tomography mode; a hemispherical array of ultrasound transducers (optimised for breast imaging) can image to depths of several cm with 250μm resolution [19], while whole-body animal imaging at 200μm resolution is possible via a single-element cylindrical scanning configuration [20], and all-optical planar detection arrays afford resolutions of tens of microns to depths of several mm [11]. However, the detection geometry will impact the accuracy of the reconstructed image. If the detection surface surrounds a point in the target such that any line through that point intersects the detection surface, the point is said to be within the 'visible region' and can be reconstructed exactly. This is illustrated by point A in FIg. 1.3. Points outside the visible region will be reconstructed with artefacts since a line through such a point would not intersect the detection surface, as illustrated by the horizontal line through point B. However, points on an edge outside the visible region with a normal that intersects the detection surface will be reconstructed perfectly, as shown by point C in Fig. 1.3.

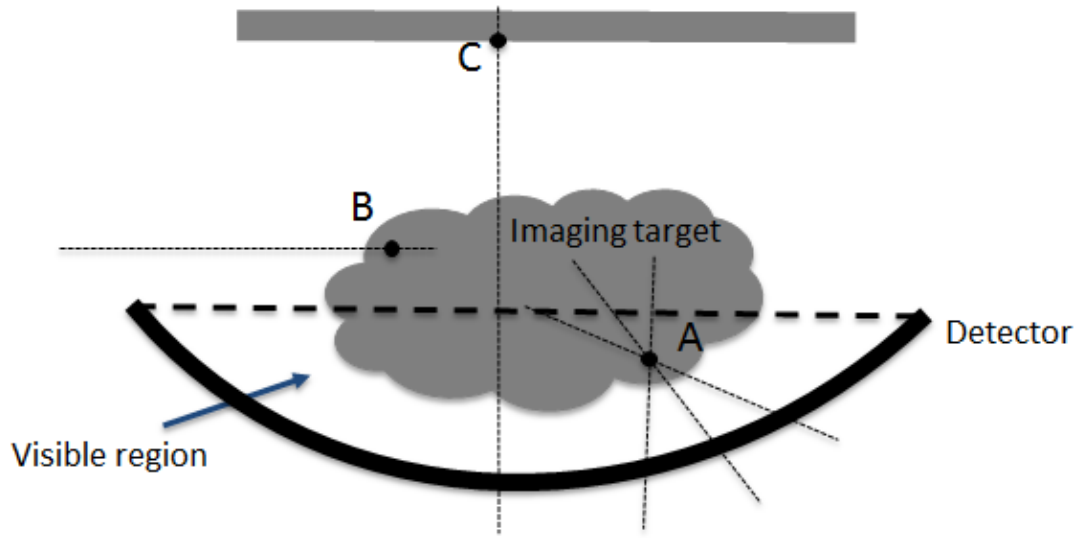

FIGURE 1.3: Diagram illustrating portions of the imaging target within the visible region for an arc-shaped sensor acoustic sensor. Point A has all lines passing through it intersecting the sensor and is therefore in the visible region; point B has lines that do not intersect sensor meaning it is outside the visible region; point C is on an edge whose normal intersects the sensor.

In AR-PAM the tissue is illuminated using wide-field or weakly-focused light but a mechanically- or electronically-steered focused acoustic detector determines both the

spatial resolution and field-of-view. The resolution obtainable with such a setup varies depending on the acoustic spot size and detector centre frequency, with penetration depth typically limited to a few mm. It has been demonstrated that the resolution at these depths can reach tens of microns for piezoelectric detectors [21, 22]. In OR-PAM, a focused transducer is used but it is not a requirement as the resolution is not determined by the acoustic spot size, but by the optical spot size. Strong focusing of the excitation beam means the signal originates from a confined, optical diffraction-limited spot whose size is limited by the strong scattering in tissue. As a result of the strong focusing penetration depth is typically less than 1mm and resolution can be in the microns range [10, 23].

In both microscopy modes, lateral localisation of the signal origin is provided by the optical or acoustic focus and axial resolution is determined by the frequency response of the acoustic detector. Reconstruction of the initial acoustic pressure in microscopy mode simply involves using an estimate of the speed of sound in tissue to determine the distance of absorbers from the detector within each A-line. The illumination, and the detector in confocal set-ups, are then scanned to form an image. This means that, of the three modes, PAT is the only one requiring acoustic reconstruction and these methods are described in the next section.

1.3.2 The acoustic inverse problem

The acoustic inversion is an inverse initial value problem that can be solved for the initial acoustic pressure distribution, p_0, using a number of methods. The solution methods to this problem can be categorised as being series solutions, back-projection formulae, time-reversal or optimisation-based. These methods are discussed below.

Series solutions involve the projection of the acoustic field into a basis such that the PDE of Eq. (1.6) simplifies to an ordinary differential equation; there are a number of options when choosing a basis depending on the detection geometry, but a common choice for planar, circular or spherical detection geometries is a Fourier basis. Decomposition of the field into the Fourier domain affords significant computational advantages, namely that

the derivatives in Eq. (1.6) can be computed using repeated Fast Fourier Transforms [24–26]. Back-projection algorithms, inspired by X-ray CT reconstruction methods, compute the inverse spherical mean Radon transform [27] in the time domain and are applicable to a number of detection geometries [28], including spherical [29, 30], cylindrical [30], planar [31] geometries. The back-projection algorithm has also been generalised to universal geometries [28].

Time-reversal exploits the fact that Eq. (1.6) is invariant under the transformation $t \rightarrow -t$ and that the acoustic field is zero after a time $t = T$ (in odd dimensions) [32, 33]. As such, the model propagates the detected pressure time-series, from the detection surface D^2, $p_d(\mathbf{x} \in D^2, t)$ backwards in time until $t = 0$. k-Wave [34] performs this calculation using a pseudospectral algorithm, performing spatial derivatives in the Fourier domain and allows for the incorporation of power law acoustic attenuation (with suitable regularisation) [35] and sound speed heterogeneities [36]. Time-reversal approaches, either in time-domain [37] or pseudospectral domain [34], provide accurate reconstruction of the initial acoustic pressure for arbitrary detection geometry, provided the imaging target is within the visible region.

More recent approaches to solving the acoustic inverse problem have been model-based; these techniques consider the inversion within an optimisation framework in order to minimise the difference between measured data and simulated data:

$$\arg\min_{x} \epsilon = \frac{1}{2}||y - f(x)||^2, \tag{1.7}$$

where ϵ is the error functional, $f(x)$ is the forward model applied to the estimate of the data and y is the measurement. The computational load and run-time associated with these methods varies quite significantly depending on the optimisation algorithm chosen and the level of convergence required; the system matrix and its derivatives can have very high memory demands for high-resolution images, while a rapid forward model can ensure the time taken per iteration is minimised [38]. Suitably fast time-domain algorithms have been used within such schemes to ensure rapid image reconstruction

[39]. This more general, model-based approach allows the introduction of complimentary information through other fidelity or penalty terms in Eq. (1.7) (e.g. smoothness priors [40]) as well as acoustic heterogeneities in the forward model [41, 42].

1.3.3 Practical considerations of acoustic reconstruction

As discussed in Section 1.3.1, a number of different scanning configurations exist, with a range of methods used to reconstruct the initial acoustic pressure from acoustic time-series. However, depending on the configuration and reconstruction algorithm, acoustic artefacts can impact the reconstructed initial acoustic pressure which in turn impact the accuracy with which chromophore concentrations can be estimated. As the acoustic inversion is outside the focus of this thesis, this section presents three common sources of artefacts and demonstrates their effect on the reconstructed pressure distribution, which in turn will impact the accuracy of the optical inversion. Artefacts can be introduced from incomplete or inaccurate modelling of the acoustic detector response, position and directionality as well as incorrect speed of sound estimation in the medium and acoustic attenuation. In this section, we consider the effect of the imaging target being partially outside the visible region, non-uniform detector frequency response and frequency-dependent acoustic attenuation in tissue. The k-Wave Toolbox [18] was used throughout this section as both a forward and inverse acoustic solver.

- **Partial data**, or the limited-view problem, arises when the photoacoustic source does not lie completely within the visible region (or has edges that are not parallel to the detection surface). This of course depends on the scanning configuration, where a circular detection scheme can acquire complete data from a circular source whereas a linear detector (unless infinite) will only acquire regions of the circle enclosed in the angle subtended by the sensor.

  We can illustrate the impact of reconstructing a photoacoustic source using partial data by placing a 0.2mm disk, representing the initial acoustic pressure, running a forward model of acoustic propagation, detecting the acoustic time series and reconstructing the initial acoustic pressure from two different source types using

time-reversal; in this example, a broadband line detector was placed at z=0mm, at the bottom of the domain, in one case, while in the other a circular detector was inscribed in the 2mm square domain. The grid spacing in the domain was 5μm with a homogeneous sound speed of 1500ms^{-1}, thus supporting a maximum frequency of 750MHz. No acoustic attenuation was present. Fig. 1.4 shows images of the 'true' photoacoustic source, full-view reconstruction, partial-view reconstruction and profiles through the source.

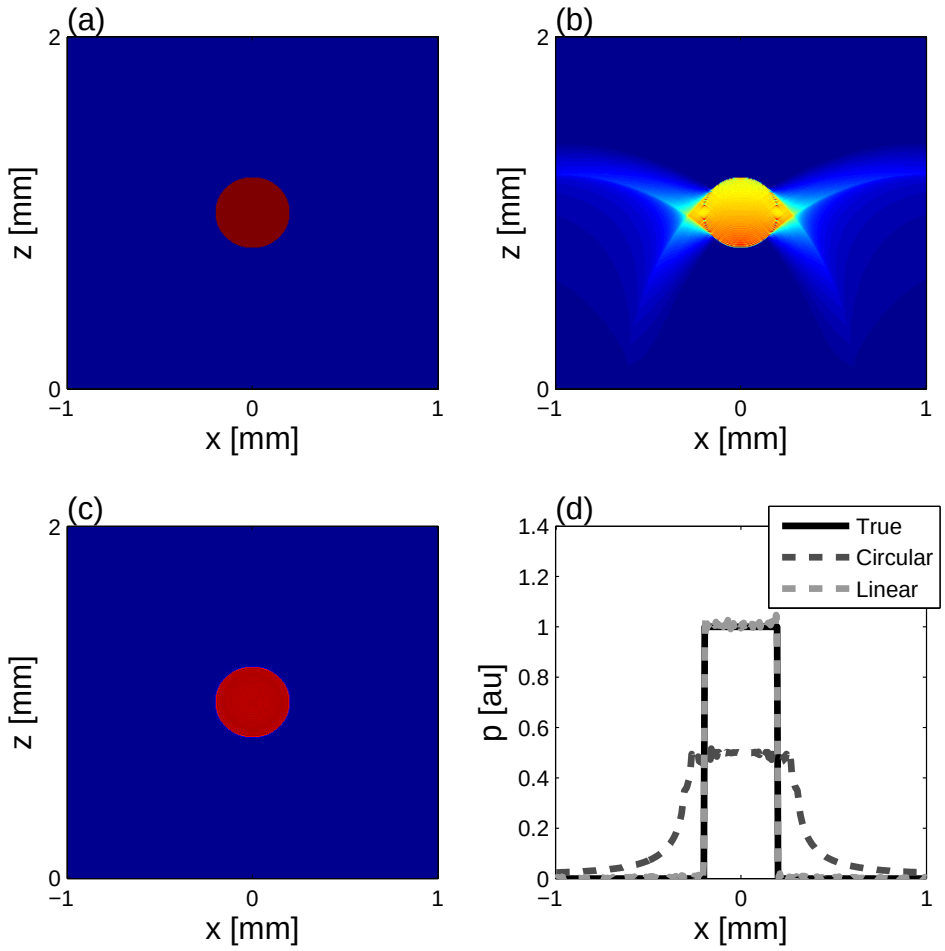

FIGURE 1.4: (a) True p_0 in 2mm squared domain; (b) Reconstructed acoustic pressure from 2mm diameter circular detector arrangement (concentric with PA source); (c) Reconstructed acoustic pressure from 2mm linear detector arrangement at z=0mm; (d) Profiles through true and reconstructed pressure distributions at z=1mm.

From Fig. 1.4 it is clear that a full-view reconstruction in (c) using, for instance, a

circular detector array is much more likely to yield accurate chromophore concentrations when performing the optical inversion as the reconstructed initial acoustic pressure is almost identical to the true value of p_0 in (a), whereas the use of a planar detector leads to significant artefacts in the reconstruction, shown in (b).

- **Bandlimited data** occurs due to limited frequency response of the acoustic sensor. Piezoelectric sensors, routinely used in PAI [21, 22], can have a high degree of sensitivity but rarely have fla frequency response as they are resonant; this results in them having nonuniform sensitivity and low sensitivity at either low or high frequency bands will result in inaccurate reconstruction of the PA source. In order to demonstrate the effect of bandlimiting of acoustic signals at the sensor, the same PA source was used with the acoustic waves propagated to the circular detector used above in Fig. 1.4(b), but the signals were then bandpass filtered using a Gaussian filter with a 3MHz centre frequency and 100% bandwidth before reconstructing the acoustic pressure.

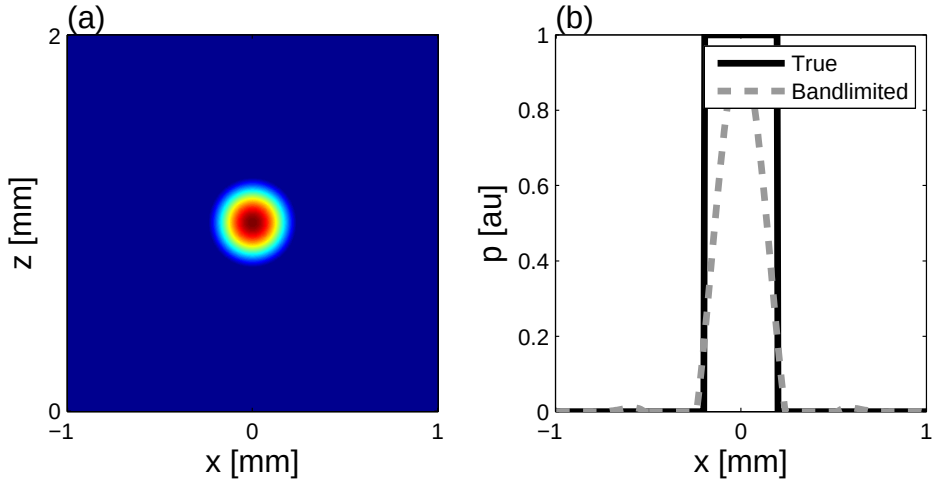

FIGURE 1.5: (a) Image of reconstructed acoustic pressure from bandpass filtered (3MHz centre frequency, 100% bandwidth) acoustic time series acquired on 2mm diameter circular detector; (b) Profiles through true and reconstructed bandlimited pressure distributions at z=1mm.

The effect of bandpass filtering the acoustic time-series is that high frequencies are lost, thus smoothing out the acoustic source. Without wideband acoustic detectors, the accuracy of quantitative estimates yielded by the optical inversion will be poor.

- **Frequency-dependent acoustic absorption** is an inherent property of matter and of relevance in PA imaging, due to the fact that signals travelling further from source to detector will have their high frequencies attenuated more. Thus, if only detecting from one side, signals originating from deeper regions in the tissue will be reconstructed with lower resolution than those originating near the tissue surface. As one-sided detection produces limtied-view artefacts, in order to demonstrate the impact that frequency-dependent attenuation has on the reconstructed image, both these effects must be included. In this set of simulations, the PA source consisted of two disks, both 0.1mm in diameter, one at 1.5mm from the sensor and the other at 0.25mm. The initial acoustic pressure is shown in Fig. 1.6(a). The simulation parameters were the same as above, except in the forward simulation of the data from which Fig. 1.6(c) was reconstructed, acoustic attenuation was modelled using a power law of the form $a = a_0\omega^u$ where $a_0 = 2$dBMHzcm^{-1} ^{u} with $u = 1.5$, which is similar to biological tissues of interest [28].

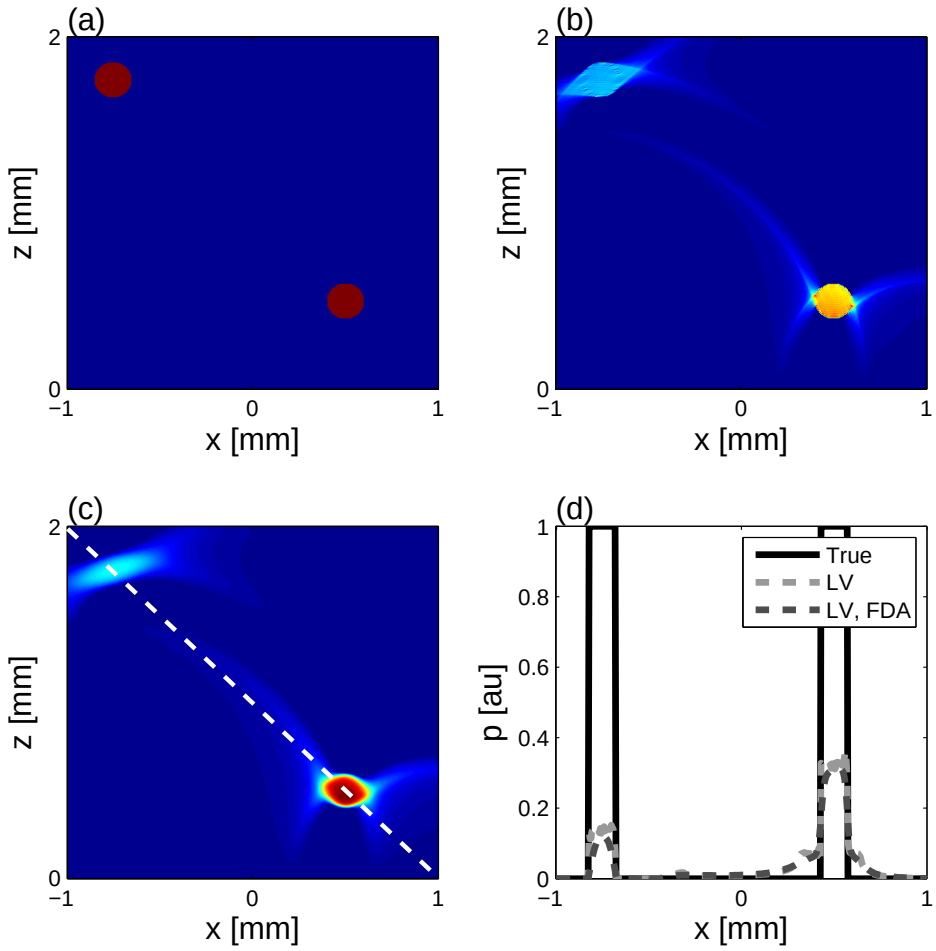

FIGURE 1.6: (a) True p_0 in 2mm square domain; (b) Reconstructed acoustic pressure from PA source); (c) Reconstructed acoustic pressure from linear detector arrangement at z=0mm with limited-view artefacts; (d) Reconstructed acoustic pressure from linear detector arrangement at z=0mm with limited-view artefacts and frequency-dependent attenuation ($a = a_0\omega^u$ where $a_0 = 2$dBMHz^u with $u = 1.5$); (d) Profiles through image diagonal in true and reconstructed pressure distributions with limited-view (LV) and frequency-dependent attenuation (FDA) artefacts.

The loss in high frequency components between Figs. 1.6(b) and (c) is evident, particularly in the source at the top of the domain, far from the detector, where higher frequencies will have undergone greater attenuation as they travelled to the sensor. Fig. 1.6(d) shows a profile through the images (denoted by the grey line in Fig. 1.6(a)) where the smoothing that has taken place due to frequency-dependent attenuation is visible in both PA sources.

In practice, the above limitations, along with limited sensor directivity and heterogeneities in the acoustic properties, would severely impact the accuracy of chromophore concentration estimates made using PAT images. It may be possible to account for unknown speed of sound [43–45] and acoustic attenuation [36], while wide band detectors have been in use for several years [11] and iterative reconstruction schemes for reducing artefacts exist. The success of QPAT relies on challenges within the acoustic inverse problem being adequately solved. The focus of this thesis however is on the optical inverse problem.

1.4 Quantitative photoacoustic imaging

As mentioned in Section 1.1, QPAT refers to the estimation of physiological properties such as blood oxygenation or chromophore concentrations from photoacoustic images following acoustic reconstruction (i.e. the top block of in Fig. 1.1). QPAT generally involves the exploitation of the wavelength-dependent absorption of various endogenous or exogenous absorbers, as the linear combination of their concentrations with their absorption spectra gives rise to the wavelength-dependent absorption coefficient, $\mu_{\mathrm{a}}(\lambda)$. Recall that the photoacoustic forward problem is given by (Eq. (1.5)):

$$p_0(\mathbf{x}, \lambda) = \Gamma(\mathbf{x})\Phi(\mathbf{x}, \lambda;\ \mu_{\mathrm{a}}, \mu_{\mathrm{s}}, g)\mu_{\mathrm{a}}(\mathbf{x}, \lambda), \tag{1.8}$$

where μ_{s} is the scattering coefficient and g describes the scattering anisotropy; expanding the absorption coefficient in terms of the chromophore concentrations c_k, their molar absorption coefficients, α_k, yields

$$p_0(\mathbf{x}, \lambda) = \Gamma(\mathbf{x})\Phi(\mathbf{x}, \lambda;\ c_k, \mu_{\mathrm{s}}, g)\sum_{k=1}^{K}\alpha_k(\lambda)c_k(\mathbf{x}) \tag{1.9}$$

It is clear from Eq. (1.9) that, given the objective in QPAT to estimate the c_k, linear unmixing to resolve concentrations from the absorption coefficient is required; however, prior to that it also necessary to remove of the impact of the fluence, $\Phi(\mathbf{x}, \lambda)$, and the Grüneisen coefficient, $\Gamma(\mathbf{x})$. A few methods for the removal of the impact of the

Grüneisen parameter have been suggested, most common of which is one that exploits the fact that Γ does not depend on wavelength. These approaches are discussed in detail in CHapter 2. The fluence – which depends on wavelength through the wavelength-dependent scattering coefficient, $\mu_{\mathrm{s}}(\mathbf{x}, \lambda)$, anisotropy factor, $g(\mathbf{x}, \lambda)$, and chromophore concentrations – cannot be divided out of Eq. (1.8) so must be modelled; in fact, the fluence cannot be modelled without knowledge of the c_k and must therefore be estimated simultaneously with it.

The internal fluence can be determined noninvasively by probing it using acoustic modulation to encode light passing through the tissue [46] but spatial resolution and depth penetration are limited using this technique, with the resolution suffering a reduction by an order of magnitude. Resolution could in principle be similar to that achievable by PAT but absorption of high acoustic frequencies will negatively impact depth penetration of the acoustic field, presenting an unfortunate trade-off using this technique. Diffuse optical imaging has also been employed [47, 48] to aid the estimation of the absorption coefficient from PAT images, but resolution is reduced to the millimetre scale. Another noninvasive technique relies on the absorption being dependent on the fluence [49]; if the absorption from a fluence-dependent chromophore were to suddenly fall to zero with increasing fluence, the fluence could be determined in the tissue. This approach is fraught with practical limitations, namely that it requires the assumption that the fluence outside a region of interest is unchanged by a change in the absorption coefficient in that region. Moreover, attempts to implement such a scheme in practice [50] used chromophores that are weakly absorbing at 488nm, a region of the spectrum in which the high absorption of blood will severely limit the penetration of light into the tissue. Invasive methods have been used for blood oxygenation measurements by simply placing a known absorber inside the tissue [51], though this method has its obvious limitations. The distribution of the fluence can also be modelled, as discussed in the next section.

1.4.1 Models of light transport

There are a number of methods by which light transport can be modelled in tissue [52, 53]. Due to the computational complexity of solving Maxwell's electromagnetic

equations for a distribution of scatterers in a large region of interest [54], a suitably accurate approximation is to use the energy-balancing, integro-differential radiative transfer equation (RTE). For a given volume, this equation balances the energy scattered from any direction into the direction of interest, and source contributions in this direction, with energy lost due to absorption and scattering away from the direction of interest. The variable in this equation is the radiance, which varies spatially and angularly. Thus, in order to simulate the RTE directly, and therefore compute the radiance, discretisation must be over space, angle and time. Note that it is not necessary to solve the time-dependent RTE for PAT because the optical part of the problem is considered instantaneous. Furthermore, the high degree of scattering in biological tissue means that it is not always necessary to accurately model the high directionality of the light field. Models such as the P_n approximations, diffusion or δ-Eddington approximation only model low directional components or scattered and unscattered components of the field to reduce computational burden of simulating the full RTE.

A range of solution methods exist for solving the RTE directly, such as the finite difference method, the finite element method or Monte Carlo methods. Finite element methods, which solve the weak formulation of the RTE, have fewer geometrical constraints compared with finite difference techniques, but both approaches have significant memory demands due to the need to discretise in angle and space. Monte Carlo methods, which simulate the fluence rather than the radiance, have limited memory demands and are inherently parallel; however, the need to simulate million or billions of photons for convergence to a sufficiently low-noise fluence estimates means the algorithms can be highly compute intensive. Approximations to the RTE, such as the diffusion equation which assumes near-diffuse propagation of light, have computationally light solution methods, but are inaccurate in superficial regions of the tissue, near the source, which is a region of interest in PAT. Therefore balancing the accuracy and computational demands of the model and its implementation are very important, particularly in their application to the optical inverse problem.

1.4.2 The optical inverse problem

The optical inverse problem in PAT refers broadly to the estimation of the absorption coefficient from a single, or multiple photoacoustic images, although the estimation of chromophore concentrations, blood oxygenation and/or the Grüneisen parameter can also be the aim of the optical inversion. Linear inversions assume that the fluence does not depend on the unknown; the unknown may be the absorption coefficient, chromophore concentrations or blood oxygenation. However, this assumption is not true in general and the inversion is complicated by the fact that the fluence depends on the unknown, as well as the heterogeneous and wavelength-dependent scattering coefficient $\mu_{\mathrm{s}}(\mathbf{x}, \lambda)$, which is also usually unknown. This yields at least two sets of unknowns in the inverse problem, which is a potential source of ill-posedness. Tomographic photoacoustic images are high-resolution 3D datasets, often containing hundreds of voxels on each edge of a 3D volume, meaning in order to perform the optical inversion, the fluence must be known at millions of positions throughout the domain; the resulting inverse problem is large-scale given that the number of unknowns is well into the millions.

The 'holy grail' of the optical inversion is one that is sufficiently general to tackle the issues raised in this section. A general method for solving the optical inverse problem is one that:

(1) is 3D

(2) can estimate the absorption coefficient accurately in every voxel in the image (SNR permitting)

(3) can estimate the scattering coefficient (sufficiently) accurately

(4) is computationally tractable, in terms of both memory and algorithmic demands

(5) can perform a full inversion in a practical amount of time

The above qualities are not independent; for instance, finite element implementations of the RTE can accurately simulate the fluence throughout the domain, but the scale of 3D simulations becomes intractable with the system matrix requiring in the range of

10^9-10^{12} Bytes to store and invert. On the other hand, modelling error of the diffusion approximation limits its application in QPAT to estimation of the absorption coefficient in deeper structures. In this thesis, Monte Carlo methods are considered as a candidate light model with the aim of demonstrating, where possible, that a Monte Carlo model satisfies the above characteristics.

1.5 Contributions of this thesis

This thesis explores three key areas of QPAT (Chapters 3, 4, 5 and 6) and the conclusions derived from these investigations make the following contributions to the field.

- Chapter 3 investigates whether blood oxygenation, a fundamental biomarker, can be estimated from PAT images without explicitly accounting for the wavelength-dependence of the fluence. Such a study has not yet been conducted in detail [55, 56] and this work demonstrates that there are in fact very few experimental scenarios where neglecting the impact of the fluence would allow accurate estimation of blood oxygenation.

- Given the need to account for the fluence, Chapter 4 presents a novel and accurate Monte Carlo model of light transport capable of computing the radiance. This is demonstrated using a Fourier basis at every position in the domain to account for the angle-dependent propagation of light.

- It has previously been demonstrated, in 2D, that the absorption and scattering coefficient can be estimated from multiple images by simulating the radiance. Chapter 5 demonstrates the calculation of gradients of the least-squares error between simulated and measured data with respect to the absorption and scattering coefficient using distributions of Fourier harmonics calculated using the radiance Monte Carlo model presented in Chapter 4. These functional gradients are then applied in inverting simulated PAT images for the absorption and/or scattering coefficient. The Monte Carlo-based inversion is the first of its kind in QPAT and presents a computationally tractable means for obtaining physiologically relevant information,

such as chromophore concentrations or oxygen saturation, from multiwavelength photoacoustic images.

Chapter 2

The optical problem in photoacoustic imaging

2.1 Radiative transfer

The most complete description of light propagation is provided by Maxwell's electromagnetic equations. However, solving such equations for light transport in tissue is unattractive for a number of reasons. Firstly, most biological tissues are highly scattering meaning the positions and shapes of all the scatterers must be known in order for the model to be accurate. This microscopic description of the tissue results in vast computations which are intractable for tissue volumes on the scale of PAT images. Instead, light transport is commonly modelled using the time-independent radiative transfer equation (RTE). The time-independent form is used as stress and thermal confinement ensure acoustic propagation only begins when the incident pulse of light has been fully absorbed (see Sec. 1.2). Nevertheless, the RTE was originally derived in time-dependent form in heuristic fashion, i.e. through the consideration of energy conservation in a small volume dV [57], but can also be derived from energy conservation in Maxwell's equations making use of a time-averaged, volume-averaged expression of the flow of energy [58, 59].

Here we present the derivation of the time-dependent form using energy balancing. Consider the net outflow of energy in time dt from a small volume dV centred around a point $\mathbf{x}$, around a small solid angle in direction $\hat{\mathbf{s}}$ and at time t; the energy is incident upon an area dA perpendicular to $\hat{\mathbf{s}}$ and travels a distance $ds = cdt$ at speed c. The net outflow is then equal to the sum of four contributions: a positive contribution from a light source, negative contributions from absorption and scattering, and a positive contribution from light being scattered into the direction of interest. This is illustrated in Fig. 2.1.

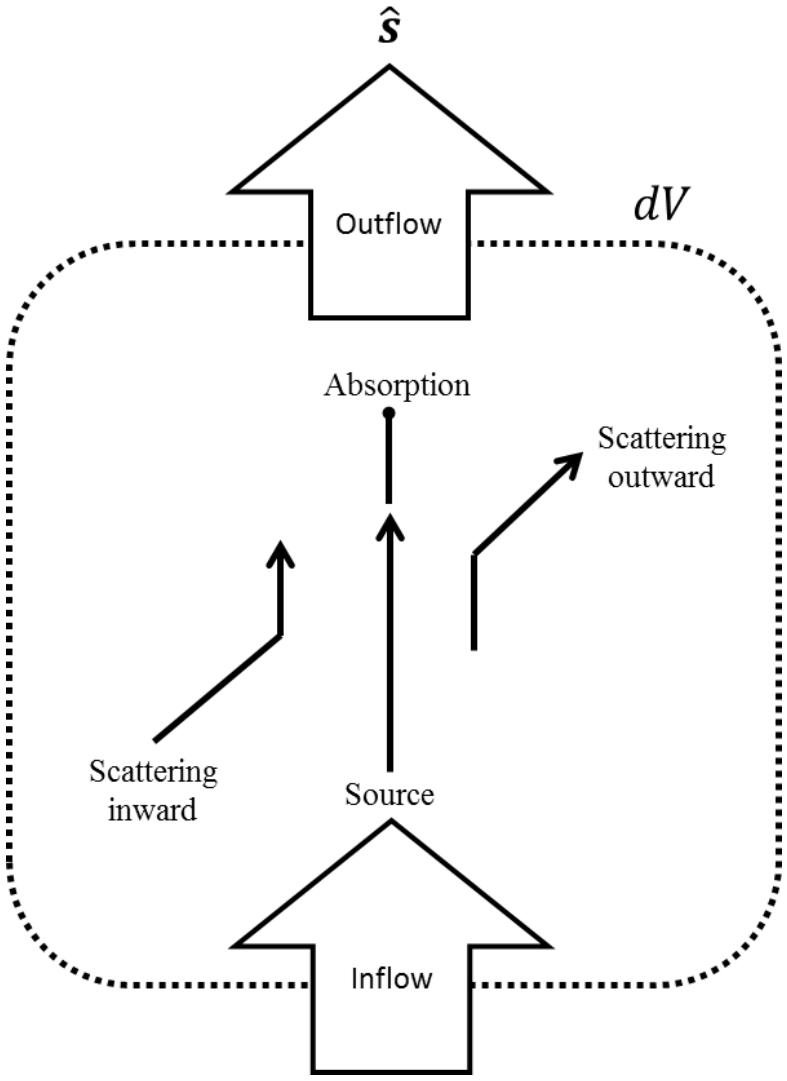

FIGURE 2.1: Illustration of absorption, outward and inward scattering in a small volume of tissue in direction $\hat{\mathbf{s}}$.

Define the time-depenent radiance, $\frac{\partial \phi}{\partial t}$, as the energy per unit area per steradian per unit time; the net change in the radiance per unit length across the volume is given by $cdtdV\frac{\partial \phi(\lambda,\mathbf{x},\hat{\mathbf{s}},t)}{\partial t}$, where dV is very small, λ is the wavelength of light, $\mathbf{x}$ is position and $\hat{\mathbf{s}}$ is the direction of propagation. Note that all nonlinearities are neglected (e.g. inelastic (Raman) scattering, fluorescence), meaning wavelength is constant. Based on the four contributions above, the change in radiance per unit length per unit time can then be written as

$$cdtdV\frac{\partial \phi(\mathbf{x},\hat{\mathbf{s}},t)}{\partial t} = cdtdVq(\mathbf{x},\hat{\mathbf{s}},t) - cdtdV(\mu_{\mathrm{a}}(\mathbf{x}) + \mu_{\mathrm{s}}(\mathbf{x}))\phi(\mathbf{x},\hat{\mathbf{s}},t) + cdtdV\mu_{\mathrm{s}}(\mathbf{x})\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi($$

(2.1)

where $P(\hat{\mathbf{s}},\hat{\mathbf{s}}')$, called the scattering phase function, represents the probability of light scattering from any direction $\hat{\mathbf{s}}'$ into $\hat{\mathbf{s}}$, and μ_{a} and μ_{s} are the absorption and scattering coefficients, respectively. The integral over $\mathcal{S}^{n-1}$ represents integrating the contributions over $n-1$ dimensions which for $n=3$ is performed over the unit sphere over 4π steradians (an analogous form exists for $n=2$). As the photoacoustic source is a function of the time-dependent fluence, we must consider Eq. (2.1) integrated over the entire illumination duration such that we are left with a time-independent equation using the substitution $ds = cdt$

$$\frac{\partial \phi(\mathbf{x},\hat{\mathbf{s}})}{\partial s} = q(\mathbf{x},\hat{\mathbf{s}},t) - (\mu_{\mathrm{a}}(\mathbf{x}) + \mu_{\mathrm{s}}(\mathbf{x}))\phi(\mathbf{x},\hat{\mathbf{s}}) + \mu_{\mathrm{s}}(\mathbf{x})\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi(\mathbf{x},\hat{\mathbf{s}}')d\hat{\mathbf{s}}'. \tag{2.2}$$

Eq. (2.2) can be re-written to give the more familiar form,

$$(\hat{\mathbf{s}}\cdot\nabla + \mu_{\mathrm{a}}(\mathbf{x}) + \mu_{\mathrm{s}}(\mathbf{x}))\phi(\mathbf{x},\hat{\mathbf{s}}) = \mu_{\mathrm{s}}(\mathbf{x})\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi(\mathbf{x},\hat{\mathbf{s}}')d\hat{\mathbf{s}}' + q(\mathbf{x},\hat{\mathbf{s}}). \tag{2.3}$$

2.1.1 Absorption

The absorption coefficient is an expression of the probability of absorption per unit length [60]. As chromophore concentrations within the tissue may change (for physiological

reasons such as blood flow or metabolism) or the concentrations themselves are of interest, the absorption coefficient is resolved into the molar absorption coefficients, $\alpha_k(\lambda)$, which describes the probability of absorption per unit length per unit concentration, and the concentrations of the respective k chromophores, c_k. This is written as a linear combination of the concentrations of the K chromophores present:

$$\mu_{\mathrm{a}}(\lambda) = \sum_{k=1}^{K} \alpha_k(\lambda) c_k. \tag{2.4}$$

In practice, this relationship may not be linear, as $\alpha_k(\lambda)$ may also depend on c_k or $c_{k'}$, as in the case of certain molecular dyes. However, this thesis assumes linearity of Eq. (2.4), which is a fair assumption in biological tissues as many chromophores are contained within a cell, meaning the concentrations of the chromophore in its local environment is fixed. Taking the heme molecule as an example, it is bound within the haemoglobin protein and therefore will not undergo changes in concentration.

In the absence of scattering and with a boundary source at $s = 0$ in direction $\hat{\mathbf{s}}$ ($q = q_0\delta(s-0)\delta(\hat{\mathbf{s}}-0)$), the RTE simplifies to

$$\frac{d\phi}{ds} = -\mu_{\mathrm{a}}\phi + q_0\delta(s-0)\delta(\hat{\mathbf{s}}-0), \tag{2.5}$$

which has the solution $\phi = q_0 \exp(-\int_{s=0}^{s} \mu_{\mathrm{a}} ds')$ or $\phi = q_0 \exp(-\mu_{\mathrm{a}} s)$ in a homogeneous medium (where q_0 is the radiance on the boundary). Since no light is scattered inwards from other directions, the solution for ϕ can be integrated over angle, giving

$$\Phi = Q_0 \exp(-\int_{s=0}^{s} \mu_{\mathrm{a}} ds'), \tag{2.6}$$

with Q_0 being the fluence on the boundary, and is known as the Beer-Lambert law.

2.1.2 Scattering

The scattering coefficient, μ_{s}, is defined in a similar manner to the absorption coefficient in that it is the probability per unit length of a scattering event [60]. Its origins are associated with refractive index changes between cells in the tissue [9].

2.1.2.1 Scattering phase function

On a microscopic level, the angular distribution of singly-scattered light may be adequately described by Rayleigh or Mie theory, but this does not appropriately model scattering by cells in tissue. In addition, multiple scattering is common in biological tissue due to the very short mean-free path between scattering events, and series solutions to scattering become computationally intractable for modelling large volumes of tissue. Instead it is common to use either the Henyey-Greenstein [61] or the delta-Eddington [62] phase functions, though a number of others exist. One expression of the Henyey-Greenstein phase function is in terms of spherical harmonics:

$$P_{HG}(\hat{\mathbf{s}}, \hat{\mathbf{s}}') = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} g^l Y_l^{m+}(\hat{\mathbf{s}}') Y_l^m(\hat{\mathbf{s}}), \tag{2.7}$$

with Y_l^m as the spherical harmonics and the superscript $+$ indicates complex conjugation. g, the anisotropy factor, is on $[-1, 1]$; scattering is isotropic when $g = 0$, forward-peaked when g approaches +1 and backward-peaked when g approaches -1. The product $Y_l^{m+}(\hat{\mathbf{s}}')Y_l^m(\hat{\mathbf{s}})$ ensures the phase function is symmetric since odd components cancel. The phase function has the attractive property that it has a closed form expression:

$$P_{HG}(\hat{\mathbf{s}}, \hat{\mathbf{s}}') = \frac{1}{4\pi} \frac{1 - g^2}{(1 + g^2 - 2g(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'))^{3/2}}. \tag{2.8}$$

With the phase function being properly normalised such that,

$$\int_{4\pi} P_{HG}(\hat{\mathbf{s}}, \hat{\mathbf{s}}')d(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}') = 1, \tag{2.9}$$

it then has the expectation

$$g = \int_{4\pi} P_{HG}(\hat{\mathbf{s}}, \hat{\mathbf{s}}')(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}')d(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'). \tag{2.10}$$

2.2 Models of light transport in tissue

The RTE is the most complete macroscopic model of light transport in tissue. However, memory demands and challenges associated with its computational implementation have led to other models being sought. Three models commonly used derive from the RTE: the $\mathrm{P_n}$ approximations, the diffusion approximation and the δ-Eddington approximation. These are described below.

2.2.1 $\mathrm{P_n}$ approximations

This class of methods expresses the radiance, source term and scattering phase function in the RTE as a series of spherical harmonics [63]:

$$\phi(\mathbf{x}, \hat{\mathbf{s}}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi}\right)^{1/2} a_l^m Y_l^m(\hat{\mathbf{s}}), \tag{2.11}$$

$$q(\mathbf{x}, \hat{\mathbf{s}}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi}\right)^{1/2} q_l^m Y_l^m(\hat{\mathbf{s}}), \tag{2.12}$$

where a_m^l and q_m^l are the weighting coefficients associated with each harmonic for the radiance and source, respectively. It is in this basis where the expansion in Eq. (2.7) is particularly useful as its substitution into the RTE in Eq. (2.5) along with the expression for the radiance and source term in Eqs. (2.11) and (2.12) allow significant

simplification, as shown below.

$$\sum_{l=0}^{\infty}\sum_{m=-l}^{l}\left(\frac{2l+1}{4\pi}\right)^{1/2}(\hat{\mathbf{s}}\cdot\nabla)a_l^m Y_l^m(\hat{\mathbf{s}})+ \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\left(\frac{2l+1}{4\pi}\right)^{1/2}a_l^m Y_l^m(\hat{\mathbf{s}}) = \mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\left(\sum_{l=0}^{\infty}\sum_{m=-l}^{l}g^l Y_l^{m+}(\hat{\mathbf{s}}')Y_l^m(\hat{\mathbf{s}})\right)\left(\sum_{l=0}^{\infty}\sum_{m=-l}^{l}\left(\frac{2l+1}{4\pi}\right)^{1/2}a_l^m Y_l^m(\hat{\mathbf{s}})\right)d\hat{\mathbf{s}}'+ \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\left(\frac{2l+1}{4\pi}\right)^{1/2}q_l^m Y_l^m(\hat{\mathbf{s}}) \tag{2.13}$$

This approach has the potential to reduce the memory demands of simulating the RTE because storing the field some distance away from directional sources will only require a few orders of spherical harmonics [64].

2.2.2 Diffusion approximation

One particular approach to solving the RTE within the $\mathrm{P_n}$ approximations is the $\mathrm{P_1}$ (which can be further simplified to give the diffusion approximation (DA)). The $\mathrm{P_1}$ assumes near-diffusive propagation by truncating the spherical harmonic series expasions for the terms in Eq. (2.13) for $l = 1$ giving [63]

$$\frac{\partial}{\partial z}a_1^0+\frac{1}{\sqrt{2}}\frac{\partial}{\partial x}\left(a_1^{-1}-a_1^1\right)+\frac{1}{i\sqrt{2}}\frac{\partial}{\partial y}\left(a_1^{-1}+a_1^1\right)+(\mu_{\mathrm{a}}+\mu_{\mathrm{s}})a_0^0=\mu_{\mathrm{s}}a_0^1+q_0^0 \tag{2.14}$$

$$\frac{1}{3}\frac{\partial}{\partial z}a_0^0+(\mu_{\mathrm{a}}+\mu_{\mathrm{s}})a_1^0=\mu_{\mathrm{s}}g^1a_1^0+q_1^0 \tag{2.15}$$

$$\frac{\sqrt{2}}{6}\left(\frac{\partial}{\partial x}+i\frac{\partial}{\partial y}\right)a_0^0+(\mu_{\mathrm{a}}+\mu_{\mathrm{s}})a_1^{-1}=\mu_{\mathrm{s}}g^1a_1^{-1}+q_1^{-1} \tag{2.16}$$

$$-\frac{\sqrt{2}}{6}\left(\frac{\partial}{\partial x}+i\frac{\partial}{\partial y}\right)a_0^0+(\mu_{\mathrm{a}}+\mu_{\mathrm{s}})a_1^1=\mu_{\mathrm{s}}g^1a_1^1+q_1^1, \tag{2.17}$$

where x, y and z are the spatial dimensions in 3D.

The fluence is equal to the isotropic, zeroth spherical harmonic a_0^0, and the flux vector $\mathbf{J} = \int_{\mathcal{S}^{n-1}} \phi(\hat{\mathbf{s}})\hat{\mathbf{s}}d\hat{\mathbf{s}}$ is given by

$$\mathbf{J} = \left(\begin{bmatrix} \frac{1}{\sqrt{2}} \left(a_1^{-1} - a_1^1 \right) \\ \frac{1}{i\sqrt{2}} \left(a_1^- + a_1^1 \right) \\ a_1^0 \end{bmatrix} \right). \tag{2.18}$$

This allows the four equations (2.14)–(2.17) to be written as two equations in $\mathbf{J}$ and Φ

$$\nabla \cdot \mathbf{J} + \mu_\mathrm{a}\Phi = q_0^0 \tag{2.19}$$

$$\frac{1}{D}\mathbf{J} + \frac{1}{3}\nabla\Phi = q_1, \tag{2.20}$$

where $D = \frac{1}{3(\mu_\mathrm{a}+\mu_\mathrm{s}(1-g))}$ and $Q_0 = q_0^0$(Eqs. (2.19) & (2.20) are known as the P_1 approximation). Re-arranging gives

$$\nabla \cdot \mathbf{J} + \mu_\mathrm{a}\Phi = Q_0 \tag{2.21}$$

$$\mathbf{J} = 3D(q_1 - \frac{1}{3}\nabla\Phi), \tag{2.22}$$

and substituting Eq. (2.22) into (2.21), using the assumption that all sources are isotropic ($q_1 = 0$),

$$-\nabla \cdot D\nabla\Phi + \mu_\mathrm{a}\Phi = Q_0, \tag{2.23}$$

which is known as the photon diffusion equation.

2.2.3 δ-Eddington approximation

The limitations of the DA in modelling non-diffusive propagation can be partially remedied using an alternative phase function that considers the directionality of the light field being composed of a diffusive part and a delta-function, known as the δ-Eddington approximation [62]:

$$P_{\delta-E}(\hat{\mathbf{s}}\cdot\hat{\mathbf{s}}') = \frac{1}{4\pi}\left[2f\delta(1-(\hat{\mathbf{s}}\cdot\hat{\mathbf{s}}')) + (1-f)(1+3g(\hat{\mathbf{s}}\cdot\hat{\mathbf{s}}'))\right], \tag{2.24}$$

where f is the fraction of light scattered in the direction of anisotropy (i.e. the direction of scattered light when $g = 1$). Like the DA, the substitution of $P_{\delta-E}(\hat{\mathbf{s}}\cdot\hat{\mathbf{s}}')$ into the RTE allows simplification to a diffusion equation for the scattered component of the fluence, while also producing a collimated component, Φ_c, which has a Beer's law type solution,

$$\Phi_c(\mathbf{x}) = \Phi_0(\mathbf{x})\exp\left(-\int_0^z(\mu_\mathrm{a}(z') + \hat{\mu_\mathrm{s}}(z'))dz'\right)\delta(1-(\hat{\mathbf{s}}\cdot\hat{\mathbf{z}})), \tag{2.25}$$

where $\Phi_0(\mathbf{x})$ is the illumination at the surface, z is depth, $\hat{\mu_\mathrm{s}}$ is a modified reduced scattering coefficient to account for the different phase function [65].

2.3 Solution methods for models of light transport in tissue

The most suitable method used to solve the RTE – and approximations to it – depends on how the optical properties are distributed over the domain, and the illumination and boundary conditions. Analytic solutions for the radiance can be obtained using an infinite space Green's function in homogeneous media [66–68], and solutions have also been extended to simple finite geometries (e.g. slabs, cylinders, spheres) [69, 70] as well as bounded [71, 72] and layered [73] media. Analytic solutions for the fluence under the diffusion and δ-Eddington approximation exist for simple geometries [75–77] and can be obtained in a similar manner. Under the diffusion approximation, in the case of

a homogeneous absorbing-scattering medium, the fluence can be computed by solving $-\nabla \cdot D\nabla G_0(\mathbf{x}-\mathbf{x}') + \mu_{\mathrm{a}}\Phi = \delta(\mathbf{x})$ using the free-space Green's function G_0 [78]:

$$G_0(\mathbf{x}-\mathbf{x}') = \frac{\exp(-\mu_{eff}\left|\mathbf{x}-\mathbf{x}'\right|^2)}{4\pi\left|\mathbf{x}-\mathbf{x}'\right|^2}, \tag{2.26}$$

where μ_{eff} the effective attenuation coefficient is given by $\mu_{eff} = \sqrt{\mu_{\mathrm{a}}/D}$ where $D = \frac{1}{3(\mu_{\mathrm{a}}+\mu_{\mathrm{s}}(1-g))}$. Analytic approaches are not available for arbitrary media, such as those imaged using PAT. More general, numerical solution methods exist either through the finite element method or Monte Carlo methods, as discussed below. (Note that finite difference schemes [79, 80] and spectral methods [81] exist, but are not commonly used).

2.3.1 Finite element methods

Finite element methods project the problem into a basis in which the weak form can be solved using a simple linear system

$$\mathbf{A}\phi_j = b_j, \tag{2.27}$$

where $\mathbf{A}$ is the system matrix, ϕ_j is the field variable and b_j is the projection of the source term into the j^{th} element in the chosen basis [82]. The construction of $\mathbf{A}$ depends on the light model and the basis [63]. Typical discretisations for the basis functions, u_j, are triangles in 2D or tetrahedra in 3D, such that the quantity $\tilde{\phi}$ of interest is expressed in the basis as

$$\phi_j = \int_{\Omega} \tilde{\phi}_j u_j d\mathbf{x}, \tag{2.28}$$

where Ω is the domain of $\mathbf{x}$. Eq. (2.28) is often subject to boundary conditions that ϕ_i and $\nabla\phi_i$, or some combination, are given on the boundary (Dirichlet and Neumann boundary conditions, respectively).

When solving for the radiance according to the RTE using the FE method, the system is projected into both spatial and angular bases in order to account for anisotropic light propagation. This becomes very computationally demanding in highly anisotropic 3D media because of the need for fine angular discretisation near the source, in both the angular and spatial domains. The discrete ordinate method of angular discretisation has been employed in 2D using FE discretisation in the spatial domain [53] where the scale of the computation remains tractable. However, the significant memory demands of 3D solutions have led to reducing the dimensionality in the angular domain by using a spherical harmonic basis [83]. Discretisation that is variable in order over the domain (i.e. near the source, spherical harmonics for many l are stored, whereas far from the source $l = 1$ is sufficient) has the potential to reduce memory demands quite significantly [64].

Finite element solutions to the the δ-Eddington approximation [84] are much lighter computationally since only two angular bins are required, while the DA requires only one angular bin [85]. Although this has significant benefits in terms of computational demand, the disadvantage is that these models break down under certain conditions; the δ-Eddington does not completely accurately model anisotropic light propagation, while the diffusion approximation is inaccurate near the tissue surface or near sources where the light is more strongly directional (which is typically within a distance $1/\mu_{\mathrm{s}}' = 1/(\mu_{\mathrm{a}} + \mu_{\mathrm{s}}(1-g))$) and in low albedo ($\mu_{\mathrm{s}} \not> \mu_{\mathrm{a}}$) regions. The former is problematic for the application of the DA as a forward model in PAT because the ballistic and near-ballistic region near the tissue surface at depths less than $1/\mu_{\mathrm{s}}'$ is a region of interest for PA imaging.

2.3.2 Monte Carlo methods

The Monte Carlo (MC) method of light transport is a stochastic approach and is considered the 'gold standard' model of the RTE because (a) it converges to a solution of the RTE [86] and (b) it can implemented straightforwardly in a range of geometries. The method involves iterative sampling of the RTE such that energy packets propagate along random walks through the domain. Each photon trajectory acts as a sample of the RTE

and the deposition of energy by the packet is calculated according to the RTE. As the number of energy packets simulated grows, the quantity of interest, typically the fluence, tends to a solution of the RTE.

The number of energy packets required in a simulation varies depending on the application, domain size, illumination and boundary conditions. In the case of PAT, where the domain is 3D and has fine spatial discretisation, the number of energy packets required can easily exceed 10^6. Even simulations using well over this number of energy packets can produce fluence estimates with excessively high variance to be useful. For this reason, the application of MC models in QPAT has been limited. However, energy packets in MC simulations can be propagated completely independently, meaning this parallelism can be exploited using high throughput computing schemes and in turn may provide a tractable means of solving the RTE in 3D. Chapters 3 and 4 describe in detail the MC method of light transport, and tackles the question of how many energy packets are sufficient in a given simulation.

2.4 The optical inverse problem

As mentioned in Chapter 1, QPAT typically refers to solving the optical inverse problem given a photoacoustic image, i.e. the reconstructed initial acoustic pressure, as the data. The range of possible objectives in QPAT is to estimate

- the absorption coefficient only;
- both the absorption and scattering coefficients;
- the absorption and scattering coefficients and the Grüneisen parameter;
- chromophore concentrations; and/or,
- blood oxygenation.

A spectrum of approaches have been presented in the literature, ranging from solving the linear inverse problem alone under many simplifying assumptions, to solving the

coupled acoustic-optical inverse problem with many parameters unknown and partial data. These are discussed below.

2.4.1 Linear approaches

A key attraction of QPAT is that it has the potential to estimate physiological parameters, such as chromophore concentrations or ratios between such concentrations. In the context of QPAT, a linear inversion aims to exploit the wavelength dependence of the chromophores' molar absorption coefficients by acquiring images at two or more wavelengths and performing an inversion under the assumption that the forward problem is linear in μ_a, i.e. that $p_0(\lambda) \propto \mu_\mathrm{a}(\lambda)$. Writing out the forward problem at a single position for the multiwavelength case for wavelengths $1...N$ and $1...M$ chromophores,

$$\begin{bmatrix} p_0(\lambda_1) \\ \vdots \\ p_0(\lambda_N) \end{bmatrix} = \Gamma \begin{bmatrix} \Phi(\lambda_1) & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Phi(\lambda_N) \end{bmatrix} \begin{bmatrix} \alpha_1(\lambda_1) & \dots & \alpha_M(\lambda_1) \\ \vdots & \ddots & \vdots \\ \alpha_1(\lambda_N) & \dots & \alpha_M(\lambda_N) \end{bmatrix} \begin{bmatrix} c_1 \\ \vdots \\ c_M \end{bmatrix}, \tag{2.29}$$

illustrates that generally two steps are required to obtain the chromophore concentrations from a position in a photoacoustic image (note that in practice these equations would be solved at multiple positions in the image simultaneously); the left-hand side must be multiplied by the reciprocal fluence term, followed by multiplication using the pseudoinverse of the matrix of molar absorption coefficients:

$$\begin{bmatrix} c_1 \\ \vdots \\ c_M \end{bmatrix} = \frac{1}{\Gamma} \begin{bmatrix} \alpha_1(\lambda_1) & \dots & \alpha_M(\lambda_1) \\ \vdots & \ddots & \vdots \\ \alpha_1(\lambda_N) & \dots & \alpha_M(\lambda_N) \end{bmatrix}^\dagger \begin{bmatrix} 1/\Phi(\lambda_1) & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 1/\Phi(\lambda_N) \end{bmatrix} \begin{bmatrix} p_0(\lambda_1) \\ \vdots \\ p_0(\lambda_N) \end{bmatrix} \tag{2.30}$$

where † is used to indicate the (pseudo)inverse and the problem in Eq. (2.29) is typically well-conditioned for $N \geq M$ depending on the $\alpha_k(\lambda)$. Under the assumption of linearity between the photoacoustic image and the absorption coefficient, Eq. (2.30) can simply

be written as

$$\begin{bmatrix} c_1 \\ \vdots \\ c_M \end{bmatrix} = \frac{1}{\Gamma} \begin{bmatrix} \alpha_1(\lambda_1) & \dots & \alpha_M(\lambda_1) \\ \vdots & \ddots & \vdots \\ \alpha_1(\lambda_N) & \dots & \alpha_M(\lambda_N) \end{bmatrix}^{\dagger} \begin{bmatrix} p_0(\lambda_1) \\ \vdots \\ p_0(\lambda_N) \end{bmatrix}. \tag{2.31}$$

This relationship is not true in general but solving the optical inverse problem in this way is attractive as it avoids having to estimate the fluence, which is challenging and often computationally intensive. Many practitioners estimate blood oxygenation, $sO_2 = c_{HbO2}/(c_{HbO2} + c_{Hb})$, using this method [20, 87–92], though their accuracy has rarely been explicitly validated. Such methods are unlikely to yield sO_2 estimates that are of a practically useful level of accuracy because the wavelength-dependence of the fluence will corrupt oxygenation estimates. The conditions under which this 'linear inversion' is accurate are investigated in Chapter 3.

Partially correcting for the depth- and wavelength- dependence of the fluence has been attempted using a 1D exponential correction factor with a wavelength-dependent decay coefficient [93]. This approach is routinely used in practice [11, 94–96] to reduce the spatial and spectral dependence of the fluence, and once this correction has been applied, the spectroscopic inversion in Eq. (3.3) is used to estimate chromophore concentrations or blood oxygenation. The implementation of this fluence correction and its impact on the accuracy of sO_2 estimates are also examined in Chapter 3.

2.4.2 Nonlinear approaches

Nonlinear approaches comprise a broad range of techniques used to recover optical and physiological parameters from PAT images. The majority of approaches involve casting the inversion as a least-squares problem. However, there are others that assume a known scattering coefficient to exploit the fact that the absorption coefficient can be expressed as a function of the PA image and unknown fluence, as performed in a fixed-point iteration. The fact that the PA image appears as a term in the diffusion approximation can also be exploited by re-casting the DA assuming a known diffusion coefficient.

2.4.2.1 Fixed-point iteration

It is possible to write an expression for the absorption coefficient in terms of the measured image and the fluence, which is unknown:

$$\mu_{\mathrm{a}} = \frac{p_0^{meas}}{\Gamma\Phi}. \tag{2.32}$$

Assuming that $\Gamma = 1$ or known allows us to cast this as a fixed-point iteration as the current estimate of the absorption coefficient $\mu_{\mathrm{a}}^{(n+1)}$ will depend on the previous estimate of the fluence, $\Phi^{(n)}$:

$$\mu_{\mathrm{a}}^{(n+1)} = \frac{p_0^{meas}}{\Phi^{(n)}(\mu_{\mathrm{a}}^{(n)}, \mu_{\mathrm{s}}) + \varrho}, \tag{2.33}$$

where ϱ is a small number that ensures the denominator is never zero, without which the method could only be applied to regions of the image wher there is sufficient SNR. The absorption coefficient is estimated by iterating over this expression using a fluence model and iteration stops when the error between the measured image and the image estimate, $\mu_{\mathrm{a}}^{(n+1)}\Phi^{(n+1)}$, falls below some threshold. This method suffers from the significant drawback that *a priori* knowledge of the scattering coefficient, which is generally unknown, is required. Accurate estimation of the absorption coefficient using this approach relies on the sensitivity of the PA image to the estimate of the scattering coefficient [97]; it has however been shown that estimates of the absorption coefficient made from noisy images and a correct guess of the background scattering coefficient is possible [98].

2.4.2.2 Re-arranging of the forward model

Banerjee et al. [99] noticed that the image divided by the Grüneisen parameter appears in the diffusion equation,

$$-\nabla \cdot D\nabla\Phi + \mu_{\mathrm{a}}\Phi = Q_0, \tag{2.34}$$

so making the substitution $p_0^{meas}/\Gamma = \mu_{\mathrm{a}}\Phi$,

$$-\nabla \cdot D\nabla\Phi = Q_0 - p_0^{meas}/\Gamma, \tag{2.35}$$

in turn allows this equation to be solved for Φ. The absorption coefficient can then be calculated via $\mu_{\mathrm{a}} = \frac{p_0^{meas}}{\Gamma\Phi}$. This however requires D to be approximated via $\hat{D} = \frac{1}{3\mu_{\mathrm{s}}(1-g)} = \frac{1}{3\mu_{\mathrm{s}}'}$, where the value of μ_{s}' is known, and of course requires Γ to be known. In the situation where μ_{s}' is unknown, the accuracy of absorption estimates is then vulnerable to inaccuracies in the estimate of the reduced scattering coefficient. For this reason, schemes which can perform simultaneous estimation of the absorption and scattering coefficients are more attractive.

2.4.2.3 Linearisation of the forward model

Linearisation-type approaches can estimate the absorption and scattering coefficients simultaneously. This class of inversion expands the photoacoustic forward problem as a Taylor series about an estimate of the absorption and scattering coefficients, $\chi^{(0)} = [\mu_a^0 \ \mu_s^0]$, which in practice can be their true value or some guess of their background value. If the photoacoustic image is formed by a product of the Grüneisen parameter and a forward operator $\mathcal{F}$ applied to χ such that

$$p_0(\mu_{\mathrm{a}}, \mu_{\mathrm{s}};\ g) = \Gamma\mathcal{F}(\chi;\ g)[\chi]. \tag{2.36}$$

Then, the expansion is given by

$$p_0(\mu_{\mathrm{a}}, \mu_{\mathrm{s}};\ g) \approx \Gamma\left[\mathcal{F}(\chi;\ g)[\chi^{(0)}] + \mathcal{F}'(\chi)[\chi - \chi^{(0)}] + \mathcal{F}''(\chi)[\chi - \chi^{(0)}]^2 + \ldots\right], \tag{2.37}$$

where $\mathcal{F}'$ and $\mathcal{F}''$ are the first- and second- order Fréchet derivatives of the forward model. Therefore, despite the forward model being nonlinear, it can be linearised in the first Fréchet derivative by considering that p_0 has been perturbed by some small changes

in μ_a and μ_s: $\Delta\chi = [\mu_a - \mu_{\mathrm{a}}^{(0)},\ \mu_s - \mu_{\mathrm{s}}^{(0)}]$

$$\Delta p_0 = \Gamma \mathcal{F}'(\chi)[\Delta\chi], \tag{2.38}$$

where $\Delta p_0 = p_0^{meas} - p_0$. The update to χ is then given by

$$\Delta\chi = \frac{1}{\Gamma}\mathcal{F}'^{\dagger}[\Delta p_0], \tag{2.39}$$

where $\dagger$ is used to indicate the pseudoinverse (in practice this matrix inverse can be regularised via a number of methods to reduce noise or enforce smoothness in $\Delta\chi$) [100]. Perturbation-type methods such as those in Eq. (2.39) are only valid for small perturbations in the model parameters μ_a and μ_s and require the background absorption and scattering coefficients to be accurately known which limits the applicability of these techniques in general [101]. The formulation of the inversion in Eq. (2.39) can also be used within a minimisation scheme as achieved by Yao et al. [102], but the significant memory and algorithmic demands of computing $\mathcal{F}'$ restricts this approach to inversions of limited scale (e.g. few unknowns, low resolution of χ).

2.4.2.4 Noniterative methods

Bal et al. [103] proposed a noniterative method to uniquely recover the absorption and diffusion coefficients using the diffusion equation as the forward model. Here, the DA for two illumination positions (and thus two images H_1 and H_2) is re-arranged to yield

$$\nabla \cdot \left(\frac{\sqrt{D}}{\mu_{\mathrm{a}}} \mathbf{B} \right) = 0, \tag{2.40}$$

where

$$\mathbf{B} = H_1 \nabla H_2 - H_2 \nabla H_1 = \mu_{\mathrm{a}}^2 \Phi_1^2 \nabla \left(\frac{\Phi_1}{\Phi_2} \right), \tag{2.41}$$

and where Φ_1 and Φ_2 are the fluences corresponding to the two illuminations. This method allows stable reconstruction of D and μ_{a} when the vector field $\mathbf{B}$ connects every position in the domain to a point on the boundary, as demonstrated in a minimisation based framework [104]. Nevertheless, the reliance of this method on the DA and sensitivity to noise due to the ratio of images would limit its application in practice.

2.4.2.5 Least-squares approaches

Minimisation-based approaches derive from the Bayesian framework (discussed in more detail later) and rely on the minimisation of an error function $\epsilon(\chi)$ with respect to the parameter set, χ,

$$\chi = \arg\min_{\chi} \epsilon(\chi) = \arg\min_{\chi} ||y - \mathcal{F}[\chi]||. \tag{2.42}$$

Note that in practice the problem in Eq. (2.42) is often subject to some regularisation, but the choice of regulariser is application-specific and is discussed in more detail where relevant.

Two classes of iterative least-squares minimisation approach are trust-region methods and linesearch methods [100]; trust-region methods use a local model of the error function at the current parameter estimate and identify a direction and step length within this region where the objective will undergo a sufficient change when the parameter set is updated. Linesearch methods, on the other hand, choose a search direction using gradient information and perform a 1D minimisation along the optimal trajectory.

Linesearch methods require a gradient to be computed to minimise Eq. (2.42). Two classes of linesearch method are: Newton and quasi-Newton methods. Newton methods

rely on calculation of the Hessian, the matrix of second-order partial derivatives of the error functional with respect to the model parameters. This approach suffers from the fact that computation of this matrix is not only time-consuming in an algorithmic sense, but memory demands scale quadratically with the number of unknowns meaning the Hessian is prohibitively large to store for large PA images which can contain millions of unknowns. In Newton methods, the descent direction is given by

$$p^{(i)} = -(\nabla^2 \epsilon(\chi)^{(i)})^{-1} \nabla \epsilon(\chi)^{(i)}, \tag{2.43}$$

where $\nabla^2 \epsilon(\chi)^{(i)}$ (the Hessian) and $\nabla \epsilon(\chi)^{(i)}$ (the gradient) are the second- and first-order derivatives of the error functional with respect to the model parameters at the i^{th} iteration. This approach has the advantage that the step-length, $\alpha^{(i)}$, is naturally 1 and has a quadratic rate of convergence. The computation of the derivatives of the error function are usually calculated using some numerical scheme, such as a finite difference algorithm (methods for gradient calculation are discussed more detail in Chapter 5) which can be computationally very costly to store and compute, particularly for the Hessian matrix.

Quasi-Newton methods only require calculation of the gradient and are therefore also termed 'gradient-based' methods. Under this scheme, the gradient is used to approximate the Hessian rather than explicitly computing it. A number of different algorithms exist for forming approximations to the Hessian matrix (e.g. Broyden-Fletcher-Goldfarb-Shanno (BFGS)) and the corresponding parameter updates. Convergence of quasi-Newton methods is slower than the Newton approach, but is well-suited to large scale inverse problems where memory demands limit the range of algorithms that can be applied. Gradient-based approaches simply require a search direction $p^{(i)}$ at the i^{th} iteration. The most intuitive of which is the gradient descent (or steepest decent) method whereby the descent direction is simply given by

$$p^{(i)}(\chi) = -\nabla \epsilon^{(i)}(\chi). \tag{2.44}$$

This scheme, unlike the Newton method, is not pre-conditioned and an appropriate step length, α^i, must be selected; α^i can either be pre-defined, obtained using a line search in the p direction, or can be found using some pre-conditioning scheme [105]. The step length must be chosen in a similar way when performing a nonlinear conjugate gradient (nCG) optimisation, which descends the error hypersurface using a series of steps along orthogonal directions with well-chosen step sizes [105] and typically requires fewer iterations to converge to the global minimum compared with the gradient descent method.

Gradient-based optimisation has been applied in conjunction with forward and adjoint models of diffusion [106] and radiative transport models [107]. As discussed previously, the diffusion model does not accurately account for the fluence in a region of interest for PAT and functional gradients $\frac{\partial \epsilon}{\partial \mu_\mathrm{a}}$ and $\frac{\partial \epsilon}{\partial D}$ are not consistent with those obtained using forward and adjoint radiative transfer equations, meaning the reconstruction is limited in accuracy due to modelling error. Reconstructions using the RTE are accurate to within a few % of the true absorption and scattering distributions in the presence of 1% Gaussian noise [107]. However, this was achieved using four illumination positions and the reconstructed PA images were free from artefacts. In practice, it is rarely possible to illuminate the target from all sides and artefacts from acoustic reconstruction, such as limited view artefacts, will result in errors in the optical inversion. Possible means of remedying this issue are discussed later in this chapter.

2.4.2.6 Single-step methods

The optimisation-based approaches above aim to obtain the optical parameters $(\Gamma, \mu_\mathrm{a}, \mu_\mathrm{s})$ by inverting the absorbed energy density, H, which is first reconstructed from multiwavelength acoustic time-series $p_d(\mathbf{x}_d, t, \lambda)$ or time-series acquired by illuminating the target from different positions $p_d(\mathbf{x}_d, t, q(\mathbf{x}_s, \hat{\mathbf{s}}))$ where $q(\mathbf{x}_s)$ are sources at $\mathbf{x}_s$ positions. This one-step reconstruction can be achieved by combining the forward optical and acoustic operators, which depend on $(\Gamma, \mu_\mathrm{a}, \mu_\mathrm{s})$ and c, respectively, into one single mapping, $\Lambda(\Gamma, \mu_\mathrm{a}, \mu_\mathrm{s};\ c)$. Ding et al. [108] employ this scheme using a FE model of diffusion for the light field and a finite difference model of acoustic propagation to reconstruct μ_a

and c, assuming Γ and the diffusion coefficient are known. The inversion is achieved using an optimisation scheme that computes gradients of Λ, $\frac{\partial \Lambda}{\partial \mu_\mathrm{a}}$ and $\frac{\partial \Lambda}{\partial c}$, using adjoint diffusion and wave equations and 8 distinct illuminations. A key benefit of this approach is that optical and acoustic parameters can be reconstructed simultaneously, however this relies on access to more data through the use of multiple illuminations in order to overcome any ill-posedness in this inversion. It was also found that run-time for single-step reconstruction of μ_a can be slower than the equivalent two-step (acoustic followed by optical) inversion [109].

2.4.2.7 Bayesian approach

Bayes relation states that given the data, y_m, the probability of parameters χ_n is proportional to the product of the prior distributions and the probability of the measurement given the parameters, or likelihood function, $P(y_m|\chi_n)$ [110]:

$$P(\chi_n|y_m) \propto P(\chi_n)P(y_m|\chi_n). \tag{2.45}$$

$P(\chi_n)$ is an expression of the prior probability of the parameter set, while $P(y_m|\chi_n)$ represents the likelihood of a measurement y_m being obtained given the parameters χ_n. Pior knowledge is included in the inversion through $P(\chi_n)$. In the presence of additive Gaussian-distributed noise, $e = y_m - \mathcal{F}(\chi_n)$, the likelihood function is given by $P_{\text{noise}}(y - \mathcal{F}(\chi_n))$, where $\mathcal{F}$ is the forward model. It follows that P_{noise} is given by

$$P_{\text{noise}}(y_m|\chi_n) \propto \exp\left(-\frac{1}{2}\left[y_m - \mathcal{F}(\chi_n)\right]^T \Gamma_e^{-1}\left[y_m - \mathcal{F}(\chi_n)\right]\right). \tag{2.46}$$

which is normally distributed with covariance Γ_e. Written in this way Eq. (2.46) illustrates the flexibility of the Bayesian approach in solving inverse problems corrupted by additive noise (it is applicable to problems affected by other noise models but this is not discussed here) because uncertainty on the priors and the likelihood function can incorporated. The Bayesian formulation of the inversion also allows derivation of the least-squares approach (Section 2.4.2.5), which can be achieved using the assumption

that Γ_e is diagonal and by taking the negative natural log of the expression for P_{noise} in Eq. (2.46) to reveal the error functional in Eq. (2.42).

In the estimation of the optimum parameter set χ_n, the maximum *a posteriori* (MAP) is often sought, i.e. the parameter set which maximises $P(y_m|\chi_n)$:

$$\chi_{MAP} = \arg\max_{\chi_n} P(y_m|\chi_n). \tag{2.47}$$

This approach has been demonstrated to remove artefacts in parameter estimates introduced by the acoustic inversion [111]. However, the memory and algorithmic demands of computing Γ_e can be significant for large datasets [112], meaning statistical inversion techniques are not always attractive for application in PAT.

In summary, there are two broad classes of inversion: those that neglect the impact of the fluence or assume it is constant with wavelength, and those that account for the fluence using some analytic or numerical model. Chapter 3 investigates the validity of the first class of inversion, while Chapters 5 and 6 consider novel ways of tackling model-based inversion and examine which optimisation scheme performs best for gradient estimates made using a MC model of light transport. Both types of inversion approach are not without their challenges, discussed in Section 2.4.3, and knowledge of these obstacles is important in maximising the accuracy of these inversion schemes.

2.4.3 Challenges in QPAT

In addition to the fact that the fluence and photoacoustic efficiency are unknown, there are a number of challenges that stand in the way of obtaining accurate estimates of chromophore concentrations. These are discussed below.

2.4.3.1 Inaccuracy of fluence models

Even when the fluence is modelled, there is no guarantee that the model used will accurately model the experimental imaging configuration. The DA, despite benefits in

terms of ease of implementation and memory demands scaling quadratically with the number of elements in the domain, it is only accurate in the diffusive regime, a distance $> 1/(\mu_{\mathrm{a}} + \mu_{\mathrm{s}}(1-g))$ away from the source, and can only accomodate isotropic sources. The RTE and Monte Carlo models of the RTE are in theory accurate throughout the domain, but boundary conditions and source implementations may not always match those used in practice and the variance in MC estimates of the fluence will be strongly dependent on the number of photons simulated, domain size and discretisation, as well as the optical properties inside the domain. These inaccuracies will inevitably result in uncertainties in parameters reconstructed from the optical inversion using a deterministic reconstruction algorithm; however, it has been demonstrated [111] that these uncertainties can be accounted for using a parameter covariance matrix, Γ_χ. The estimation of Γ_χ does not come without its challenges but can at least provide some estimate of the uncertainty on parameter estimates.

2.4.3.2 Nonuniqueness

The nonuniqueness problem in QPAT refers to the inability to simultaneously accurately estimate a combination of the Grüneisen parameter, absorption coefficient and/or scattering coefficient due to a lack of data, either from only using a single illumination or a single wavelength. Strategies to overcome nonuniqueness often rely on obtaining images using multiple wavelengths or illumination positions. Even using multiple images, each acquired using a different illumination position, only two of the three quantities above can be recovered uniquely provided the positions are well-chosen (as demonstrated by Bal et al. [104]). Aside from multiple illumination positions [107, 113, 114], it is also possible to overcome the absorption-scattering nonuniqueness by using multiple illumination patterns [103] or multiple wavelengths over which the scattering coefficient's dependence is known [115]; by parameterising the reduced scattering, $\mu_{\mathrm{s}}' = \mu_{\mathrm{s}}(1-g)$, as a function of wavelength, $\mu_{\mathrm{s}}'(\mathbf{x},\lambda) = \mu_{\mathrm{s}}(\mathbf{x})''\lambda^{-b}$, with μ_{s}'' being the spatial dependence of the scattering, the optimisation minimises the objective function with respect to the parameters μ_{s}'' and chromophore concentrations [116].

Spatial priors can also help overcome the ill-posedness in Γ, μ_{a} and μ_{s}. Bal et al. [117] use multiple illuminations and require smooth distributions of the absorption and scattering coefficients in order for accurate and unique reconstruction. Malone et al. [118] treat the estimation of the absorption and scattering coefficients as a combined reconstruction-classification problem. This involves classifying regions in the domain according to an estimate of their absorption and scattering coefficients, given prior estimates of μ_{a} and μ_{s}' for each class. The algorithm alternates between reconstruction and classification steps upon each iteration allowing significant improvements on the accuracy of estimates of the scattering coefficient made using this technique.

One approach to remove the non-uniqueness between absorption, scattering and the Grüneisen parameter which does not require prior knowledge of their interrelationship is to exploit the fact that Γ in general does not depend on wavelength and to formulate a new forward problem requiring inversion [104]:

$$\frac{p_0(\mathbf{x},\lambda_1)}{p_0(\mathbf{x},\lambda_2)} = \frac{\Phi(\mathbf{x},\lambda_1;c_k,\mu_s,g)\sum_{k=1}^{K}\alpha_k(\lambda_1)c_k(\mathbf{x})}{\Phi(\mathbf{x},\lambda_2;c_k,\mu_s,g)\sum_{k=1}^{K}\alpha_k(\lambda_2)c_k(\mathbf{x})}, \tag{2.48}$$

such that the Grüneisen parameter cancels. The inversion of Eq. (2.48) for absorption and scattering would require at least three distinct wavelengths or multiple illumination positions; in the three-wavelength case a composite error function can be formed which sums over quadratic errors between differences of ratios of PA images:

$$\epsilon = \frac{1}{2}\left|\left|\frac{p_0^{meas}(\mathbf{x},\lambda_1)}{p_0^{meas}(\mathbf{x},\lambda_2)} - \frac{p_0(\mathbf{x},\lambda_1)}{p_0(\mathbf{x},\lambda_2)}\right|\right|^2 + \frac{1}{2}\left|\left|\frac{p_0^{meas}(\mathbf{x},\lambda_2)}{p_0^{meas}(\mathbf{x},\lambda_3)} - \frac{p_0(\mathbf{x},\lambda_2)}{p_0(\mathbf{x},\lambda_3)}\right|\right|^2. \tag{2.49}$$

2.4.3.3 Elliptic sensitivity to absorption and scattering

In addition to absorption-scattering non-uniqueness, the initial pressure and therefore the error functional differ in their sensitivity to perturbations in absorption and scattering – often by several orders of magnitude. If we consider the error function between a measured image, p_0^{meas}, and the absorbed energy density at an estimate of the absorption

and scattering coefficients, $p_0(\mu_\mathrm{a}, \mu_\mathrm{s})$,

$$\epsilon = \frac{1}{2} \left|\left| p_0^{meas} - p_0(\mu_\mathrm{a}, \mu_\mathrm{s}) \right|\right|^2 , \tag{2.50}$$

which with $\Gamma = 1$ can be written as

$$\epsilon = \frac{1}{2} \left|\left| p_0^{meas} - \mu_\mathrm{a}\Phi(\mu_\mathrm{a}, \mu_\mathrm{s}) \right|\right|^2 . \tag{2.51}$$

Derivatives with respect to μ_a and μ_s gives

$$\frac{\partial \epsilon}{\partial \mu_\mathrm{a}} = -(p_0^{meas} - \mu_\mathrm{a}\Phi(\mu_\mathrm{a}, \mu_\mathrm{s})) \left[\Phi(\mu_\mathrm{a}, \mu_\mathrm{s}) - \mu_\mathrm{a} \frac{\partial \Phi(\mu_\mathrm{a}, \mu_\mathrm{s})}{\partial \mu_\mathrm{a}} \right] \tag{2.52}$$

$$\frac{\partial \epsilon}{\partial \mu_\mathrm{s}} = -(p_0^{meas} - \mu_\mathrm{a}\Phi(\mu_\mathrm{a}, \mu_\mathrm{s})) \left[-\mu_\mathrm{a} \frac{\partial \Phi(\mu_\mathrm{a}, \mu_\mathrm{s})}{\partial \mu_\mathrm{s}} \right] . \tag{2.53}$$

The term inside the square brackets in Eq. (2.52) will in general be much larger than the term in the square brackets in Eq. (2.53) because Eq. (2.52) contains the large fluence term. As such, $\frac{\partial \epsilon}{\partial \mu_\mathrm{a}}$ will typically be much greater than $\frac{\partial \epsilon}{\partial \mu_\mathrm{s}}$. The consequence of this is that when absorption and scattering are being estimated simultaneously in an inversion, convergence in the scattering is slow and inaccurate due to the error surface exhibiting elliptic sensitivity to μ_a and μ_s, with much larger gradients in the absorption axis [98]. There are preconditioning schemes, whereby the error surface is 'sphered', in order to make the gradients along the absorption and scattering axes more similar in magnitude, but the stronger dependence of the error functional on μ_a is inherent as the absorbed energy density depends on the absorption coefficient directly, whereas it only depends on the scattering coefficient through the fluence.

There has been some investigation into the extent to which the accuracy of the absorption estimate – which is generally the quantity of interest given that it is often closely linked with physiology – depends on the accuracy of the scattering estimate. It has been shown that significant inaccuracies in the estimate of the scattering distribution yields highly

inaccurate estimates of the absorption coefficient; however, if the mean of the scattering coefficient is accurately estimated, then the absorption coefficient can be recovered accurately as well [98].

2.4.3.4 Large scale

The inverse problem in PAT is referred to as being 'large scale' due to the number of unknowns. PAT images are high resolution, volumetric datasets that can be cubic centimeters in size with resolution in the tens to hundreds of microns; this leads to images with several million voxels. The optimisation approaches described in Section 2.4 require gradient information, which can either be a functional gradient or an image Jacobian. Functional gradients contain the same number of elements as the domain, while the (full) image Jacobian can contain the number of elements in the domain squared (in practice, prior classification of the image into different regions can reduce the number of unknowns). While the former may only require MB of computational memory for typical high resolution PAT images, the latter can require well into the TB range of memory to store for the same size/resolution image. As such, memory limitations play a crucial role in selecting suitable optimisation algorithms for use in the optical inversion (discussed in Chapter 4). It is worth noting that this problem is inherent to QPAT as the objective is to image structures, which may have differing chromophore concentrations, at very high resolution, which in turn results in a large scale inverse problem when estimating these concentrations.

Chapter 3

Estimating blood oxygen saturation from multiwavelength photoacoustic images

Blood has a fundamental role in the body, key of which is transporting oxygen toward tissues where the partial pressure of oxygen is low and carbon dioxide along the opposite gradient. As such, blood oxygen saturation, also termed blood oxygenation or sO_2 , is a key indicator of tissue function and pathology and is defined as the ratio of oxygenated haemoglobin concentration to total haemoglobin concentration (i.e. the sum of oxygenated and deoxygenated haemoglobin concentrations). The ability to measure blood oxygenation is therefore of significant interest in clinical and pre-clinical settings because it has potential usefulness in characterising and staging cancers [119, 120], as well as monitoring brain activation [121–123].

Measuring blood oxygenation noninvasively is nontrivial; a century ago oxygen partial pressure could be measured invasively via catheterisation of a blood vessel and using gas extraction [124]. Contemporary noninvasive methods such as BOLD-MRI (Blood Oxygenation Level Dependent Magnetic Resonance Imaging) [125] and optical methods such as diffuse optical tomography (DOT) suffer from significant limitations; as well as having limited spatial resolution, the BOLD signal is sensitive to changes in both

blood volume and/or venous deoxy-haemoglobin concentration [126] and therefore cannot measure oxygenation directly, while purely optical methods like DOT [127], NIRS [128] and pulse-oximetry [129] can measure oxygenation accurately but have limited penetration depth and resolution due to the high degree of optical scattering in tissue.

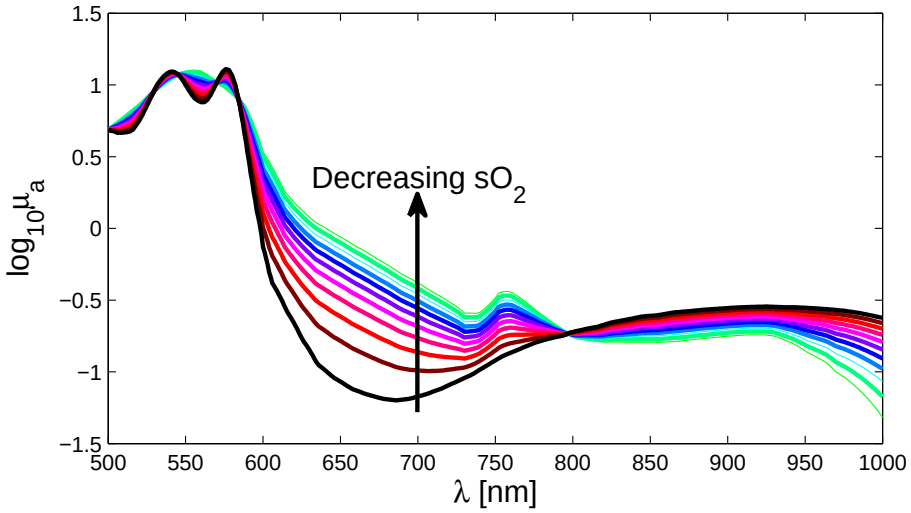

FIGURE 3.1: Absorption spectrum of blood (150gl^{-1} total haemoglobin concentration) for oxygenation levels between 10% and 100% as a function of wavelength between 500-1000nm.

Fig. 3.1 shows the dependence of the absorption coefficient of blood (150gl^{-1} total haemoglobin concentration) on oxygenation level as a function of wavelength between 500-1000nm. The sensitivity of PAI to changes in oxygenation through changes in the absorption coefficient means that it is in principle possible to recover sO_2 from PA measurements made at one or more wavelengths. This can be achieved via arterial blood sampling for calibration of photoacoustic signals in the single wavelength case [130], however this method is not readily extendable to *in vivo* applications. Calibration-free methods are desirable as they do not require prior knowledge of the true oxygenation against which the PA measurement must be calibrated. Such approaches require additional information, which can for instance be a measurement at another wavelength or oxygenation level because this helps remove system-dependent (and wavelength-independent) parameters and is sometimes assumed to remove the impact of the fluence, which in general depends on wavelength.

Approaches to remove system-dependent and fluence effects have been attempted by assuming that the internal fluence distribution does not change with oxygenation in a vessel of interest [131] or that the fluence is simply independent of wavelength [20, 87–92]. These assumptions require in-depth evaluation as they are unlikely to apply in general and the circumstances in which they do apply are of interest, given the challenges associated with modelling the fluence, as described in Chapter 2 and in subsequent chapters of this thesis. This chapter primarily investigates the spectroscopic case, i.e. the latter of the two above assumptions, through a series of simulations. Two numerical phantoms are considered: a homogeneous absorbing-scattering medium in which a blood-filled tube is submerged as well as a highly vascularised block of tissue. Section 3.0.4 describes the relationship between the photoacoustic image and the fluence, and how oxygen saturation can be estimated from multiwavelength photoacoustic images. The model of light transport used to simulate the PA images is then provided. It is first demonstrated that using several wavelengths across the 500-1000nm range yields inaccurate sO_2 estimates, while using pairs of wavelengths in this range can significantly improve the accuracy of oxygenation estimates; it is observed that the choice of wavelengths depends on the oxygenation of the blood in the tube. These observations are largely consistent with those made regarding the tissue phantom simulations considered in Section 3.2.

3.0.1 Estimating blood oxygenation from photoacoustic images

In Chapter 1, the photoacoustic forward model for a single voxel and N wavelengths and M chromophores was presented. We can re-write this equation under the assumption that the only chromophores present at the given position that are optically absorbing at wavelengths $\lambda_1 ... \lambda_N$ are oxy- and deoxy-haemoglobin:

$$\begin{bmatrix} p_0(\lambda_1) \\ \vdots \\ p_0(\lambda_N) \end{bmatrix} = \kappa \begin{bmatrix} \Phi(\lambda_1) & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \Phi(\lambda_N) \end{bmatrix} \begin{bmatrix} \alpha_{HbO2}(\lambda_1) & \alpha_{Hb}(\lambda_1) \\ \vdots & \vdots \\ \alpha_{HbO2}(\lambda_N) & \alpha_{Hb}(\lambda_N) \end{bmatrix} \begin{bmatrix} c_{HbO2} \\ c_{Hb} \end{bmatrix}. \quad (3.1)$$

where κ represents a detection system calibration factor and the photoacoustic efficiency, and c_{HbO2} and c_{Hb} are the concentrations of oxy- and de-oxyhemoglobin respectively. c_{HbO2} and c_{Hb} can be estimated from the photoacoustic amplitudes by inverting Eq. (3.1) with the full inversion given by

$$\begin{bmatrix} c_{HbO2} \\ c_{Hb} \end{bmatrix} = \frac{1}{\kappa} \begin{bmatrix} \alpha_{HbO2}(\lambda_1) & \alpha_{Hb}(\lambda_1) \\ \vdots & \vdots \\ \alpha_{HbO2}(\lambda_N) & \alpha_{Hb}(\lambda_N) \end{bmatrix}^{\dagger} \begin{bmatrix} 1/\Phi(\lambda_1) & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 1/\Phi(\lambda_N) \end{bmatrix} \begin{bmatrix} p_0(\lambda_1) \\ \vdots \\ p_0(\lambda_N) \end{bmatrix} \quad (3.2)$$

where † is used to indicate the pseudoinverse. In a 'linear inversion', the fluence matrix is treated as being proportional to the identity matrix, i.e. constant, $\Phi \propto \mathcal{I}$. Thus, the linear inversion is given by

$$\begin{bmatrix} c_{HbO2} \\ c_{Hb} \end{bmatrix} \propto \begin{bmatrix} \alpha_{HbO2}(\lambda_1) & \alpha_{Hb}(\lambda_1) \\ \vdots & \vdots \\ \alpha_{HbO2}(\lambda_N) & \alpha_{Hb}(\lambda_N) \end{bmatrix}^{\dagger} \begin{bmatrix} p_0(\lambda_1) \\ \vdots \\ p_0(\lambda_N) \end{bmatrix}, \quad (3.3)$$

and demonstrates that this inversion exlicitly assumes proportionality of the PA amplitude to the absorption coefficient.

As oxygenation is defined as a ratio,

$$sO_2 = \frac{c_{HbO2}}{c_{HbO2} + c_{Hb}}, \quad (3.4)$$

the photoacoustic efficiency and system-dependent parameters, i.e. the constants of proportionality, cancel out in Eq. (3.4), as they are independent of wavelength, and this assumption allows the fluence also to cancel. It is this assumption, that $\Phi(\lambda) = constant$, that is investigated in this chapter.

Another assumption requiring consideration given its prevalence in the literature [11, 93–96], is that the fluence can be approximated by a 1D or Beer-Lambert type correction

factor. Solutions for the fluence in homogeneous slab geometry, illuminated at its centre in the z-direction and integrated along the x- and y- directions, follow an exponential relationship as the loss in fluence per incremental length dz is proportional to the light fluence at that position:

$$\frac{d\Phi(z)}{dz} \propto -\Phi(z), \tag{3.5}$$

with $\Phi(z) = \int\int \Phi dxdy$ which can be re-expressed as an equality with a constant of proportionality μ,

$$\frac{d\Phi(z)}{dz} = -\mu\Phi(z). \tag{3.6}$$

Solutions for $\Phi(z)$ of this first-order ODE are exponential in form:

$$\Phi(z) = \Phi(z=0)\exp(-\mu z). \tag{3.7}$$

The value of μ depends on the optical properties of the slab; under the diffusion approximation, i.e. far from the slab surface at $z = 0$ and in a high albedo medium ($\mu_{\mathrm{s}} \gg \mu_{\mathrm{a}}$), μ is a function of the absorption coefficient, μ_{a}, the scattering coefficient, μ_{s}, and the anisotropy factor, g; for a purely absorbing medium, $\mu = \mu_{\mathrm{a}}$. In a real imaging scenario the optical properties are heterogeneous, so slab geometry does not apply, and are unknown; nevertheless, photoacoustic images obtained using one-sided illumination exhibit decreasing intensity with increasing depth, due to wavelength-dependent absorption and scattering of light in the tissue. Thus, we can attempt to compensate for the wavelength-dependent depth-dependence of the fluence using a 1D exponential model given in Eq. (3.7). As we do not have prior knowledge of the value of μ we can extract it by fitting a straight line to the log of the initial acoustic pressure integrated over the x- and y- directions as a function of depth z, with μ^{fit} as the decay coefficient:

$\Phi(\lambda) = \exp(-\mu^{\text{fit}}(\lambda)z)$. The accuracy of this model when used in forming sO_2 estimates is the second assumption investigated in this chapter.

The objective of this chapter is to evaluate the impact of the unknown fluence on the accuracy of a linear inversion and to identify situations in which spectral colouring, i.e. the variation of the fluence's spectrum with absorption, is minimised. It is important to note that throughout this chapter it was assumed that $\kappa = 1$, which comes at no loss of generality as it cancels in Eq. (3.4) because the focus is to determine the effect the fluence's wavelength-dependence has on the accuracy of oxygenation estimates. For the same reason, it was assumed that the initial acoustic pressure was reconstructed perfectly from boundary acoustic measurements, without any artefacts from, for instance, limited sensor aperture or bandlimited detection of acoustic frequencies. If an inversion is not accurate under these ideal circumstances, it is unlikely to perform more accurately with bandlimited or partial data.

3.1 Modelling Light Transport in Tissue: Monte Carlo Simulations

Light can be modelled in tissue using a variety of techniques. We require a model that is accurate for heterogeneous, anisotropic turbid media. Finite element implementations of the RTE satisfy these conditions, but due to significant computational demands of spatial and angular discretisation, this method becomes intractable for this study, while the diffusion approximation is not sufficiently accurate in the near-field. Monte Carlo models (discussed in more depth in Chapter 4) have relatively modest memory requirements which scale linearly with the number of voxels in the grid. Simulations involve propagating 'photon packets' or 'energy packets' along random walks, depositing 'weight' as a function of of the absorption coefficient. 'Photon packets' in the simulation are non-interacting meaning that each packet (hereinafter referred to as 'photon') can be simulated independently, which can be efficiently run using parallel computing architectures such as GPGPUs. The code used here was MCX[132], which has been validated

extensively and is GPU-accelerated, offering relatively short computation times even for the large number of photons required for convergence.

Monte Carlo models are stochastic in nature so, despite them being considered the 'gold standard' [133] model of light transport, it must be appreciated that the fluence in each voxel will exhibit some variance, and convergence to a solution of the RTE requires the number of photons simulated to tend to infinity; as this is practically impossible, we can require that the variance in a given voxel (the 'MC noise') falls below some threshold. In this particular study, this threshold must be low enough to ensure that oxygenation estimates made using initial pressure distributions obtained using MC exhibit a negligible level of variance due to uncertainties in fluence estimates (i.e. the variance in the sO_2 estimates due to MC noise falls below the precision with which we are trying to estimate oxygenation).

For a two-wavelength inversion (Eq. (3.2) with $M = 2$), the standard deviation in oxygenation in a particular voxel, σ_{sO_2}, can be estimated from the standard deviation in the initial acoustic pressure simulated for that voxel, $\sigma_{p_0(\lambda_i)}$, by assuming that the initial pressure at the two wavelengths, $p_0(\lambda_1)$ and $p_0(\lambda_2)$, are uncorrelated [134] and using the standard formula

$$\sigma_{\text{sO}_2} = sO_2 \sqrt{\left(\frac{\partial sO_2}{\partial p_0(\lambda_1)}\right)^2 (\sigma_{p_0(\lambda_1)})^2 + \left(\frac{\partial sO_2}{\partial p_0(\lambda_2)}\right)^2 (\sigma_{p_0(\lambda_2)})^2}, \tag{3.8}$$

where the partial derivatives of sO_2 with respect to initial pressure can be computed from Equations (3.3) and (3.4). The standard deviation $\sigma_{p_0(\lambda_i)}$ was estimated from multiple simulations in the *same* domain with the same number of photons, taking care to re-seed the random number generator to ensure different photon trajectories from simulation to simulation. Running the simulation multiple times at a pair of wavelengths allowed σ_{sO_2} to be calculated. As this quantity is domain-dependent, this parameter was calculated in each phantom examined in this chapter and this is addressed in the relevant sections.

3.1.1 Digital phantom

The first phantom used to examine the assumptions outlined in Section 3.0.4 was composed of a cubic domain 5mm×5mm×5mm in size – of a similar order to *in vivo* studies from literature [135] – with isotropic voxel size of 25μm (similar scale to what can be achieved in PAT) and index-matched, non-scattering boundaries. The illumination was a collimated beam of light, 5mm diameter (top-hat profile), which is representative of illumination from a fibre tip in real-life experiments. The convergence analysis described in Section 3.1 was performed by running the MC model at 500nm and 800nm (chosen due to the significant difference in the absorption of blood at these wavelengths). The maximum relative uncertainty (coefficient of variation) in the fluence between the two wavelengths and all 10 model runs was 0.16% which, propagating to the absolute uncertainty in sO_2 in the tube using Eq. (3.8), gives a maximum $\sigma_{\mathrm{sO_2}}$ value in the tube of 0.34%, indicating that uncertainties in sO_2 estimates due to MC noise are negligible. Note that this was corroborated by directly computing the standard deviation in the sO_2 estimates themselves.

A blood-filled tube, 250μm in diameter, was embedded in the domain at a depth of 1mm, parallel to the y-direction. The blood within the tube had total haemoglobin concentration, c_{HbT}, of 150gl^{-1} and oxygen saturation of 90% ($c_{HbO2} = 135$gl^{-1} and $c_{Hb} = 15$gl^{-1}). The scattering coefficient in the tube followed the relationship $21.3\mathrm{mm}^{-1}\left(\frac{\lambda}{500\mathrm{nm}}\right)^{-1.2}$. The arrangement is shown in Figure 3.2.

In the first instance, in order to carefully assess the impact of spectral colouring on sO_2 estimates made in the blood-filled tube, two seperate cases were considered: one where the background had no spectral dependence, and one where the background was composed of 55% water and 45% blood with a total haemoglobin concentration, c^{bg}_{HbT}, of 5.63gl^{-1} and 60.7% oxygenation ($c_{HbO2} = 3.42$gl^{-1} and $c_{Hb} = 2.11$gl^{-1}) and was used to represent the capillary bed. The wavelength-dependent scattering coefficient of the background was the same as that of the tube. (The optical properties of the background were chosen in order to be consistent with simulations later in this paper). The Henyey-Greenstein phase function was used throughout this paper, with the background

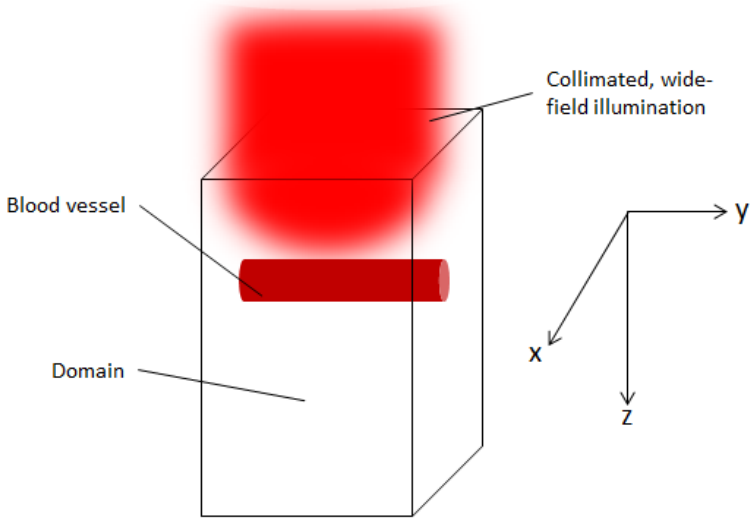

FIGURE 3.2: Schematic of 5mm×5mm×5mm domain with blood vessel and collimated wide-field illumination.

anisotropy factor consistently set at $g = 0.9$ and that of blood being 0.9945. All of the above optical properties are consistent with values quoted in literature [9].

Figure 3.3 shows examples of the initial acoustic pressure and fluence distributions at 580nm and 600nm with a spectrally-dependent background. Both are maximum intensity projections (MIP) in the y-direction, and are normalised by the maximum value of acoustic pressure or fluence, respectively.

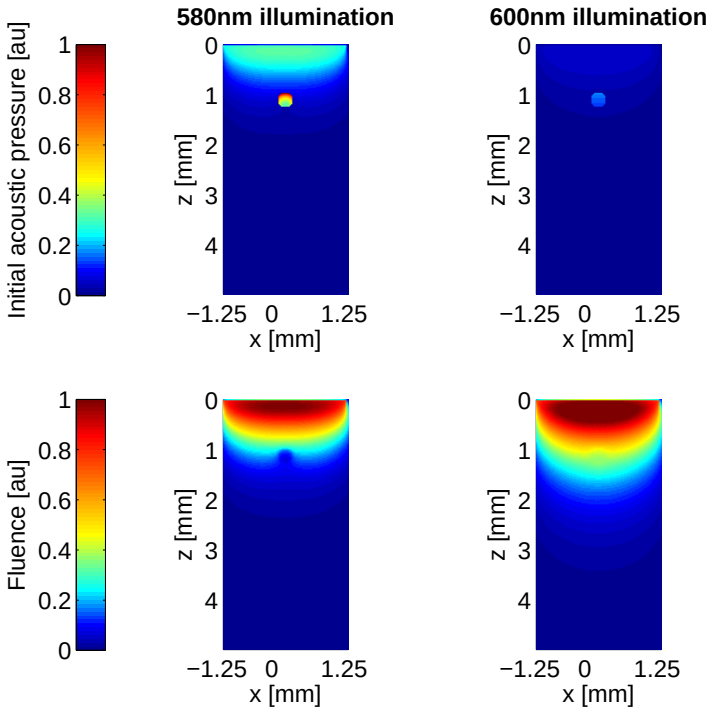

FIGURE 3.3: Top row: MIP of initial acoustic pressure in tube phantom with spectrally-dependent background illuminated with top hat profile at 580nm and 600nm, normalised by maximum acoustic pressure value. Bottom row: MIP of fluence in tube phantom with spectrally-dependent background illuminated with top hat profile at 580nm and 600nm, normalised by maximum fluence value.

The key observation that can be made regarding Fig. 3.3 is that despite the fact that absorption coefficient in the tube is homogeneous, the initial acoustic pressure across the tubes is heterogeneous due to the heterogeneity of the fluence. In addition, the fluence distribution is different at each wavelength, meaning oxygen saturation cannot be estimated accurately without knowledge of the fluence. This is demonstrated in the subsequent sections.

3.1.2 Estimates using 26 wavelengths over the 500-1000nm range

The domain was illuminated at wavelengths in the 500nm to 1000nm range at 20nm intervals and sO_2 was computed according to Eqs. (3.3) and (3.4), assuming $\Phi^{-1} = \mathcal{I}$, the identity matrix, using all wavelengths at each voxel. Defining the absolute error in oxygenation as $|sO_2^{\text{true}} - sO_2|$ (a definition used consistently throughout this thesis), this

quantity is shown as a MIP in the x-z plane half way along the y-axis of the domain in Fig. 3.4 and histograms showing the oxygenation estimate at various depths, for both the spectrally-independent and spectrally-dependent backgrounds. The histograms represent the distribution of sO_2 estimates in voxels contained in x-y slices through the tube at four depths (a, b, c, d), where each x-y slice is orthogonal to the x-z slice shown on the left $sO_2^{\text{mean}}(\text{i})$ (i=[a,b,c,d]) represent the mean of the sO_2 estimates of the voxels contained within the tube for each x-y slice. The global mean of sO_2, shown on the right-hand side of Figure 3.4 represents the mean of the sO_2 estimates in all voxels enclosed within the entire tube.

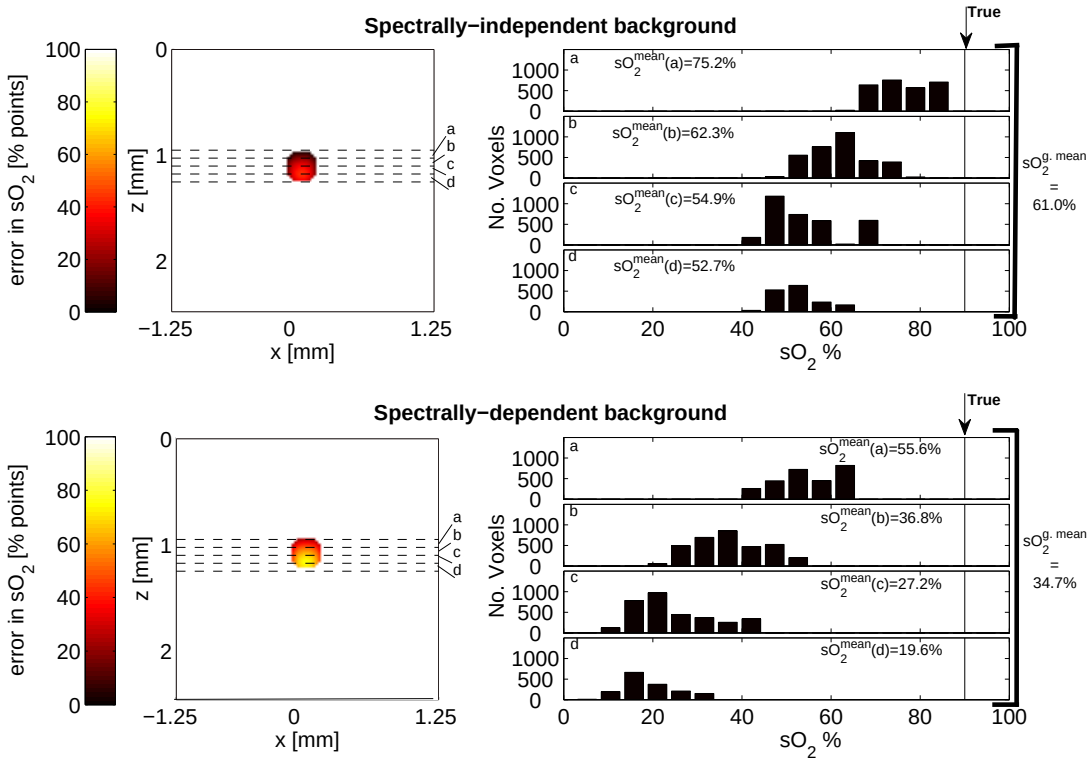

FIGURE 3.4: Oxygenation estimates made in 250μm tube using 26 equally spaced wavelengths in 500-1000nm range. Top row: x-z slice of error in oxygenation mid way along tube's length when the background is spectrally-independent and the error originates from spectral colouring of the fluence in the blood-filled tube. Dashed lines used to indicate four x-y slices (a-d) from which histograms of sO_2 are formed (true value of 90% indicated by black line). Bottom row: x-z slice of error in oxygenation mid way along tube's length when the background properties vary with wavelength. Dashed lines used to indicate four x-y slices (a-d) from which histograms of sO_2 are formed (true value of 90% indicated by vertical black line). $sO_2^{\text{mean}}(\text{a})$, $sO_2^{\text{mean}}(\text{b})$, $sO_2^{\text{mean}}(\text{c})$ and $sO_2^{\text{mean}}(\text{d})$ represent the mean of the sO_2 estimates of the voxels contained within the tube in slices a, b, c and d, respectively. The global mean sO_2 estimate averaged over all voxels enclosed within the entire tube is shown to the right of the histograms.

It can be seen from the x-z slices in the error in sO_2 is minimised near the tube surface, near the surface where the illumination was applied. The minimum error, obtained at the most central region in the tube in the x-y plane of slice 'a', for the spectrally-independent background was about 6% and increased to almost 27% with a spectrally-dependent background. With increasing depth the minimum error increased to 27% and 59% for the spectrally-independent and spectrally-dependent backgrounds, respectively. The significant errors in the oxygenation estimates is the result of spectral colouring which describes the nonlinear, wavelength-dependent relationship between $p_0(\lambda)$ and $\mu_a(\lambda)$; this nonlinearity is highlighted in Fig. 3.5, where the initial acoustic pressure (normalised by its maximum value) is plotted against the normalised absorption coefficient.

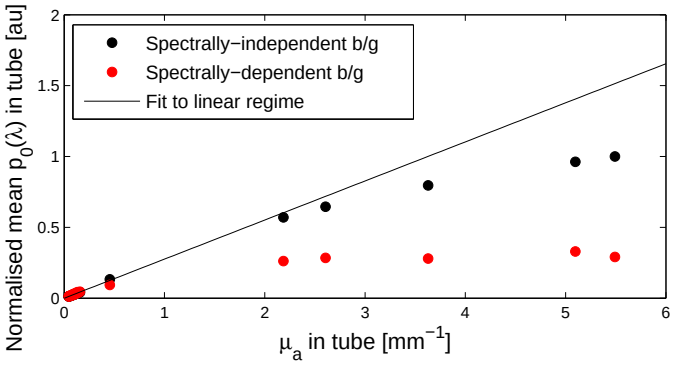

FIGURE 3.5: Plot of average initial acoustic pressure in 250μm tube as a function of absorption coefficient in tube. Fluence simulated using Monte Carlo with 10^{12} photons, collimated, illuminated 2.5mm in diameter at wavelengths in 500-1000nm range. Black line fit to linear region of acoustic pressure vs μ_a.

As the background contained blood, this would in part have been responsible for spectral colouring of the fluence. For this reason, the poor level of accuracy in estimates made in the when the background was spectrally-independent may be somewhat surprising; however, the back-scattered light from within the tube was coloured as it travelled through the tube, which was spectrally-dependent. This effect was also present in the phantom where the background was spectrally-dependent, but the effect of spectral colouring is much more severe here, resulting in significantly more inaccurate oxygenation estimates in the tube.

It can be seen in Fig. 3.4 that there was a tendency to underestimate oxygenation relative to the true value of 90%. This effect arises due to weighting of the inversion

by the fluence at wavelengths where it undergoes large changes, e.g. between 500nm and 600nm. This is illustrated in Fig. 3.6 which shows the spectrum of the absorption coefficient and the fluence averaged over the volume of the tube. As one might expect, large variations in fluence are correlated with significant changes in absorption; when wavelengths in the 500-600nm range were omitted and the inversion was computed using the data at the remaining 20 wavelengths (>600nm), sO_2 was overestimated relative to 90%.

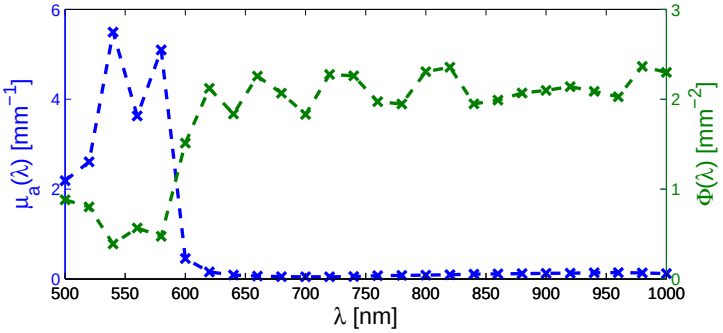

FIGURE 3.6: Plot of absorption spectrum in tube (blue) for 90% oxygenated blood and average fluence spectrum (green) in tube. Fluence simulated using Monte Carlo with 10^{12} photons, collimated, illuminated 2.5mm in diameter at wavelengths in 500-1000nm range and averaged over tube volume.

In order to determine whether oxygenation estimates converge to the true value when spectral colouring of the fluence is reduced through the application of the fluence correction, the initial acoustic pressure at each wavelength was averaged in the x- and y-dimensions and a straight line was fit to the log of the image as a function of z to compute $\mu^{\text{fit}}(\lambda)$. Each voxel in $p_0(\lambda)$ was then divided by the exponential model of Eq. (3.7) and the linear inversion was performed as described above. The resulting estimates are displayed in Figure 3.7.

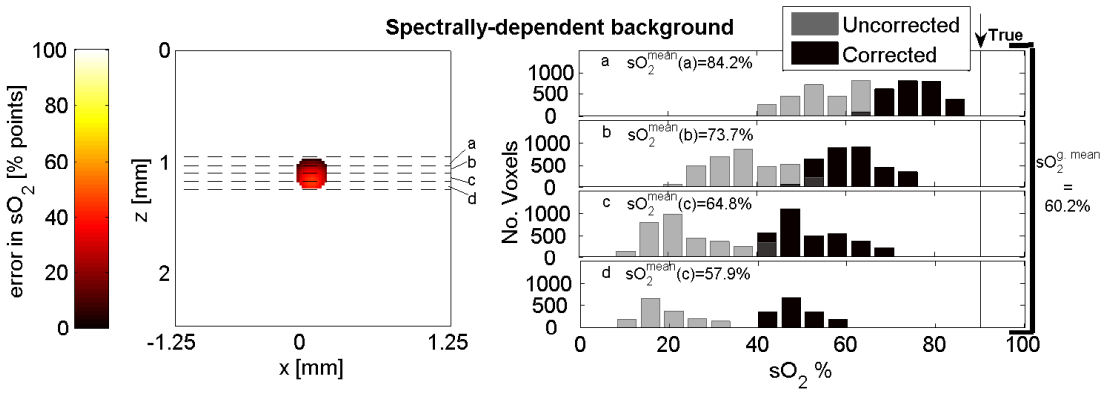

FIGURE 3.7: Oxygenation estimates made in 250μm tube using 26 wavelengths in 500-1000nm range, with initial pressure distributions corrected by 1D expoential fluence correction. x-z slice of error in oxygenation for fluence-corrected initial pressure distirbutions in spectrally-dependent background. Dashed lines used to indicate four x-z slices (a-d) from which histograms of sO_2 are formed (true value of 90% indicated by vertical black line). sO_2^{mean}(a), sO_2^{mean}(b), sO_2^{mean}(c) and sO_2^{mean}(d) represent the mean of the sO_2 estimates of the voxels contained within the tube in slices a,b,c and d respectively. The global mean sO_2 estimate averaged over all voxels enclosed within the entire tube is shown to the right of the histogram.

Figs. 3.7 and 3.4 demonstrate that the fluence correction improves oxygenation estimates made in the tube immersed in the spectrally-dependent background by between 20% and 30%. Despite these improvements, the error remains significant, at 7% at shallow regions in the tissue and up to 30% in deeper regions of the tube. It is traditionally believed that using wavelengths across a wide band of the NIR-visible yields more accurate oxygenation estimates [136], possibly applying the thinking that a better fit is obtained when fitting to more data. This may not be applicable to a linear inversion in QPAT and, in fact, the inclusion of data obtained at wavelengths where the absorption coefficient of blood is high ($\lambda < 620$nm) may reduce the accuracy of the inversion because spectral colouring is more marked at these wavelengths [137]. The validity of this assertion is examined in the next section.

3.1.3 Determining optimal wavelengths in the 500-1000nm range

It was shown in the previous section that a linear inversion using many wavelengths across the 500-1000nm range yielded inaccurate oxygenation estimates in a tube immersed in an absorbing-scattering background when the background was spectrally-independent as well as when it was spectrally-dependent. This was found to be due to spectral colouring of the fluence, whereby the fluence's spectrum is coloured as it passes through the medium which has wavelength-dependent absorption and scattering. Given a small change in the absorption coefficient, $\delta\mu_a$, the resulting change in fluence and therefore initial acoustic pressure can be written as,

$$\delta p_0 = \kappa \left[\delta\mu_a \Phi + \mu_a \delta\Phi + \delta\mu_a \delta\Phi\right], \tag{3.9}$$

where $\mu_a = \mu_a(\lambda_1)$, $\Phi = \Phi(\lambda_1)$ and κ is a factor that takes into account the Grüneisen coefficient and the detection sensitivity of the system (which cancels out in Eq. (3.4)). Dividing through by p_0 yields,

$$\frac{\delta p_0}{p_0} = \frac{\delta\mu_a}{\mu_a} + \frac{\delta\Phi}{\Phi} + \frac{\delta\mu_a}{\mu_a}\frac{\delta\Phi}{\Phi}. \tag{3.10}$$

Using the notation, $\hat{\delta x}$ to denote $\frac{\delta x}{x}$, we require $\left|\frac{\hat{\delta\Phi}}{\hat{\delta\mu_a}}\right| \ll 1$ and $\hat{\delta\mu_a}\hat{\delta\Phi}$ to be negligible in order for Equation (3.10) to show

$$\hat{\delta p_0} \approx \hat{\delta\mu_a}, \tag{3.11}$$

i.e. when the dependence of the fluence on the absorption coefficient is minimised, the change in the initial pressure approaches that of the absorption coefficient. This means for the relative change in acoustic pressure to be predominantly due to a change in the

absorption coefficient, the relative change in the absorption coefficient must far exceed the relative change in the fluence, which may only be true for particular wavelength pairs.

A linear inversion was performed for all pairs of the 26 wavelengths in the 500-1000nm range and considering the sO_2 estimates for which $\left|\frac{\delta\hat{\Phi}}{\delta\hat{\mu}_a}\right|$ is minimised. sO_2 was calculated for each pair of the 26 wavelengths, giving ${}^{26}C_2 = 325$ sets of oxygenation estimates; with the estimates for each pair of wavelengths averaged over the tube volume, the most accurate were those obtained at wavelengths between 620nm and 1000nm, for both the spectrally-independent and spectrally-dependent backgrounds; for the spectrally-dependent case, the 1D fluence correction, discussed in Section 3.0.4, was applied to the initial pressure distribution. Figure 3.8 shows histograms of sO_2 estimates in each voxel in the tube, obtained using a pair of wavelengths for the two background types and for estimates made with the 1D fluence correction, when $\left|\delta\hat{\Phi}\right| \leq 0.01 \left|\delta\hat{\mu}_a\right|$ and when $\left|\delta\hat{\Phi}\right| > 0.01 \left|\delta\hat{\mu}_a\right|$. Note that the wavelengths chosen (620/920nm and 700/820nm, respectively) yielded the most accurate sO_2 estimates on average over the tube in each case.

It can be seen in Fig. 3.8(a) that sO_2 is accurately estimated (to within 0.5%) when the background is spectrally-independent, and that estimates made under the condition that the relative change in the absorption coefficient far exceeds that of the fluence (in red) are more accurate. The red band of the most accurate oxygenation estimates in Fig. 3.8(a) is composed of the most superficial voxels in the tube, at the centre of the illumination in the x- and y- directions. For the spectrally-dependent background in Fig. 3.8(b), it can be seen that oxygenation estimates are less accurate with 1-2% errors relative to the true value of 90% and the variance of estimates is much higher. Applying the threshold $\frac{\delta\hat{\Phi}}{\delta\hat{\mu}_a} < 0.01$, no voxels satisfy this relationship due to spectral colouring; however, the application of the fluence correction shifts the oxygenation estimates toward the true value in Fig. 3.8(c), with the red bars in the histogram representing the values which satisfy the above inequality.

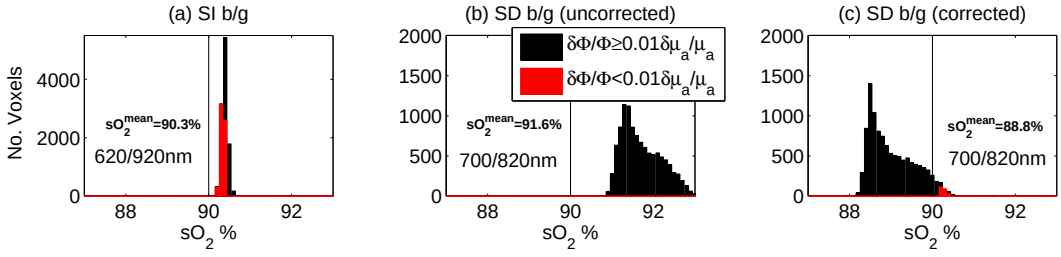

FIGURE 3.8: Histograms of sO_2 estimates in each voxel within the tube for (a) spectrally-independent (SI) with mean sO_2 of 90.3%, (b) spectrally-dependent (SD) background with mean sO_2 of 91.6%, and (c) spectrally-dependent with 1D fluence correction (SD background corrected, see Section 3.0.4) with mean sO_2 of 88.8% obtained using linear inversion at 620/920nm, 700/820nm, 700/820nm, respectively; sO_2 estimates for which $\frac{\delta\hat{\Phi}}{\delta\hat{\mu}_a} < 0.01$ shown in red and $\frac{\delta\hat{\Phi}}{\delta\hat{\mu}_a} \geq 0.01$ in black. Notes: vertical line indicates true oxygenation of 90%; there are no sO_2 estimates for which $\frac{\delta\hat{\Phi}}{\delta\hat{\mu}_a} < 0.01$ for SD background (uncorrected) due to severity of spectral colouring of fluence; note the scale change on vertical axis between SI background to SD background. These wavelength pairs yielded the most accurate sO_2 estimates of all pairs examined. sO_2^{mean} represents the mean of the sO_2 estimates all voxels enclosed within the entire tube and is termed the global mean sO_2.

The above result is in line with the prediction that the accuracy of oxygenation estimates can be improved by ensuring the change in fluence is small compared to the change in absorption. In practice this is not a useful guide as $\frac{\delta\hat{\Phi}}{\delta\hat{\mu}_a}$ is unknown (and therefore cannot be used directly as an indicator of the accuracy of the sO_2 estimate). A more useful result, and one that could be more applicable in general, would be to find which – if any – wavelengths pairs are best-suited to recovering sO_2. Fig. 3.9 shows the error in oxygenation, averaged over the tube, obtained at wavelengths λ_1 and λ_2 when the background was spectrally-dependent. Four different values of oxygenation in the tube, sO_2^{tube}, were considered: 70%, 80%, 90% and 100%. In each plot, the bottom-left triangle in the matrix shows values obtained without the 1D fluence correction, while the top-right triangle contains values obtained from data using this correction.

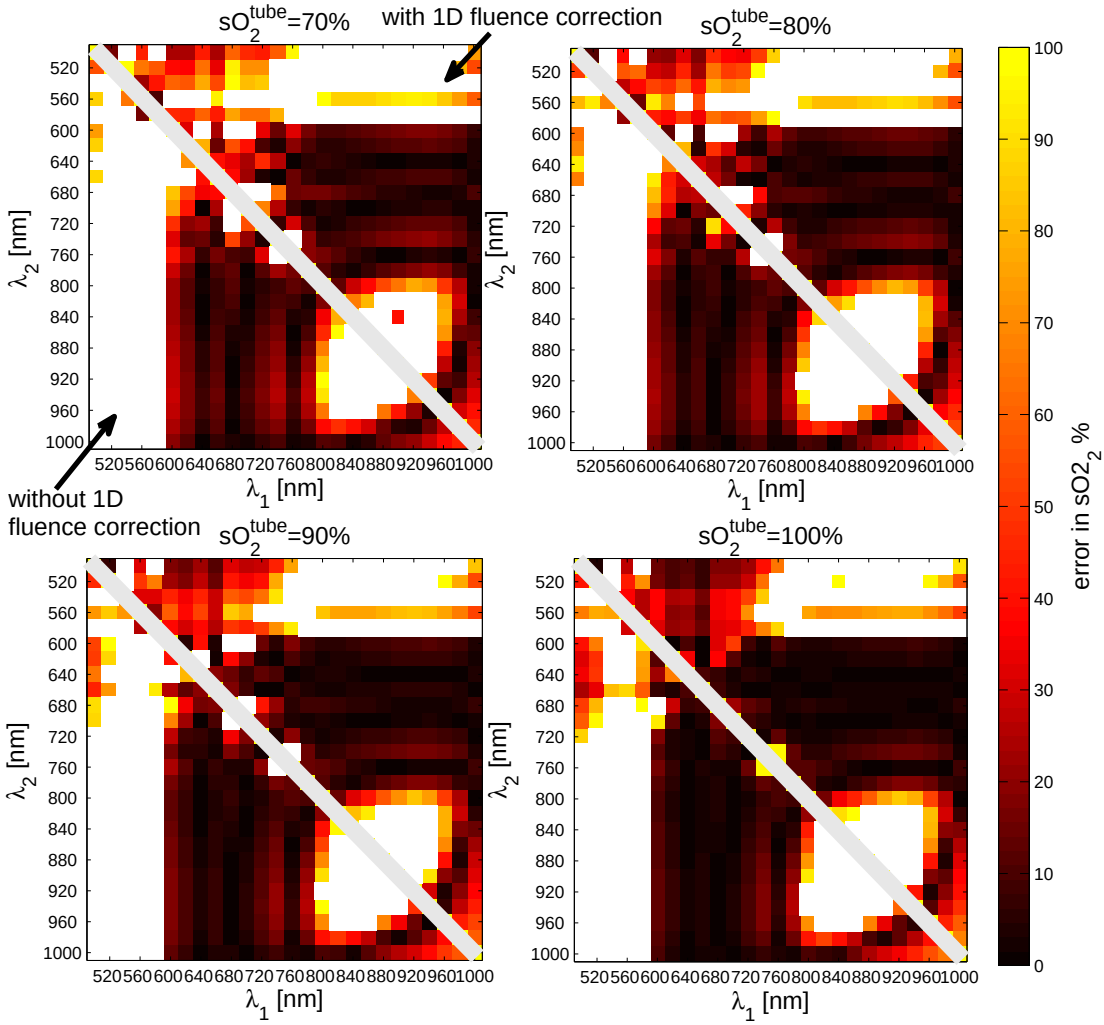

FIGURE 3.9: Matrix of average error in oxygenation for 26 wavelength pair combinations in 500-1000nm range for different values of oxygenation in the tube, sO_2^{tube}. In each figure, the bottom-left triangle represents sO_2 estimates made with a spectrally-dependent background and the top-right triangle represent sO_2 estimates made with a spectrally-dependent background using a 1D fluence correction. The white pixels represents estimates with an error $>$100%.

A number of observations can be made regarding Figure 3.9. Comparing the bottom left- and top right-hand triangles, it can be observed that if wavelengths between 600nm and 1000nm are chosen, the fluence correction has little impact on the accuracy of results. On the other hand, the effect of the correction is significant when one wavelength lies in the 500-580nm range where spectral colouring is more severe. The practical implication of this is that the spectral colouring may yield more accurate sO_2 estimates in applications where high SNR is required.

With or without the fluence correction, the most accurate oxygenation estimates are obtained in the 620-1000nm range, regardless of the sO_2 value in the tube. However, few wavelength pairs in this range yielded a reasonable degree of accuracy ($\leq$10% error) and those that did in the case where $sO_2^{\text{tube}} = 100\%$, produced lower accuracy estimates when $sO_2^{\text{tube}} = 70\%$. This is due to blood absorption, and its wavelength dependence, at wavelengths below the isobestic point at 800mn increasing with decreasing sO_2. As such, the number of wavelength pairs yielding reasonable accuracy also decreases with decreasing sO_2^{tube}.

Based on these findings, it is of interest whether the use of more than two wavelengths in the 620-920nm range affords more accurate oxygenation estimates, as it is in this range that the absorption coefficient undergoes minor changes, meaning the corresponding change in fluence is also likely to be small (see Fig. 3.6). Oxygenation was computed in each voxel for every possible combination of 2 to 26 wavelengths and then averaged over the tube volume, giving on the order of 67 million ($\sum_{r=2}^{26} {}^{26}C_r$=6.71$\times$10^7) unique estimates. The majority of the most accurate estimates were obtained using combinations of wavelengths in the 620-1000nm range. Nevertheless, it is not possible to determine optimal wavelength sets without prior knowledge of the internal fluence distributions and the true oxygenation. Thus, wavelengths between 620nm and 1000nm are *most likely* to provide accurate oxygenation values. This rule-of-thumb depends on a number of factors such as tissue geometry and composition and will vary from one imaging target to another, as well as depending on the oxygenation state of the target (discussed in more detail in Section 3.1.4). Moreover, it was only possible to obtain this rule-of-thumb using explicit knowledge of internal fluence distributions.

3.1.4 Dependence of accuracy of sO_2 estimates on background oxygenation

Given the sensitivity of sO_2 estimates made in the tube to the tube oxygenation, it is of interest to determine how strong the dependence of the accuracy of estimates on the background oxygenation, sO_2^{bg}, is. A separate set of experiments was conducted to examine how physiologically realistic changes in background oxygenation, ranging from

60% to 90%, affect the accuracy of oxygenation estimates with $sO_2^{\text{tube}} = 90\%$. Due to the extended computation times of running the MC model for 26 illumination wavelengths and 4 values of sO_2^{bg}, Toast++ [138], an open-source FE model of the DA was used instead. The domain was again 5mm$\times$5mm$\times$5mm in size but was composed of mesh elements with edge lengths of approximately 0.056mm, rather than 0.025mm, due to the significant memory demands of dense meshes in FE simulations. The simulations used top-hat square illumination on the top surface of the domain (rather than top-hat circular illumination in the MC simulations) and the boundaries were index-matched as above. The fluence correction was applied in a similar manner as before (i.e. by fitting to the log of the image averaged over the x- and y-directions to obtain a decay coefficient, μ^{fit}), with the data being interpolated to a regular 3D grid with 100μm voxel edge length.

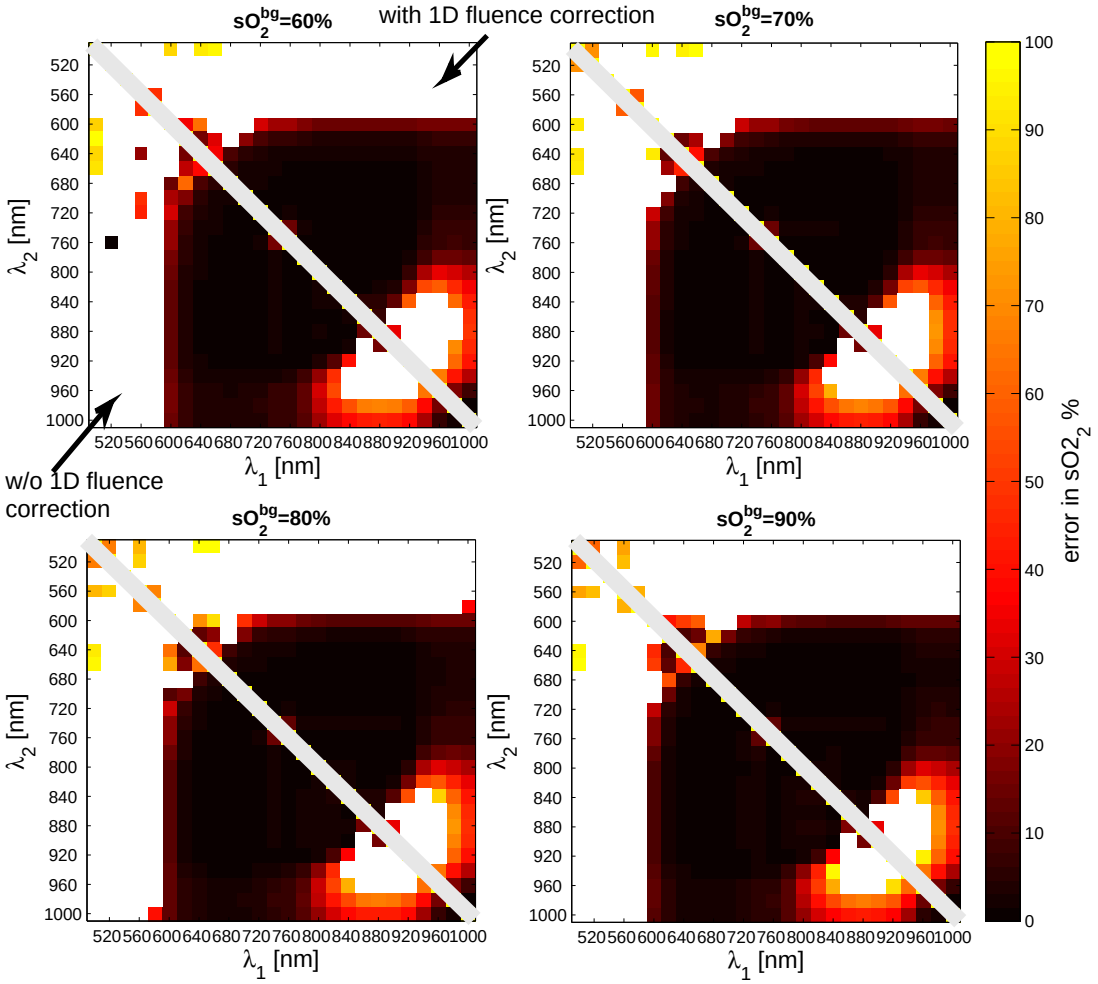

FIGURE 3.10: Matrix of average error in oxygenation for 26 wavelength pair combinations in 500-1000nm range for different values of oxygenation in the background, sO_2^{bg}. In each figure, the bottom-left triangle represents sO_2 estimates made with a spectrally-dependent background and the top-right triangle represent sO_2 estimates made with a spectrally-dependent background using a 1D fluence correction. The white pixels represent estimates with an error $>$100%.

Comparing Fig. 3.10 with Fig. 3.9, there are only minor discrepancies in the distribution of the error in sO_2 estimates. The most significant differences occur at the shorter wavelengths (500-600nm), which is most likely due to the fact that the albedo throughout the domain is much smaller at these wavelengths, thus resulting in modelling error of the DA relative to MC. In the dark regions between 600nm and 1000nm, the error estimates in Fig. 3.10 are within $\sim$3% of those in Fig. 3.9 obtained using the MC model. Looking at the variation in accuracy in this region with changing sO_2^{bg}, there is very little variation overall. The most notable change is a reduction in accuracy at wavelengths

straddling 800nm, the isosbestic point of blood, with increasing sO_2^{bg}. This is likely to be due to the difference in absorption coefficient at wavelengths surrounding the isosbestic point increasing with increasing sO_2^{bg} (see Fig. 3.1), resulting in a greater degree of spectral colouring. The reason for the relatively weak dependence of the the accuracy of sO_2 estimates made in the tube to the background oxygenation value is likely to be due to the absorption coefficient in the background being 12-15$\times$ smaller than in the tube, meaning a large proportion of spectral colouring in these simulations occurs in the tube, rather than outside it. This is consistent with observations made when comparing the accuracy of inversions with and without spectral dependence of the background in Section 3.1.2.

3.1.5 Pessimal wavelengths

Thus far, the wavelength pairs that appear to be optimal in terms of accurately estimating sO_2 are those in the 600-1000nm range and the optimal wavelengths are largely independent of sO_2^{bg} and sO_2^{tube}, though more accurate estimates are generally obtained at these wavelengths when sO_2^{tube} is nearer 100%. Nevertheless, from Figs. 3.9 and 3.10 it is clear that there are certain wavelength pairs that are far from optimal; these are represented by the white regions in Figs. 3.9 and 3.10 below 600nm and between approximately 800-920nm. The former is, for the most part, the result of severe spectral colouring at the short wavelengths as already described in this chapter. The latter occurs in a region of the spectrum where spectral colouring is minimal due to low absorption and a relatively flat spectrum of blood; the poor accuracy in the estimates made in this small region is due to the lack of linear independence of the spectra of oxy- and deoxy-haemoglobin, which means that α, the matrix of molar absorption coefficients (see Eq. (3.3)), is ill-conditioned for those wavelength pairs. This is confirmed by considering a plot of the log_{10} of the condition number of α as a function of the two wavelength pairs as shown in Fig. 3.11.

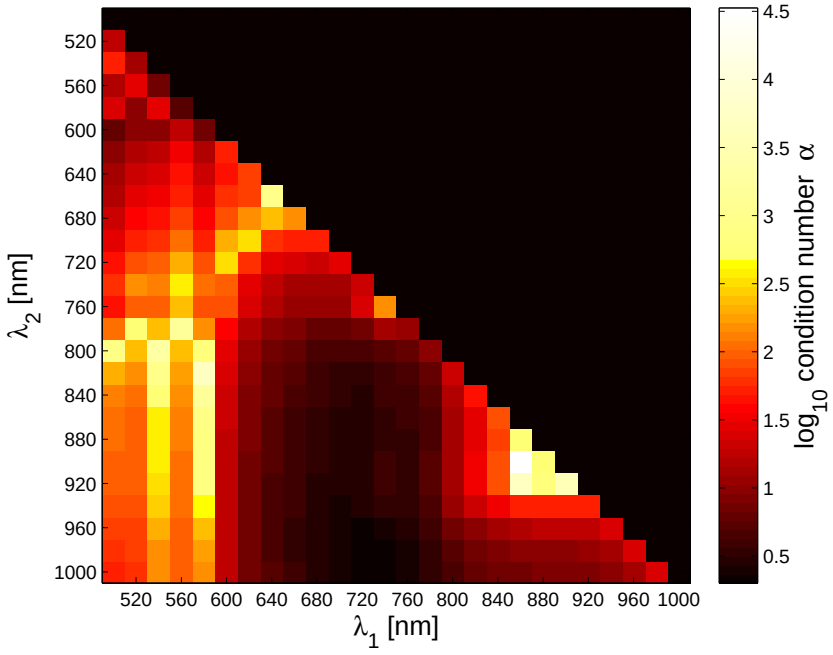

FIGURE 3.11: Matrix of $\log_{10}$ of condition number on α for wavelength pairs in 500-1000nm range.

Novel wavelength selection techniques that exploit independence of the spectra in the target have been proposed [95, 139]. As observed above, the condition number on α must be kept sufficiently low to prevent the inversion becoming ill-conditioned and blowing up any noise in the data, otherwise oxygenation estimates can have errors $>$100%. Nevertheless, the studies by Luke et al. ignore the impact of spectral colouring which can result in equivalently large errors in sO_2 estimates, evidenced by the white regions between 500nm and 680nm in Fig. 3.9 which correspond to relatively low condition numbers of 6-50, as shown in Fig. 3.11. The above analysis has demonstrated that optimal wavelengths for imaging a blood-filled tube are those between 600-1000nm, which avoid regions of the spectrum where spectral colouring is greatest; however, regions of the spectrum where α_{HbO2} and α_{Hb} change little with wavelength must also be avoided. These conclusions were reached with the PA image data being completely noise-free. With this in mind, the presence of noise may have significant implications on the accuracy of sO_2 estimates made using the guidelines above, given that the absorption

coefficient of blood is much lower in the 600-1000nm range compared with those in the 500-600nm range. The effects of noise are explored below.

3.1.6 Trade-off between Signal-to-Noise and accurate sO_2 estimation

It was found earlier in this chapter that oxygenation estimates made using a linear inversion will be most accurate when spectral colouring is minimised, which is equivalent to stating that the relative change in the absorption coefficient must far exceed the corresponding change in fluence, or that the relative change in PA amplitude must be predominantly due to a relative change in absorption ($\hat{\delta p} \approx \delta\hat{\mu}_{\mathrm{a}}$). This relationship poses the significant issue that if the change in PA amplitude falls below the noise floor, the inversion will become inaccurate or ill-posed. This is likely to be the case when using wavelengths above 620nm, where blood absorption is two orders of magnitude smaller than in the 500-600nm range, meaning we are faced with a choice: using shorter wavelengths affords sufficient SNR at the expense of a greater degree of spectral colouring of the fluence, or SNR is sacrificed in favour of accuracy in oxygenation estimates. This trade-off is made evident by considering that we require that the change in the initial pressure, as a result of a change in oxygenation state or illumination wavelength, exceeds the noise equivalent pressure of the detection system, η: $|\delta p_0| > \eta$. Substituting Equation (3.11),

$$|\delta\mu_a| > \sqrt{2}\frac{\eta}{\Gamma\Phi}. \tag{3.12}$$

η, the noise equivalent pressure of the detection system can be measured, while the Grüneisen parameter is in this case assumed to be unity and Φ is in general unknown. Note that the $\sqrt{2}$ factor on the right-hand side of Equation (3.12) is obtained by assuming that $p_0(\lambda_1)$ and $p_0(\lambda_2)$ are uncorrelated and have equal SNR such that their variance is additive [140].

The 325 wavelength pairs in the 500-1000nm range were used to form sO_2 estimates averaged over the tube volume were used to test the inequality in Eq. (3.12). Typical

values were used for the parameters in Eq. (3.12): $\eta = 0.2$kPa, $\Gamma = 0.15$ [141]. In order to obtain fluence distributions in realistic units, the MC-generated fluence distribution was normalised by the maximum value atthe centre of the most superficial surface and then scaled by 20mJcm^{-2}, the maximum permissible exposure for tissue. The resulting SNR from these parameters were in the range of 26.5dB-45.7dB, averaged over the entire image volume, depending on wavelength and the value of the tube sO_2; these SNRs are significantly higher than those of high quality *in vivo* photoacoustic images of superficial blood vessels [142, 143] and thus provide favourable conditions for the inversion. Fig. 3.9 was re-plotted 'rejecting' sO_2 estimates made at wavelengths that did not satisfy the inequality in Eq. (3.12); the estimates that were rejected are coloured in light blue as shown in Fig. 3.12.

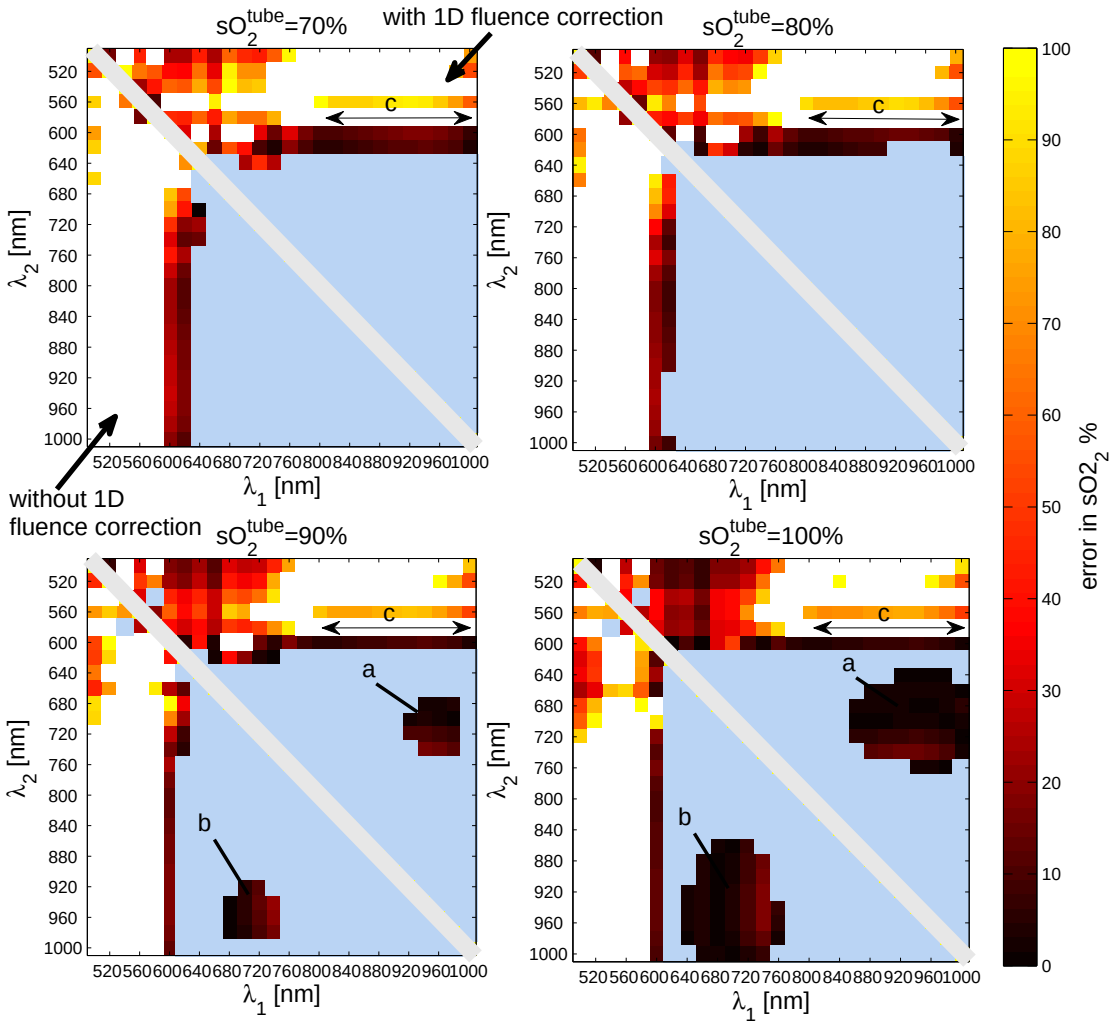

FIGURE 3.12: Matrix of average error in oxygenation for 26 wavelength pair combinations in 500-1000nm range for tube oxygenation values of 70%, 90% and 100%. Bottom-left triangles: sO_2 estimates made with spectrally-dependent background. Top-right triangles: sO_2 estimates made with spectrally-dependent background using 1D fluence correction. The white pixels represent estimates with an error >100%. Light blue regions indicate wavelengths that do not satisfy the condition in Equation (3.12).

It is immediately evident that with realistic levels of noise in the images, the number of wavelength pairs that can be used to obtain a *reasonable* sO_2 estimates is much reduced. This is due to the fact that many wavelength pairs do not yield sufficiently large $|\delta\mu_a|$ to satisfy Eq. (3.12). In addition, it is noticeable that the 1D fluence correction often provides a clear improvement in accuracy, particularly for low values of sO_2^{tube} and for wavelengths in the 500-620nm range where SNR is typically high. Thus, while spectral

colouring is aggravated when using these shorter wavelengths, it can be moderated by the application of the fluence correction.

Many regions in Fig. 3.12 appear to yield 'reasonably' accurate estimates. The most notable regions are marked 'a', 'b' and 'c' for reference. For $sO_2^{\text{tube}} = 100\%$ and $sO_2^{\text{tube}} = 90\%$, the regions labelled 'a' and 'b' provide accurate estimates with errors predominantly below approximately 10%. However, as sO_2^{tube} decreases, the number of wavelengths yielding sufficient SNR decay away because, as oxygenation decreases, the absorption coefficient of blood also decreases. Thus, these spectral bands may be useful when estimating sO_2 in vessels where it is known *a priori* that oxygenation is high (>90%); however, as demonstrated in the figures for $sO_2^{\text{tube}} = 70\%$ and $sO_2^{\text{tube}} = 80\%$, using wavelengths in these bands may result in highly erroneous estimates when oxygenation is low.

Given that the spectral range labelled 'c' is present, at least for λ_2=600nm, in for the tube sO_2 values between 70-100%, this may be a more useful range for implementation in practice. With one wavelength set to 600nm and the other between 800nm and 1000nm, the chosen wavelength pair will not only satisfy Eq. (3.12) but will also yield reasonable accuracy for all the sO_2^{tube} values represented in Fig. 3.12. This spectral band emerges because blood absorption at 600nm is high so μ_a is large enough to satisfy Eq. (3.12) but not so high as to introduce excessive spectral colouring. Despite this band appearing to be a panacea, it should be noted that estimates made made with $sO_2^{\text{tube}} = 100\%$ have very low errors between 2.8% and 8.10%, depending on the choice of λ_1, but errors grow to between 7.7% and 18% when $sO_2^{\text{tube}} = 70\%$. The wavelength pair that yielded the best accuracy across all four sO_2 values in Fig. 3.12 was found to be 600/800nm with an error ranging from 9.8% for $sO_2^{\text{tube}} = 70\%$ to 3.6% for $sO_2^{\text{tube}} = 100\%$.

It may seem as though estimates with <10% error can be obtained using a wavelength pair of 600/800nm. However, the simulations in this study represent the 'best case scenario', in which it is assumed that the images are perfectly reconstructed and the only heterogeneity in the domain is the tube; in practice neither apply so lower accuracy than that predicted by these simulations can be expected. In addition, the typical SNR achieved *in vivo* is likely to be significantly lower than used in these simulations. Not

only will this have the effect of reducing the number of wavelengths that satisfy Eq. (3.12), but for the wavelengths that do satisfy it, noise in $p_0(\lambda)$ will corrupt the accuracy of estimates. This is compounded by the fact that the accuracy of the estimate made at this seemingly 'optimal' wavelengths also depends on the value of sO_2^{tube}. Thus, in practice, not only will accuracy be compromised by lower SNRs and image reconstruction artefacts but there is no way of knowing *a priori* which wavelengths will provide even these compromised estimates since they depend strongly on the sO_2^{tube} which is unknown. This presents a formidable twofold challenge.

3.2 Estimating sO_2 in a Realistic Tissue Phantom

The tube phantom in Section 3.1.1 is not representative of many PAT imaging geometries, as the imaging target often consists of a highly vascularised volume of tissue rather than just a single vessel. Attempts to obtain absolute or relative measures of oxygenation in the brain using PAI are common in the literature [20, 92, 130, 144–152]. This reason, combined with the fact that anatomically the brain is representative of many highly vascularised tissues, is why a mouse brain phantom is used for this portion of the study, which aims to address the following questions:

- Do the simple guidelines for wavelength selection obtained from analysing data from the tube experiments help us obtain more accurate sO_2 estimates in this tissue-realistic phantom?
- Does a simple exponential correction to the fluence improve oxygenation estimates when the tissue is both multilayered and contains a complex network of blood vessels?

3.2.1 Digital phantom design

The phantom was generated from a μCT image of a mouse brain vascular cast, which has high contrast of the vasculature with 2.5μm isotropic resolution. The image volume was

down-sampled to 25μm isotropic voxel size and segmented using two methods. First, thresholding was used to remove noise; the second step used connected-component analysis to identify the connected portions of the vasculature. The vessel of interest (VOI) was selected by choosing the largest component produced from connected-component analysis and the smaller components were determined to form the background vasculature.

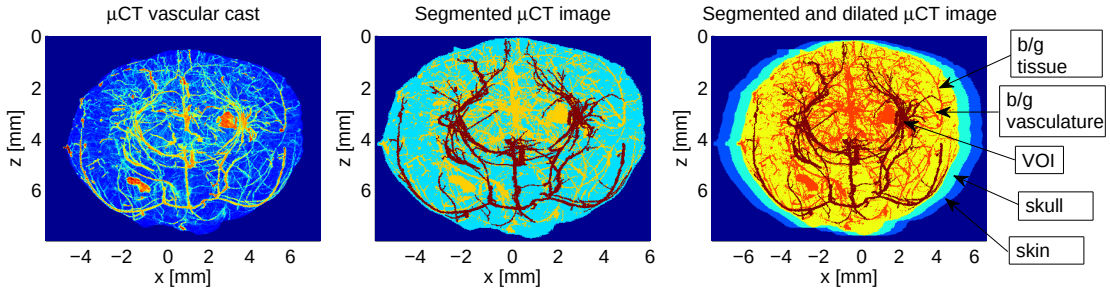

FIGURE 3.13: MIPs (coronal view) of μCT vascular cast (left), segmented μCT vascular cast (centre) and segmented and dilated (via Minkowski addition) μCT vascular cast (right) with skin and skull layers. Background tissue and vasculature, skin and skull, and vessel of interest (VOI) denoted.

The total haemoglobin concentration in vasculature in the brain was 150gl^{-1} and background tissue was assumed to consist of a 55% water volume fraction and 45% haemoglobin volume fraction with total haemoglobin concentration $c_{HbT}^{bg} = 5.63$gl^{-1} and $sO_2 = 60.7\%$ oxygenation – values taken from rat brain cortical tissue [153]. The lack of skin and skull layers in the original dataset meant that these had to be introduced via post-processing. The enveloping skin and skull layers were generated by dilating the segmented image volume (via Minkowski addition using an isotropic cubic 500μm kernel) by 500μm in all directions, twice sequentially, producing two layers, variable in thickness, that were approximately 500μm thick on average (realistic values from literature [154, 155] for mice). Of course, in practice, the skull does not adhere to the cortical surface and is separated by cerebrospinal fluid (CSF). The impact of CSF on the fluence distribution can be significant, but effect on the fluence's spectrum in the brain is likely to be quite small because of the low-absorbing, low-scattering nature of the CSF. The skin layer had $c_{HbT}^{skin} = 0.48$gl^{-1} at 98.5% oxygenation, 0.87% melanosome volume fraction with the absorption coefficient following the trend $\mu_{a,\text{mel}} = 0.452\left(\frac{\lambda}{500nm}\right)^{-3}$, and 21.4% water volume fraction [9]. The skull had wavelength-dependent absorption and scattering

coefficients between 0.025mm^{-1} and 0.036mm^{-1}, for μ_a, and 25mm^{-1} and 45mm^{-1}, for μ_s, where the scattering coefficient was obtained by fitting to scattering data of cortical bone [156].

The scattering coefficient of the background tissue and skin regions varied as a function of wavelength according to $\mu_s' = A\left(\frac{\lambda}{500nm}\right)^{-B}$ [9]. The brain and blood regions of the phantom had $A = 2.14$mm^{-1} and $B = 1.20$, the skin region had $A_{\text{skin}} = 2.97$mm^{-1} and $B_{\text{skin}} = -0.705$. The anisotropy factor was heterogeneously distributed with $g = 0.9$ for the tissue background and skin and $g = 0.9945$ for blood vessels.

Simulation of the fluence inside the phantom was achieved using MCX [132] with the phantom illuminated using a top-hat profile at $z = 0$mm in the +z direction at wavelengths in the 500-1000nm range at 20nm intervals. The illumination was collimated in the z-direction, 11.7mm in diameter, and consisted of 10^{14} photons (deemed sufficient by performing a convergence test in Section 3.1).

3.2.2 Estimates using 26 wavelengths in the 500-1000nm range

As with the tube phantom, a common approach for estimating blood oxygenation is to perform a linear inversion using many wavelengths across a wide band of the NIR-visible spectrum [136] as it is believed that the inclusion of more data leads to a more accurate fit. All 26 images at wavelengths in the 500-1000nm range were used in the linear inversion described described by Eq. (3.3). Fig. 3.14 shows the maximum intensity projection (MIP) in the y-direction of the oxygenation estimate made in the VOI, overlayed on the anatomy of the brain in greyscale. The oxygenation estimates in the VOI are of very limited accuracy and only a few voxels near the surface of the brain, around 1mm in depth, contain accurate oxygenation values to within 20% of the true value of 90% (absolute error calculated as $|sO_2^{\text{true}} - sO_2|$). The impact of the 1D fluence correction is very small and only marginally improves the accuracy of sO_2 estimates made using 26 wavelengths. The oxygenation estimates obtained with this phantom are of equivalently poor accuracy compared with those obtained in the tube phantom.

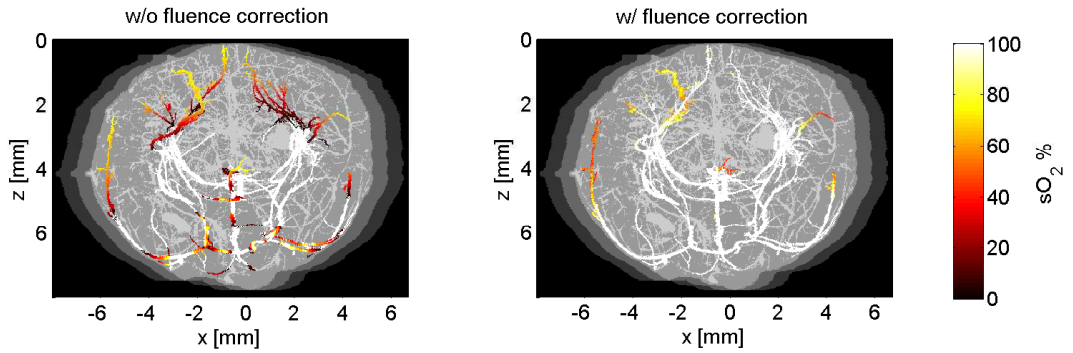

FIGURE 3.14: MIPs of oxygenation estimates in VOI made using 26-wavelength linear inversion with wavelengths in range 500-1000nm without fluence correction (left) and with fluence correction (right), both overlayed on greyscale phantom. True oxygenation was 90%. White indicates sO_2 estimates outside valid range of 0-100%.

3.2.3 Estimates using two or more wavelengths

Since similar observations can be made when using wavelengths from a broad spectral range for the brain phantom compared with the tube phantom, it is worthwhile testing multiple pairs of wavelengths between 500nm and 1000nm. The decay coefficient, μ_a^{fit}, used to apply the fluence correction to the images was obtained in a slightly different manner compared with the tube experiments in Section 3.1.1; due to the presence of skin and skull layers, the fit was performed to the data between the first maximum in the image averaged over x and y (at approximately 500-800μm below the surface) and the end of the most highly absorbing region. The data omitted is circled in grey in Fig. 3.15, which shows the decay in the PA image (averaged over the x- and y- directions) at 840nm with the straight line (in blue) fit to the data between 0.675mm and 7.88mm along the z-axis. The oxygenation estimates made using each wavelength pair was then averaged over the entire vessel-of-interest (VOI) and plotted in a matrix for each wavelength pair in Figure 3.16; the bottom-left triangle represents estimates made using a linear inversion without fluence correction and the top-right triangle represents estimates made with the fluence correction.

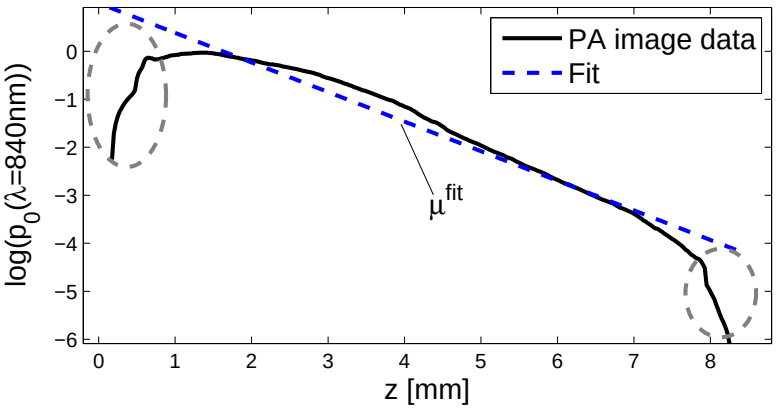

FIGURE 3.15: Plot of log initial acoustic pressure at 840nm averaged along x- and y-directions as a function of z. Straight line fit to data between 0.675mm and 7.88mm with slope μ^{fit}. Grey lines indicate skin and skull regions omitted from fitting of straight line to obtain μ^{fit}

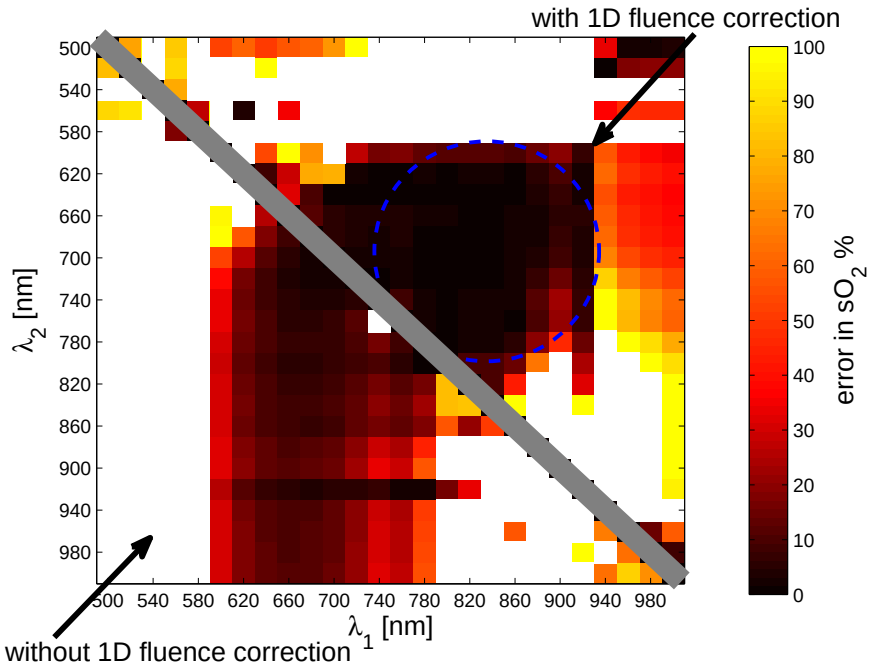

FIGURE 3.16: Matrix of average error in oxygenation in vessel of interest in tissue phantom for 26 wavelength combinations in 500-1000nm range. Bottom-left triangle: sO_2 estimates made without depth-dependent fluence correction. Top-right triangle: sO_2 estimates made with depth-dependent fluence correction. Blue dashed circle indicates wavelength pairs which yielded substantial improvement in accuracy of sO_2 estimates due to application of 1D fluence correction.

It can be seen that using uncorrected images, many wavelengths paired with those in the 600-1000nm range yield average oxygenation estimates over the VOI well within 30%

of the true value of 90%. This is a slightly narrower wavelength range to that observed with the tube phantom, but overall estimates are of lower accuracy in the mouse brain phantom. Given that the wavelengths that yielded the more accurate sO_2 estimates are different from those in the tube phantom, the question 'are the optimum wavelengths dependent on the target geometry?' can be asked. However, swapping the VOI and the background vasculature (see Figure 3.13) did not have a significant effect on the optimal wavelength pairs (i.e. 600-1000nm remained the optimal range). This suggests therefore that in highly vascularised media, the specific geometry may not determine which subset of wavelengths in the 600-1000nm range is most accurate; nevertheless, using wavelengths outside this range will lead to highly inaccurate oxygenation estimates everywhere except in the most superficial regions.

The 1D fluence correction reduces the range of wavelengths that yield the most accurate sO_2 estimates from 600-1000nm to 600-920nm while also improving the overall accuracy of sO_2 estimates; a much greater number of elements in the top-right triangle in the matrix show an error of less than 30% compared with the estimates obtained without the fluence correction. The most accurate average sO_2 estimate was obtained using 640nm and 840nm, which gave a minimum average error over the VOI of 2.9%, compared with 7.6% when the inversion is performed without correcting the PA images. Figure 3.17 shows oxygenation estimates made in the VOI using the optimal 640/840nm wavelength pair with and without a 1D fluence correction.

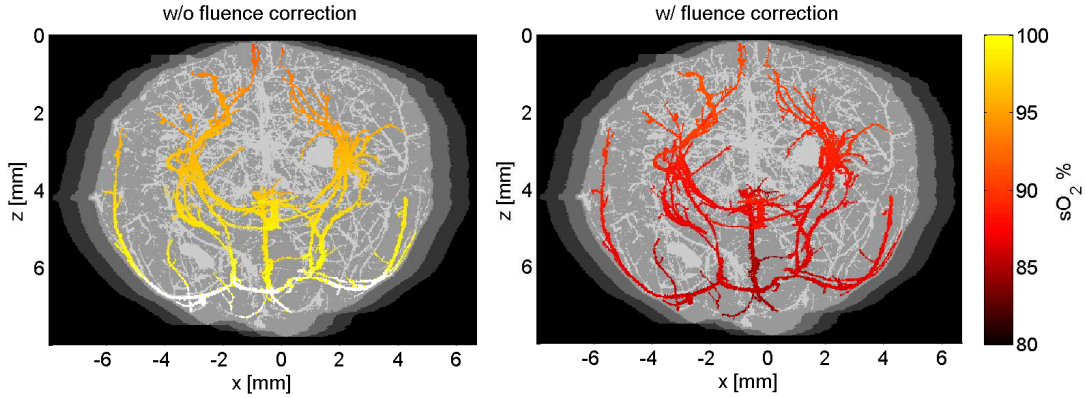

FIGURE 3.17: MIPs of oxygenation estimates in VOI made using two-wavelength linear inversion at 640/840nm without fluence correction (left) and with fluence correction (right), both overlayed on greyscale phantom. True oxygenation 90%.

Without the fluence correction, oxygenation estimates are predominantly above 95%, while when the image is divided through by $\exp(-\mu^{\text{fit}}z)\ \forall x, y$, the values are centered around the true value of 90% plus-or-minus a few % and show less variation with depth. In practice, however, without knowing the *true* oxygenation, it is impossible to choose the optimal wavelengths, and the image will be corrupted with some level of noise. The wavelengths, 640nm and 840nm, provide little signal at depth in practice, which may compromise the ability to accurately estimate sO_2. In Section 3.1.6, it was found that the change in PA amplitude from one wavelength to another must exceed the NEP of the system in order for the inversion to yield a plausible oxygenation estimate. Using this relationship it was derived that under a wavelength change $\lambda_1 \leftrightarrow \lambda_2$ the corresponding change in absorption must exceed $\sqrt{2}\frac{\eta}{\Gamma\Phi(\lambda_1)}$. Typical values were selected for these parameters: $\eta = 0.21$KPa (selected to be consistent with image SNRs obtained experimentally [142, 143]), $\Gamma = 0.15$ [141]; the fluence in this expression was computed using the MC-generated fluence used in the inversion, and was then normalised by its maximum value and scaled by 20mJcm^{-2}, the maximum permissible exposure for tissue. Fig. 3.17 was then reproduced, but only showing the sO_2 estimates for which this condition is satisfied, shown in Fig. 3.18.

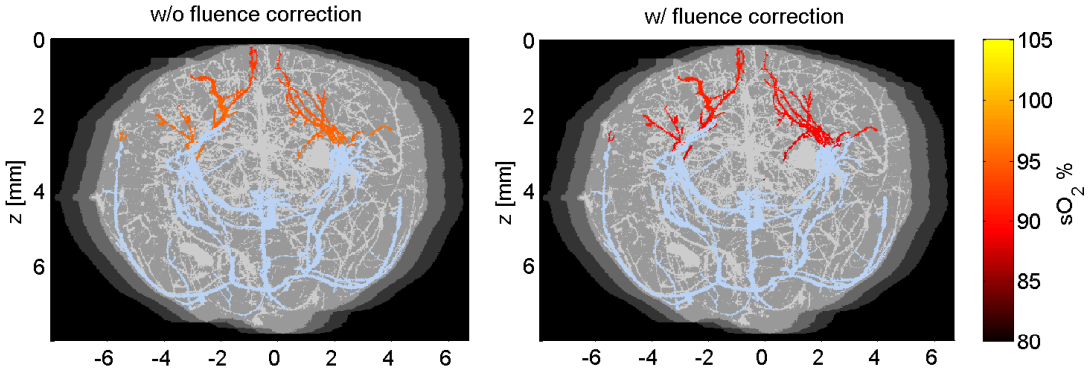

FIGURE 3.18: MIPs of oxygenation estimates in VOI made using two wavelength linear inversion at 640/840nm without fluence correction (left) and with fluence correction (right), both overlayed on greyscale phantom. Only values satisfying condition $|\delta\mu_a| > \sqrt{2}\frac{\eta}{\Gamma\Phi(\lambda_1)}$ shown. True oxygenation 90%.

The application of this condition demonstrates that oxygenation can only be estimated at less than 3mm at these wavelengths. At depths greater than 3mm, the SNR is insufficient and will yield spurious estimates of sO_2, demonstrating again that there is

an inherent trade-off between using wavelengths that are optimal for estimating sO_2 and wavelengths that yield sufficient SNR over the entire imaging volume. Moreover, accurate estimates to a depth of 3mm were only possible because optimal wavelengths, which required knowledge of the true oxygenation, were chosen; without access to the ground truth value, it would be impossible to know how accurate estimates are and the presence of noise in an experimental setting would lead to further inaccuracies in sO_2 values. This suggests that a linear inversion, with or without a simple fluence correction, may be of limited value in *in vivo* PA imaging beyond very superficial depths.

3.3 Summary

In this chapter, the accuracy with which oxygenation estimates obtained using simple linear spectroscopic inversions was investigated. This was carried out in two separate phantoms; the first consisted of a tube submerged in an absorbing-scattering background, while the second was a realistic phantom of vascularised tissue. It was also examined whether oxygenation estimates are more accurate when the PA image is divdided through by a simple 1D exponential decay, with the decay coefficient obtained by fitting a straight line to the log of the PA image averaged over the (transverse) x- and y- directions. It was found in both phantoms that using several wavelengths across a wide band of the NIR-visible spectrum yields inaccurate sO_2 values, except at the most superficial regions of the vessel of interest. Moreover, there was a tendency to underestimate oxygenation when wavelengths in the 500-600nm range were used, due to the fact that the absorption coefficient, and therefore the fluence, undergoes large changes at these wavelengths.

Based on these observations, *pairs* of wavelengths were tested which allowed trends of which regions of spectrum are better suited to estimating oxygenation using a linear inversion to be identified. It was found that spectral colouring of the fluence is minimised at wavelengths in the 620-1000nm range and that the simple 1D fluence correction could be used to improve estimates, provided optimal wavelengths were chosen. However, wavelengths in the range deemed optimal do not provide as high SNR as those in the shorter range ($\leq$600nm). Thus, the key conclusion was that wavelengths more

likely to yield accurate sO_2 estimates do not provide sufficient SNR to obtain accurate estimates beyond a few mm in depth, as oxygenation estimates become corrupted with measurement noise beyond this depth. Even selecting wavelengths found to provide the most accurate oxygenation values (on average in the complex geoemtry of the mouse brain phantom) did not yield accurate sO_2 values at a depth greater than 3mm. Furthermore, the determination of the optimal wavelengths required prior knowledge of the ground truth sO_2 value. Thus, there are very few scenarios in practice in which a linear inversion (with or without 1D fluence correction) is likely to yield accurate oxygenation estimates.

The above trends may not extend to dynamic sO_2 estimation, discussed in Section 3.1.2, which relies on the background absorption coefficient not changing under oxygenation changes [131], and photoacoustic microscopy, which is an inherently superficial imaging technique. As mentioned in Chapter 1, PAM exists in acoustic-resolution (AR-PAM) and optical-resolution modalities (OR-PAM), each of which may find a means of minimising the effect of spectral colouring on sO_2 estimation. Guo et al. [157] argue that in AR-PAM the fluence within a diffusely illuminated blood vessel will be a radial function of absorption of the form $\exp(-\mu_{\mathrm{a}} r)$ with r the radius; this would be a viable method for estimation of the absorption coefficient in the vessel (from which sO_2 can be computed) if the fluence were not coloured by wavelength-dependent absorption in the interstitial diffusing medium, which in practice is not likely to be the case as has been demonstrated [158]. OR-PAM involves focusing of the light field, so if the optical focus is unaffected due to wavelength dependence of the external optics (e.g. laser, optical fibre, lenses) which in practice can be measured anyway, spectral colouring of the fluence in the tissue is minimal due to the limited penetration depth; whether the distance travelled by the light into the tissue is sufficiently small to have a negligible impact on the fluence at the focus requires further investigation.

While it is possible to measure the fluence incident on the tissue surface, and thereby correct for system wavelength-dependent factors, it is not possible to make internal measurements. Thus, for PAT, to obtain estimates that are accurate beyond a few mm, modelling of the internal fluence distribution is a necessity. One model of particular

interest, due to its parallelisability, is the Monte Carlo method, which is described in detail in Chapter 4.

3.4 Future work

In this chapter it was observed that the accuracy of certain wavelength pairs was better than others due to better image SNR, reduced spectral colouring and/or better conditioning of α. The degree to which these factors impact the accuracy of the oxygenation estimate are unknown *a priori* and will depend on the experimental imaging conditions (e.g. tissue geometry and composition, illumination conditions, detection scheme, amongst many other experimental features). This of course makes optimal wavelength selection impossible without prior modelling of the internal fluence distribution. The investigation proposed herein is to acquire images at as many wavelengths across the NIR-visible spectrum as possible (laser linewidth and acquisition time permitting) and perform linear inversions for every pair of wavelengths; the expectation is that while certain pairs of wavelengths will yield wholly inaccurate – and even physiologically implausible – estimates, a number of wavelength pairs will ensure that sufficiently low spectral colouring of the fluence has taken place thus resulting in fairly accurate oxygenation estimates. Initial investigations were performed using the tube phantom of Chapter 3 with 19 illumination wavelengths spaced by 20nm in the 520-920nm range. The light model used here was TOAST++ [138]. Linear inversions were applied to the initial acoustic pressure at the 153 wavelength pairs. This was carried out for tube oxygenationation values, sO_2^{tube}, of 70%, 80%, 90% and 100%. Histograms of the sO_2 estimates averaged over all the voxels in the tube are plotted for each value of sO_2^{tube}. Note that only estimates between 0% and 100% were included when the averaging was performed, i.e. voxels with estimates outside this range were rejected.

It can be seen from Fig. 3.19 that this sO_2 estimation technique yields a spread of sO_2 estimates that appear to be centred near the true oxygenation value, sO_2^{tube}, indicated by the black vertical line. The most likely value in the histograms occurs nearly at the true value, possibly suggesting this approach may yield accurate oxygenation estimates in

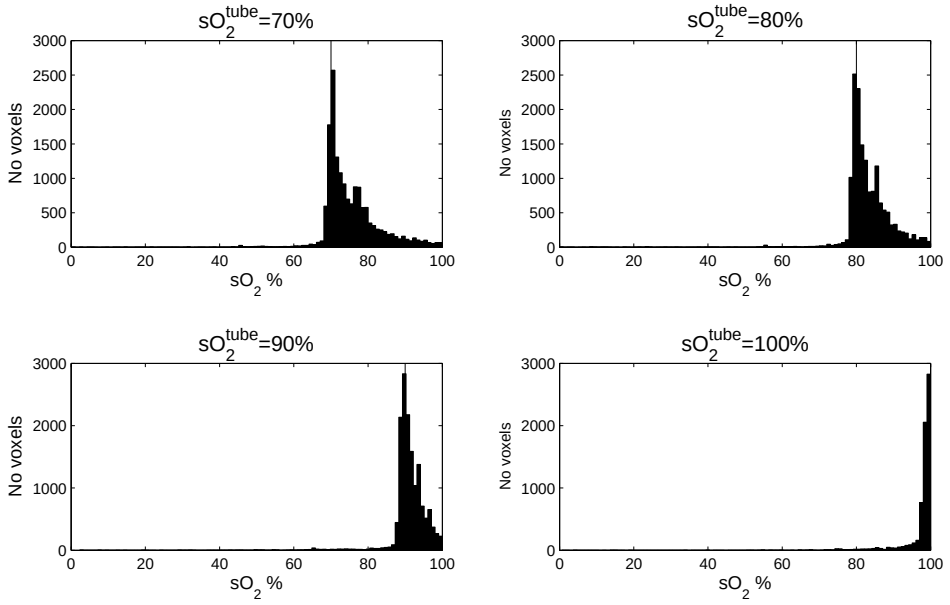

FIGURE 3.19: Histograms of sO_2 estimates averaged over all voxels in the tube obtained using every combination of 26 wavelengths in the 500-1000nm range for different values of sO_2^{tube}.

other imaging scenarios. This of course comes with the caveat that this technique may no longer yield reasonable estimates in the presence of noise; however, as the noise in $p_0(\lambda_1)$ and $p_0(\lambda_2)$ is uncorrelated, the effect of noise in the images is likely only to increase the spread of the distribution of sO_2 estimates, rather than to introduce a systematic error that shifts the position of the largest bin.

The method described above was applied to *in vivo* data. A dataset containing reconstructed images of a mouse flank at wavelengths between 680nm and 910nm with 10nm separation were used. Here, a blood vessel was manually segmented and the two-wavelength inversion procedure was applied to every voxel in the segmented vessel giving 351 estimates for each position in the vessel. However, due to very few wavelength pairs providing valid sO_2 estimates in the 0-100% range (sometimes only a pair of wavelengths provided a valid estimate), the mean of the values was used. Fig. 3.20(a) shows a maximum intensity projection of the blood vessel in the z-direction obtained at 900nm and Fig. 3.20(b) shows a maximum intensity projection in the z-direction of the mask used to segment the voxels in the vessel. One adjustment to the method described above had to be made due to the fact that the acquisition at each wavelength takes on the order of 2 minutes, meaning the animal moved due to breathing. Upon

each inversion, therefore, one image was registered to the other using Matlab's rigid registration tool. The voxel-wise distribution of sO_2 values obtained using this method is shown in the histogram in Fig. 3.20(c) with the estimates then plotted spatially inside the vessel in Fig. 3.20(d) as an MIP.

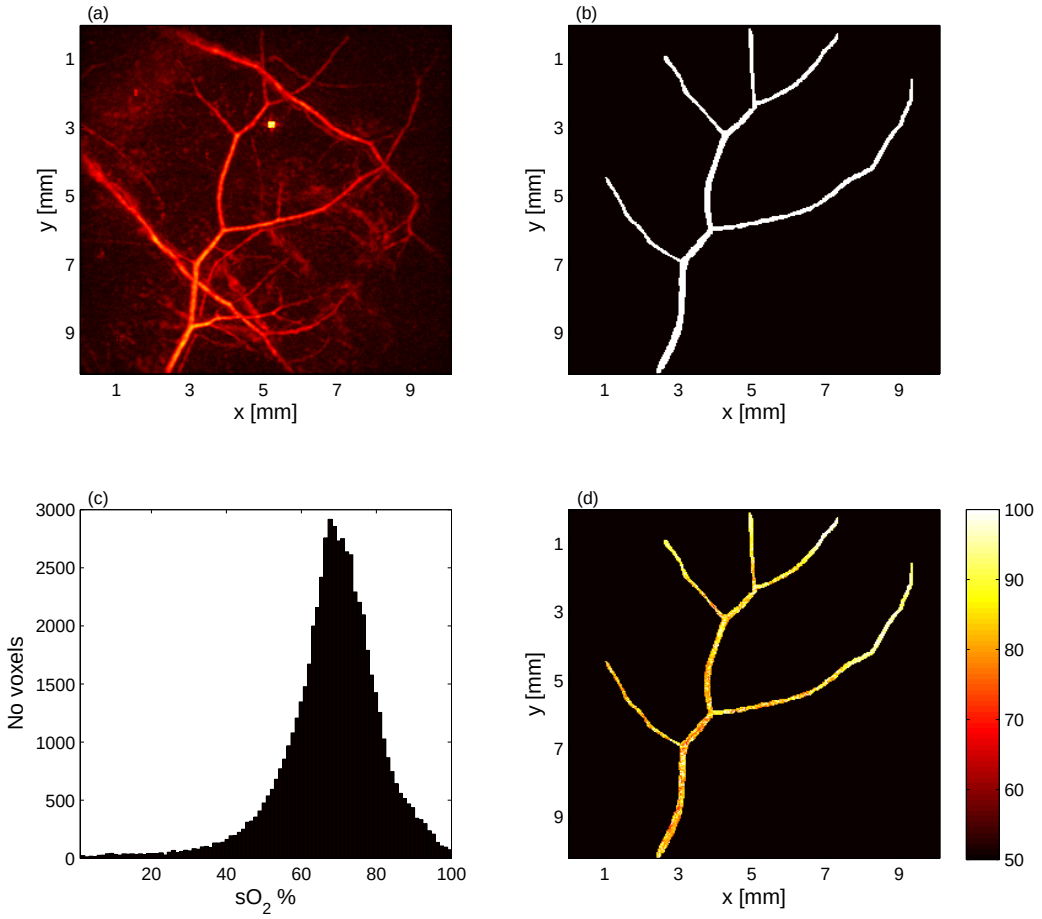

FIGURE 3.20: (a) MIP of PA image of mouse flank; (b) MIP of manually segmented mask of large vessel shown in (a); (c) Histograms over sO_2 estimates in each voxel in the mask in (b); (d) MIP of sO_2 estimates obtained using linear inversions using all possible combinations of illuminations plotted in the vessel.

The distribution of the values in each voxel in Fig. 3.20(c) demonstrates quite promising results given that the majority of the values lie in the physiologically expected range for a vein. Although it is unknown whether the segmented vessel is a vein, consideration of the images on either side of the isosbestic point demonstrated a decrease in the average

image intensity of the segmented vessel with increasing wavelength. This is supported by the fact that nearby vessels, thought to be arteries, increased with increasing wavelengths on either side of the isosbestic point. Although the mean value of 68% for the entire vessel appears to be low, the mouse was under isoflurane which is linked to decreased systemic oxygenation [159]. Nevertheless, the mouse was ventilated using 100% O_2, which is likely to have increased systemic sO_2 . Despite these confounding factors, the oxygenation estimates are within a reasonable range for a superficial vein.

By using multiple two-wavelength inversions to construct a distribution of sO_2 estimates, the preliminary experiments above demonstrate that this approach has the capability to estimate sO_2 *in vivo*. The vein in the example above was very shallow, at a depth of less than 750μm, but provides a compelling argument for further investigation into this approach.

Chapter 4

Monte Carlo modelling of light transport

The Monte Carlo method is a statistical technique that uses the law of large numbers to approximate the expectation value of a given quantity. The method has application in numerical integration, whereby the integral is approximated by summing the function of interest evaluated at randomly sampled points in the function's domain, as well as solving integral equations in physics and chemistry. The application of the Monte Carlo method discussed here is that of the transport of light energy through biological tissue. The aim is to simulate the propagation of packets of energy through the domain to form the quantity of interest – often termed a 'scored quantity' – which, in PAT, is typically the fluence or absorbed energy density.

The basic algorithm can be derived from either the integral [160] or derivative form of the RTE, but in this thesis only the derivation from the latter form is presented, treated in Section 4.1. The Monte Carlo method for light transport involves propagating an energy packet a distance s before depositing a fraction of its energy, or weight W, and being scattered by some angle, determined by sampling the scattering phase function; this process is repeated until the packet's weight has decayed below some threshold value. Note that in the rest of this thesis the term 'photon' is used to refer to an energy packet or ensemble of photons rather than a distinct photon. The distinction however is important

as the absorption of individual photons occurs, resulting in an overall reduction of the packet energy. An illustration of a photon packet's trajectory is shown in Fig. 4.1. In a Monte Carlo simulation, as the number of photons simulated tends to infinity, the quantity of interest tends to its true value; i.e. in the case of the fluence, it would tend to the value predicted by the RTE. A simulation using a finite number of photons will therefore produce a result that is not completely accurate because it will be a sample of a field with some variance. The degree variance in the estimate is a function of many parameters, including spatial discretisation in the domain and the number of photons. Despite this drawback, the variance in fluence estimates reduces sufficiently quickly with increasing number of photons so as to be useful as a forward model in PAT, as discussed in Chapter 3. This chapter presents a novel Monte Carlo model that simulates the radiance rather than the fluence and is hence termed radiance Monte Carlo, or RMC.

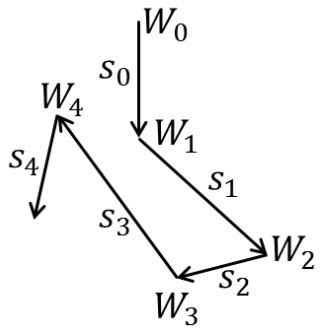

FIGURE 4.1: Illustration of photon packet trajectory, with starting weight is W_0. Photon packet travels a series of steps s_n, scattering between steps and undergoing a reduction in weight W_n upon each stop.

4.1 Derivation of MC method for light transport

The calculation of step length, the weight deposited and scattering angle can all be derived from the RTE by considering trajectories aligned with direction $\hat{\mathbf{s}}$ [57], as described in the following sections.

4.1.1 Absorption and scattering

In this section, the absorption of energy and the step length between scattering events are derived. Writing the time-independent RTE along the $\hat{\mathbf{s}}$ direction,

$$(\hat{\mathbf{s}} \cdot \nabla + \mu_{\mathrm{a}}(\mathbf{x}) + \mu_{\mathrm{s}}(\mathbf{x}))\phi(\mathbf{x}, \hat{\mathbf{s}}) = \mu_{\mathrm{s}}(\mathbf{x}) \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}, \hat{\mathbf{s}}')\phi(\mathbf{x}, \hat{\mathbf{s}}')d\hat{\mathbf{s}}' + q(\mathbf{x}, \hat{\mathbf{s}}), \tag{4.1}$$

allows us to write two separate transport equations for two sources of photons; the first (Eq. (4.2)) only models the transport of photons emitted by q anywhere along the trajectory $\hat{\mathbf{s}}$, while the second (Eq. (4.3)) only models the propagation of photons scattered inward from directions $\hat{\mathbf{s}}'$:

$$\left(\frac{\partial}{\partial s} + \mu_{\mathrm{a}}(s) + \mu_{\mathrm{s}}(s)\right) \phi_{sources}(s;\ \hat{\mathbf{s}}) = q(s;\ \hat{\mathbf{s}}) \tag{4.2}$$

$$\left(\frac{\partial}{\partial s} + \mu_{\mathrm{a}}(s) + \mu_{\mathrm{s}}(s)\right) \phi_{scattered}(s;\ \hat{\mathbf{s}}) = \mu_{\mathrm{s}}(s) \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}, \hat{\mathbf{s}}')\phi(\mathbf{x}, \hat{\mathbf{s}}')d\hat{\mathbf{s}}', \tag{4.3}$$

where $\phi_{sources}$ and $\phi_{scattered}$ are the radiances due to the sources and inward scattered light, respectively. The total radiance along the trajectory in direction $\hat{\mathbf{s}}$ and position s (using the convention that $\mathbf{s} = |\mathbf{s}|\,\hat{\mathbf{s}} = s\hat{\mathbf{s}}$) is equal to the sum of the radiance due to photon sources and inward-scattered photons: $\phi(s) = \phi_{sources}(s) + \phi_{scattered}(s)$.

The solution to Eq. (4.2), a first-order differential equation in s [160], is obtained by considering that contributions to ϕ in the $-\hat{\mathbf{s}}$ direction, e.g. due to back-scattering or sources emitting in this direction, are treated by writing Eqs. (4.2) and (4.3) for the $-\hat{\mathbf{s}}$ direction explicitly. In the case of discrete sources, q_{n_s}, separated by a distance Δs, then the light at position $S = N_s \Delta s$ along the trajectory $\hat{\mathbf{s}}$ is given by

$$\begin{aligned}\phi_{sources}(S = N_s\Delta s) = q_0 \exp\left(-\int_0^{N_s\Delta s} \mu_t(s)ds\right) \\ +q_1 \exp\left(-\int_{\Delta s}^{N_s\Delta s} \mu_t(s)ds\right) + \ldots + q_{N_s},\end{aligned} \tag{4.4}$$

where $\mu_t = \mu_{\mathrm{a}} + \mu_{\mathrm{s}}$. Exploiting the fact that a product of exponentials is equivalent to a sum of their exponents, the attenuation over the entire trajectory can be factorised out, provided source terms after the zeroth one are attenuated over the correct path

length. This is achieved by changing the sign and limits on the relevant integrals, as shown below,

$$\phi_{sources}(S = N_s\Delta s) = \exp\left(-\int_0^{N_s\Delta s} \mu_t(s)ds\right)\left[q_0 + q_1 \exp\left(\int_0^{\Delta s} \mu_t(s)ds\right) + \ldots + q_{N_s} \exp\left(\int_0^{N_s\Delta s} \mu_t(s)ds\right)\right] \tag{4.5}$$

$$\phi_{sources}(S = N_s\Delta s) = \exp\left(-\int_0^{N_s\Delta s} \mu_t(s)ds\right)\left[\sum_{n_s=0}^{N_s} q_{n_s} \exp\left(\int_0^{n_s\Delta s} \mu_t(s)ds\right)\right]. \tag{4.6}$$

When the source distribution, q, is a continuous function of s, the summation inside the square brackets in Eq. (4.6) becomes an integral over s between 0 and the position at which the radiance is calculated, which in this case is $N_s\Delta s$,

$$\phi_{sources}(S = N_s\Delta s) = \exp\left(-\int_0^{N_s\Delta s} \mu_t(s)ds\right)\left[\int_0^{N_s\Delta s} q(s) \exp\left(\int_0^{s} \mu_t(\zeta)d\zeta\right) ds\right], \tag{4.7}$$

where ζ is an integration variable representing a distance along $\hat{\mathbf{s}}$ between 0 and S.

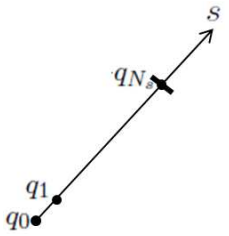

FIGURE 4.2: Diagram of trajectory s with source points q_{n_s} present.

Eq. (4.3) can be solved in a similar manner, except now the contributions to $\phi_{scattered}$ are simply the result of light scattered from the trajectories $\hat{\mathbf{s}}'$. Let us define a new source term to account for inward-scattered light $q_{scattered,\ n_s}(s) = \mu_{s,\ n_s} \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}, \hat{\mathbf{s}}')\phi(\mathbf{x}, \hat{\mathbf{s}}')d\hat{\mathbf{s}}'$. Similar to sources of photons on the trajectory, inward-scattered photons are attenuated

between the position at which they were in-scattered and $N_s\Delta s$ as shown in Fig. 4.3,

$$\phi_{scattered}(S = N_s\Delta s) = q_{scattered,\ 0} \exp\left(-\int_0^{N_s\Delta s} \mu_t(s)ds\right) + q_{scattered,\ 1} \exp\left(-\int_{\Delta s}^{N_s\Delta s} \mu_t(s)ds\right) + \ldots + q_{scattered,\ N_s}. \tag{4.8}$$

The contributions from light scattered inwards at the positions q_{n_s} must be summed or integrated along the trajectory from 0 to $N_s\Delta s$ such that the radiance in direction $\hat{\mathrm{s}}$ at position s is given by

$$\phi_{scattered}(S = N_s\Delta s) = \exp\left(-\int_0^{N_s\Delta s} \mu_{\mathrm{t}}(s)ds\right)\left[\int_0^{N_s\Delta s} q_{scattered}(s) \exp\left(\int_0^{s} \mu_{\mathrm{t}}(\zeta)d\zeta\right) ds\right]. \tag{4.9}$$

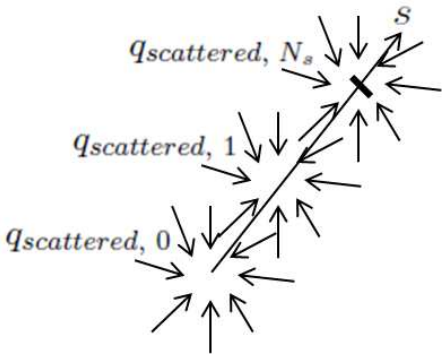

FIGURE 4.3: Diagram of trajectory s with points at which light is scattered inwards, $q_{scattered,\ n_s}$, present.

We now have a model for how light is accumulated through sources and inward scattering along the trajectory as well as how light is lost due to absorption and scattering. Writing the complete expression for the radiance as a sum of the contributions from sources and inward scattered light at a point S along trajectory $\hat{\mathrm{s}}$,

$$\phi(S) = \exp\left(-\int_0^{S} \mu_{\mathrm{t}}(s)ds\right)\left[\left[\int_0^{S} q_{scattered}(s) \exp\left(\int_0^{s} \mu_{\mathrm{t}}(\zeta)d\zeta\right) ds\right] + \left[\int_0^{S} q(s) \exp\left(\int_0^{s} (\mu_t(\zeta))d\zeta\right) ds\right]\right]. \tag{4.10}$$

Consider a source of photons at the origin of a trajectory, $q_0\delta(s)$, with uniform absorption and scattering coefficients and photons being scattered inward at a position $s = S/2$, then Eq. (4.10) simplifies to

$$\phi(S) = \exp\left(-\int_0^S \mu_{\mathrm{t}}(s)ds\right)\left[\left[\int_0^S q_{scattered}\delta(s - S/2)\exp\left(\int_0^s \mu_{\mathrm{t}}(\zeta)d\zeta\right)ds\right] + \left[\int_0^S q_0\delta(s)\exp\left(\int_0^s (\mu_t(\zeta))d\zeta\right)ds\right]\right]. \tag{4.11}$$

The second term in the square brackets obviously collapses to q_0, while the first term yields $q_{scattered}\exp(-\mu_{\mathrm{t}}S/2)$. Multiplying in the attenuative portion outside of the brackets gives

$$\phi(S) = q_{scattered}\exp(-\mu_{\mathrm{t}}S/2) + q_0\exp(-\mu_{\mathrm{t}}S) \tag{4.12}$$

$$= q_{scattered}\exp(-\mu_{\mathrm{t}}S/2) + q_0\exp(-\mu_{\mathrm{a}}S)\exp(-\mu_{\mathrm{s}}S). \tag{4.13}$$

The energy from the source at the origin will have been attenuated due to absorption according to $\exp(-\mu_{\mathrm{a}}S)$ and scattering according to $\exp(-\mu_{\mathrm{s}}S)$. The probability per unit length of a photon travelling a distance S before being scattered is

$$P(S) = \mu_{\mathrm{s}}\exp(-\mu_{\mathrm{s}}S). \tag{4.14}$$

In order to determine S, the cumulative probability density function, $cPDF(s)$ is sampled using a uniform random variable on $[0,1]$, $\mathcal{U}([0,1])$. The $cPDF$ is formed by integrating (4.14) between 0 and S

$$cPDF(s) = \mu_{\mathrm{s}}\int_0^S \exp(-\mu_{\mathrm{s}}\zeta)d\zeta = 1 - \exp(-\mu_{\mathrm{s}}S). \tag{4.15}$$

Since the domain of a cumulative probability density function is $[0,1]$, the right-hand side of Eq. (4.15) can be set equal to $\mathcal{U}([0,1])$ and solved for S (using the fact that $\mathcal{U}([0,1]) \equiv 1 - \mathcal{U}([0,1])$):

$$S = -\frac{ln(\mathcal{U}([0,1]))}{\mu_{\mathrm{s}}}. \tag{4.16}$$

It is useful at this point to introduce the concept of 'weight'. Weight is a unitless quantity representing the amount of energy of a photon packet. This allows photon weight as a function of distance S to be written as

$$W = W_0 \exp\left(-\mu_{\mathrm{a}} S\right) \tag{4.17}$$

where W_0 is the weight of the photon packet at the start of the trajectory which, if the photon trajectory starts from the source is just equal to q_0 and S is the length of the trajectory obtained using Eq. (4.16). The weight deposited after a distance S is therefore

$$\Delta W = W_0 \left(1 - \exp(-\mu_{\mathrm{a}} S)\right), \tag{4.18}$$

for the absorbed energy density, and

$$\Delta W = \frac{W_0}{\mu_{\mathrm{a}}} \left(1 - \exp(-\mu_{\mathrm{a}} S)\right), \tag{4.19}$$

for the fluence. The above absorption-scattering scheme is described by Sassaroli et al. [161] as the microscopic Beer-Lambert Law (mBLL) and has been shown to converge to other weight deposition schemes; in fact, the mBLL approach has been demonstrated [161] to produce lower variance fluence estimates for a given number of photons compared with other weight deposition schemes, such as the albedo-weight method proposed by Prahl [162], and for this reason it was selected as the scheme in the RMC algorithm.

Note that the mBLL method can be readily extended to 'white Monte Carlo' (WMC) [163] implementations whereby rather than storing the weight deposited in each voxel, complete photon trajectories are stored. This involves storing one or more photon positions between scattering events and the photon weight is modelled by computing the distances between the positions: $W = W_0 \exp\left(-\sum_{n_s} \mu_{\mathrm{a}}(\mathbf{r}_{n_s}) S_{n_s}\right)$, where $\mathbf{r}_{n_s}$ is the photon's position at step n_s, after travelling a distance S_{n_s}. Thus, although this method is not of any particular value if used once as a forward model, rapid updating of the distribution of the absorption coefficient allows scored quantities like the fluence and/or absorbed energy density to be quickly modified. This type of implementation is suitable when the distribution of the scattering coefficient is known; however this is not

typically the case in PAT and, thus, photon trajectories must be modelled for each set of scattering coefficient distributions. Furthermore, the potential application of variance reduction techniques when source-detector pairs are sparsely distributed makes WMC more applicable to the quantitative problem in DOT, while in QPAT source-detector pairs exist for every voxel in the domain resulting in excessive memory demands of storing billions of photon trajectories.

4.1.2 Deflection of photon packet

Following the absorption event, after travelling a distance S to a new position such that $\mathbf{x}' = \mathbf{x} + S\hat{\mathbf{s}}$, the photon is scattered. The phase function, $P(\hat{\mathbf{s}}, \hat{\mathbf{s}}')$, is typically chosen to be the Henyey-Greenstein phase function [61] because it is a good approximation for tissue and exists in closed form, making it straightforward to integrate; in order to determine the scattering angle, the cumulative integral of $P(\hat{\mathbf{s}}, \hat{\mathbf{s}}')$ must be evaluated:

$$\text{cPDF}_{3D}(\hat{\mathbf{s}}, \hat{\mathbf{s}}') = \int_{-1}^{\cos(\theta_{3D})} \frac{1}{2} \frac{1 - g^2}{(1 + g^2 - 2g(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'))^{3/2}} d(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'). \tag{4.20}$$

Eq. (4.20) can now be solved for the deflection angle between $\hat{\mathbf{s}}$ and $\hat{\mathbf{s}}'$ using $\theta_{3D} = \arccos(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}')$ by substituting another uniform random variable

$$\theta_{3D} = \arccos\left(\frac{1}{2g}\left[1 + g^2 - \left(\frac{1 - g^2}{1 - g + 2g\mathcal{U}([0,1])}\right)^2\right]\right), \tag{4.21}$$

for $g \neq 0$. When $g = 0$, scattering is isotropic and the deflection angle is selected from a uniform random variable over $[0, 2\pi]$

$$\theta_{3D} = \arccos\left(2\mathcal{U}([0,1]) - 1\right). \tag{4.22}$$

In 3D, there is an additional scattering angle, ψ, which rotates $\hat{\mathbf{s}}'$ isotropically, as shown in Fig. 4.4.

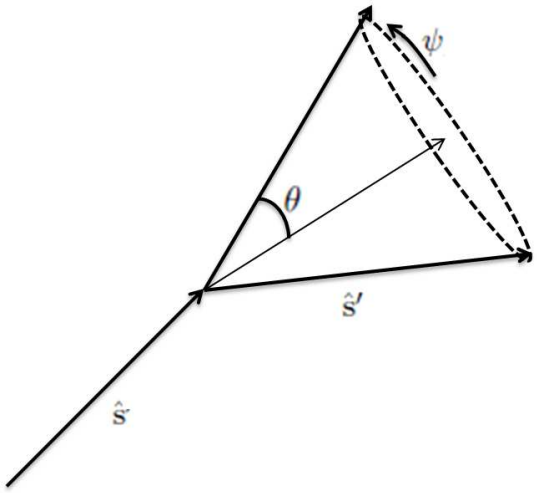

FIGURE 4.4: Diagram of photon scattering from $\hat{\mathrm{s}}$ to $\hat{\mathrm{s}}'$. The deflection angle, θ, is selected from the phase function and ψ rotates the new trajectory about the original trajectory.

As such it is simply chosen using another uniform random variable according to

$$\psi = 2\pi\mathcal{U}([0,1]). \tag{4.23}$$

Then, the updates to the photon direction are given by Peplow et al. [164] and can be derived by performing a coordinate transform on the scattering angles and rotating the original photon direction

$$\mu_x \leftarrow \mu_x \cos(\theta_{3D}) + (\mu_x\mu_z \cos(\psi) - \mu_y \sin(\psi))\frac{\sqrt{1-\cos^2(\theta_{3D})}}{\sqrt{1-\mu_z^2}} \tag{4.24}$$

$$\mu_y \leftarrow \mu_y \cos(\theta_{3D}) + (\mu_y\mu_z \cos(\psi) + \mu_x \sin(\psi))\frac{\sqrt{1-\cos^2(\theta_{3D})}}{\sqrt{1-\mu_z^2}} \tag{4.25}$$

$$\mu_z \leftarrow \mu_z \cos(\theta_{3D}) - \cos(\psi)\sqrt{1-\cos^2(\theta_{3D})}\sqrt{1-\mu_z^2}, \tag{4.26}$$

where μ_x, μ_y and μ_z are the direction cosines of the incident photon, and the arrow is used to indicate replacement. Eq. (4.24)–(4.26) collapse to

$$\mu_x \leftarrow \sin(\theta_{3D})\cos(\psi) \tag{4.27}$$

$$\mu_y \leftarrow \sin(\theta_{3D})\sin(\psi) \tag{4.28}$$

$$\mu_z \leftarrow \cos(\theta_{3D}), \tag{4.29}$$

when μ_z is very close to unity.

In 2D, there is only one deflection angle, $\theta_{2D} = \arccos(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}')$, and a suitable choice of scattering phase function to model scattering in biological tissue is analogous to the Henyey-Greenstein phase function [165]:

$$P_{2D}(\hat{\mathbf{s}}, \hat{\mathbf{s}}') = \frac{1}{2\pi} \frac{1 - g^2}{(1 + g^2 - 2g(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'))} d(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'), \tag{4.30}$$

and the cumulative integral can be sampled according to

$$\theta_{\text{2D}} = 2 \arctan\left(\frac{1 - g}{1 + g} \tan(\pi \mathcal{U}([0, 1]))\right). \tag{4.31}$$

Based on the scattering angle, the new photon direction is obtained by

$$\mu_x' = \cos(\theta_{2D})\mu_x + \sin(\theta_{2D})\mu_z \tag{4.32}$$

$$\mu_z' = \cos(\theta_{2D})\mu_z - \sin(\theta_{2D})\mu_x. \tag{4.33}$$

RMC was only implemented in 2D meaning it was this scheme that was used.

4.2 Implementation of MC

There are a variety of implementational aspects of the basic MC algorithm for light transport, including spatial discretisation, source implementation, boundary conditions as well as parallelisation. These are discussed below.

4.3 Spatial discretisation and basis

The quantity of interest, whether it be the fluence or the absorbed energy density, must be scored in a particular region or subdomain. This is achieved by splitting the domain into subdomains u_j, where $\Omega = \cup_j^J u_j$, over which there are distributed basis functions governing the energy deposited in a set of subdomains. Given that N_p photons traverse a subdomain, the fractions of weight they deposit ΔW_{n_p} must be summed over the number

of photons, $\sum_{n_p=1}^{N_p} \Delta W_{n_p}$, and then integrated over the basis function. In the case of a piecewise constant basis, the integral is performed over the subdomain to form the total weight deposited: $W_{tot,\,j} = \int_{u_j} \sum_{n_p=1}^{N_p} \Delta W_{n_p} dV_j$ where dV_j is the volume of the j^{th} subdomain, u_j, (in 2D, the integral is performed over an area). This is consistent with definitions provided by Boas et al. [166] excluding normalisation, which is addressed in the next section.

A simple model of tissue is to consider it being composed of a number of stacked semi-infinite slabs with different optical properties. MCML [167] models the tissue using multiple homogeneous layers in the xy-directions. As such, this geometry has cylindrical symmetry and, using a piecewise constant basis, scored quantities are obtained by integrating over the radius and z-axis between z_n to $z_n + \Delta z$, for layer thickness Δz and the n^{th} slab. This way of simplifying tissue geometry is not applicable to most geometries imaged using PAT; for this reason, this model has recently been extended to include inclusions, first, of spherical tissues [168] and later spheroidal and cylindrical inclusions [169], which may be suitable for modelling light transport for lymph node or tumour imaging.

In order to model light transport in arbitrarily complex geometries, a voxelised grid or mesh is required. Grids offer the advantage that they are simple to navigate computationally for weight deposition, achieved through simple truncation of the photon's physical coordinates. A mesh requires the photon's coordinates to be mapped to the barycentric coordinates associated with the mesh element, which is more computationally expensive. However, a grid composed of cuboid elements cannot accurately – or at least not straightforwardly – account for internal index of refraction mismatch [170] (unless these occur between cuboid objects), and can only model mismatch at the domain boundaries [132, 162]; on the other hand, the use of tetrahedral or higher order polyhedral elements can be used in such a way that surfaces are accurately modelled, thus allowing regions where there is a refractive index change (e.g. between CSF and the cortex) to be accurately modelled using ray-polyhedra intersections [133].

The intersection tests required for the refraction of light via index mismatch are also useful for updating the step length of Eq. (4.16) due to changes in the scattering coefficient

[162]

$$S \leftarrow (S - d)\frac{\mu_{s,1}}{\mu_{s,2}}, \tag{4.34}$$

where the arrow is used to indicate replacement, d is the distance from the photon's current position to the nearest polyhedron face that it crossed that has a value of $\mu_{s,2} \neq \mu_{s,1}$, where $\mu_{s,1}$ is the scattering coefficient in the current element (note that with alternative weight deposition schemes $\mu_t = \mu_\mathrm{a} + \mu_\mathrm{s}$ is used instead of just μ_s [161]). It is desirable to avoid the computationally expensive ray-polyhedra tests, which was achieved by Fang et al. [132] by propagating the photon in increments, typically less than a voxel edge length dS, and updating the total step length S according to Eq. (4.34) except with $d = n_S dS$ where n_S is the number of voxel edges traversed. Of course, this is likely to yield a small error, ΔS, in the updated step length as it has travelled this distance in $\mu_{s,2}$ rather than $\mu_{s,1}$ as illustrated in Fig. 4.5; the erroneous portion of the expression, $\Delta\mu_S\frac{\mu_{s,2}}{\mu_{s,1}}$, is very small for changes in media where the scattering coefficient is similar in magnitude (true for many biological tissues) and the error will also tend to zero as $dS \to 0$, i.e. as the photon is then stepped in small fractions along the trajectory. RMC uses the stepping scheme for benefits in terms of speed.

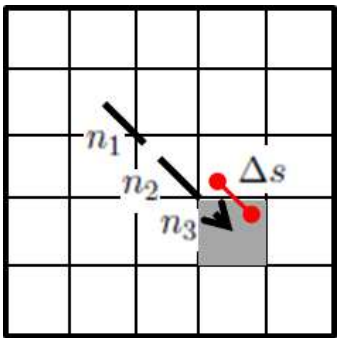

FIGURE 4.5: Diagram of photon trajectory traversing multiple voxels, illustrating error in path length, ΔS, when there is a change in the scattering coefficient $\mu_{s,2}$ in the grey region.

This stepping approach is extendable to unstructured grids such as an octree geometry. Discretisation in an octree consists of a grid with a maxmimum voxel edge length L_{max} and all other voxels are either this size or are the result of subdivision of a parent voxel n_v times, yielding a family of voxels with edge length $L_{max}/2^{n_v}$. This offers a great deal

of flexibility in terms of coarseness of discretisation and has significant benefits in terms of memory storage of the medium properties and scored quantities; however, in order for the error in S to remain small, the photon would have to step the minimum $L_{max}/2^{n_v}$ associated with each parent voxel, which is compute intensive.

4.3.1 Normalisation and sources

Quantities output by MC simulations require normalisation so that they are representative of the energy input and are not scaled by the number of photons simulated. The latter can be removed by simply dividing the result by the number of photons, while the former requires the data to be multiplied by the total energy input; e.g. for a 20mJ pulse of energy, the fluence output from MC must just be multiplied by this number to have the correct units [171].

Source implementation in grid geometry is straightforward as it simply involves launching photons from a particular position with a given direction. An isotropic point source, either inside the domain or on the boundary, simply involves launching photons from the same starting position and distributing their direction randomly over 4π Steradians:

$$\theta_{in} = 2\pi\mathcal{U}([0,1]) \tag{4.35}$$

$$\psi_{in} = \arccos(2\mathcal{U}([0,1]) - 1), \tag{4.36}$$

where θ_{in} and ψ_{in} are the polar and azimuthal angles, respectively (in 2D only Eq. (4.35) is required for an isotropic source). A pencil beam is achieved by launching all the photons in the same direction, whereas a collimated broadbeam can be achieved by launching all photons in the direction of collimation and adding uniformly distributed values in the x- or xy-directions about the central position. Similarly a Gaussian beam can be used by shifting the launching positions of photons about their central position using normally distributed random variables and in homogeneous media this can be realised using convolution with a Gaussian kernel [171].

4.3.2 Boundary conditions

Treatment of photons exiting the domain often involves ray-polyhedra intersections, except when the internal and external indices of refraction, n_i and n_e, are matched giving a reflection coefficient of zero. This can be generalised to the situation where there is a non-zero reflection coefficient at the boundary and photons are reflected specularly: $\theta_i = \theta_r$ i.e. the angle of incidence to the normal is equal to the reflected angle and the new photon direction is given by $\hat{\mathbf{s}} \leftarrow 2(\hat{\mathbf{n}} \cdot \hat{\mathbf{s}})\hat{\mathbf{n}} - \hat{\mathbf{s}}$, where $\hat{\mathbf{n}}$ is the normal to the reflecting surface. Under these conditions the reflected photon weight is scaled by R where

$$R = \left(\frac{n_i - n_e}{n_i + n_e} \right)^2, \tag{4.37}$$

which may be representative of many situations in both *in vivo* and *ex vivo* PAT imaging, as the target is often immersed in water or acoustically coupled to the detector using ultrasound gel. Thus, in the case where $n_i = n_e$, the boundary is index matched and $R = 0$, meaning no light travels inward from the boundary (except from sources on the boundary).

4.3.3 Photon termination and variance reduction

Photons that escape the domain are terminated, but those that remain in the domain must be terminated using a threshold condition as their weight approaches zero asymptotically and can survive indefinitely without the application of a threshold, W_T. It is typical to propagate a photon as long as its weight is greater than W_T. However, this is often coupled with a variance reduction scheme called 'roulette'. Roulette prevents the propagation of a large number of low-weight photons which do not contribute significantly to the result and instead performs unbiased selection of a single photon to propagate the weight of others; this is achieved by allowing the survival of $1/C_T$ photons whose weight is below W_T, if a uniform random variable $\mathcal{U}([0, 1]) < C_T$; the surviving photon's weight is then scaled to be $C_T W$ [52]; a typical choice for C_T is $1/10$.

Certain variance reduction techniques involve *importance sampling* which encourages photon trajectories that are more likely to interact with a light detector to be sampled more frequently than those that do not yield a detection event [172]. Of course, in PAT, the optical detector is effectively distributed over the entire domain as light absorption anywhere in the tissue will yield a PA signal. Domain splitting techniques [173] can also help reduce variance in fluence estimates, but given that this would involve compromising resolution in PAT, this is unattractive; nevertheless, this is akin to the octree method of spatial discretisation described in Section 4.3 and may be useful in variance reduction away from regions of interest in the inverse problem. A number of other variance reduction methods exist [174, 175] but are not either applicable to photon transport using the RTE for a single wavelength/time gate or when optical detection is distributed over the illuminated volume and isotropic as in PAT.

4.3.4 Parallelisation and hardware

Photons in MC simulations are propagated completely independently, meaning each photon can be treated using an independent thread with minimal cross-talk between threads. This has long been an area of reserach [176] in biomedical applications of MC light transport simulations; however, recent advances in computing hardware have facilitated highly parallelised MC implementations. The availability of multicore CPU processors have allowed multithreading of C/C++ MC codes providing approximately an order of magntidue increase in speed compared with serial implementations, while the accessibility of high memory (several GB) GPGPUs [177] have allowed the development of high-speed MC codes such as MCX [132], MMC [133] and a GPU-accelerated version of MCML, CUDAMCML [178], which provide two to three orders of magnitude speed up over their serial counterparts. RMC was implemented in 'Julia' [179], a high-level language, which allowed straightforward parallelisation over multiple CPU cores on a single node.

4.4 Monte Carlo modelling of the radiance

As discussed in Chapter 2, there are a number of methods that can be used to solve the RTE for the radiance. However, given that analytic solutions in homogeneous or layered media are not applicable in QPAT, numerical techniques must be used. FE modelling of the radiance is a popular approach to simulating the radiance which requires some discretisation over angle, the most intuitive of which, although not necessarily optimal, is to take the unit sphere, in 3D, or the unit circle, in 2D, and divide it up into segments. This piecewise constant set of basis functions over angle may not be optimal for describing light propagation in PAT; PAT images have a region near the tissue surface in which ballistic or near-ballistic propagation is prevalent meaning many orders of discretisation are required to capture the high directionality of the radiance. This results in excessive memory demands in FE implementations of the RTE [53, 180] as realistic levels of anisotropy in tissue (e.g. $g \geq 0.9$) require greater than 64 angles per node, which equates to more than 64^2=4096 elements per node in the system matrix. This fine discretisation is then poorly utilised deep into the tissue as the light field is diffusive in this region and therefore dense in this basis.

A number of schemes have been proposed to overcome these limitations when simulating the fluence. One approach uses a model accurate for directional light propagation, e.g. the RTE, near the optical source, and the DA using internal sources for the rest of the domain [53, 181]. A similar approach has been attempted using MC in the optical near-field and the DA in the far-field [182] but was demonstrated to be of limited accuracy compared with pure MC simulations and is challenging to implement due to the combination of computational meshes and the placement of diffuse sources for the DA simulation.

However, the increasing availability of highly parallel computing platforms means that the parallelism of the MC algorithm should be exploited; this is especially relevant as the method is inherently capable of simulating the angularly-dependent radiance because the directional information is only lost when the implicit integral over angle is performed in calculating the total weight deposited in a given voxel; with the introduction of some

angular discretisation, the radiance could be simulated without a substantial increase on the already modest memory demands.

4.4.1 Angular discretisation and basis

For the reasons described above, it is not attractive to discretise in angle by dividing the angular domain into segments (also known as the discrete ordinate method). Taking inspiration from the P_n approximation [63], the field at every position can be expanded as a series of spherical harmonics [183, 184]. The direct analogue in 2D is to use a Fourier expansion. The expansion for the radiance in these bases is therefore given by

$$\phi(\mathbf{x}, \hat{\mathbf{s}}') = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi} \right)^{1/2} a_l^m(\mathbf{x}) Y_l^m(\mathbf{x}, \hat{\mathbf{s}}'), \tag{4.38}$$

in 3D, where $Y_l^m(\mathbf{x}, \hat{\mathbf{s}}')$ are the real spherical harmonics, expressed as

$$Y_l^m(\hat{\mathbf{s}}') = \begin{cases} \left(\left(\frac{2l+1}{4\pi} \right) \left(\frac{l-|m|!}{l-|m|!} \right) \right)^{1/2} (-1)^{1/2(m+|m|)} \cos(m\psi) P_l^{|m|}(\cos\theta) & \text{for } m \geq 0 \\ \left(\left(\frac{2l+1}{4\pi} \right) \left(\frac{l-|m|!}{l-|m|!} \right) \right)^{1/2} (-1)^{1/2(m+|m|)} \sin(m\psi) P_l^{|m|}(\cos\theta) & \text{for } m < 0 \end{cases} \tag{4.39}$$

In this case, the a_l^m are then given by $a_l^m(\mathbf{x}) = \sum_{n_p=1}^{N_p} dW_{n_p} \int_{\mathcal{S}^2} \delta(\hat{\mathbf{s}} - \hat{\mathbf{s}}'_{n_p}) Y_l^{m+}(\mathbf{x}, \hat{\mathbf{s}}) d\hat{\mathbf{s}}$, where $+$ indicates complex conjugation; in the case of real spherical harmonics, the coefficients are also real, meaning this is equivalent to $a_l^m(\mathbf{x}) = \sum_{n_p=1}^{N_p} dW_{n_p} \int_{\mathcal{S}^2} \delta(\hat{\mathbf{s}} - \hat{\mathbf{s}}'_{n_p}) Y_l^m(\mathbf{x}, \hat{\mathbf{s}}) d\hat{\mathbf{s}}$. The 3D model was however not implemented, and only the 2D model is discussed in subsequent chapters of this thesis.

In 2D, the Fourier expansion is simply given by

$$\phi(\mathbf{x}, \theta') = \frac{1}{2\pi} a_0(\mathbf{x}) + \frac{1}{\pi} \sum_{n=1}^{N=\infty} a_n(\mathbf{x}) \cos(n\theta') + \frac{1}{\pi} \sum_{n=1}^{N=\infty} b_n(\mathbf{x}) \sin(n\theta'), \tag{4.40}$$

where

$$a_0 = \sum_{n_p=1}^{N_p} dW_{n_p} \int_{\mathcal{S}^1} \delta(\theta - \theta') d\theta = \sum_{n_p=1}^{N_p} dW_{n_p} \tag{4.41}$$

$$a_n = \sum_{n_p=1}^{N_p} dW_{n_p} \int_{\mathcal{S}^1} \delta(\theta - \theta'_{n_p}) \cos(n\theta) d\theta = \sum_{n_p=1}^{N_p} dW_{n_p} \cos(n\theta'_{n_p}) \tag{4.42}$$

$$b_n = \sum_{n_p=1}^{N_p} dW_{n_p} \int_{\mathcal{S}^1} \delta(\theta - \theta'_{n_p}) \sin(n\theta) d\theta = \sum_{n_p=1}^{N_p} dW_{n_p} \sin(n\theta'_{n_p}), \tag{4.43}$$

for every position $\mathbf{x}$ in the domain. Thus, at every grid/mesh element, the total weight deposited, W_{tot}, is scaled by the coefficient for the associated spherical/Fourier harmonic.

This method of discretisation has two potential benefits over segmenting the unit sphere/-circle. Firstly, memory demands are reduced as few orders ($<$3) are required to store diffuse light fields, although many more are required near the source to adequately represent directional propagation. The impact of running simulations using a limited number of Fourier/spherical harmonics is discussed in Sections 4.4.4 and 5.3.3. Secondly, every photon contributes to the weight deposited in the n^{th} harmonic, while when the unit sphere/circle is segmented, a photon only deposits weight in the segment corresponding to the photon's direction, meaning a greater number of photons would be required to achieve the same variance. This novel MC model calculates the fluence as it is simply equal to $a_0(\mathbf{x})$, and the directional portion of the field is represented by the higher order terms. This can be illustrated by plotting the energy deposited in each harmonic from an isotropic point source. The domain was 4mm$\times$4mm with 0.1mm pixel size and up to second order Fourier harmonics were stored. Figure 4.6 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 4.6.

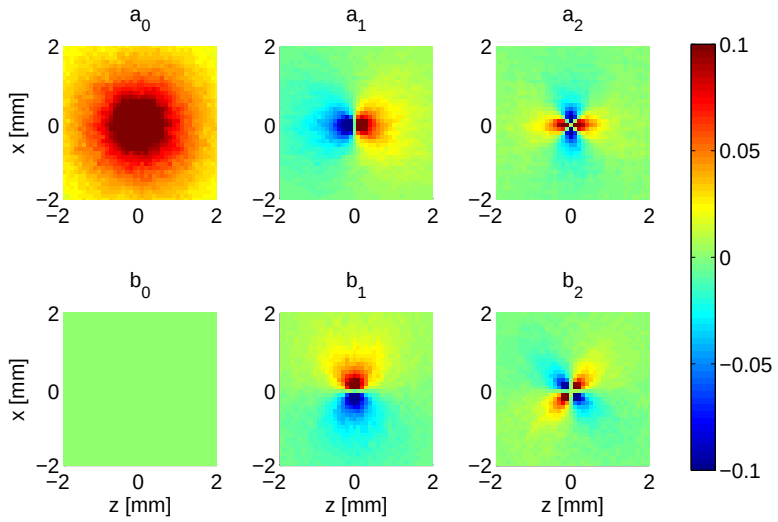

FIGURE 4.6: Fourier harmonics a_n (top row) and b_n (bottom row), for $n \in [0, 2]$, simulated using MC model with 10^6 photons in 4mm×4mm domain and isotropic point source at centre. Images normalised by maximum value of a_0.

4.4.2 Program flow chart

Other than the deposition of weight, the program flow of both the 2D and 3D radiance Monte Carlo models is the same as that presented by Fang [132] and Prahl [162], as displayed in Fig. 4.7.

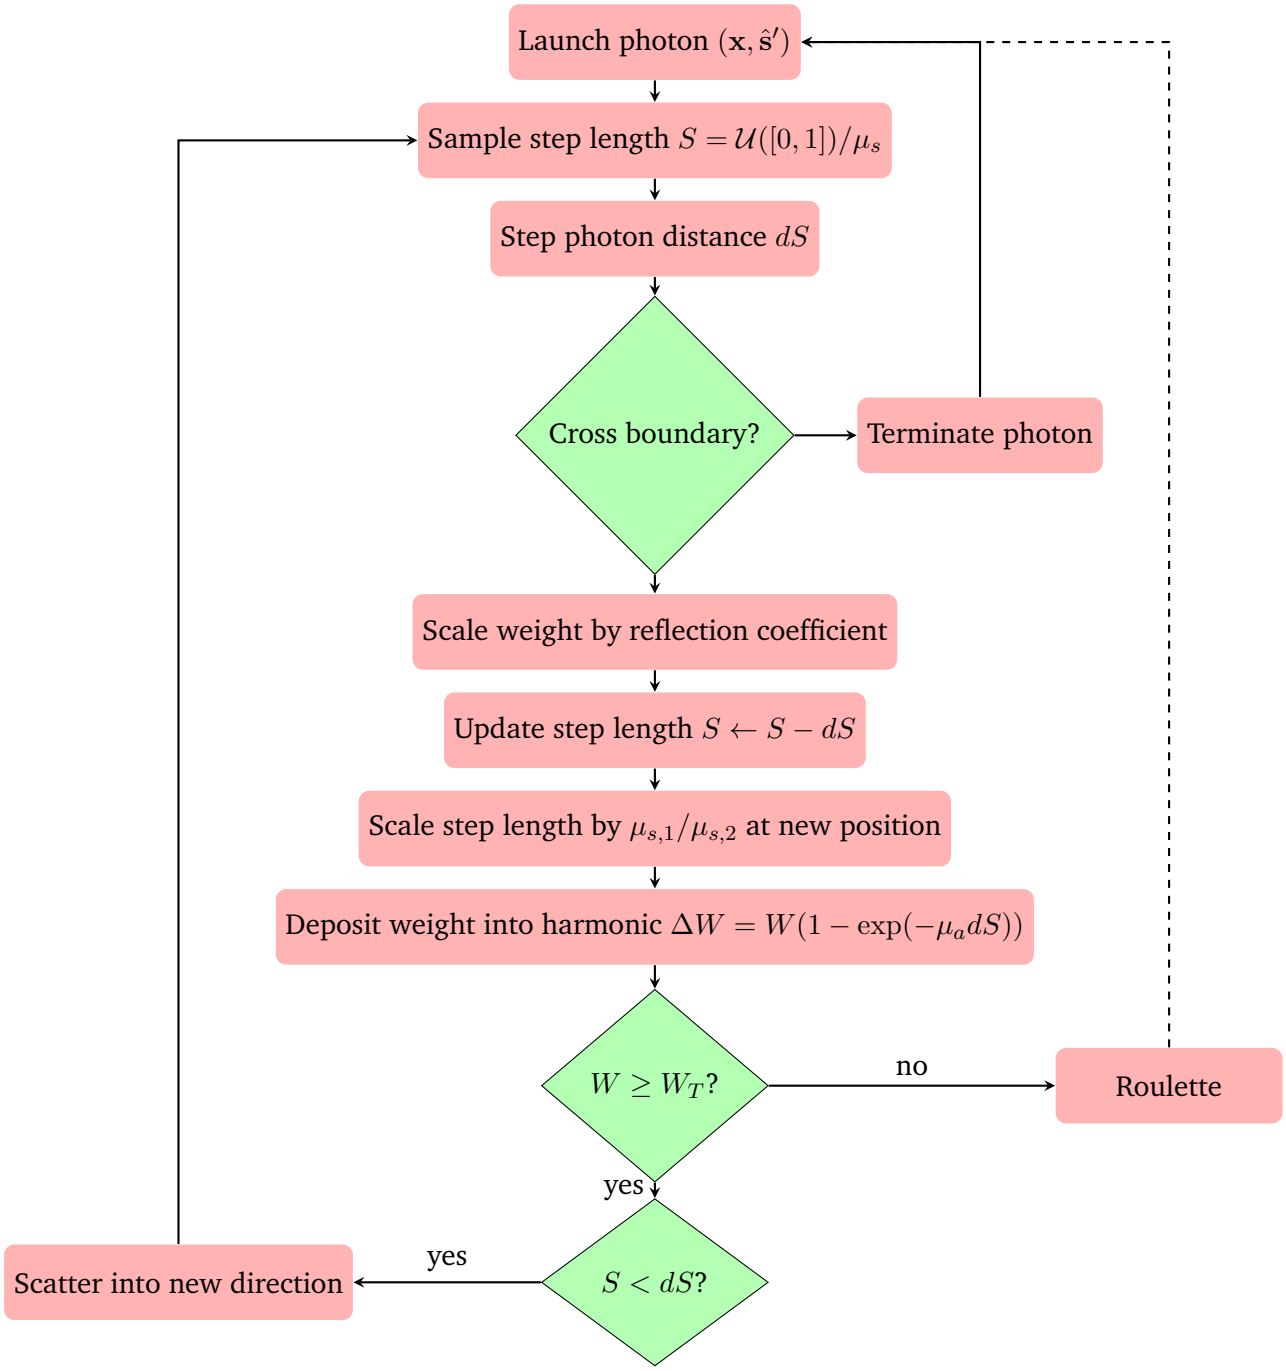

FIGURE 4.7: Flowchart of regular grid radiance Monte Carlo algorithm.

4.4.3 Validation of RMC (2D radiance Monte Carlo model)

It is common to validate other models of light transport against Monte Carlo models as MC is often considered the gold standard. This is not an option here as MC models

that compute the radiance are not available. There is a simple analytic solution for the fluence in a homogeneous 2D half-space illuminated by a pencil beam. This model indicates that in the far-field, at distances $z \gg 1/\mu_{\mathrm{s}}'$ into the domain, the fluence decays exponentially as $\exp(-\mu_{eff} z)$ [167]. This model derives from the DA in 2D and yields the effective attenuation coefficient $\mu_{eff} = \sqrt{2\mu_{\mathrm{a}}(\mu_{\mathrm{a}} + \mu_{\mathrm{s}}(1-g))}$. In order to compare the 2D radiance MC code against this simple diffusion model in 2D, a MC simulation was run with a domain 30mm×30mm containing 600×600 pixels illuminated at the centre of the edge at z=0mm with a collimated pencil beam consisting of 10^7 photons. The medium properties were chosen to be typical of tissue and consistent with those used in the validation performed by Wang et al. [167]: $\mu_{\mathrm{a}} = 0.01$mm^{-1}, $\mu_{\mathrm{s}} = 10$mm^{-1} and $g = 0.9$ and the boundaries were index-matched. Comparing the 2D radiance MC code in the far-field against this diffusion model by fitting a line to the log fluence integrated in the transverse direction as a function of z, it was found that a sufficiently accurate far-field fluence estimate is obtained as the predicted value of μ_{eff} was 0.1421mm^{-1} and the MC simulation yielded 0.1424mm^{-1}, i.e. an error of 0.2%.

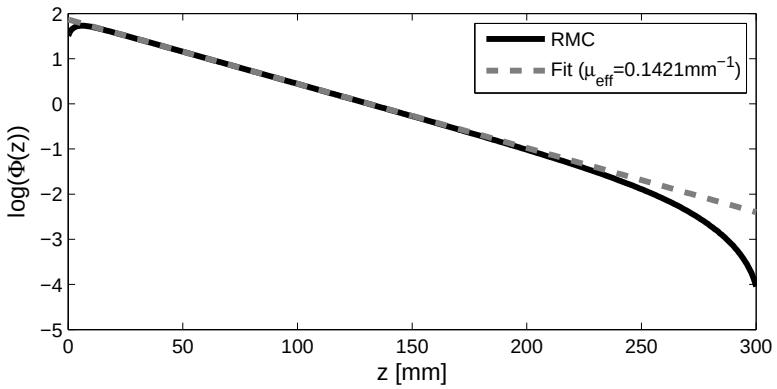

FIGURE 4.8: Plot of the log fluence in a homogeneous 30mm×30mm domain illuminated by a pencil beam at the centre of the edge at z=0mm, integrated over the x-axis and plotted as a function of depth, z. A straight line was fit to the region of the black curve (RMC) between 20mm and 160mm, which yielded a slope of 0.1421mm^{-1}, which is equal to μ_{eff}. The rapid decrease in the black curve after a distance of approximately 220mm is due to light escaping the domain at the index-matched boundary at 300mm.

However, the accuracy of the model in the near-field, particularly the accuracy of radiance estimates, is also of interest. There are few methods by which the radiance

can be calculated in 2D which have been independently validated. Finite element solutions to the RTE in 2D exist, but these have only undergone basic validation for angle-independent quantities such as the fluence [180], exitance [53] or transmittance [185] and have not demonstrated the validity of angle-resolved radiance estimates. Analytic solutions for the radiance also exist [67, 72, 73, 186] but these suffer from inaccuracies near the source when evaluated numerically [74] and the methods themselves have only been validated for semi-infinite or bounded circular media [72]. Nevertheless, good agreement of the angle-resolved radiance a few mm from the source would provide some confidence in the accuracy of the radiance MC model proposed in this chapter. Due to a lack of analytic solutions for rectangular bounded media, the comparison was performed using a large, bounded 15mm$\times$15mm square domain for the MC model with an isotropic point source at its centre, which emitted 10^8 photons, and 5 Fourier harmonics were recorded. The analytic solution was for an infinite domain also illuminated by an isotropic point source [67]. The radiance was measured as a function of angle at 2mm and 3mm away from the source in the positive x-direction and is plotted in Fig. 4.9.

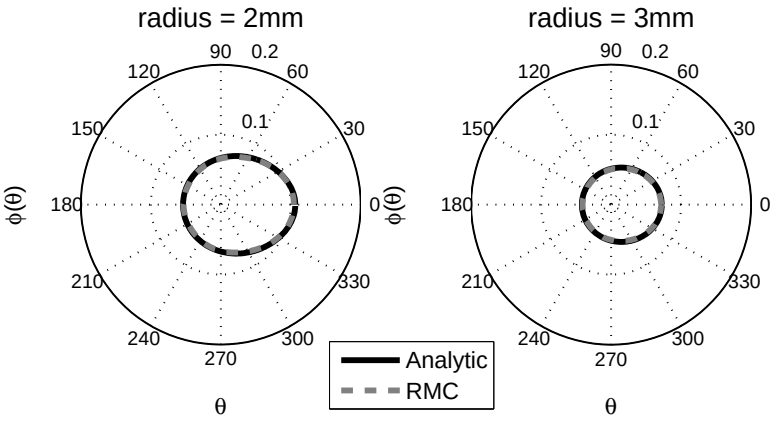

FIGURE 4.9: Polar plots of the angle-resolved radiance due to an isotropic point source in a homogeneous domain with $\mu_{\mathrm{a}} = 0.01\mathrm{mm}^{-1}$, $\mu_{\mathrm{s}} = 10\mathrm{mm}^{-1}$ and $g = 0.9$. Results from an analytic method (infinite domain) and RMC simulations (15mm$\times$15mm square domain) shown.

Very good agreement is observed between the two sets of radiance profiles plotted. There is a very minor discrepancy at zero and π radians (parallel to the horizontal axis), where the MC model very slightly underestimates the estimate made using the analytic solution. This discrepancy is more clearly illustrated in Fig. 4.10, which plots the same data as

in Fig. 4.9 but the log of the angle-resolved radiance. The MC model underestimating the radiance compared with the analytic solution is most likely due to the fact that the MC simulation was run in a finite domain meaning the light field decays slightly more quickly with distance from the source compared with the analytic solution which is valid for an infinite domain. While this effect is hardly noticeable near the source, at a distance of 2mm, it is more prominent at a distance of 3mm. Note that comparisons were also performed using 10, 20 and 40 harmonics in the MC simulations; however, these were only run using 10^6 photons due to extended computation times. Good agreement was observed between the MC model and the analytic solutions, but the high frequency components were very prominent in the plots because an insufficient number of photons were simulated for convergence of the MC solution.

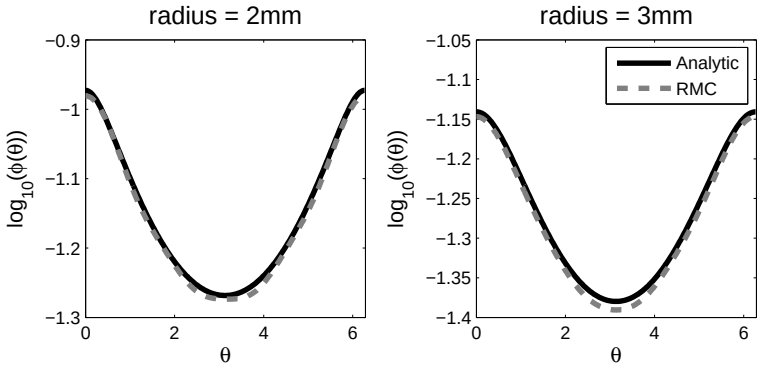

FIGURE 4.10: Plots of the angle-resolved log-radiance due to an isotropic point source in a homogeneous domain with $\mu_{\mathrm{a}} = 0.01\mathrm{mm}^{-1}$, $\mu_{\mathrm{s}} = 10\mathrm{mm}^{-1}$ and $g = 0.9$. Results from an analytic method (infinite domain) and RMC simulations (15mm$\times$15mm square domain) shown.

4.4.4 Alternative bases

Since a collimated illumination profile is a useful approximation to a beam of laser light, the capability to model this situation is desirable. The directional portion of the source term associated with this illumination type is a delta-function, which is not accurately represented using small N in the corresponding Fourier or spherical harmonic series. It may therefore be worthwhile to consider an amendment to these bases whereby a δ-function in the source direction, $\hat{\mathbf{s}}_0$ or θ_0, is included explicitly. These alternative series

can be written as

$$\phi(\mathbf{x},\hat{\mathbf{s}}') = \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\left(\frac{2l+1}{4\pi}\right)^{1/2} a_l^m(\mathbf{x})Y_l^m(\mathbf{x},\hat{\mathbf{s}}') + a_\delta\delta(\hat{\mathbf{s}}-\hat{\mathbf{s}}_0), \tag{4.44}$$

and

$$\phi(\mathbf{x},\theta') = \frac{1}{2\pi}a_0(\mathbf{x}) + \frac{1}{\pi}\sum_{n=1}^{N=\infty} a_n(\mathbf{x})\cos(n\theta') + \frac{1}{\pi}\sum_{n=1}^{N=\infty} b_n(\mathbf{x})\sin(n\theta') + a_\delta\delta(\theta'-\theta_0), \tag{4.45}$$

where a_δ is the coefficient associated with directional component of the field. Of course, a photon travelling in the direction $\hat{\mathbf{s}}_0$ will deposit energy into all the angular components in a given voxel; however, this calculation is unnecessary and the field can be split into scattered and unscattered components – also referred to as the collided and uncollided components [62] – with unscattered energy deposited into a_δ which then acts as a source term for the scattered field [185]. This method however was not implemented in this thesis and, as will be observed later, only 10-20 Fourier harmonics are required when applying RMC in 2D to the inverse problem in QPAT.

4.5 Summary

This chapter presented fundamental approaches to the development of MC light transport algorithms and demonstrated that, simply by decomposing the incident photon direction at each voxel into Fourier harmonics scaled by the photon weight, the angularly-dependent radiance could be simulated. The particular approach applied in solving the RTE for the radiance was to use 'frequency'-domain angular discretisation at every voxel, which were Fourier and spherical harmonics in 2D and 3D, respectively. The 2D code was validated against analytic solutions for the fluence, in the diffuse regime, and the angle-resolved radiance was shown to match analytic solutions a few mm away from an isotropic point source. There are two important benefits of having codes that are capable of computing the radiance accurately in heterogenous media. Firstly, it can be used as a forward model of the radiance and offers more compact storage of the radiance

through the use of Fourier or spherical harmonic angular discretisation. The second, more crucial, advantage of this radiance MC code in the context of QPAT is that it can be used in conjunction with an adjoint model to compute gradients of an error functional for application in the optical inverse problem. This is investigated in the next chapter.

Chapter 5

Computing functional gradients using an adjoint Monte Carlo model of radiance

In Chapter 4, 2D and 3D Monte Carlo models capable of modelling the radiance were presented. The 2D model used a Fourier representation for the angular part of the radiance, while the 3D model used spherical harmonics. Forward models of the radiance on their own are of limited applicability because measuring the radiance is not straightforward and PA images depend on the fluence. The key motivation for computing the radiance however is that the inverse problem in QPAT is both large scale and nonlinear, meaning the capability to quickly and accurately estimate gradients of the error between measured and simulated data is desirable. Monte Carlo models of the radiance are a promising candidate as they are highly parallelisable so simulations can be run on distributed computing clusters or GPGPUs with minimal cross-talk between nodes or cores. In this chapter, the radiance Monte Carlo code is extended to compute the adjoint radiance, which can be used to assist the computation of gradients of an error functional. The derivation of expressions for functional gradients is demonstrated and validated against the gold standard method of calculating gradients: finite differencing.

5.1 Least-squares approach to inverting photoacoustic images

The estimation of the absorption and/or scattering coefficient from PAT images can be achieved using least-squares minimisation of some error between simulated and measured data as shown in Eq. (5.1).

$$\underset{[\mu_\mathrm{a},\mu_\mathrm{s}]}{\arg\min}\,\epsilon = \frac{1}{2}\left|\left|p_0^{meas} - p_0(\mu_\mathrm{a},\mu_\mathrm{s})\right|\right|^2. \tag{5.1}$$

This inversion approach was chosen here due to it being mature within the field of QPAT and due to the availability of a number of optimisation routines. In order for the absorption and scattering coefficients to converge to their true values, i.e. those that yielded the measured image p_0^{meas}, gradients of the error functional with respect to the absorption and scattering coefficients are required for the majority of optimisation schemes that are applicable to the large scale inverse problem in QPAT. As there is no closed-form expression for a Monte Carlo model, analytic derivatives of the absorbed energy density, and hence the error functional, with respect to μ_a and/or μ_s cannot be calculated. Instead, finite differencing using a perturbation in μ_a or μ_s can be applied to approximate the derivative numerically.

5.1.1 Finite difference algorithm for gradient computation

Finite differencing computes the gradient of the error functional about a point χ_0 (e.g. μ_a) by perturbing the value of χ and measuring the respective change in ϵ. This typically involves computing the difference between the perturbed and unperturbed values of ϵ over the perturbation magnitude; it is worth noting that this is referred to as a forward-difference (shown in Eq. (5.2)), but other first-order differencing schemes exist, such as central differences, as well as higher-order schemes, all deriving from Taylor's theorem.

$$\left.\frac{\partial\epsilon(\mathbf{x})}{\partial\chi}\right|_{\chi_0} = \lim_{\delta\chi\to 0}\left(\frac{\epsilon(\mathbf{x},\chi_0+\delta\chi)-\epsilon(\mathbf{x},\chi_0)}{\delta\chi}\right), \tag{5.2}$$

where the derivative is being performed with respect to χ (μ_a and/or μ_s in the case of the inverse transport problem) at position $\mathbf{x}$, χ_0 is the current estimate of the parameter

and $\delta\chi$ is the perturbation. The value of the perturbed error functional is calculated in the following way

$$\epsilon(\mathbf{x}, \chi_0 + \delta\chi) = \frac{1}{2} \left|\left| p_0^{meas} - p_0(\mathbf{x}, \chi_0 + \delta\chi) \right|\right|^2, \tag{5.3}$$

which in a piecewise constant basis, in a structured or unstructured grid, can be expressed discretely at the r^{th} voxel as

$$\epsilon_r(\chi_0 + \delta\chi) = \frac{1}{2} \langle (p_0^{meas} - p_0(\chi_{r,0} + \delta\chi)) dV, p_0^{meas} - p_0(\chi_{r,0} + \delta\chi) \rangle, \tag{5.4}$$

where dV is a diagonal matrix containing grid element volumes.

Eq. (5.2) must be evaluated using a sufficiently small value of $\delta\chi$ in order for the functional gradient to converge to its true value. A typical approach therefore is to test the convergence of the gradient estimate by progressively decreasing $\delta\chi$ until there is no noticeable change in $\frac{\partial\epsilon}{\partial\chi}$. However, due to the finite variance in quantities estimated using a MC model, reducing the perturbation magnitude to a very small fraction of χ_0 will result in noise dominating the gradient estimate. MC noise in $p_0(\chi_0)$ and $p_0(\chi_0 + \delta\chi)$ is additive when calculating a difference through Eq. (5.4), resulting in greater noise levels in $\frac{\partial\epsilon}{\partial\chi}$ than in just the fluence or absorbed energy density. The consequence of the noise will become apparent later in this chapter.

The finite difference scheme relies on the perturbation $\delta\chi$ to be applied independently at every position r where the gradient is to be computed. For large PAT images, which can contain millions of unknowns, as many model evaluations may be required, each of which is expensive given the low noise, and therefore large number of photons, required for accurate finite difference calculation. Thus, in order to obtain an estimate of the gradient over several voxels in the domain and ensure that the estimate has converged, the model must be run at several positions, for several values of $\delta\chi$ and for a large number of photons. Despite the fact that this process introduces extremely large computation times for accurate 3D transport models, finite difference schemes remain a useful tool in the validation of alternative gradient computation approaches.

5.2 Adjoint Monte Carlo model of radiance

The principal attraction of adjoint-assisted gradient calculation compared with finite differencing is that the gradient can be obtained using two model evaluations per N_{vox} unknowns where N_{vox} is the number of elements in the domain. The computation of the adjoint radiance, ϕ^*, is similar to that of the forward radiance, except for the fact that sources are distributed internally in the domain, emit photons isotropically and the source strength is scaled by $\mu_{\text{a}}\left(H^{meas} - H(\mu_{\text{a}}, \mu_{\text{s}})\right)$, the adjoint source term. In order to write down the RTE in adjoint form, reciprocity of this equation is exploited, which can be thought of as propagating photons along $-\hat{\mathbf{s}}$ rather than $\hat{\mathbf{s}}$, and similarly using the symmetry of the scattering phase function since $P(\hat{\mathbf{s}}' \cdot \hat{\mathbf{s}}) \equiv P(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}')$:

$$(-\hat{\mathbf{s}} \cdot \nabla + \mu_{\text{a}}(\mathbf{x}) + \mu_{\text{s}}(\mathbf{x}))\phi^*(\mathbf{x}, \hat{\mathbf{s}}) = \mu_{\text{s}}(\mathbf{x}) \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}', \hat{\mathbf{s}})\phi^*(\mathbf{x}, \hat{\mathbf{s}}')d\hat{\mathbf{s}}' + q^*(\mathbf{x}, \hat{\mathbf{s}}). \tag{5.5}$$

The particular choice of source term in the adjoint RTE, $q^*(\mathbf{x}, \hat{\mathbf{s}})$, is discussed in Section 5.3, while the implementation of q^* is discussed below.

5.2.1 Source implementation

Because the absorption of energy in a particular voxel does not depend on angle, and since the adjoint model represents the probability of energy in the forward model being absorbed in a given voxel, the emission of photons from adjoint sources must be isotropic. This was implemented by launching photons from the centre of source voxels $\mathbf{x}_{\mathbf{source}}$ with the position distributed randomly and uniformly over the voxel: $\mathbf{x}_{\mathbf{source}} + \frac{dx}{2}\mathcal{U}[0, 1]$, where dx is the voxel edge length and $\mathcal{U}$ is a uniform random variable. Distributing the starting positions of photons over the voxel in this way is only valid for the piecewise constant basis used in the adjoint model and other bases would require photons' starting positions to be distributed differently. The isotropic initial directions were sampled in the usual way using Eqs. (4.35) and (4.36).

Since the source strength at a given voxel is given by $\mu_{\mathrm{a}}^{\mathbf{r}}\left(H^{meas} - H(\mu_{\mathrm{a}}^{\mathbf{r}}, \mu_{\mathrm{a}}^{\mathbf{r}})\right)$, rather than photons being launched with unitary weight, photon weight must be scaled according to the source strength with normalisation of the output quantity (i.e. radiance, absorbed energy density, harmonic, etc.) being N_p. It is worth noting that certain voxels will of course contain negative sources, meaning photon weight is initially negative and weight deposition is also negative. This has the implication that photon termination occurs when the absolute value of the photon weight falls below the threshold value: $|W_p| < W_T$ (see Section 4.3.3).

5.2.2 Validation of the adjoint model

In order to determine whether the implementation of the adjoint MC model is genuinely the adjoint of the forward model, the definition of the adjoint operator was applied:

$$\langle \mathcal{L}\mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathcal{L}^{*}\mathbf{y} \rangle, \tag{5.6}$$

where $\mathcal{L}$ and $\mathcal{L}^{*}$ are the forward and adjoint operators of the RTE and $\mathbf{x}$ and $\mathbf{y}$ represent source/measurement operator applied to the resultant field. Specifically, the three cases that require testing are:

$$\text{Case 1} \quad \mathbf{x}_1 = \tfrac{1}{2\pi}\delta(\mathbf{r} - \mathbf{r}_s) \quad \mathbf{y}_1 = \delta(\mathbf{r} - \mathbf{r}_d) \tag{5.7}$$

$$\text{Case 2} \quad \mathbf{x}_2 = \tfrac{1}{2\pi}\delta(\mathbf{r} - \mathbf{r}_s) \quad \mathbf{y}_2 = \delta(\mathbf{r} - \mathbf{r}_d)P_d(\mathbf{r}, \hat{\mathbf{s}}) \tag{5.8}$$

$$\text{Case 3} \quad \mathbf{x}_3 = \delta(\mathbf{r} - \mathbf{r}_s)P_s(\hat{\mathbf{s}}) \quad \mathbf{y}_3 = P_d(\mathbf{r}), \tag{5.9}$$

where $\mathbf{r}_s$ and $\mathbf{r}_d$ are the positions of the source and virtual detector, respectively, $P_d(\mathbf{r}, \hat{\mathbf{s}})$ is the spatial and angular sensitivity of the 'detector' and $P_s(\mathbf{r}, \hat{\mathbf{s}})$ is the spatial and angular sensitivity of the 'source'. Thus, a summary of the cases being considered are an isotropic point source-detector pair (Case 1), an isotropic point source and an anisotropic point detector (Case 2) and an anistropic point source and an isotropic detector distributed over the domain (Case 3). Substituting the expressions for $\mathbf{x}_{1,2,3}$ and $\mathbf{y}_{1,2,3}$ into Eq. (5.6)

yields

$$\int_{\Omega}\int_{2\pi}\phi_1(\mathbf{r},\hat{\mathbf{s}})\delta(\mathbf{r}-\mathbf{r}_d)d\mathbf{r}d\hat{\mathbf{s}} = \int_{\Omega}\int_{2\pi}\phi_1^*(\mathbf{r},\hat{\mathbf{s}})\delta(\mathbf{r}-\mathbf{r}_s)d\mathbf{r}d\hat{\mathbf{s}}, \tag{5.10}$$

$$\int_{\Omega}\int_{2\pi}\phi_2(\mathbf{r},\hat{\mathbf{s}})\delta(\mathbf{r}-\mathbf{r}_d)P_d(\mathbf{r},\hat{\mathbf{s}})d\mathbf{r}d\hat{\mathbf{s}} = \int_{\Omega}\int_{2\pi}\phi_2^*(\mathbf{r},\hat{\mathbf{s}})\delta(\mathbf{r}-\mathbf{r}_s)d\mathbf{r}d\hat{\mathbf{s}}, \tag{5.11}$$

$$\int_{\Omega}\int_{2\pi}\phi_3(\mathbf{r},\hat{\mathbf{s}})P_s(\mathbf{r})d\mathbf{r}d\hat{\mathbf{s}} = \int_{\Omega}\int_{2\pi}\phi_3^*(\mathbf{r},\hat{\mathbf{s}})\delta(\mathbf{r}-\mathbf{r}_s)P_s(\hat{\mathbf{s}})d\mathbf{r}d\hat{\mathbf{s}}, \tag{5.12}$$

where $\phi_{1,2,3}$ and $\phi^*_{1,2,3}$ are the forward and adjoint radiances from computing $\mathcal{L}\mathbf{x}_{1,2,3}$ and $\mathcal{L}^*\mathbf{y}_{1,2,3}$, respectively. Eqs. (5.10), (5.11) and (5.12) simplify to

$$\Phi_1(\mathbf{r}_d) = \Phi_1^*(\mathbf{r}_s), \tag{5.13}$$

$$\int_{2\pi}\phi_2(\mathbf{r}_d,\hat{\mathbf{s}})P_d(\mathbf{r}_d,\hat{\mathbf{s}})d\hat{\mathbf{s}} = \Phi_2^*(\mathbf{r}_s), \tag{5.14}$$

$$2\pi\int_{\Omega}\Phi_3(\mathbf{r})P_d(\mathbf{r})d\mathbf{r} = \int_{2\pi}\phi_3^*(\mathbf{r}_d)P_s(\hat{\mathbf{s}})d\hat{\mathbf{s}}. \tag{5.15}$$

It can be seen from Eq. (5.13) that the case where a pair of isotropic δ-functions are used for $\mathbf{x}_1$ and $\mathbf{y}_1$, that we expect the resulting fluence values at their respective positions, $\Phi_1(\mathbf{r}_d)$ and $\Phi_1^*(\mathbf{r}_s)$, to be equal. This is an intuitive result given the reciprocity of the RTE and the angular indpendence of the source-detector combination.

Simulations were performed using a domain 40mm$\times$40mm in size with 101 pixels along each dimension; 10 Fourier harmonics were used and each source distribution emitted 10^6 photons. $\mathbf{r}_s$ was at the centre of the domain with $\mathbf{r}_d$ moved along the x-direction across the domain at y=3.6mm from the boundary. The left- and right-hand sides of Eq. (5.13) are plotted in Fig. 5.1 as a function of $r_{d,x}$ and it can be seen that indeed the fluences are equal for the range of $\mathbf{r}_d$ considered. Note that the plots for $\mathcal{L}^*\mathbf{y}_{1,2,3}$ are consistently noisier than for $\mathcal{L}\mathbf{x}_{1,2,3}$ because $\mathcal{L}^*\mathbf{y}_{1,2,3}$ must be evaluated for several positions $\mathbf{r}_d$ meaning the MC model must be re-run for each position. In subsequent cases a drop in SNR is noticeable which is due to a decrease in signal from introducing angular dependence in the source/detector.

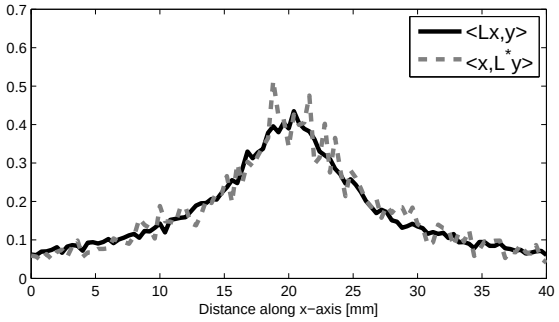

FIGURE 5.1: Plot of $\langle \mathcal{L}\mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathcal{L}^*\mathbf{y} \rangle$ for validation of adjoint model. Plot was produced with $\mathbf{x}$ and $\mathbf{y}$ as two isotropic point sources, with $\mathbf{x}$ at the centre of the domain and $\mathbf{y}$ translated across the domain at y=23.6mm from the boundary.

However, it is attractive to demonstrate that the equality in Eq. (5.6) holds when an anisotropic detector is used. The functions chosen for $\mathbf{x}_2$ and $\mathbf{y}_2$ were: $\mathbf{x}_2 = \delta(\mathbf{r} - \mathbf{r}_s)$ and $\mathbf{y}_2 = \delta(\mathbf{r} - \mathbf{r}_s)\frac{1}{\pi}\sin^2(2\theta)$. The angular portion of $\mathbf{y}_2$ is plotted in Fig. 5.2(a). $\mathbf{r}_s$ was again placed at the centre of the domain and with $\mathbf{r}_d$ moved along the x-direction across the domain at y=23.6mm. The two sides of the equality are plotted as a function of $r_{d,x}$ in Fig. 5.2(b), where again good agreement is obtained, suggesting the MC model is correctly modelling $\mathcal{L}^*$.

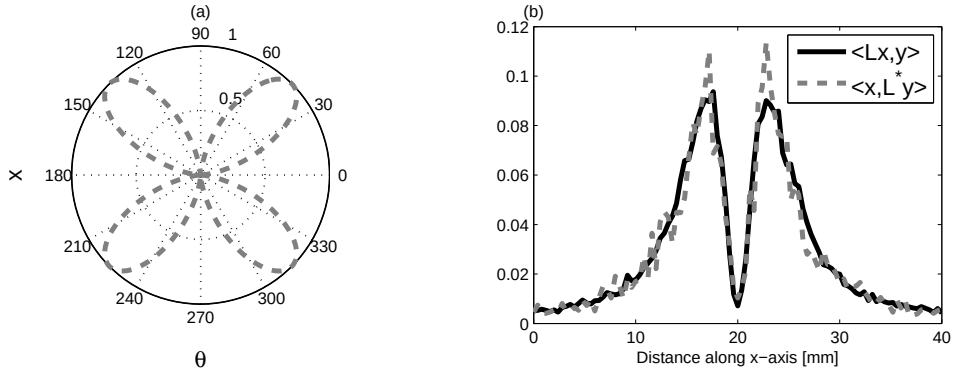

FIGURE 5.2: (a) Polar plot of source distribution for $\mathbf{y}_2 = \delta(\mathbf{r} - \mathbf{r}_s)\frac{1}{\pi}\sin^2(2\theta)$; (b) Plot of $\langle \mathcal{L}\mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathcal{L}^*\mathbf{y} \rangle$ for validation of adjoint model. Plot was produced with $\mathbf{x}$ as an isotropic point source at the centre of the domain and $\mathbf{y}$ as a point source with angular dependence following $\frac{1}{\pi}\sin^2(2\theta)$. $\mathbf{y}$ was translated across the domain along a line 4.0mm from the centre of the domain.

In practice, while correct modelling of the angular portion of $\mathcal{L}^*$ is important, the computation of functional gradients is performed using a source which is isotropic and

distributed in space over the domain as $P_s(\mathbf{r})$. $P_s(\mathbf{r})$ is plotted in Fig. 5.3(a). The 'detector' was the same as that above, and was again translated across the domain, with the resulting left- and right-hand sides of Eq. (5.15) plotted in Fig. 5.3(b); good agreement is observed once again.

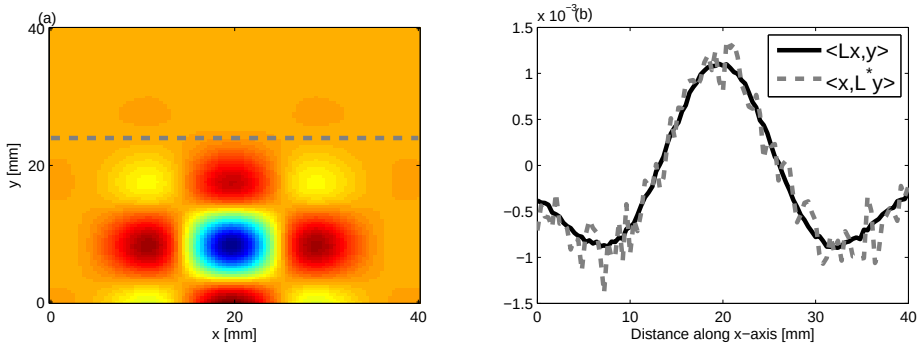

FIGURE 5.3: (a) Plot of the isotropic detector distribution $\mathbf{y}_3 = P_d(\mathbf{r})$; (b) Plot of $\langle \mathcal{L}\mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathcal{L}^*\mathbf{y} \rangle$ for validation of adjoint model. Plot was produced with $\mathbf{x}$ as anisotropic point source emitting light over angle following $\frac{1}{\pi}\sin^2(2\theta)$ and $\mathbf{y}$ was a distributed-source emitting light isotropically. $\mathbf{x}$ was translated along a line across the domain at y=4.0mm from the centre of the domain, as shown by the grey line dashed line in (a).

It is worth noting that good agreement was obtained in Case 1 and Case 2, without any scaling of the left- and right-hand sides, as the angular portion of the source and detector were properly normalised and spatially they were just δ-functions, meaning the consideration of the basis functions in space, u_j, and in angle, s_j, in the integrals was not necessary. However, in Case 3, $\mathbf{y}_3$ is a function of space, meaning that the integral over Ω of the basis functions requires proper consideration. Substituting the fields projected into their respective bases into Eq. (5.12) yields

$$\sum_{i,j,k,l} \int_{\Omega} \int_{2\pi} \phi_3(\mathbf{r}, \hat{\mathbf{s}})_{i,j} u(\mathbf{r})_i s(\mathbf{r})_j P_d(\mathbf{r})_{k,l} u(\mathbf{r})_k s(\mathbf{r})_l d\mathbf{r} d\hat{\mathbf{s}}$$
$$= \sum_{i,j,k,l} \int_{\Omega} \int_{2\pi} \phi_3^*(\mathbf{r}, \hat{\mathbf{s}})_{i,j} u(\mathbf{r})_i s(\mathbf{r})_j \delta(\mathbf{r} - \mathbf{r}_s) P_s(\hat{\mathbf{s}}) u(\mathbf{r})_k s(\mathbf{r})_l d\mathbf{r} d\hat{\mathbf{s}}. \quad (5.16)$$

The integral over space of the basis functions, recalling that a structured grid was used, produces a factor of (number of elements) $\times\, |d\mathbf{r}|$. The element area $|d\mathbf{r}|$ cancels as it appears on both sides in Eq. (5.16), but $N_{vox} =$ (number of elements) only appears on

the left-hand side and does not cancel, as shown in Eq. (5.17)

$$2\pi N_{vox} \sum_{i,j} \int_{\Omega} \phi_3(\mathbf{r}, \hat{\mathbf{s}})_i P_d(\mathbf{r})_i u(\mathbf{r})_i d\mathbf{r}$$
$$= 2\pi \sum_{i,j} \int_{2\pi} \phi_3^*(\mathbf{r}_s, \hat{\mathbf{s}})_{i,j} P_s(\mathbf{r}_s)_{i,j} \int_{\Omega} u(\mathbf{r})_i s(\mathbf{r})_j d\mathbf{r} d\hat{\mathbf{s}}, \tag{5.17}$$

A fourth case, which most closely simulates the conditions present in the inversions in Chapter 6, involves collimated illumination with a line source on one side of the domain and an isotropic internal adjoint source distribution. Under these conditions, $\mathbf{x}$ consists of a δ-function in angle, $\mathbf{x}_4 = P_s(x)\delta(\theta)$, while $\mathbf{y}_4$ is just a distribution of isotropic point sources through the domain, $\mathbf{y}_4 = P_d(\mathbf{r})$. The issue with the δ-function in angle is that very few photons have exactly this direction at the source position on the boundary; this means that estimates of $\langle \mathbf{x}_4, \mathcal{L}^* \mathbf{y}_4 \rangle$ are very noisy as few photons reaching the line source are travelling at exactly $\theta = 0$ (of course this is only problematic when evaluating Eq. (5.6), which is not carried out when evaluating functional gradients). In order to improve the SNR in this case many orders of magnitude more photons would need to be simulated when applying $\mathcal{L}^*$ to $\mathbf{y}_4$. Initial simulations of this fourth case showed too much noise to determine whether there was good agreement between $\langle \mathbf{y}_4, \mathcal{L}\mathbf{x}_4 \rangle$ and $\langle \mathbf{x}_4, \mathcal{L}^* \mathbf{y}_4 \rangle$, even with the acceptance angle of $\mathbf{x}_4$ expanded by ± 5 degrees. This data is therefore not presented here, but Cases 1–3 are an adequate demonstration that $\mathcal{L}^*$ is indeed the adjoint of $\mathcal{L}$.

5.3 Adjoint-assisted gradient computation

Functional gradients can be calculated using the adjoint radiance by differentiating the expression for the error functional with respect to the model parameters $\chi = (\mu_\mathrm{a}, \mu_\mathrm{s})$ [63, 107]. Starting with the definition in Eq. (5.1), this function can be re-expressed

with an integral over the domain Ω and Γ set to one with no loss of generality:

$$\epsilon = \frac{1}{2} \left|\left| H^{meas} - H(\mu_{\mathrm{a}}, \mu_{\mathrm{s}}) \right|\right|^2 \tag{5.18}$$

$$= \frac{1}{2} \int_{\Omega} \left(H^{meas} - H(\mu_{\mathrm{a}}, \mu_{\mathrm{s}}) \right)^2 d\mathbf{r}. \tag{5.19}$$

Taking the derivative with respect to μ_{a} in the r^{th} voxel gives

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} = - \int_{\Omega} \frac{\partial H}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} \left(H^{meas} - H \right) d\mathbf{r}, \tag{5.20}$$

which can be re-written using the product rule and the definition $H = \mu_{\mathrm{a}} \Phi$,

$$\frac{\partial H}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} = \frac{\partial \mu_{\mathrm{a}}}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} \Phi + \frac{\partial \Phi}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} \mu_{\mathrm{a}}. \tag{5.21}$$

Substituting Eq. (5.21) into (5.20) and using the approximation that the basis function derivative in the voxel-wise basis is just a delta-function, such that $\int_{\Omega} \frac{\partial \mu_{\mathrm{a}}}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} \Phi d\mathbf{r} = \delta(\mathbf{x} - \mathbf{x}^{\mathbf{r}}) \Phi$, yields

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} = \Phi \delta(\mathbf{x} - \mathbf{x}^{\mathbf{r}}) \left(H^{meas,r} - H^{\mathbf{r}} \right) - \int_{\Omega} \frac{\partial \Phi}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} \left(H^{meas} - H \right) d\mathbf{r}. \tag{5.22}$$

The sensitivity of the fluence to changes in the absorption is generally unknown, but can be estimated by differentiating the RTE in Eq. (5.23) with respect to the absorption coefficient

$$(\hat{\mathbf{s}} \cdot \nabla + \mu_{\mathrm{a}} + \mu_{\mathrm{s}}) \phi(\hat{\mathbf{s}}) = \mu_{\mathrm{s}} \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}', \hat{\mathbf{s}}) \phi(\hat{\mathbf{s}}') d\hat{\mathbf{s}}' + q(\hat{\mathbf{s}}). \tag{5.23}$$

$$(\hat{\mathbf{s}} \cdot \nabla + \mu_{\mathrm{a}} + \mu_{\mathrm{s}}) \frac{\partial \phi(\hat{\mathbf{s}})}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} - \mu_{\mathrm{s}} \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}', \hat{\mathbf{s}}) \frac{\partial \phi(\hat{\mathbf{s}}')}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} d\hat{\mathbf{s}}' = -\phi(\hat{\mathbf{s}}) \delta(\mathbf{x} - \mathbf{x}^{\mathbf{r}}). \tag{5.24}$$

At this point, the adjoint model is introduced by calculating the product of ϕ^* with Eq. (5.24) and subtracting the product of $-\partial\phi/\partial\mu_{\mathrm{a}}^{\mathbf{r}}$ with Eq. (5.5), where ϕ^* is the solution

to the adjoint RTE in Eq. (5.5) with source term $\mu_{\mathrm{a}}\left(H^{meas}-H(\mu_{\mathrm{a}},\mu_{\mathrm{s}})\right)$:

$$\begin{aligned}
&\phi^*(\hat{\mathbf{s}})(\hat{\mathbf{s}}\cdot\nabla)\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{\mathbf{r}}}+\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{\mathbf{r}}}(\hat{\mathbf{s}}\cdot\nabla)\phi^*(\hat{\mathbf{s}})\\
-\phi^*(\hat{\mathbf{s}})\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\frac{\partial\phi(\hat{\mathbf{s}}')}{\partial\mu_{\mathrm{a}}^{\mathbf{r}}}d\hat{\mathbf{s}}'+\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{\mathbf{r}}}\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}')d\hat{\mathbf{s}}'&\\
=-\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}})\delta(\mathbf{x}-\mathbf{x}^{\mathbf{r}})-\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{\mathbf{r}}}\mu_{\mathrm{a}}\left(H^{meas}-H(\mu_{\mathrm{a}},\mu_{\mathrm{s}})\right).&\qquad(5.25)
\end{aligned}$$

Eq. (5.25) can be simplified by noting that the left-hand side is equal to zero. To show this, begin by integrating Eq. (5.25) over angle and space,

$$\begin{aligned}
&\int_\Omega\int_{\mathcal{S}^{n-1}}\phi^*(\hat{\mathbf{s}})(\hat{\mathbf{s}}\cdot\nabla)\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{\mathbf{r}}}d\hat{\mathbf{s}}d\mathbf{r}+\int_\Omega\int_{\mathcal{S}^{n-1}}\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{r}}(\hat{\mathbf{s}}\cdot\nabla)\phi^*(\hat{\mathbf{s}})d\hat{\mathbf{s}}d\mathbf{r}\\
&\quad-\int_\Omega\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\phi^*(\hat{\mathbf{s}})\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\frac{\partial\phi(\hat{\mathbf{s}}')}{\partial\mu_{\mathrm{a}}^{r}}d\hat{\mathbf{s}}'\,\mathrm{d}\hat{\mathbf{s}}d\Omega\\
&\quad+\int_\Omega\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{r}}\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}')d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega\\
&\quad=-\int_\Omega\delta(\mathbf{x}-\mathbf{x}^{\mathbf{r}})\int_{\mathcal{S}^{n-1}}\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}})d\hat{\mathbf{s}}d\Omega\\
&\quad-\int_\Omega\mu_{\mathrm{a}}(H^{meas}-H)\int_{\mathcal{S}^{n-1}}\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{r}}d\hat{\mathbf{s}}d\Omega.\qquad(5.26)
\end{aligned}$$

Using an alternative form of the divergence theorem

$$\int_\Omega ab\cdot\nabla cd\Omega+\int_\Omega cb\cdot\nabla ad\Omega=\int_{\partial\Omega}b\cdot\hat{n}acd\Omega,\qquad(5.27)$$

with

$$a=\phi^*(\hat{\mathbf{s}}),\quad b=\hat{\mathbf{s}}\quad\text{and}\quad c=\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^{r}},$$

allows the simplification of Eq. (5.26) to

$$\begin{aligned}
&\int_{\partial\Omega}\int_{\mathcal{S}^{n-1}}(\hat{\mathbf{s}}\cdot\hat{\mathbf{n}})\phi^*(\hat{\mathbf{s}})\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^r}d\hat{\mathbf{s}}d\Omega\\
-&\int_{\Omega}\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\phi^*(\hat{\mathbf{s}})\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\frac{\partial\phi(\hat{\mathbf{s}}')}{\partial\mu_{\mathrm{a}}^r}d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega\\
+&\int_{\Omega}\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^r}\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}')d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega\\
&=-\int_{\Omega}\delta(\mathbf{x}-\mathbf{x}^{\mathbf{r}})\int_{\mathcal{S}^{n-1}}\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}})d\hat{\mathbf{s}}d\Omega\\
&-\int_{\Omega}\mu_{\mathrm{a}}(H^{meas}-H)\int_{\mathcal{S}^{n-1}}\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^r}d\hat{\mathbf{s}}d\Omega. \qquad (5.28)
\end{aligned}$$

The boundary conditions state that the adjoint radiance, ϕ^*, and derivative of the radiance, $\partial\phi^*/\partial\mu_{\mathrm{a}}^{\mathbf{r}}$, go to zero on the boundary meaning the first term in Eq. (5.28) is equal to zero, giving

$$\begin{aligned}
-&\int_{\Omega}\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\phi^*(\hat{\mathbf{s}})\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\frac{\partial\phi(\hat{\mathbf{s}}')}{\partial\mu_{\mathrm{a}}^r}d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega\\
+&\int_{\Omega}\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^r}\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}')d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega\\
&=-\int_{\Omega}\delta(\mathbf{x}-\mathbf{x}^{\mathbf{r}})\int_{\mathcal{S}^{n-1}}\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}})d\hat{\mathbf{s}}d\Omega\\
&-\int_{\Omega}\mu_{\mathrm{a}}(H^{meas}-H)\int_{\mathcal{S}^{n-1}}\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^r}d\hat{\mathbf{s}}d\Omega. \qquad (5.29)
\end{aligned}$$

The aforementioned symmetry in the phase function (see Section 5.2) means that

$$\int_{\Omega}\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}})\frac{\partial\phi(\hat{\mathbf{s}}')}{\partial\mu_{\mathrm{a}}^r}d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega \qquad (5.30)$$

$$=\int_{\Omega}\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\int_{\mathcal{S}^{n-1}}P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}')\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^r}d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega. \qquad (5.31)$$

and therefore

$$\int_{\Omega}\mu_{\mathrm{a}}(H^{meas}-H)\int_{\mathcal{S}^{n-1}}\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{a}}^r}d\hat{\mathbf{s}}d\Omega=-\int_{\Omega}\delta(\mathbf{x}-\mathbf{x}^{\mathbf{r}})\int_{\mathcal{S}^{n-1}}\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}})d\hat{\mathbf{s}}d\Omega. \qquad (5.32)$$

Recalling the definition that $\Phi = \int_{\mathcal{S}^{n-1}} \phi(\hat{\mathbf{s}}) d\hat{\mathbf{s}}$, allows re-arranging of Eq. (5.32) to

$$\begin{aligned}\int_\Omega \frac{\partial \Phi}{\partial \mu_{\mathrm{a}}^{\mathbf{r}}} \mu_{\mathrm{a}} (H^{meas} - H) d\mathbf{r} &= -\int_\Omega \delta(\mathbf{x} - \mathbf{x}^{\mathbf{r}}) \int_{\mathcal{S}^{n-1}} \phi^*(\hat{\mathbf{s}}) \phi(\hat{\mathbf{s}}) d\hat{\mathbf{s}} d\Omega \\ &= \int_{\mathcal{S}^{n-1}} \phi^{*,r}(\hat{\mathbf{s}}) \phi^{\mathbf{r}}(\hat{\mathbf{s}}) d\hat{\mathbf{s}}, \end{aligned} \tag{5.33}$$

where $\phi^{\mathbf{r}}$ and $\phi^{*,r}$ are the radiance and adjoint radiance at position r, respectively. Substitituing into Eq. (5.20) yields an expression for the functional gradient with respect to absorption at any position in the domain:

$$\frac{\partial \varepsilon}{\partial \mu_{\mathrm{a}}} = -\Phi(H^{meas} - H) + \int_{\mathcal{S}^{n-1}} \phi^*(\hat{\mathbf{s}}) \phi(\hat{\mathbf{s}}) d\hat{\mathbf{s}}. \tag{5.34}$$

Using the derivative of the error function as the starting point, the sensitivity of this function with respect to the scattering coefficient can be obtained using a similar procedure as for the absorption coefficient,

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{s}}^{\mathbf{r}}} = -\int_\Omega \frac{\partial H}{\partial \mu_{\mathrm{s}}^{\mathbf{r}}} (H^{meas} - H) \, d\mathbf{r}. \tag{5.35}$$

Differentiating the RTE with respect to μ_{s} gives

$$(\hat{\mathbf{s}} \cdot \nabla + \mu_{\mathrm{a}} + \mu_{\mathrm{s}}) \frac{\partial \phi(\hat{\mathbf{s}})}{\partial \mu_{\mathrm{s}}^{\mathbf{r}}} - \mu_{\mathrm{s}} \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}', \hat{\mathbf{s}}) \frac{\partial \phi(\hat{\mathbf{s}}')}{\partial \mu_{\mathrm{s}}^{\mathbf{r}}} d\hat{\mathbf{s}}' = -\phi(\hat{\mathbf{s}}) \delta(\mathbf{x} - \mathbf{x}^{\mathbf{r}}), \tag{5.36}$$

which, in conjunction with the adjoint equation in Eq. (5.37),

$$\begin{aligned}&(-\hat{\mathbf{s}} \cdot \nabla + \mu_{\mathrm{a}}(\mathbf{x}) + \mu_{\mathrm{s}}(\mathbf{x})) \phi^*(\mathbf{x}, \hat{\mathbf{s}}) \\ &= \mu_{\mathrm{s}}(\mathbf{x}) \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}', \hat{\mathbf{s}}) \phi^*(\mathbf{x}, \hat{\mathbf{s}}') d\hat{\mathbf{s}}' + \mu_{\mathrm{a}} \left(H^{meas} - H(\mu_{\mathrm{a}}, \mu_{\mathrm{s}}) \right), \end{aligned} \tag{5.37}$$

allows us to evaluate $\phi^* \times$(5.36) $-\partial\phi/\partial\mu_{\mathrm{s}}^{\mathbf{r}} \times$(5.5)

$$\begin{aligned}
&\phi^*(\hat{\mathbf{s}})(\hat{\mathbf{s}}\cdot\nabla)\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r} + \frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}(\hat{\mathbf{s}}\cdot\nabla)\phi^*(\hat{\mathbf{s}}) \\
-\phi^*(\hat{\mathbf{s}})\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\frac{\partial\phi(\hat{\mathbf{s}}')}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}' &+ \frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}\mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}')d\hat{\mathbf{s}}' \\
= \left(\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}}')d\hat{\mathbf{s}}' - \phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}})\right)&\delta(\mathbf{x}-\mathbf{x}^{\mathbf{r}}) - \qquad (5.38)\\
&\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}\mu_{\mathrm{a}}(H^{meas}-H). \qquad (5.39)
\end{aligned}$$

Integrating both sides over angle, $\hat{\mathbf{s}}$, and the domain, Ω,

$$\begin{aligned}
&\int_\Omega\int_{\mathcal{S}^{n-1}} \phi^*(\hat{\mathbf{s}})(\hat{\mathbf{s}}\cdot\nabla)\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}d\Omega + \int_\Omega\int_{\mathcal{S}^{n-1}} \frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}(\hat{\mathbf{s}}\cdot\nabla)\phi^*(\hat{\mathbf{s}})d\hat{\mathbf{s}}d\Omega \\
&\quad - \int_\Omega \mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}} \phi^*(\hat{\mathbf{s}})\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\frac{\partial\phi(\hat{\mathbf{s}}')}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega \\
&\quad + \int_\Omega \mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}} \frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}')d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega \\
&= \int_\Omega\int_{\mathcal{S}^{n-1}} \left(\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}}')d\hat{\mathbf{s}}' - \phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}})\right)d\hat{\mathbf{s}}\ \delta(\mathbf{x}-\mathbf{x}^{\mathbf{r}})d\Omega \\
&\quad - \int_\Omega \mu_{\mathrm{a}}(H^{meas}-H)\int_{\mathcal{S}^{n-1}} \frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}d\Omega. \qquad (5.40)
\end{aligned}$$

Using the statement of the divergence theorem in Eq. (5.27) to write

$$\begin{aligned}
&\int_\Omega\int_{\mathcal{S}^{n-1}} \phi^*(\hat{\mathbf{s}})(\hat{\mathbf{s}}\cdot\nabla)\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}d\Omega + \int_\Omega\int_{\mathcal{S}^{n-1}} \frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}(\hat{\mathbf{s}}\cdot\nabla)\phi^*(\hat{\mathbf{s}})d\hat{\mathbf{s}}d\Omega \\
&\quad = \int_{\partial\Omega}\int_{\mathcal{S}^{n-1}} (\hat{\mathbf{s}}\cdot\hat{\mathbf{n}})\phi^*(\hat{\mathbf{s}})\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}d\Omega. \qquad (5.41)
\end{aligned}$$

we find that the first two terms of Eq. (5.40) are equal to zero as the fluence (and its derivative with respect to μ_{s}) are zero on the boundary. Thus, Eq. (5.40) reduces to

$$\begin{aligned}
&\int_\Omega \mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}})\frac{\partial\phi(\hat{\mathbf{s}}')}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}'\ \mathrm{d}\hat{\mathbf{s}}d\Omega \\
&\quad - \int_\Omega \mu_{\mathrm{s}}\int_{\mathcal{S}^{n-1}}\int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}}')\frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}'d\hat{\mathbf{s}}d\Omega \\
&= \int_\Omega\int_{\mathcal{S}^{n-1}} \left(\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}}) - \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}})\phi(\hat{\mathbf{s}}')d\hat{\mathbf{s}}'\right)d\hat{\mathbf{s}}\ \delta(\mathbf{x}-\mathbf{x}^{\mathbf{r}})d\Omega \\
&\quad + \int_\Omega \mu_{\mathrm{a}}(H^{meas}-H)\int_{\mathcal{S}^{n-1}} \frac{\partial\phi(\hat{\mathbf{s}})}{\partial\mu_{\mathrm{s}}^r}d\hat{\mathbf{s}}d\Omega. \qquad (5.42)
\end{aligned}$$

Once again exploiting the symmetry of the phase function, the left-hand side of Eq. (5.42) is equal to zero and

$$
\begin{aligned}
&- \int_{\Omega} \mu_{\mathrm{a}}(H^{meas} - H) \int_{\mathcal{S}^{n-1}} \frac{\partial \phi(\hat{\mathbf{s}})}{\partial \mu_{\mathrm{a}}^{r}} d\hat{\mathbf{s}} d\Omega \\
&= \int_{\Omega} \int_{\mathcal{S}^{n-1}} \left(\phi^{*}(\hat{\mathbf{s}}) \phi(\hat{\mathbf{s}}) - \int_{\mathcal{S}^{n-1}} P(\hat{\mathbf{s}}, \hat{\mathbf{s}}') \phi^{*}(\hat{\mathbf{s}}) \phi(\hat{\mathbf{s}}') d\hat{\mathbf{s}}' \right) d\hat{\mathbf{s}}\ \delta(\mathbf{x} - \mathbf{x}^{\mathbf{r}}) d\Omega. \qquad (5.43)
\end{aligned}
$$

It follows that

$$
\begin{aligned}
&- \int_{\Omega} \frac{\partial \Phi}{\partial \mu_{\mathrm{s}}^{r}} \mu_{\mathrm{a}}(H^{meas} - H) d\Omega \qquad (5.44) \\
&= \int_{\mathcal{S}^{n-1}} \phi^{*,r}(\hat{\mathbf{s}}) \phi^{\mathbf{r}}(\hat{\mathbf{s}}) d\hat{\mathbf{s}} - \int_{\mathcal{S}^{n-1}} \int_{\mathcal{S}^{n-1}} \phi^{*,r} P(\hat{\mathbf{s}}, \hat{\mathbf{s}}') \phi^{\mathbf{r}}(\hat{\mathbf{s}}') d\hat{\mathbf{s}}' d\hat{\mathbf{s}}, \qquad (5.45)
\end{aligned}
$$

which can be substituted into Eq. (5.35) to obtain the expression for the gradient of the error function with respect to the scattering coefficient at any position:

$$
\frac{\partial \varepsilon}{\partial \mu_{\mathrm{s}}} = \int_{\mathcal{S}^{n-1}} \phi^{*}(\hat{\mathbf{s}}) \phi(\hat{\mathbf{s}}) d\hat{\mathbf{s}} - \int_{\mathcal{S}^{n-1}} \int_{\mathcal{S}^{n-1}} \phi^{*}(\hat{\mathbf{s}}) P(\hat{\mathbf{s}}, \hat{\mathbf{s}}') \phi(\hat{\mathbf{s}}') d\hat{\mathbf{s}}' d\hat{\mathbf{s}}. \qquad (5.46)
$$

5.3.1 Expressions using Fourier series

Recall that the MC code presented in Chapter 4 directly computes Fourier coefficients of the radiance, rather than the angularly resolved radiance required in the expressions for the functional gradients with respect to absorption and scattering obtained in the previous section. As such, we wish to evaluate Eqs. (5.34) and (5.46) using the expressions for the radiance and adjoint radiance in the Fourier domain:

$$
\phi(\theta') = \frac{1}{2\pi} a_0 + \frac{1}{\pi} \sum_{n=1}^{N=\infty} a_n \cos(n\theta') + \frac{1}{\pi} \sum_{n=1}^{N=\infty} b_n \sin(n\theta'), \qquad (5.47)
$$

and

$$
\phi^{*}(\theta') = \frac{1}{2\pi} a_0^{*} + \frac{1}{\pi} \sum_{m=1}^{M=\infty} a_m^{*} \cos(m\theta') + \frac{1}{\pi} \sum_{m=1}^{M=\infty} b_m^{*} \sin(m\theta'). \qquad (5.48)
$$

where a_0^*, a_n^* and b_n^* are the Fourier coefficients of the adjoint radiance. Substituting these expressions into the equations for the functional gradients (Eqs. (5.34) & (5.46)),

$$\begin{aligned}\frac{\partial \epsilon}{\partial \mu_{\mathrm{a}}} &= -\Phi(H^{meas} - H0) + \int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}})\phi^*(\hat{\mathbf{s}})d\hat{\mathbf{s}} \\ &= -a_0(H^{meas} - \mu_{\mathrm{a}}a_0) + \int_{2\pi}\left[\frac{1}{4\pi^2}a_0a_0^* + \frac{1}{2\pi^2}a_0\sum_{m=1}^{\infty}a_m^*\cos(m\theta') + \frac{1}{2\pi^2}a_0\sum_{m=1}^{\infty}a_m^*\sin(m\theta')\right. \\ &+\frac{1}{2\pi^2}a_0^*\sum_{n=1}^{\infty}a_n\cos(n\theta') + \sum_{n=1}^{\infty}\sum_{m=1}^{\infty}a_na_m^*\cos(n\theta')\cos(m\theta') + \sum_{n=1}^{\infty}\sum_{m=1}^{\infty}a_nb_m^*cos(n\theta')\sin(m\theta') + \\ &\left.\frac{1}{2\pi^2}a_0^*\sum_{n=1}^{\infty}b_m\cos(m\theta') + \sum_{n=1}^{\infty}\sum_{m=1}^{\infty}a_m^*b_n\sin(n\theta')cos(m\theta')\sum_n\sum_m b_nb_m^*\sin(n\theta')\sin(m\theta')\right]d\theta',\end{aligned} \tag{5.49}$$

noting that since $\Phi = \int_{\mathcal{S}^1}\phi(\hat{\mathbf{s}}')d\hat{\mathbf{s}}'$ and that integrals of the sine and cosine terms over angle equal zero, the fluence is equal to the zeroth order Fourier term of the radiance: $\Phi = a_0$. Using orthogonality, all terms for which $n \neq m$ integrate to zero and Eq. (5.49) reduces to

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{a}}} = -a_0(H^{meas} - \mu_{\mathrm{a}}a_0) + \int_{2\pi}\left[\frac{1}{4\pi^2}a_0a_0^* + \frac{1}{\pi^2}\sum_{n=1}^{\infty}a_na_n^*\cos^2(n\theta') + \frac{1}{\pi^2}\sum_{n=1}^{\infty}b_nb_n^*\sin^2(n\theta')\right]d\theta',$$

which is simply

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{a}}} = -a_0(H^{meas} - \mu_{\mathrm{a}}a_0) + \frac{1}{2\pi}a_0a_0^* + \frac{1}{\pi}\sum_{n=1}^{\infty}a_na_n^* + \frac{1}{\pi}\sum_{n=1}^{\infty}b_nb_n^*. \tag{5.50}$$

This closed form expression of the second term in Eq. (5.34) is computationally straightforward to evaluate due to the fact that we are simply summing over products of Fourier coefficients already loaded in memory, meaning numerical integration is not necessary. This is a significant advantage of having a 'frequency'-domain representation of the radiance because the evaluation of integrals numerically can be very expensive computationally. The current RMC implementation has undergone limited parallelisation meaning the forward and adjoint MC simulations are the bottleneck in gradient computation. However, the ability to distribute RMC over many cores and multiple nodes means that the MC simulations can be accelerated significantly, at which point the run-time bottleneck then becomes the actual computation of the gradients. However, as gradients

$\left(\frac{\partial \epsilon}{\partial \mu_a}\right)^{n_n}$ computed on $n_n = 1, ..., N_n$ nodes can simply be averaged over each node, the gradient computation can be performed locally (i.e. by each node) and transferred to a host node for averaging. This has the benefit that all memory- and compute-intensive operations of the gradient calculation are performed in a distributed manner. This benefit extends to the gradient computations using spherical harmonics as well as the calculation of the functional gradient with respect to scattering using Fourier/spherical harmonics dervied below.

The expression for the gradient with respect to scattering in Eq. (5.46) contains two terms; the first term is the same as the second term in Eq. (5.34), for which we already have an expression using Fourier coefficients. The second term in Eq. (5.46) is

$$\int_{\mathcal{S}^{n-1}} \int_{\mathcal{S}^{n-1}} \phi^*(\hat{\mathbf{s}}) P(\hat{\mathbf{s}}, \hat{\mathbf{s}}') \phi(\hat{\mathbf{s}}') d\hat{\mathbf{s}}' d\hat{\mathbf{s}},$$

which contains the phase function. The analogue to the Henyey-Greenstein phase function in 2D can be expanded using a Fourier series in powers of g, the anisotropy factor, [165]:

$$P_\theta(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'; g) = \frac{1}{2\pi} + \frac{1}{\pi} \sum_{l=1}^{\infty} g^l \cos(l\Delta\theta), \tag{5.51}$$

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

$$\int_{\mathcal{S}^1} \int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}}') P_\theta(\hat{\mathbf{s}}, \hat{\mathbf{s}}') \phi^*(\hat{\mathbf{s}}) d\hat{\mathbf{s}} d\hat{\mathbf{s}}' = \int_{\mathcal{S}^1} \int_{\mathcal{S}^1} \left[\frac{1}{2\pi} a_0 + \frac{1}{\pi} \sum_{n=1}^{\infty} a_n \cos(n\theta') + \frac{1}{\pi} \sum_{n=1}^{\infty} b_n \sin(n\theta') \right]$$
$$\left[\frac{1}{2\pi} + \frac{1}{\pi} \sum_{l=0}^{\infty} g^l \cos(l(\theta - \theta')) \right] \left[\frac{1}{2\pi} a_0^* + \frac{1}{\pi} \sum_{m=1}^{\infty} a_m^* \cos(m\theta) + \frac{1}{\pi} \sum_{m=1}^{\infty} b_m^* \sin(m\theta) \right] d\theta d\theta',$$

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

$$\int_{\mathcal{S}^1}\int_{\mathcal{S}^1}\phi(\hat{\mathbf{s}}')P_\theta(\hat{\mathbf{s}},\hat{\mathbf{s}}')\phi^*(\hat{\mathbf{s}})d\hat{\mathbf{s}}d\hat{\mathbf{s}}' = \frac{1}{2\pi}a_0a_0^* + \frac{1}{\pi}\sum_{n=1}^{\infty}a_na_n^*g^n + \frac{1}{\pi}\sum_{n=1}^{\infty}b_nb_n^*g^n, \quad (5.52)$$

Substituting this expression into Eq. (5.46), we can write the full expressions for the functional gradients with respect to the absorption and scattering coefficients:

$$\begin{aligned}\frac{\partial\epsilon}{\partial\mu_s} &= \frac{1}{2\pi}a_0a_0^* + \frac{1}{\pi}\sum_{n=1}^{\infty}a_na_n^* + \frac{1}{\pi}\sum_{n=1}^{\infty}b_nb_n^* \\ &-\frac{1}{2\pi}a_0a_0^* + \frac{1}{\pi}\sum_{n=1}^{\infty}a_na_n^*g^n + \frac{1}{\pi}\sum_{n=1}^{\infty}b_nb_n^*g^n \quad (5.53)\\ &= \frac{1}{\pi}\sum_{n=1}^{\infty}\left[a_na_n^* + b_nb_n^*\right](1-g^n) \quad (5.54)\end{aligned}$$

Apart from the simplicity of the expression for the gradient in Eq. (5.54), this way of calculating the functional gradient avoids the need for expensive numerical integration and can be achieved in a completely distributed manner, where the gradients from multiple nodes must only be averaged on a host node, with the actual computations being performed on a cluster.

5.3.2 Expressions using spherical harmonic series

The radiance in 3D can be expanded in a similar fashion to that in 2D, except using a series of spherical harmonics. Recall the expression from Chapter 4 (Eq. (5.56)) for the radiance,

$$\phi(\hat{\mathbf{s}}') = \sum_{l=0}^{\infty}\sum_{m=-l}^{l}\left(\frac{2l+1}{4\pi}\right)^{1/2}a_l^mY_l^m(\hat{\mathbf{s}}'), \quad (5.55)$$

then the adjoint radiance is given by

$$\phi^*(\hat{\mathbf{s}}') = \sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}\left(\frac{2l'+1}{4\pi}\right)^{1/2}a_{l'}^{m'*}Y_{l'}^{m'}(\hat{\mathbf{s}}'), \quad (5.56)$$

where a_l^m and $a_{l'}^{m'*}$ are the weights associated with the spherical harmonics in the forward and adjoint fields, respectively. These definitions can be substituted into the expressions for the functional gradients with respect to absorption and scattering, as above. Starting with the gradient with respect to absorption, noting that $\Phi = \int_{\mathcal{S}^2} \phi(\hat{\mathbf{s}}')d\hat{\mathbf{s}}' = a_0^0$,

$$\begin{aligned}\frac{\partial \epsilon}{\partial \mu_{\mathrm{a}}} &= -a_0^0(H^{meas} - \mu_{\mathrm{a}} a_0^0) + \int_{\mathcal{S}^1} \phi(\hat{\mathbf{s}})\phi^*(\hat{\mathbf{s}})d\hat{\mathbf{s}} \\ &= -a_0^0(H^{meas} - \mu_{\mathrm{a}} a_0^0)\end{aligned} \tag{5.57}$$

$$\int_{\mathcal{S}^2} \left[\sum_{l=0}^{\infty}\sum_{m=-l}^{l}\left(\frac{2l+1}{4\pi}\right)^{1/2} a_l^m Y_l^m(\hat{\mathbf{s}})\right]\left[\sum_{l'=0}^{\infty}\sum_{m'=-l'}^{l'}\left(\frac{2l'+1}{4\pi}\right)^{1/2} a_{l'}^{m'*} Y_{l'}^{m'}(\hat{\mathbf{s}})\right] d\hat{\mathbf{s}}. \tag{5.58}$$

Spherical harmonics satisfy the orthogonality relationship $\int_{\mathcal{S}^2} Y_l^m Y_{l'}^{m'+} d\hat{\mathbf{s}} = \delta_{mm'}\delta_{ll'}$; as such, cross-products that form as a result of expanding the summations over m, m', l and l' will integrate to zero when these indices are not equal. However, the integral over angle of products of terms with $m = -m'$ are non-zero because $Y_l^{m+} = Y_{-m}^l$ and $\int_{\mathcal{S}^2} Y_l^m Y_l^{m+} d\hat{\mathbf{s}} = \delta_{ll'}$. This leads to the gradient with respect to absorption being

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{a}}} = -a_0^0(H^{meas} - \mu_{\mathrm{a}} a_0^0) + \sum_{l=0}^{\infty}\sum_{m=0}^{l}\left(\frac{2l+1}{4\pi}\right) a_l^m a_l^{-m*}. \tag{5.59}$$

As observed in the gradient with respect to scattering in the Fourier domain, the two terms in this gradient can be expressed using sum of products of spherical harmonic coefficients. The first term is $\sum_{l=0}^{\infty}\sum_{m=0}^{l}(-1)^m\left(\frac{2l+1}{4\pi}\right) a_l^m a_l^{-m*}$ as in Eq. (5.59), while the second is obtained by substituting the spherical harmonic series expressions for the forward and adjoint radiances, and the spherical harmonic expansion, in powers of g, for the Henyey-Greenstein phase function [165]:

$$P(\hat{\mathbf{s}}, \hat{\mathbf{s}}') = \sum_{l''=0}^{\infty}\sum_{m''=-l''}^{l''} g^{l''} Y_{l''}^{m''+}(\hat{\mathbf{s}}')Y_{l''}^{m''}(\hat{\mathbf{s}}), \tag{5.60}$$

where g is the anisotropy factor and the superscript $+$ is used to indicate complex conjugation. Thus, the second term in Eq. (5.46) is given by

$$\begin{aligned}
&\int_{\mathcal{S}^{n-1}} \int_{\mathcal{S}^{n-1}} \phi^*(\hat{\mathbf{s}}) P(\hat{\mathbf{s}}, \hat{\mathbf{s}}') \phi(\hat{\mathbf{s}}') d\hat{\mathbf{s}}' d\hat{\mathbf{s}} \\
&= \left[\sum_{l'=0}^{\infty} \sum_{m'=-l'}^{l'} \left(\frac{2l'+1}{4\pi} \right)^{1/2} a_{l'}^{m'*} Y_{l'}^{m'}(\hat{\mathbf{s}}) \right] \\
&\quad \left[\sum_{l''=0}^{\infty} \sum_{m''=-l''}^{l''} g^{l''} Y_{l''}^{m''+}(\hat{\mathbf{s}}') Y_{l''}^{m''}(\hat{\mathbf{s}}) \right] \\
&\quad \left[\sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi} \right)^{1/2} a_l^m Y_l^m(\hat{\mathbf{s}}') \right].
\end{aligned} \tag{5.61}$$

The integrals above can be performed separately if we define an integral I_1

$$I_1(\hat{\mathbf{s}}') = \int_{\mathcal{S}^2} I_2(\hat{\mathbf{s}}') \left[\sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi} \right)^{1/2} a_l^m Y_l^m(\hat{\mathbf{s}}') \right] d\hat{\mathbf{s}}', \tag{5.62}$$

where

$$I_2 = \int_{\mathcal{S}^2} \left[\sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi} \right)^{1/2} a_{l'}^{m'*} Y_{l'}^{m'}(\hat{\mathbf{s}}) \right] \left[\sum_{l''=0}^{\infty} \sum_{m''=-l''}^{l''} g^{l''} Y_{l''}^{m''+}(\hat{\mathbf{s}}') Y_{l''}^{m''}(\hat{\mathbf{s}}) \right] d\hat{\mathbf{s}}. \tag{5.63}$$

Due to orthogonality of spherical harmonics, I_2 becomes

$$I_2(\hat{\mathbf{s}}') = \sum_{l'=0}^{\infty} \sum_{m'=-l'}^{l'} (-1)^{m'} \left(\frac{2l'+1}{4\pi} \right)^{1/2} g^{l'} a_{l'}^{m'*} Y_l^m(\hat{\mathbf{s}}'), \tag{5.64}$$

where the identity $Y_l^{m+} = (-1)^m Y_l^{-m}$ was used. Substituting I_2 back into the expression for I_1 gives

$$\begin{aligned}
I_1 &= \int_{\mathcal{S}^2} \left[\sum_{l'=0}^{\infty} \sum_{m'=-l'}^{l'} (-1)^{m'} \left(\frac{2l'+1}{4\pi} \right)^{1/2} g^{l'} a_{l'}^{m'*} Y_l^m(\hat{\mathbf{s}}') \right] \left[\sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi} \right)^{1/2} a_l^m Y_l^m(\hat{\mathbf{s}}') \right] d\hat{\mathbf{s}}' \\
&= \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi} \right) g^l a_l^m a_l^{-m*},
\end{aligned} \tag{5.65}$$

Thus, the full expression for the gradient with respect to scattering is given by

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{s}}} = \sum_{l=0}^{\infty} \sum_{m=0}^{l} \left(\frac{2l+1}{4\pi} \right) a_l^m a_l^{-m*} - \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \left(\frac{2l+1}{4\pi} \right) g^l a_l^m a_l^{-m*}$$
$$= \sum_{l=0}^{\infty} \sum_{m=0}^{l} \left[\left(\frac{2l+1}{4\pi} \right) a_l^m a_l^{-m*} \right] \left(1 - g^l \right). \tag{5.66}$$

This expression mirrors that for the gradient with respect to scattering in 2D quite closely; the expression in 2D there is not contribution from the zeroth terms, i.e. the forward and adjoint fluence, and although not immediately obvious, this is also the case in 3D. The zeroth term has $l = 0$, which results in $\left(1 - g^0\right)$ in Eq. (5.66) and this equation could therefore just be written as

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{s}}} = \sum_{l=1}^{\infty} \sum_{m=0}^{l} \left[\left(\frac{2l+1}{4\pi} \right) a_l^m a_l^{-m*} \right] \left(1 - g^l \right).$$

5.3.3 Selecting the number of harmonics in 2D

When simulating the radiance due to a strongly directional or δ source, it is of course beneficial to include a large number of harmonics. The field is inherently less directional when the illumination is formed by internally distributed isotropic sources, meaning a large number of harmonics is not crucial as the contributions from the higher harmonics may make little difference to the radiance estimate. This section considers the question of how many harmonics are required in the forward and adjoint RMC simulations. This is worthwhile as the number of harmonics simulated affects the cost both algorithmically and in terms of computational memory; the memory costs can be significant (particularly in 3D where for every l there are $2l + 1$ harmonics; even for just $l = 10$, this leads to 120 additional harmonics needing to be calculated and stored which leads to tens of GB in storage for domains similar in size to high-resolution 3D PAT images!)

The question of how many harmonics are required to accurately represent radiance estimates is of particular relevance to the functional gradients. This is only treated here in 2D, but a similar argument applies for the 3D case using spherical harmonics. The gradient with respect to absorption in 2D in Eq. (5.50) is dominated by the first term

which depends on the difference between the measured data and the simulated data, so harmonics beyond the first are unlikely to have a significant impact on the gradient. On the other hand, the gradient with respect to scattering in Eq. (5.54) is not dominated by a single term and the decay rate is governed by the factor $(1 - g^n)$. Plotting the value of this factor against n for different values of g in Fig. 5.4, it can be observed that when $g > 0$ and tends towards unity, the low harmonics carry less weight whilst the higher ones maintain a large contribution to the gradient. However, when g is close, or equal, to zero harmonics are weighted more equally. The harmonics themselves however decay with increasing n as higher frequencies are attenuated rapidly due to the diffusive nature of scattering; when the light field is completely diffuse, the only non-zero coefficients remaining are those at $n = 0, 1$.

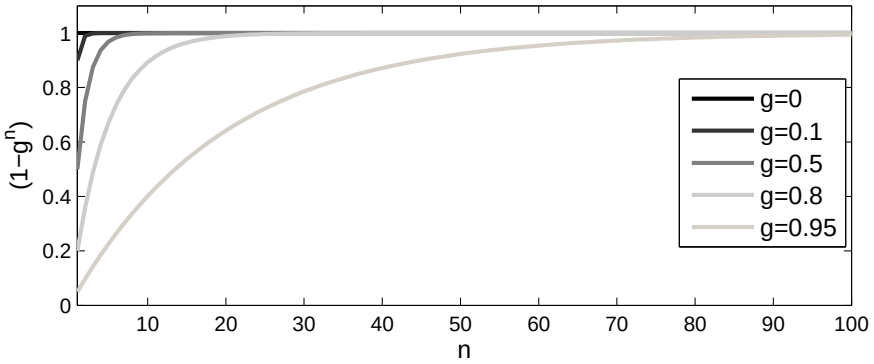

FIGURE 5.4: Plot of $(1 - g^n)$ as a function of n for 5 different values of g.

Thus, although it may seem from Fig. 5.4 that accurate gradient estimates in a domain with high anisotropy require a high number of harmonics to be simulated (possibly $\geq$100), this is not the case in general, at least far from boundaries and from the location of the source used in the forward simulation. This can be demonstrated by considering a plot of the harmonic a_n as a function of n (normalised to one for ease visualisation) on the same plot as the curve for $(1 - g^n)$. The simulation involved a 4mm$\times$4mm domain, containing 40 pixels along each dimension, with μ_a and μ_s homogeneously distributed and set to 0.01mm^{-1} and 5mm^{-1}, respectively, with $g = 0.95$. The illumination was a pencil beam at the centre of one axis consisting of 10^7 photons. Fig. 5.5 shows a plot of the absorbed energy density with arrows indicating the positions at which a_n, the $(1 - g^n)$ term and their product $a_n(1 - g^n)$ are plotted as a function of n.

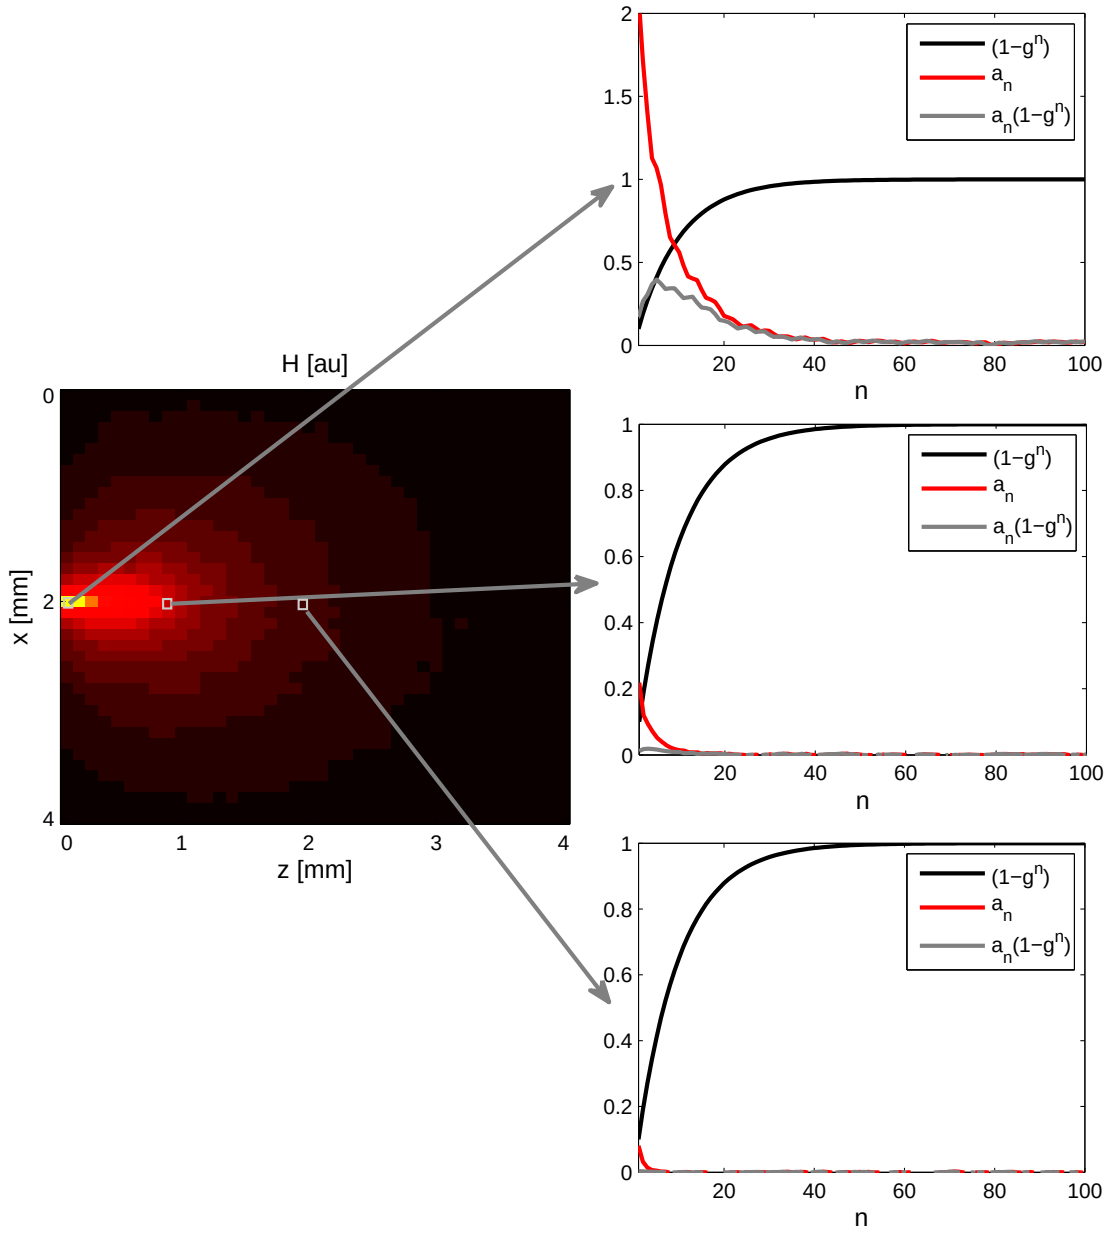

FIGURE 5.5: Absorbed energy distribution from a pencil beam illuminating a homogeneous medium. Plots of Fourier coefficients a_n, $(1-g^n)$ term and their product as a function of n at positions in H as indicated by the arrows showing that fewer harmonics are required with increasing depth when computing the functional gradient with respect to scattering.

As might be expected, the magnitude of a_n at $n = 0$ is greatest at the voxel nearest the source and it decays more slowly with increasing n than at the positions further from the source. This is due to there not only being more light near the source, as it will not have

yet been attenuated through absorption and scattering, but also frequency content of the radiance field will be highest near the source. This is particularly true for a_n compared with b_n due to the alignment of the Fourier series with the source direction ($\theta' = 0$). Despite the fact that in practice the gradient contains the term $a_n a_n^*(1-g^n)$, $a_n(1-g^n)$ can be used here for illustration since a_n^* contains predominantly low frequencies anyway. In the voxel nearest the source, $a_n(1-g^n)$ has fallen to 37% of the maximum value of a_n by $n = 20$ and only 9% of the maximum value of a_n remains when $n = 40$. In the voxel 1mm away from the source the contribution of $a_n(1-g^n)$ for $n = 20$ is only 0.7% of the maximum value and the contribution at 2mm away from the source is negligble. This is a compelling demonstration of the fact that few harmonics (estimated to be $\sim$3 from the second plot in Fig. 5.5) are required for accurate gradient calculation far from sources and boundaries ($\sim$1mm). Although only $n = 10$ was used in the validation of the gradient with respect to scattering above, this was sufficient in all regions apart from the pixels nearest the source given that the product of a_n and a_n^* will decay more quickly with increasing n than the two quantities individually.

5.3.4 Validation of functional gradients against finite differences in 2D

In this section, adjoint-assisted functional gradients in 2D are validated against gradients obtained using finite differences. In order for more accurate estimates of the gradients computed using finite differences, a central finite difference (FD) was used,

$$\left.\frac{\partial \epsilon}{\partial \chi}\right|_{\chi_0} \approx \frac{(\epsilon(\chi_0+\delta\chi)-\epsilon(\chi_0))-(\epsilon(\chi_0)-\epsilon(\chi_0-\delta\chi))}{2\delta\chi}. \tag{5.67}$$

5.3.4.1 Validation of the functional gradient with respect to the absorption coefficient

In this section, the word 'true' is used to refer to quantities that are either 'measured', such as the data H^{true}, or parameters used to form the data, such as $\mu_{\text{a}}^{\text{true}}$ and $\mu_{\text{s}}^{\text{true}}$. Two arbitrary distributions were chosen for the true absorption and scattering coefficients, as shown in the top row of Fig. 5.6. The RMC model was then used to simulate the data

using a line source at z=0mm consisting of 10^9 photons and just the zeroth harmonic, as shown in Fig. 5.6(c). H^0 is shown in Fig. 5.6(d) and was calculated assuming μ_s^{true} was known and a homogeneous initial guess of the absorption of 0.01mm^{-1}, which is referred to as μ_a^0. 10 harmonics were used in simulating H^0, and this number is used in the subsequent adjoint-assisted gradient calculations.

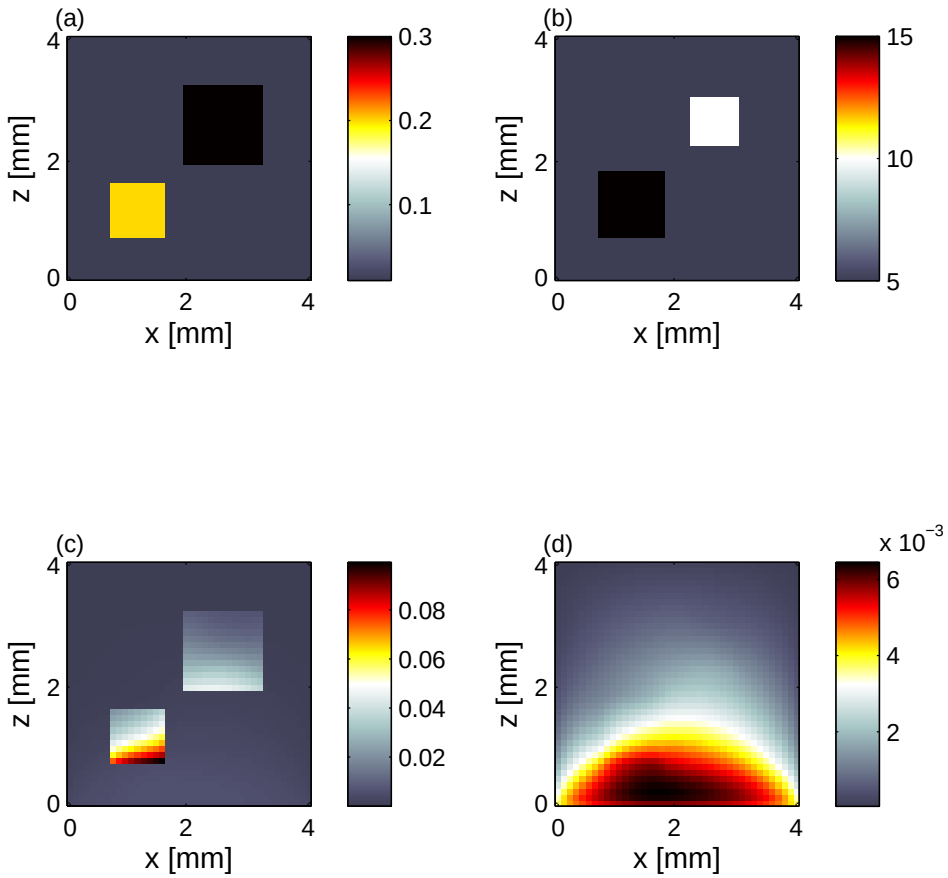

FIGURE 5.6: (a) True absorption coefficient, μ_a^{true}, with values between 0.01mm^{-1} and 0.3mm^{-1}; (b) True scattering coefficient, μ_s^{true}, with values between 5mm^{-1} and 15mm^{-1}; (c) 'Measured' data, H^{true}, in arbitrary units formed using RMC with 1 harmonic and 10^9 photons using true absorption and scattering coefficients; (d) Initial estimate of data, H^0, in the same units as (c), formed using RMC with 10 harmonics and 10^9 using homogeneous absorption coefficient of 0.01mm^{-1} and true scattering coefficient.

The adjoint source was calculated as $\mu_{\mathrm{a}}^{0}\left(H^{\mathrm{true}}-H^{0}\right)$ and is shown in Fig. 5.7(a), with the resulting fluence distribution, calculated using the adjoint RMC in (b). The adjoint model used 10 Fourier harmonics and 10^9 photons.

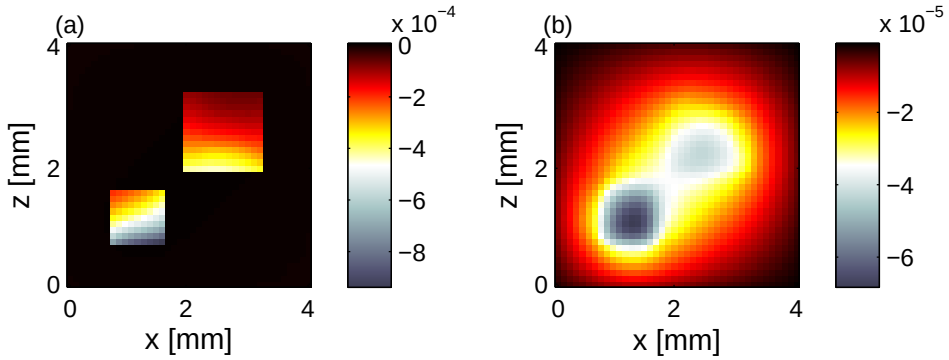

FIGURE 5.7: (a) Adjoint source, $q^{*}=\mu_{\mathrm{a}}^{0}(H^{\mathrm{true}}-H^{0})$, used for gradient calculation for absorption using RMC; (b) adjoint fluence from RMC, Φ^{*}, used to compute adjoint-assisted gradient with respect to absorption.

The most prominent feature in q^{*} in Fig. 5.7(a) is the inclusion in μ_{a} that was clearly visible in H^{true}, demonstrating a strong inherent sensitivity to differences in the absorption coefficient in the data. The resulting adjoint-assisted gradient is shown in Fig. 5.8(a). The FD calculation was run for all z between 0-4mm for x=1.5mm using RMC, where 10^9 photons and just the zeroth harmonic were simulated. The magnitude of the perturbation in the FD calculation, $\delta\mu_{\mathrm{a}}$, was 15% of the background μ_{a} value (0.01mm^{-1}). A profile through the domain at x=1.5mm in the adjoint-assisted gradient with respect to absorption was plotted with the FD gradient, as shown in Fig. 5.8(b). The computation time required for the adjoint-assisted gradient was $\sim$1 hour and 25 minutes on a Dell 2U R820 server with 32 cores, compared with 12 hours and 30 minutes for the FD calculation on the same system. All gradient calculations were run using UCL's Legion cluster.

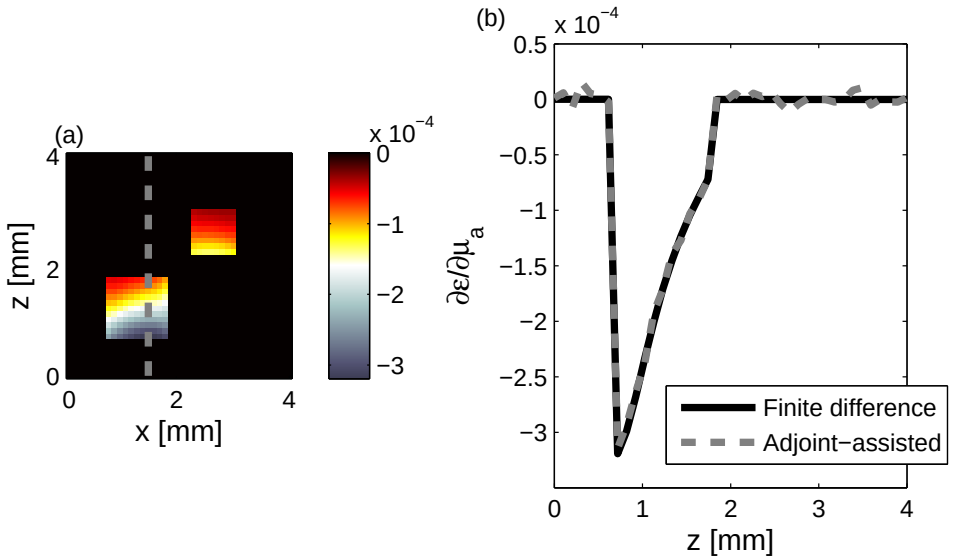

FIGURE 5.8: (a) Adjoint-assisted gradient with respect to absorption computed using forward and adjoint radiance MC simulations with grey line indicating profile plotted in (b); (b) profile through adjoint-assisted gradient with respect to absorption plotted at x=1.5mm for all z positions between 0-4mm; also plotted is gradient calculated using FD method for same positions and a value of $\delta\mu_{\mathrm{a}}$ equal to 15% of the background absorption.

Very good agreement is observed between the two gradient calculation methods. However, close inspection of the FD gradient shows some noise, particularly in regions where the gradient is close to zero. This is the result of noise in the fluence estimates made using RMC being propagated to the gradient estimates. However, the strong dependence of ϵ on μ_{a} means SNR is high in the gradient with respect to absorption. The much weaker dependence of ϵ on μ_{s} means that we can expect significantly more noise in the gradient with respect to scattering when applying the FD method, as is observed in the next section. Furthermore, the second-order dependence of the fluence on μ_{s} means a smaller perturbation magnitude, $\delta\mu_{\mathrm{s}}$, in the FD approach may be required to accurately estimate the functional gradient; however, reducing $\delta\mu_{\mathrm{s}}$ will significantly decrease SNR in the gradient estimate. This problematic trade-off is examined in the next section.

5.3.4.2 Validation of the functional gradient with respect to the scattering coefficient

The same domain and true absorption and scattering cofficients were used initially to evaluate the accuracy of the functional gradient with respect to scattering. However, here it was assumed that $\mu_{\mathrm{a}}^{\mathrm{true}}$ was known and the initial guess of μ_{s} was a homogeneous estimate of 5mm^{-1}, equal to the background value in $\mu_{\mathrm{s}}^{\mathrm{true}}$. The true absorption and scattering coefficients are shown in Fig. 5.9(a) and (b) with the true absorbed energy density and an estimate of the absorbed energy density, H^0.

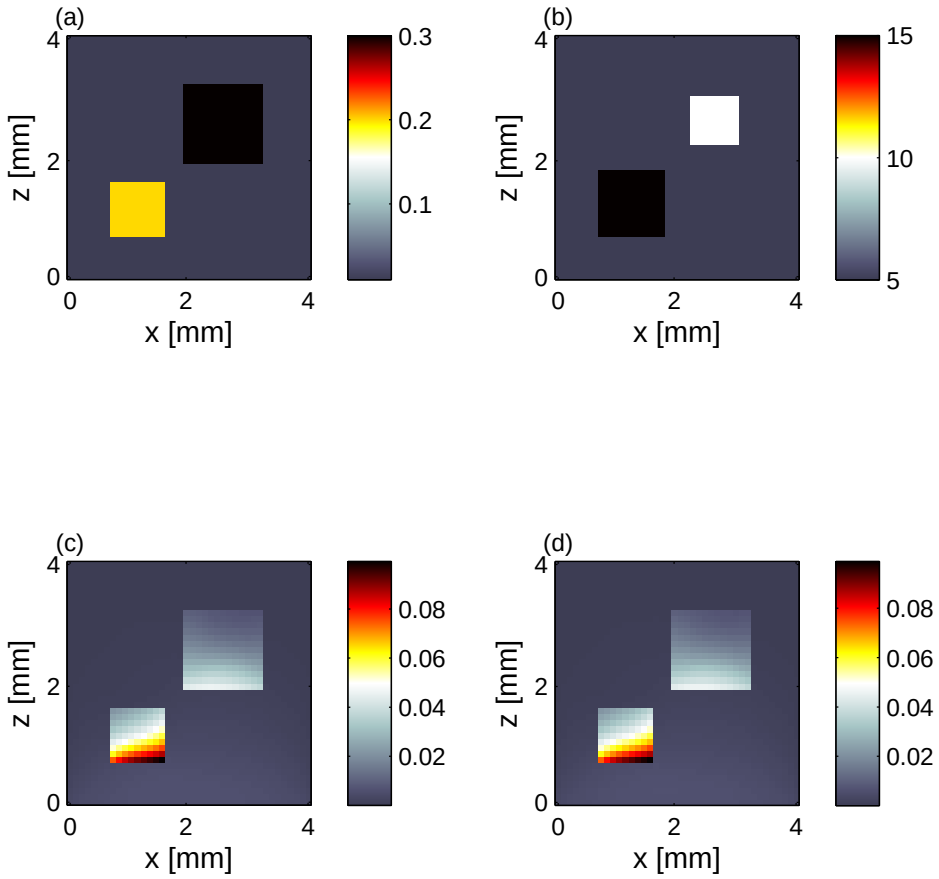

FIGURE 5.9: (a) True absorption coefficient, μ_a^{true}, with the background absorption coefficient set to 0.01mm^{-1}, containing two inclusions with the absorption coefficient between 0.2mm^{-1} and 0.3mm^{-1}; (b) True scattering coefficient, μ_s^{true}, with the background scattering coefficient set to 5mm^{-1}, containing two inclusions with the scattering coefficient between 10mm^{-1} and 15mm^{-1}; (c) 'Measured' data, H^{true}, in arbitrary units formed using RMC with 1 harmonic and 10^9 photons using true absorption and scattering coefficients; (d) Initial estimate of data, H^0, in arbitrary units formed using RMC with 10 harmonics and 10^9 photons using homogeneous scattering coefficient of 5mm^{-1} and true absorption coefficient.

The difference in H^0 between Fig. 5.6(d) and 5.9(d) is clear: with the true absorption known, the inclusions in absorption are much more prominent than the inclusions in scattering. This is more apparent in the adjoint source, plotted in Fig. 5.10(a), which is expressed as a product of the residual, $\left(H^{true} - H^0\right)$, and $\mu_a^0 = \mu_a^{true}$. This demonstrates

a large inherent sensitivity of the error functional to the absorption coefficient, which is to be expected given the strong dependence of the image on the absorption coefficient (in diffuse optical imaging this dependence is weaker as the image does not depend to first-order on μ_a). Furthermore, the adjoint fluence in Fig. 5.10(b) resembles the distribution of the absorption coefficient more than that of the scattering coefficient.

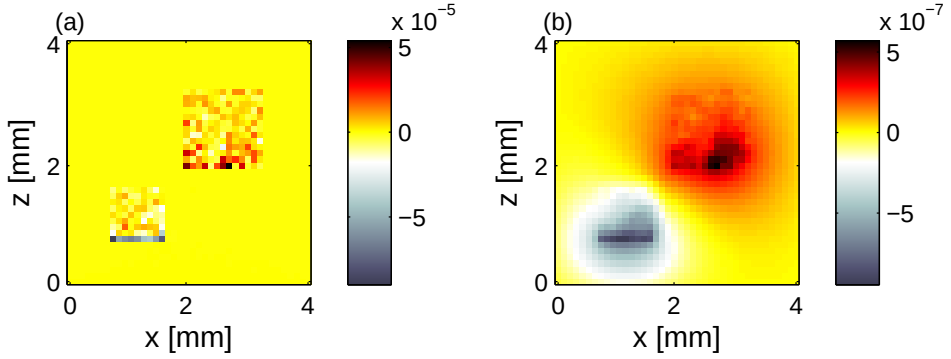

FIGURE 5.10: (a) Adjoint source, $q^* = \mu_a^0(H^{\text{true}} - H^0)$, used for gradient calculation for scattering using RMC; (b) adjoint fluence from RMC, Φ^*, used to compute adjoint-assisted gradient with respect to scattering.

The adjoint-assisted gradient yields the largest magnitude of the gradient in the region where the absorption is greatest, demonstrated in Fig. 5.11(a). Taking the profile through the adjoint-assisted gradient at x=1.5mm and comparing with the FD calculation along the same line, which was performed using a perturbation of 15% times the background μ_s value (5mm^{-1}), yields a significant discrepancy in regions where the gradient is non-zero as shown in Fig. 5.11(b). The large negative portion of the two curves is located between 0.61mm and 1.64mm, which is the region where the inclusion in μ_a^{true} is located. Meanwhile, the inclusion in μ_s^{true} is between 0.61mm and 1.85mm in the z-direction, suggesting that the gradient estimate using both methods is dominated by the presence of the region of large μ_a.

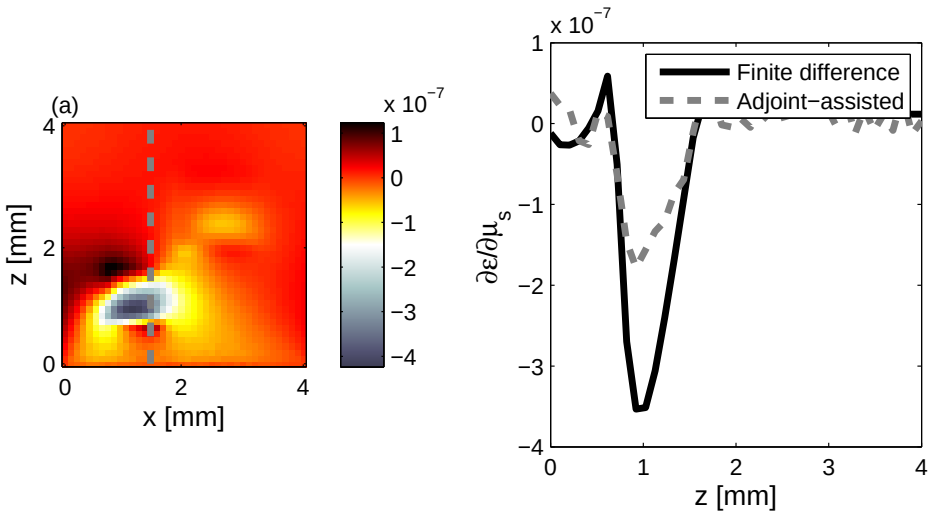

FIGURE 5.11: (a) Adjoint-assisted gradient with respect to scattering computed using forward and adjoint radiance MC simulations with grey line indicating profile plotted in (b); (b) profile through adjoint-assisted gradient with respect to scattering plotted at x=1.5mm for all z positions between 0-4mm; also plotted is gradient calculated using FD method for same positions and a value of $\delta\mu_s$ equal to 15% of the background absorption.

This issue can be more clearly demonstrated using another domain where there is no overlap in regions of large μ_a and μ_s, as shown in Fig. 5.12. The illumination conditions were the same as in the simulations above and the number of harmonics used in the adjoint-assisted gradient calculation was set to 10. The distribution of the scattering coefficient was fixed as shown in Fig. 5.12(b). Three different magnitudes of the square inclusions in μ_a visible in Fig. 5.12(a) were tested. Starting with 5.12(a) the inclusion magnitude in μ_a was between 0.1mm^{-1} and 0.2mm^{-1}, this was then decreased to be between 0.05mm^{-1} and 0.08mm^{-1} in Fig. 5.12(c), and subsequently to 0.01mm^{-1} in Fig. 5.12(d) (i.e. a homogeneous μ_a).

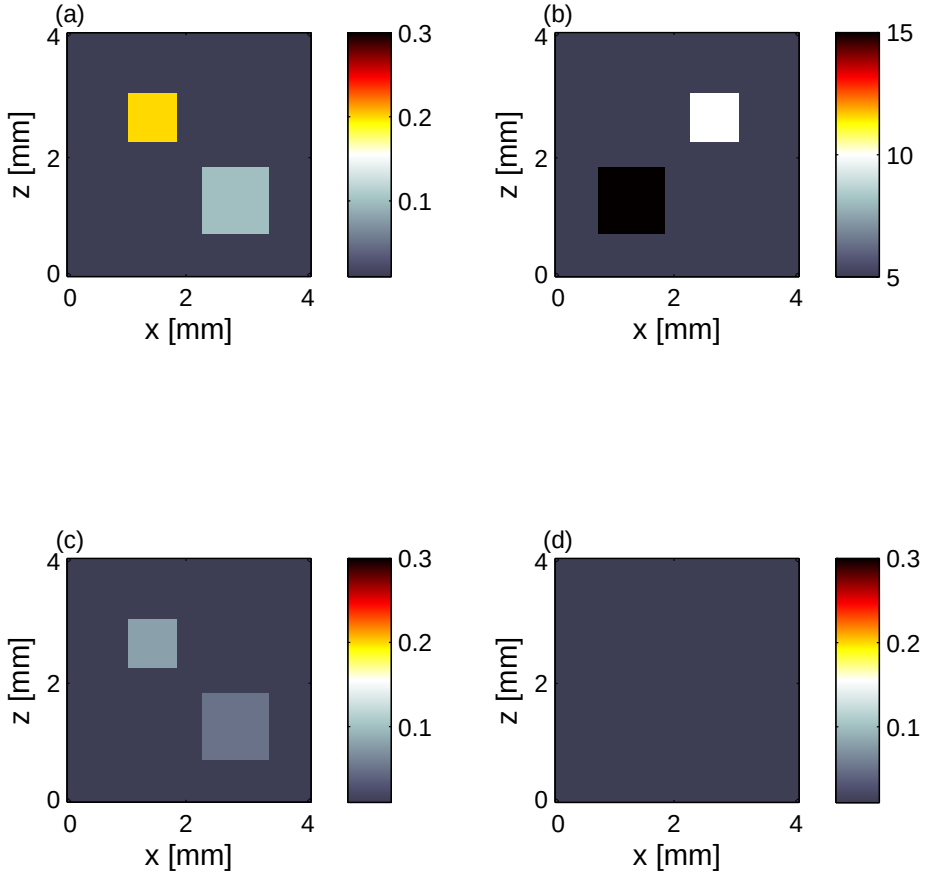

FIGURE 5.12: (a) True absorption coefficient, μ_a^{true}, with background absorption coefficient set to 0.01mm^{-1}, containing two inclusions with the absorption coefficient between 0.1mm^{-1} and 0.2mm^{-1}; (b) True scattering coefficient, μ_s^{true}, with values between 5mm^{-1} and 15mm^{-1}; True absorption coefficient, μ_a^{true}, with background absorption coefficient set to 0.01mm^{-1}, containing two inclusions with the absorption coefficient between 0.05mm^{-1} and 0.08mm^{-1}; (d) True absorption coefficient, μ_a^{true}, with homogeneous absorption coefficient of 0.01mm^{-1}.

The FD algorithm and adjoint-assisted gradient computation were run for each pair of μ_a and μ_s above with the zeroth estimate of the absorption coefficient being equal to μ_a^{true} and the scattering coefficient set to a homogeneous value of the background (5mm^{-1}). The FD was run for 40 z-positions for x=1.5mm. The adjoint-assisted gradient was

plotted for the same positions on the same plot as the corresponding FD gradient for each of the three pairs of absorption and scattering coefficients, in Figs. 5.13(a)-(c).

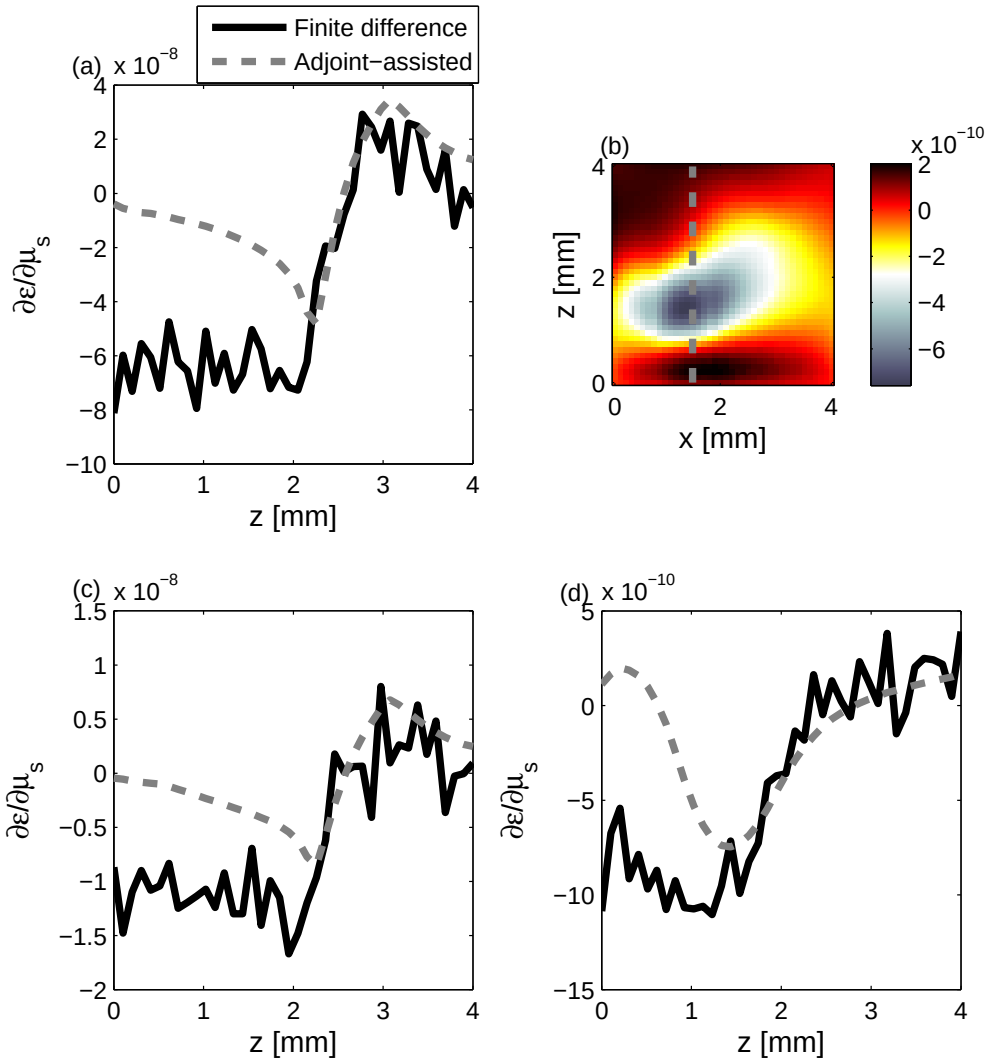

FIGURE 5.13: (a) Plot of adjoint-assisted and FD gradients with respect to scattering in domain containing two inclusions with the absorption coefficient between 0.1mm^{-1} and 0.2mm^{-1} and heterogneous μ_s (see Fig. 5.12(a) & (b)); (b) Plot of adjoint-assisted gradient with respect to scattering in domain for homogeneous μ_a and heterogneous μ_s (see Fig. 5.12(d) & (b)). Grey line indicates profile plotted in (d); (c) Plot of adjoint-assisted and FD gradients with respect to scattering in domain containing two inclusions with the absorption coefficient between 0.05mm^{-1} and 0.08mm^{-1} and heterogneous μ_s (see Fig. 5.12(c) & (b)); (d) Plot of adjoint-assisted and FD gradients with respect to scattering in domain for homogeneous μ_a and heterogneous μ_s (see Fig. 5.12(d) & (b)).

Figs. 5.13(a), (c) and (d) show that the region of the gradient obtained using the FD method is greatest in magnitude between 0mm and 2.8mm. This coincides partly with

the inclusion in μ_{s}, which is located between z=0.8mm and 1.8mm. It may seem therefore that the FD method is correctly estimating the gradient with respect to scattering, regardless of whether there are inclusions in μ_{a}. However, comparing the FD and adjoint-assisted gradients in Fig. 5.13(d) indicates that the FD method is not estimating the gradient accurately near the source, because an almost zero gradient is expected between 0mm and 0.8mm where the scattering coefficient is equal to the background value. Moreover, the adjoint-assisted gradient when μ_{a} was homogeneous, shown in Fig. 5.13(b), matches more closely the distribution of the scattering coefficient in Fig. 5.12(b). There are therefore two effects being observed in Figs. 5.13(a), (c) and (d). Firstly, coupling between μ_{a} and μ_{s} in the gradients means the gradient with respect to μ_{s} will always in some way be dominated by the presence of inclusions in μ_{a}, as illustrated by comparing the adjoint-assisted gradients in Figs. 5.13(a),(c) with (d). Second, the FD method has a tendency to overestimate the magnitude of the gradient near the source, evidenced by the discrepancy between the FD and adjoint-assisted method between z=0.8mm to 1.8mm.

The reason for the FD method overestimating the functional gradient with respect to scattering is discussed later in this chapter; however, in order to decouple the gradient with respect to μ_{s} from heterogeneity in μ_{a}, a simpler domain is used below. Here, μ_{a} is homogeneously distributed (shown in 5.14(a)) and only a single inclusion is present in μ_{s}. The distribution of the scattering coefficient, with the inclusion now at the centre of the domain, is shown in Fig. 5.14(b). The simulation parameters (illumination, number of photons, number of harmonics and $\delta\mu_{\mathrm{s}} = 0.15\mu_{\mathrm{s}}^{bg}$) were the same as in the cases examined above.

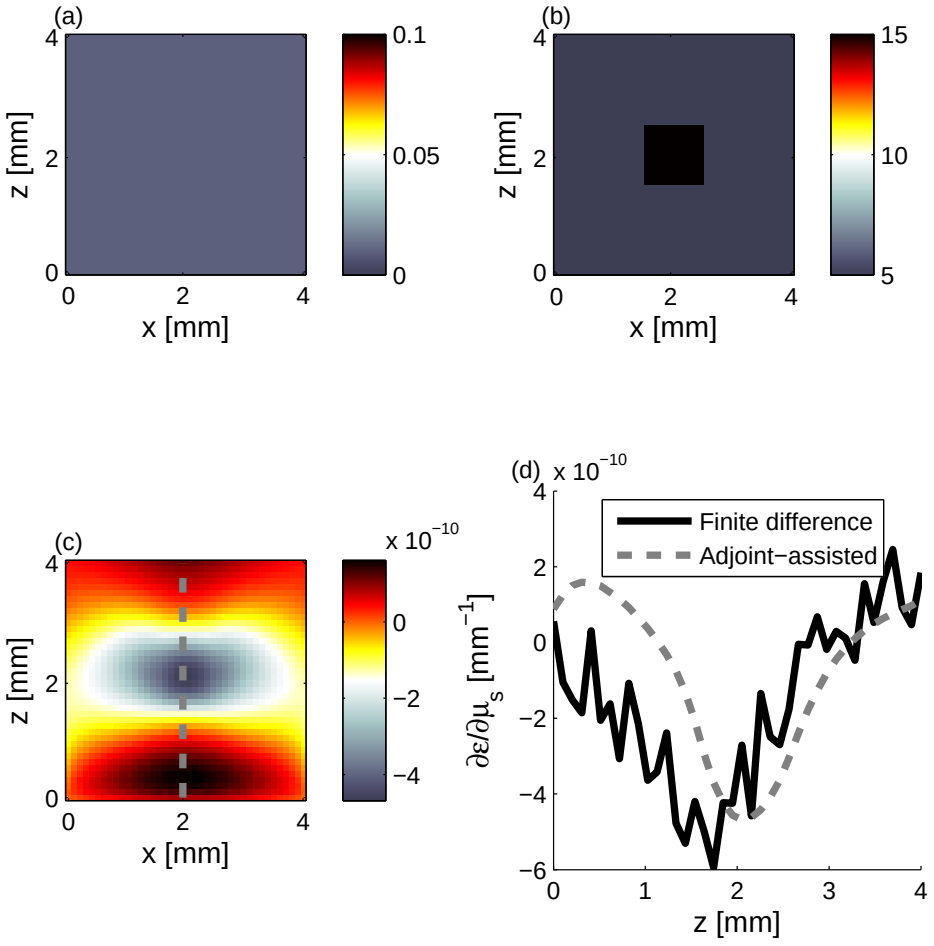

FIGURE 5.14: (a) Adjoint-assisted gradient with respect to scattering computed using forward and adjoint radiance MC simulations with grey line indicating profile plotted in (b); (b) profile through adjoint-assisted gradient with respect to scattering plotted at x=2.0mm for all z positions between 0-4mm; also plotted is gradient calculated using FD method for same positions and a value of $\delta\mu_s$ equal to 15% of the background absorption.

The distribution of the adjoint-assisted gradient with respect to scattering coefficient is displayed in Fig. 5.14(c). The grey line is used to indicate the profile through the adjoint-assisted gradient plotted in grey in Fig. 5.14(d) alongside the gradient obtained using FD. Much better agreement is observed between the two curves when the distribution of μ_a is homogeneous, however there is still a significant discrepancy for z<2mm, near the source. Consider two additional profiles, now for all x positions at z=0.5mm (outside

the inclusion in μ_s) and z=2.5mm (inside the inclusion in μ_s) shown in Figs. 5.15(a) and (b). The trend observed in the profiles in the z direction are observed once again: near the source there is a large discrepancy, with the FD gradient underestimating the adjoint-assisted gradient, and this discrepancy decreases with increasing distance from the source.

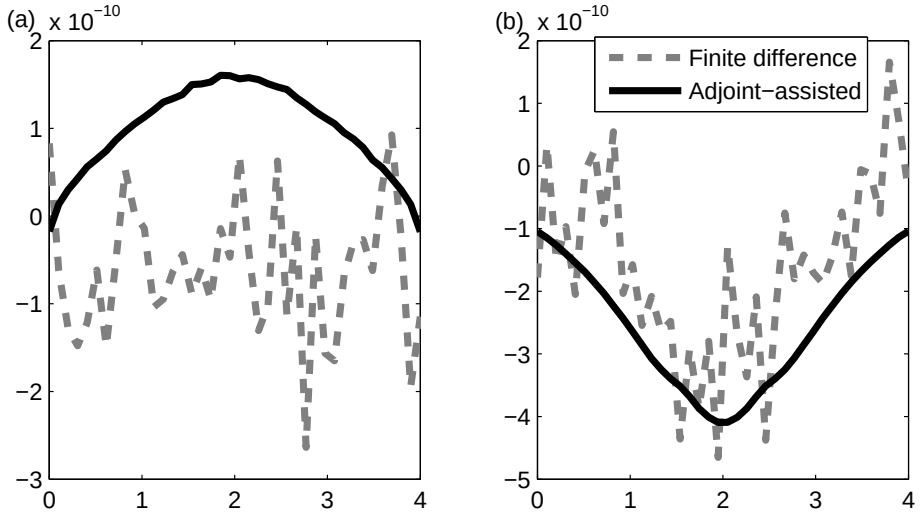

FIGURE 5.15: (a) Profiles through the domain of the gradient with respect to scattering at all x positions for z=0.5mm (does not traverse inclusion in μ_s) obtained using adjoint-assisted and FD methods; (b) Profiles through the domain of the gradient with respect to scattering at all x positions for z=2.4mm (traverses inclusion in μ_s) obtained using adjoint-assisted and FD methods.

One possible reason for the FD method underestimating the gradient near the source is that the fluence is much more sensitive to changes in the scattering coefficient in this region compared with large distances from the source. Recall that the perturbation magnitude used, $\delta\mu_s$, was 15% of the background scattering value. Due to the strongly nonlinear relationship between the fluence and the scattering coefficient, particularly near the source, this perturbation magnitude may result in gradient estimates with significant inaccuracies. The next section discusses why the use of a smaller value of $\delta\mu_s$ introduces significant computational challenges.

5.3.5 Convergence of the FD method in estimating $\partial\epsilon/\partial\mu_{\mathrm{s}}$

Earlier in this chapter it was stated that estimating the gradient with respect to scattering using the FD approach is subject to a trade-off. Correct application of the FD method involves calculating $\partial\epsilon/\partial\mu_{\mathrm{s}}$ and iteratively decreasing $\delta\mu_{\mathrm{s}}$ until there is no observable change in the functional gradient, i.e. reduce $\delta\mu_{\mathrm{s}}$ until the gradient has converged. This however was not possible using the RMC model as even decreasing $\delta\mu_{\mathrm{s}}$ by a factor of 3 to 5% of the background μ_{s} resulted in the shape of the gradient no longer being visible due to the high degree of noise (this result is not presented here). This limitation presents quite a significant dilemma: validation of the adjoint-assisted gradient is not possible without reducing the perturbation size in the FD calculation; however, the knock-on effect of this is that many more photons would be required in the FD calculation. This leads to the questions of what perturbation size is required for adequate convergence of the FD calculation in order to be able to validate the adjoint-assisted gradient? And, how many photons would be required in the MC simulations for the FD calculation to ensure sufficiently low variance in the gradient estimate?

It would only be possible to answer the above questions using the RMC model with access to quite substantial computing resources; the gradient calculations above were run using a 32-core system and computation times were similar to those quoted above for the gradient with respect to absorption: 1-2 hours for the adjoint-assisted method and 12-13 hours for the FD approach. Thus, without access to many more nodes or a very lengthy access period to a 32-node machine, it is not possible to use the RMC model. Instead TOAST++ [138], a FE model of the DA, was used. Due to the relatively short run-times of TOAST++, the FD computation could be run at every position throughout the domain.

Direct comparison between gradients computed using the RMC model and TOAST++ is not possible due to MC being a model of the RTE and TOAST++ simulating the DA. Under the diffusion approximation the adjoint-assisted gradients can be obtained by taking derivatives of the forward and adjoint models with respect to μ_{s}', similar to the method used in the transport case in Section 5.3. The adjoint-assisted gradient with

respect to scattering, provided by Cox et al. [115], is given by

$$\frac{\partial \epsilon}{\partial \mu_{\mathrm{s}}'} = -3D^2 \nabla \Phi^* \cdot \nabla \Phi, \tag{5.68}$$

where Φ and Φ^* are the fluence and adjoint fluence, respectively. Thus, instead of attempting to validate the gradient obtained in the transport case using MC against those calculated using the DA, the aim is to determine for which value of $\delta\mu_{\mathrm{s}}'$ the gradient obtained using the FD method converges to the adjoint-assisted gradient.

A circular domain, 40mm in diameter, containing 4370 elements was used for this part of the investigation. The absorption coefficient was set to a homogneous value of 0.01mm^{-1} and the background reduced scattering was set to 0.5mm^{-1}. There was a 10mm diameter circular inclusion in μ_{s}' with μ_{s}' equal to 1.5mm^{-1}. The source illuminated one side along a 6mm arc and was Gaussian in profile (the source is on the right-hand side of the plots in Fig. 5.16). The boundaries were index-matched. The gradient was computed using the adjoint-assisted approach, $\left(\frac{\partial \epsilon}{\partial \mu_s'}\right)^{\mathrm{AA}}$, and using the FD method, $\left(\frac{\partial \epsilon}{\partial \mu_s'}\right)^{\mathrm{FD}}$, with perturbation magnitudes of 15%, 1.5% and 0.15%. The gradients were plotted alongside the percent error between these two quantitaties, $100\% \times \frac{\left(\frac{\partial \epsilon}{\partial \mu_s'}\right)^{\mathrm{AA}} - \left(\frac{\partial \epsilon}{\partial \mu_s'}\right)^{\mathrm{FD}}}{\left(\frac{\partial \epsilon}{\partial \mu_s'}\right)^{\mathrm{FD}}}$.

Fig. 5.16 shows plots of the $\left(\frac{\partial \epsilon}{\partial \mu_s'}\right)^{\mathrm{AA}}$, in the first column, $\left(\frac{\partial \epsilon}{\partial \mu_s'}\right)^{\mathrm{FD}}$, in the second column, and the percentage error between them in the third column. In each row in Fig. 5.16, the FD gradient was calculated with a different value of $\delta\mu_{\mathrm{s}}$, starting with 15%, and decreasing through 1.5%, to 0.15%.

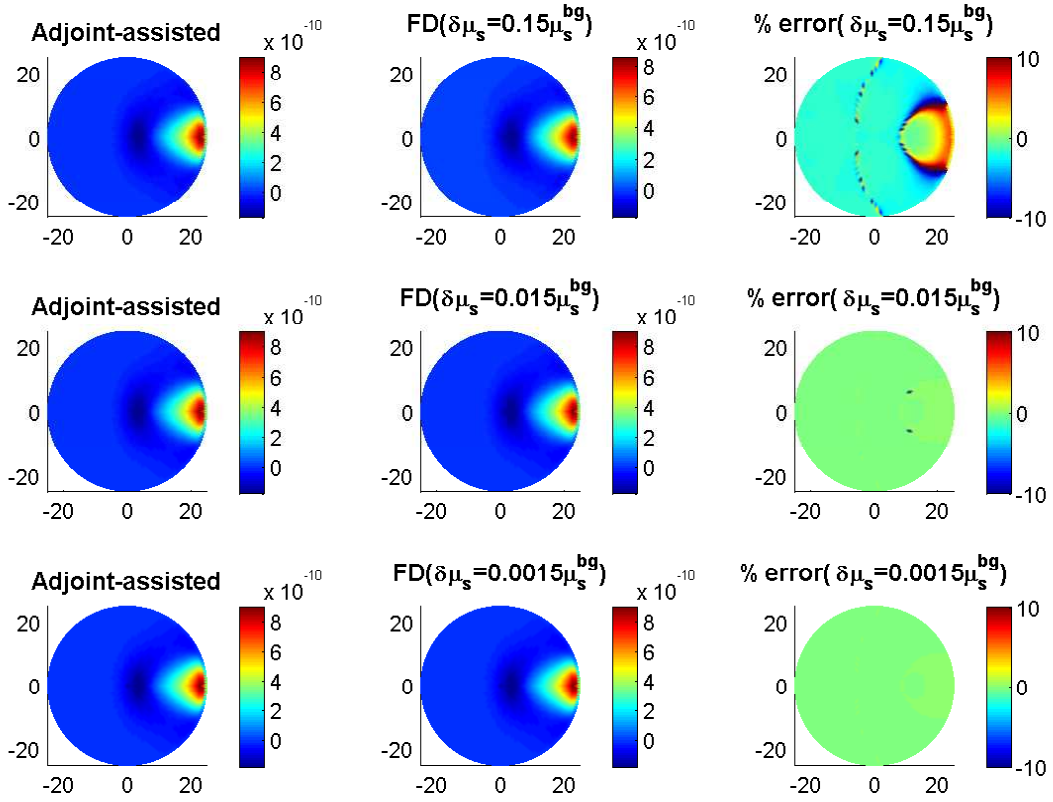

FIGURE 5.16: Functional gradients with respect to scattering obtained using TOAST++ in a 40mm diameter circular domain, illuminated by diffuse sources on the boundary. The absorption coefficient was set to a homogeneous value of 0.01mm^{-1} and the background reduced scattering coefficient was set to 0.5mm^{-1}, with a 10mm diameter circular inclusion in μ_s' equal to 1.5mm^{-1}. Left-hand column represents the functional gradient obtained using the adjoint-assisted method (see Eq. 5.68), centre column represents functional gradients obtained using the FD method and right-hand column shows the percentage error between the adjoint-assisted and FD methods. Each row represents a difference perturbation magnitude, $\delta\mu_s$, being used in the FD calculation: 15% of the background μ_s (top row), 1.5% of the background μ_s (middle row), 0.15% of the background μ_s (bottom row). The percentage error between the adjoint-assisted and FD methods is significant with a 15% perturbation, while it drops significantly when the perturbation is $\leq$1.5%.

Although it is not immediately clear from looking at the left and middle columns of Fig. 5.16, there is a significant difference in the gradients calculated using each method for the largest value of $\delta\mu_s$. This discrepancy is greatest near the source on the right side of the plots and decays to zero with increasing distance from the source in the -x direction, which is consistent with observations made regarding Fig. 5.15 when the gradients were

obtained using RMC. With an order of magnitude decrease in $\delta\mu_{\mathrm{s}}$, the % error decreases quite significantly and is almost zero everywhere and, with a further order of magnitude decrease in the perturbation, the FD and adjoint-assisted gradients are almost equal everywhere throughout the domain. This suggests that indeed the value of $\delta\mu_{\mathrm{s}}$ used for the FD calculation for the gradient with respect to scattering using the MC model was too large, and that a decrease of 1-2 orders of magnitude in $\delta\mu_{\mathrm{s}}$ would most likely result in convergence of the FD method to the adjoint-assisted approach. However, in order to decrease $\delta\mu_{\mathrm{s}}$ and maintain sufficient SNR, many more photons would have to be simulated. This is practically intractable with the CPU-parallelised MC code described in Chapter 4; assuming Poisson statistics apply, at least 10^2-100^2 more photons would be required for convergence, meaning for the FD calculation to complete execution in the same amount of time ($\sim$13 hours), 10^2-10^4 more cores would be required which is infeasible using the computing resources available. This problem would be more severe in 3D where the photon density is inherently lower. Nevertheless, highly optimised 3D GPU codes, which can compute thousands of threads simultaneously, are currently in existence [132] and would make this calculation tractable but such fast codes are not available for 2D simulations.

5.4 Summary

This chapter demonstrated that through the use of an accurate adjoint model of the RTE, functional gradients with respect to the absorption and scattering coefficients can be computed given a PAT image. The gradients were calculated using the radiance Monte Carlo model presented in Chapter 4, which stores the radiance in Fourier or spherical harmonics, and therefore allows straightforward calculation of the functional gradients as products of harmonic coefficients. The functional gradients in 2D were validated against the central finite difference method for gradient calculation, which yielded very good agreement for the gradient with respect to asborption; however, due to noise in the forward and adjoint radiance fields computed using RMC, it was not possible to obtain the same level of agreement between the two methods of gradient calculation. Nevertheless, the discrepancy between the two methods was demonstrated

to be the result of too large a perturbation used in the finite difference calculation, which was shown to be consistent with the discrepancy between finite difference and adjoint-assisted approaches using the diffusion approximation. It was also demonstrated that 10-20 Fourier orders are required for accurate gradient computation, which lies in contrast to algorithms that implement a discrete ordinate method of angular discretisation that, even in 2D, require upwards of 60 for the desired level of accuracy in highly anisotropic media such as tissue.

Computation of functional gradients using a MC model of light transport that only computes the fluence is only possible via the finite difference method. It was observed in Section 5.3.5 that convergence of the gradient calculated using the finite difference method would require enormous computing resources. On the other hand, the RMC-based adjoint-assisted approach to computing functional gradients presented in this chapter was found to be accurate, and has the potential to be computationally tractable. Since MC models are accurate in both the ballistic and diffusive regimes, their application in the optical inverse problem is highly attractive, but this has not been possible previously because functional gradients have only been accessible via the finite difference approach. The method of computing gradients presented in this chapter provides a new tool for solving the optical inversion, which is demonstrated in the next chapter.

Chapter 6

Inverting for the absorption and scattering coefficients from PAT images

In Chapter 5, the calculation of functional gradients with respect to absorption and scattering was demonstrated using forward and adjoint radiance Monte Carlo models in 2D and the accuracy of these gradients was verified against a finite difference calculation, to the extent that this is possible. In this chapter, the gradients with respect to absorption and scattering are applied in the inversion of simulated PAT images. A number of questions arise in this process, key of which is the choice of optimisation algorithm as this will determine both the rate of convergence of the inversion as well as its stability. This is of particular importance because exact inversion of PAT images simulated using a MC forward model, which contain intrinsic MC noise, is impossible. Furthermore, the convergence of the inversion is complicated by the presence of noise in the functional gradients and the error functional due to the variance in the radiance estimates from the MC simulation. These issue are discussed in detail in this chapter.

6.1 Choice of optimisation algorithm

The presence of variance in quantities estimated using MC models, such as ϵ and the functional gradients, means that these parameters are simply samples of their true value, which have some finite variance. As such, one approach to solving the inverse problem in the presence of uncertainty on parameters of interest is to consider these parameters as random variables and treat the inversion as a problem of statistical inference. This approach comes with added computational load because the distribution on each parameter must be calculated and it is therefore unclear to what extent this method would yield a more stable or rapid inversion, if at all. For this reason, the approach to solving the inversion used here starts by ignoring uncertainties due to MC modelling and treats the inversion as a least-squares minimisation problem:

$$\arg\min_{\chi} \epsilon = \frac{1}{2} \left|\left| H^{meas} - H(\chi) \right|\right|^2 , \tag{6.1}$$

where χ are the quantities being estimated, i.e. μ_{a} and/or μ_{s}. The large scale nature of the inverse problem in QPAT limits the choice of optimisation in the least-squares framework; linesearch methods are one well-documented approach.

Linesearch methods perform a 1D minimisation, the 'linesearch', in the direction $p^{(i)} = -\nabla\epsilon^{(i)}$ where $\nabla\epsilon^{(i)}$ is the gradient of the error functional with respect to χ. In general, $p^{(i)}$ is given by

$$p^{(i)} = -B^{(i)\ -1} \nabla\epsilon^{(i)}, \tag{6.2}$$

and is a descent direction if $p^{(i)T}\nabla\epsilon^{(i)} < 0$. In the gradient descent (GD) or steepest descent method, $B^{(i)}$ is just $\mathcal{I}$, the identity matrix, while quasi-Newton methods approximate the Hessian matrix of second-order derivatives, $B^{(i)}$. A popular choice of quasi-Newton method is the BFGS or l-BFGS (limited-memory) minimisation procedure due to their superior rate of convergence over the gradient-descent approach and the ability to self-correct when a bad estimate of the Hessian has been used for a step [100]. The limited-memory implementation discards some information about the curvature of ϵ by only storing $B^{(i)}$ for a limited number of previous iterates, making it suitable for large

scale optimisation problems such as the optical inversion in PAT. The update in l-BFGS is given by

$$\Delta\chi = -\alpha^{(i)} B^{(i)} \nabla\epsilon^{(i)}(\chi), \tag{6.3}$$

where $\alpha^{(i)}$ is the step length and $B^{(i)}$ is an approximation to the Hessian and is given by

$$B^{(i+1)} = \left(\mathcal{I} - \frac{y^{(i)} s^{(i)T}}{y^{(i)T} s^{(i)}}\right)^T B^{(i)} \left(\mathcal{I} - \frac{y^{(i)} s^{(i)T}}{y^{(i)T} s^{(i)}}\right) + \frac{s^{(i)} s^{(i)T}}{y^{(i)T} s^{(i)}} \tag{6.4}$$

with $s^{(i)} = \chi^{(i+1)} - \chi^{(i)}$ and $y^{(i)} = \nabla\epsilon^{(i+1)} - \nabla\epsilon^{(i)}$.

Both methods employ a linesearch that requires two conditions to be satisfied, collectively known as the Wolfe conditions: the sufficient decrease condition and the curvature condition [100]. The sufficient decrease condition stipulates that

$$\epsilon(\chi + \alpha^{(i)} p^{(i)}) \leq \epsilon(\chi) + \nu_1 \alpha^{(i)} \nabla\epsilon^{(i)T} p^{(i)}, \tag{6.5}$$

where $\nu_1 \in (0, 1)$, which is typically chosen to be quite small. However, this condition is not sufficient to ensure the linesearch is efficient because the condition will be satisfied for arbitrarily small values of $\alpha^{(i)}$ (provided, of course, that $p^{(i)}$ is a descent direction). The second condition considers whether the slope of the function at a position $\alpha^{(i)}$ is large and negative as this indicates further minimisation can be performed through the linesearch. This condition is expressed as

$$\nabla\epsilon(\chi + \alpha^{(i)} p^{(i)})^T p^{(i)} \geq \nu_2 \nabla\epsilon^{(i)T} p^{(i)}, \tag{6.6}$$

where ν_2 is chosen to be between ν_1 and unity.

Although convergence of quasi-Newton methods is in general faster than the GD method, the presence of noise in first-order derivatives of the error functional will result in an even larger degree of noise in second-order derivatives. In the context of the BFGS algorithm, as there is noise present from the MC model in $y^{(i)}$ (due to noisy gradients, $\nabla\epsilon^{(i)}$), the noise will be propagated and amplified when calculating $B^{(i+1)}$, which in turn is likely to result in a search direction $p^{(i)}$ that is sub-optimal. In addition, the self-correcting

properties of the BFGS algorithm may be compromised if $y^{(i)T}s^{(i)}$ becomes too small (but remains positive), which also may occur in the presence of large amounts of noise in the functional gradients. Noise in $H(\chi^{(i)} = [\mu_\mathrm{a}, \mu_\mathrm{s}]^{(i)})$, the estimate of the absorbed energy density, will also result in a non-smooth error functional in the search direction as at every point $\chi^{(i)}$, the error function will be corrupted by some different noise σ: $\epsilon(\chi^{(i)} = [\mu_\mathrm{a}, \mu_\mathrm{s}]^{(i)}) = \left|\left|H^{meas} - H(\chi = [\mu_\mathrm{a}, \mu_\mathrm{s}])(1 + \sigma^{(i)})\right|\right|^2$. This may be problematic within the linesearch, particularly when using the l-BFGS scheme; Nocedal et al. [100] state that the self-correcting properties of l-BFGS are less effective if the linesearch does not capture appropriate curvature information of the error functional. Given the noise levels observed in Chapter 5 in quantities estimated using the RMC model, it may in many cases be more effective to use the GD approach, rather than the l-BFGS algorithm as discussed in subsequent sections.

6.2 Noise in H^{meas} and the inverse crime

Rather than adding noise to the measured image, H^{meas}, in this chapter noise in the data is simply the result of noise from the MC forward model. The noise from MC simulations follows a Poisson distribution for high number of expected occurences, which can be approximated to be Gaussian. The variance of estimates made using MC simulations depends on domain size, the number of photons simulated, illumination type, the optical properties of the domain and boundary conditions; however, the primary factors in the magnitude of variance are the domain size and discretisation, and the number of photons simulated. In order to estimate typical noise levels in H^{meas}, the forward model was run using a line source illuminating a 4mm$\times$4mm domain containing 40 pixels along each edge; the optical properties in the domain were homogeneously distributed with μ_a, μ_s and g set to 0.01mm^{-1}, 5mm^{-1} and 0.9, respectively. The optical properties, domain size and discretisation are consistent with the gradient computations run in Chapter 5 and the reconstructions carried out below. The simulation used a single Fourier harmonic and 10^5 photons. This was repeated 1000 times in order to estimate the noise, which was achieved by computing three times the standard deviation of H^{meas} from the 1000 model runs, divided by the mean. This yielded a spatially-varying noise level of 6-21%

with the most common value being 8%. Repeating this procedure using 10^6 photons yielded a noise level between 2-8% with most of the pixels having 3% noise. Assuming Poisson statistics, to estimate the noise as a function of the number of photons simulated, the model $\sigma = aN_p^{-1/2} + c$ where N_p is the number of photons simulated, and a and c are constants, was fitted. This showed that simulating H^{meas} using 10^8 photons would yield a most probable noise value in the domain of approximately 0.9%, which is similar to the levels of extreneously added Gaussian noise in H^{meas} in other inversion approaches presented in literature [107].

6.3 Inversions for the absorption coefficient

The reconstruction of the absorption coefficient only was performed on the same domain used to demonstrate the gradients in Chapter 5; the true absorption and scattering coefficients are displayed in Fig. 6.2(a) and (b). The measured data was calculated using 10^8 photons, and the forward and adjoint RMC simulations also used 10^8 photons. Note that reconstructions in this chapter are performed under the assumptions that the initial acoustic pressure is perfectly reconstructed and that the Grüneisen parameter was known. The inversion was performed using Julia's in-built optimisation library which implements both the l-BFGS minimisation and the GD algorithm. The default linesearch in Julia's optimisation toolbox is that described by Hager and Zhang [187] (HZ-linesearch) and utilises a modified version of the second Wolfe condition; the modifications however have little impact in the presence of noisy estimates of the error functional. The termination condition for both l-BFGS and GD optimisation algorithms was

$$\left|\epsilon^{(i)} - \epsilon^{(i-1)}\right| / \left|\epsilon^{(i)}\right| < 10^{-9} \tag{6.7}$$

and 10 iterations of the approximation to the Hessian were stored when using the l-BFGS algorithm.

Fig. 6.1 shows the log_{10} value of the error functional (normalised by the value at the zeroth iteration) as a function of iteration number for the l-BFGS and GD method. The choice to normalise the error functionals was simply due to the fact that their value was

not equal at the zeroth iteration due to MC noise in the estimate of $H^{(0)}$. Comparing the two curves, it can be seen that their convergence rates are quite similar initially with a larger change in the value of the error functional at the first iteration by the l-BFGS algorithm; however, the GD approach then outperforms l-BFGS as each step results in a significant reduction in the value of the error functional where as l-BFGS takes many steps that do not result in any significant change in ϵ. Nocedal et al. [100] write that self-correcting of $B^{(i)}$ takes place over a few iterations which may explain the trend of the red curve in Fig. 6.1 to take a large descent step after several small or non-descent steps. The poor convergence of the l-BFGS optimisation suggests that noise in the Hessian approximation rarely allows the descent condition, $(-B^{(i)\ -1}\nabla\epsilon^{(i)})^T\nabla\epsilon^{(i)} < 0$, to be satisfied. This has the consequence that the HZ-linesearch is inefficient; computation time for the l-BFGS approach, which completed 34 iterations, was $\sim$15 hours compared with $\sim$11 hours for 39 iterations for the GD algorithm. This indicates that the l-BFGS method used many more runs of the forward model while the optimisation converged to a much larger value of ϵ. Note that all reconstructions were carried out using a Dell 2U R820 32-core server on Legion.

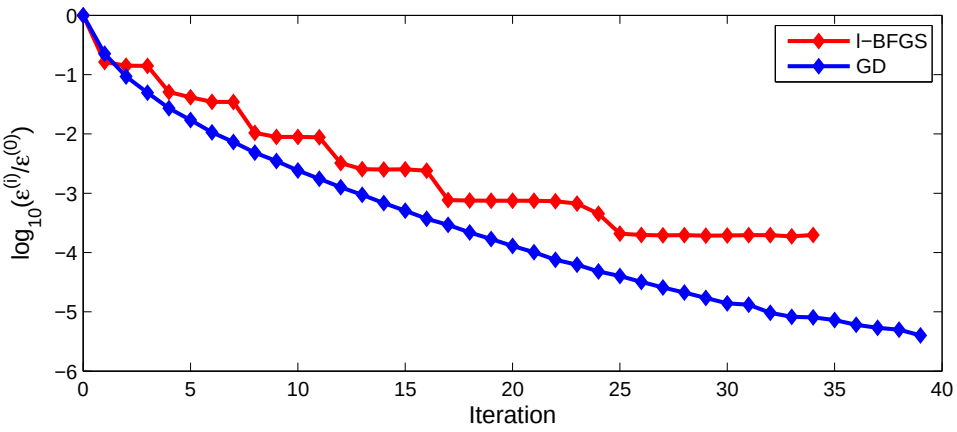

FIGURE 6.1: Plots of the error functional from l-BFGS (which approximated the Hessian using information from the 10 previous iterations) and GD optimisations as a function of iteration when inverting for the absorption coefficient using adjoint-assisted functional gradients computed using RMC using 10^8 photons.

The reconstructed absorption coefficient from the GD-based inversion is shown in Fig. 6.2(c), with profiles through the true and reconstructed absorption coefficient at x=1.5mm for all z-positions shown in (d). It can be seen that the reconstruction of the absorption

coefficient, μ_{a}^{est}, is very accurate throughout the domain. Note that the average relative error in the reconstructed μ_{a}, $\left|\mu_{\mathrm{a}}^{true} - \mu_{\mathrm{a}}^{est}\right| / \mu_{\mathrm{a}}^{true}$, was 1.9%. Although not presented here, the spatial heterogeneity of the fluence remained visible in reconstruction of μ_{a} using the l-BFGS algorithm.

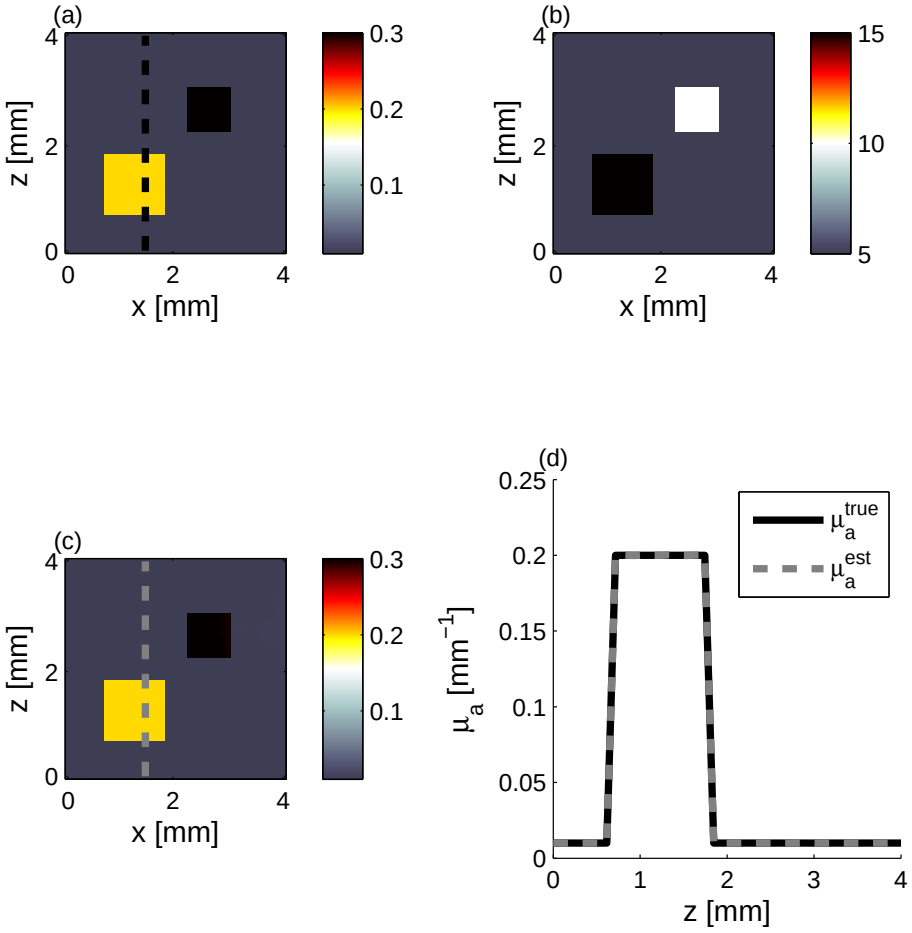

FIGURE 6.2: (a) True absorption coefficient with a background absorption coefficient of 0.01mm^{-1} and inclusions where the absorption coefficient is equal to 0.2mm^{-1} and 0.3mm^{-1}; (b) True scattering coefficient with a background scattering coefficient of 5mm^{-1} and inclusions where the scattering coefficient is equal to 10mm^{-1} and 15mm^{-1}; (c) Distribution of reconstructed absorption coefficient using GD algorithm after 39 iterations; (d) Profiles through reconstructed and true absorption coefficients at x=1.5mm for all z-positions.

Convergence of the optimisation is strongly dependent on the number of photons simulated in the forward and adjoint simulations, N_p, as this will determine the level of noise in the gradient as well as in the error functional. Fig. 6.3 shows a plot of log_{10} of the minimum value of the error functional once the termination condition was satisfied against the log_{10} of the number of photons used in the forward and adjoint models. It can be seen that the final value of the error functional decreases exponentially with increasing N_p, indicating that performing the inversion with greater values of N_p would allow further reduction of the terminal value of ϵ. Note that the reduction in the final value of the error functional is not simply due to a reduction in σ in the expression $||H^{meas} - H(\chi)(1+\sigma)||^2$, but is predominantly due to the fact that greater N_p means the Wolfe conditions have a higher probability of being satisfied for larger values of α in the the linesearch, thus allowing more progress to be made in the inversion, and prevents premature termination of the optimisation.

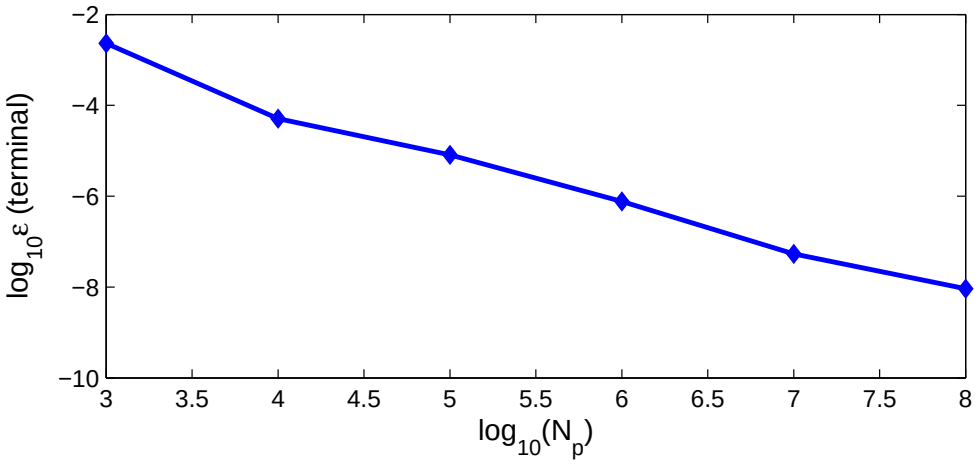

FIGURE 6.3: Plot of the log_{10} terminal value of the error functional (termination condition given in Eq. (6.7)) as a function of $log_{10}(N_p)$, where N_p is the number of photons used in the forward and adjoint RMC simulations.

It can be seen that the terminal value of the error functional continues to decrease with increasing number of photons, as might be expected. The termination of the optimisation for each value of N_p was the result of the termination condition in Eq. (6.7) being met, which was triggered by a non-descent step. A non-descent step is likely to be due to noise in the gradient resulting from an insufficient value of N_p in the forward and adjoint MC simulations, which in turn results in an upward step being selected from the linesearch.

Nevertheless, the iteration at which this occurs is implicitly a function of N_p which might mean that it is not necessary for the MC model to use a large number of photons from the start of the optimisation as it may progress initially even with low values of N_p. This idea is based on the fact that at the start of the inversion, the descent direction is well defined (i.e. the descent condition is close to -1) because the gradient is large, and for this reason the direction can be resolved even using MC estimates that have a high variance (i.e. from low values of N_p being simulated). As the optimisation progresses, the magnitude of the gradient decreases and therefore requires more photons to estimate the gradient accurately.

This approach was used in the estimation of the absorption coefficient in Fig. 6.2(a) with a homogeneous initial guess of 0.03mm^{-1} and the scattering coefficient shown in Fig. 6.2(b). The optimisation was started using 10^5 photons in the calculation of the forward and adjoint fields, which completed 12 iterations before meeting the termination condition in Eq. (6.7); the optimisation was then restarted using the terminal estimate of the absorption coefficient from the first inversion but using 10^6, completing 10 iterations, and then 10^7, completing 2 iterations. The inversion then finished after 2 iterations using 10^8 photons. The benefit of this approach was that the value of the error functional after the first 20 iterations in the optimisation using 10^8 photons, which took 5.5 hours, could be achieved in around 45 minutes using a variable number of photons, thereby significantly reducing the total time required to perform the inversion. An undesired effect of using a variable number of photons in the inversion was that the optimisation terminated at a value 6.3×10^{-7} compared with 1.3×10^{-8} for the inversion using a fixed value of 10^8 photons. This suggests that noise introduced into the inversion from the earlier high-variance inversions may have affected the progress of the inversions for larger N_p. The final estimate of the absorption coefficient is shown in Fig. 6.4(b) with the value of the error functional over the optimisations run using 10^5-10^8 photons shown in Fig. 6.4(a).

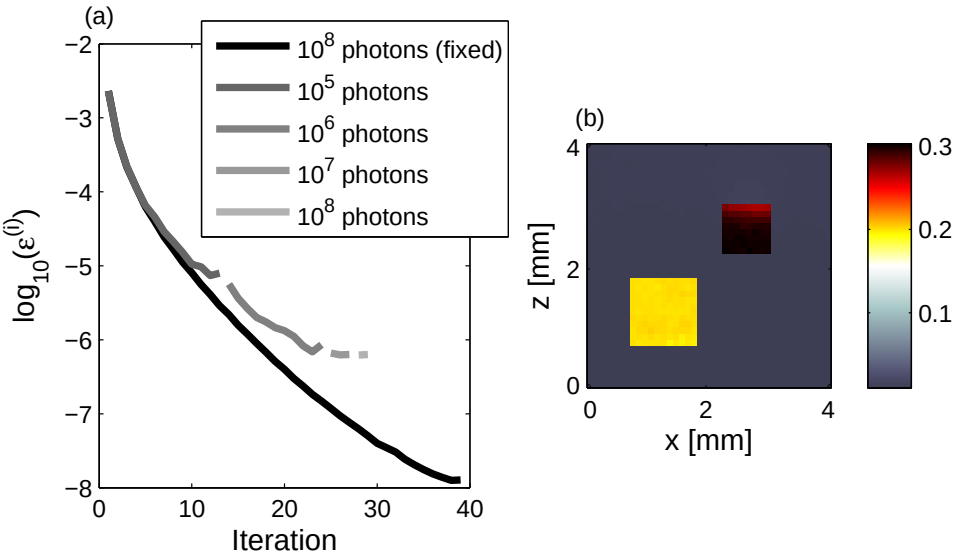

FIGURE 6.4: (a) Plot of error functional as a function of number of photons used in inversion for each set of iterations (grey lines). Note that the grey lines do not join up as the starting value of ϵ depends on N_p. The value of the error functional at each iteration in which the number of photons was fixed at 10^8 is plotted in black; (b) Estimate of absorption coefficient after 26 iterations reconstructed using a variable number of photons and and known scattering coefficient using GD optimisation.

The convergence of the inversion using a variable number of photons, plotted using a series of grey lines in Fig. 6.4(a), is much worse than the inversion in which the number of photons was fixed to 10^8, plotted in black. This is reflected in the reconstructed absorption coefficient using the method in which N_p is varied in Fig. 6.4(b). The inclusion nearest the source, in yellow, is reconstructed fairly accurately, with a visible spread around the true value of 0.2mm^{-1}; however, the inclusion far from the source, in red, has significant errors on the side furthest from the source deviating by up to 11% in some pixels relative to the true value of 0.3mm^{-1}. Divergence from the variable N_p optimisations from the black curve occurs initially at around the 10^{th} iteration but the inversion, using 10^5 photons, progresses until iteration number 20. At the 20^{th} iteration, progress can no longer be made, even when 10^8 photons are used in the inversion. The failure of the inversion to converge when N_p was increased to 10^8 photons may be due to a local minimum in the error functional having been found. This is supported by the fact that the inversion using a fixed N_p of 10^6 photons converged to 4.6×10^{-7}, which is lower than the value of ϵ that the inversion with variable N_p terminated at which was 6.3×10^{-7}.

6.4 Inversions for the scattering coefficient

Reconstruction of the scattering coefficient is typically more challenging than inverting for μ_{a} because the image depends on μ_{a} to first order, whereas this is not the case for μ_{s}. A consequence of this, as observed in Chapter 5, is that the functional gradient with respect to scattering is significantly smaller than the gradient with respect to absorption. This also has the effect that the error functional is much more sensitive to MC noise in $H^{(i)}$ which impacts the accuracy and efficiency of the linesearch. Recalling the Wolfe conditions, which require a sufficient decrease in the error function for a step length $\alpha^{(i)}$ and that the gradient at this step length is less than that at the current position, in the presence of noise in the error functional many values of $\alpha^{(i)}$ would have to be tested to satisfy simultaneously the two Wolfe conditions.

The inversion was tested in the same domain in which the gradients were validated in Section 5.3.4.2. The distribution of μ_{s} is reproduced in Fig. 6.5(a) and the true absorption coefficient was set to a homogeneous value of 0.03mm^{-1}. The data was simulated using 10^8 photons, and simulation of the forward and adjoint radiance fields was also achieved using 10^8 photons. The GD optimisation was once again used. It was found that the optimisation spent several hours within the linesearch on the first iteration testing candidate values of $\alpha^{(i)}$. Recalling that inversions for the absorption coefficient completed one iteration every 20 minutes on average, this indicates that the HZ-linesearch performs highly inefficiently in the presence of noise in the error functional and the small magnitude of the functional gradient. Since requiring that both Wolfe conditions are satisfied negatively impacts the performance of the optimisation, it may be worthwhile considering a backtracking linesearch [100], which requires only the sufficient decrease condition to be satisfied for a given $\alpha^{(i)}$. The backtracking linesearch starts with a large candidate step length; a suitable choice is $\alpha^{(i,j=0)} = 1/\max(\nabla\epsilon^{(i)})$. The sufficient decrease condition (Eq. (6.5)) is then tested for $\alpha^{(i,j)}\tau$ whilst iterating over j, with $\tau \in [0, 1]$.

The gradient descent optimisation using the backtracking linesearch was used to invert the image for the scattering coefficient shown in Fig. 6.5(a). The value of ν_1 and τ were set to 0.2 and 0.5 as this yielded positive results. The true absorption coefficient

was homogeneous and was known when inverting for the scattering coefficient. The measured data was computed using a line source across all x positions at z=0mm. The forward and adjoint RMC simulations used 10^7 photons and 10 Fourier harmonics. The initial step, $\alpha^{(1)}$, was chosen to be 10^{12} because this produced an update that would make a substantial change to the scattering coefficient. Subsequent steps were obtained using the backtracking linesearch with a starting α of 10^{12}, iteratively decreased by an order of magnitude. However, to prevent the linesearch from iterating for an extended period of time and producing an update that was negligble in magnitude, a minimum step length was set to 10^6. Note that positivity of the scattering coefficient was enforced by setting all values of $\mu_{\mathrm{s}} < 0$ to zero, which is consistent with the definition of the scattering coefficient. The optimisation was run on a Dell 2U R820 32-core server on Legion and took about 13.6 hours. The value of the error functional is plotted as a function of iteration number in Fig. 6.5(c) and the final estimate of the scattering coefficient is shown in Fig. 6.5(b), with profiles through the true and final estimate of the scattering coefficent shown in Fig. 6.5(d) at x=2mm for all z positions, indicated by the vertical black and grey lines in (a) and (b).

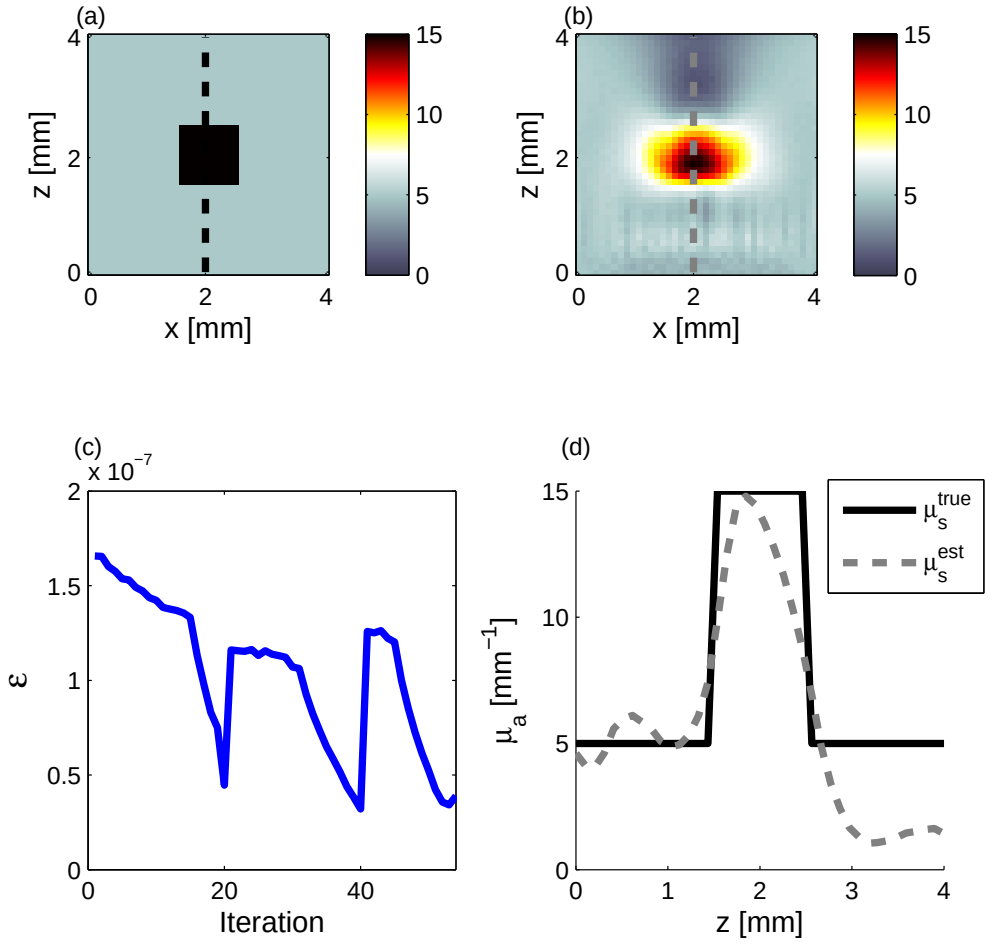

FIGURE 6.5: (a) True scattering coefficient with a background absorption coefficient of 5mm^{-1} and inclusions where the scattering coefficient is equal to 10mm^{-1} and 15mm^{-1}; (b) Estimate of scattering coefficient estimated with known absorption coefficient and GD optimisation using 10^7 photons; (c) Value of error functional plotted as a function of iteration during GD optimisation for estimating μ_s; (d) Profiles through the true and reconstructed scattering coefficient at x=2mm for all z-positions (positions shown by the black and grey lines in (a) and (b)).

From the estimate in Fig. 6.5(b), it is clear that edges are not well reconstructed; this is in part due to the diffusive nature of scattering, but may also be due to early termination of the optimisation. Reconstruction of the scattering coefficient may benefit from some regularisation, which has been demonstrated in the literature [107, 109, 113, 188] to improve the quality of the scattering estimate. Specifically, given the

piecewise constant distribution of μ_{a} and μ_{s} used here, total variation regularisation (TVR) [189] may help stabilise the solution. TVR may also be able to suppress the overestimation of μ_{s} on the left- and right-hand sides of the square inclusion as well as prevent the underestimation of μ_{s} in the deeper region, behind the inclusion, through TVR's edge preserving properties. Nevertheless, the overall accruacy of the estimate of μ_{s} in the inclusion was quite good and improved with each iteration as demonstrated by Fig. 6.6 which shows a plot of the maximum and average estimate of the scattering coefficient in the inclusion converging to their respective true values. Thus, perhaps with regularisation and if the optimisation were allowed to iterate for longer, a more accurate reconstruction would have been obtained.

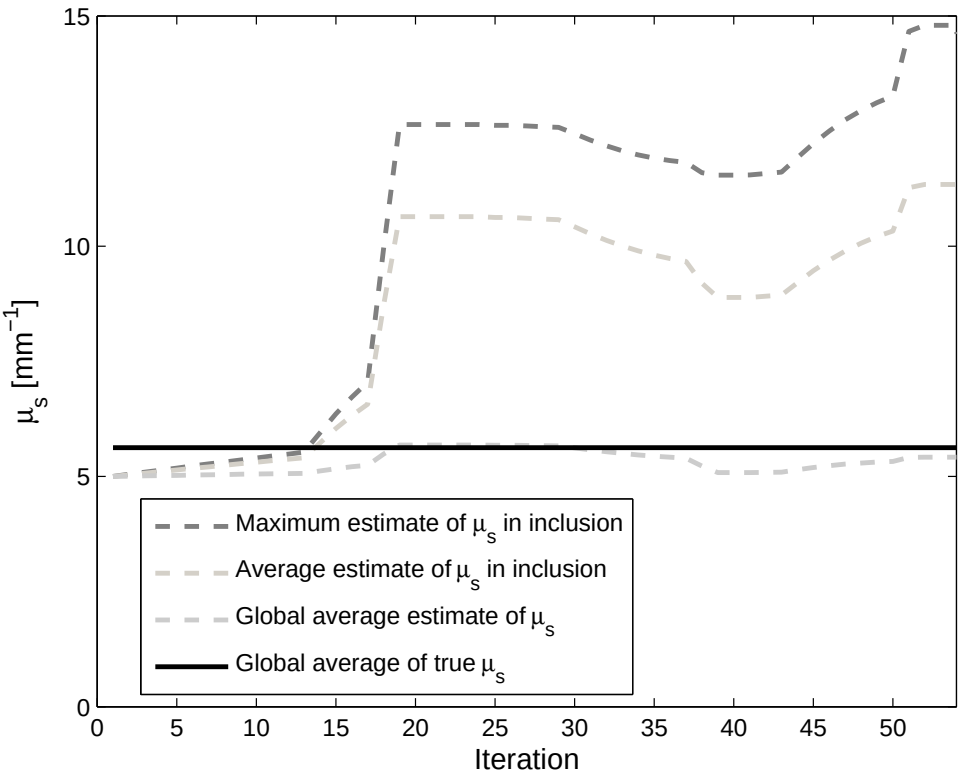

FIGURE 6.6: Plot of maximum and average estimate of μ_{s} in the inclusion in Fig. 6.5(a) as a function of iteration, and global estimate of μ_{s} as a function of iteration. Plots obtained from GD optimisation used to reconstruct μ_{s} using 10^7 photons in the forward and adjoint simulations. The true value of μ_{s} in the inclusion was 15mm^{-1} and the true global average of μ_{s}, 5.63mm^{-1}, is plotted in black.

The limited accuracy of the inversions for the scattering coefficient may compromise

the accuracy of reconstruction of the absorption coefficient when these are estimated simultaneously. Nevertheless, the primary objective of QPAT is to estimate the absorption coefficient, given that it is closely linked to physiology through chromophore concentrations. Thus, if the accuracy of the estimate of μ_{s} in Fig. 6.5(b) is sufficient to allow accurate estimation of the absorption coefficient then gradient-based inversion using RMC is a promising approach for use in practice. It has been shown [98] that a homogeneous average or background estimate of the scattering coefficient can yield estimates of the absorption coefficient that are of a useful level of accuracy. Fig. 6.6 also contains a plot of the estimate of μ_{s} averaged over the entire domain for each iteration in the optimisation, with the horizontal black line showing the average of the true scattering coefficient over all pixels. Despite the average estimate of μ_{s} at the 54^{th} iteration being lower in accuracy than the average estimate at the 18^{th} and 30^{th} iterations, it may still be sufficiently accurate to yield an accurate estimate of the absorption coefficient when inverting for μ_{a} and μ_{s} simultaneously. This type of inversion is examined in the next section.

6.5 Inversions for the absorption and scattering coefficients simultaneously

In order to overcome the nonuniqueness problem in the inversion when estimating μ_{a} and μ_{s} simultaneously, at least one additional PAT image is required. As discussed in Chapter 2, a few approaches are possible, but obtaining images using illuminations from multiple directions is the scheme used here as it is straightforward to implement in the MC model. The absorption and scattering coefficients used were the same as when reconstructing the absorption coefficient and are plotted in Fig. 6.8(a) and (b). One image was simulated using the same illumination as before: a collimated line source on the boundary at z=0mm. A second image was simulated using a collimated line source on the adjacent boundary at x=4mm. The pair of images, H_{ξ}^{meas}, for each illumination ξ, is shown in Fig. 6.7 with the position of the line source indicated by the white line. Each image was simulated using the RMC model using 10^8 photons.

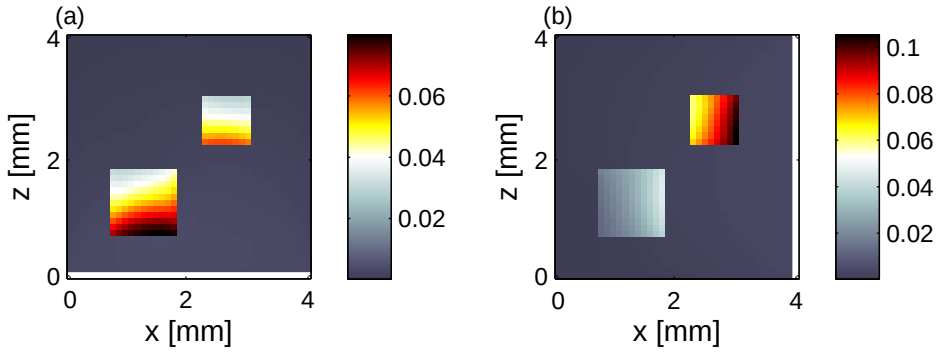

FIGURE 6.7: (a) $H_{\xi=1}^{meas}$, the true absorbed energy density from line source illumination at z=0mm, indicated by white line; (b) $H_{\xi=2}^{meas}$, the true absorbed energy density from line source illumination at x=4mm, indicated by white line.

A new error functional must now be formed with a summation over the illumination index ξ:

$$\epsilon^{(i)} = \frac{1}{2}\sum_{\xi=1}^{2}\left|\left|H_{\xi}^{meas} - H_{\xi}^{(i)}\right|\right|^2, \tag{6.8}$$

and the functional gradients with respect to absorption and scattering must also be obtained using the same summation over illumination:

$$\frac{\partial\epsilon}{\partial\mu_{\mathrm{a}}} = \sum_{\xi=1}^{2}\left(\frac{\partial\epsilon}{\partial\mu_{\mathrm{a}}}\right)_{\xi} \tag{6.9}$$

$$\frac{\partial\epsilon}{\partial\mu_{\mathrm{s}}} = \sum_{\xi=1}^{2}\left(\frac{\partial\epsilon}{\partial\mu_{\mathrm{s}}}\right)_{\xi}. \tag{6.10}$$

As observed earlier in this chapter as well as in Chapter 5, the magnitude of the gradient with respect to μ_{s} is much smaller than the gradient with respect to μ_{a}. The consequence of this is that a linesearch in the optimisation will yield step lengths, $\alpha^{(i)}$, that would produce a significant change in $\mu_{\mathrm{a}}^{(i)}$, with almost no change in $\mu_{\mathrm{s}}^{(i)}$, because of the increased sensitivity of ϵ to μ_{a}. 'Sphering' of the error surface is therefore performed by

using a change of variables

$$\hat{\mu}_a = \frac{\mu_\mathrm{a}}{\mu_\mathrm{a}^{(0)}} \tag{6.11}$$

$$\hat{\mu}_s = \frac{\mu_\mathrm{s}}{\mu_\mathrm{s}^{(0)}}, \tag{6.12}$$

where $\mu_\mathrm{a}^{(0)}$ and $\mu_\mathrm{s}^{(0)}$ are the initial guesses of the absorption and scattering coefficients, respectively. With this change of variables, the functional gradients become

$$\frac{\partial \epsilon}{\partial \hat{\mu}_a} = \mu_\mathrm{a}^{(0)} \frac{\partial \epsilon}{\partial \mu_\mathrm{a}} \tag{6.13}$$

$$\frac{\partial \epsilon}{\partial \hat{\mu}_s} = \mu_\mathrm{s}^{(0)} \frac{\partial \epsilon}{\partial \mu_\mathrm{s}}, \tag{6.14}$$

which will be more similar in magnitude provided we choose $\mu_\mathrm{a}^{(0)} < 1 < \mu_\mathrm{s}^{(0)}$.

The inversion was performed using 10^7 photons and 10 Fourier harmonics for each forward and adjoint RMC simulation and for each illumination. $\mu_\mathrm{a}^{(0)}$ and $\mu_\mathrm{s}^{(0)}$ were set to be homogeneous and equal to the background values in the true absorption and scattering coefficients, respectively. The same GD optimisation with backtracking linesearch ($\nu_1 = 0.2$, $\tau = 0.5$) was used, as before when reconstructing the scattering coefficient, but the range of $\alpha^{(i)}$ had to be adjusted to account for the different magnitude in $\frac{\partial \epsilon}{\partial \hat{\mu}_s}$. The step length was chosen to be between 10^4 and 10^6 in the linesearch as this resulted in sufficiently rapid convergence without producing extremely large estimates of μ_a and μ_s that are outside of the physiological range. Positivity in μ_a and μ_s were also enforced. Inversions in this section were again run on a Dell 2U R820 32-core server on Legion. The top row of Fig. 6.8 shows the true absorption and scattering coefficients with the reconstructed μ_a and μ_s after 300 iterations shown in the bottom row. 300 iterations took 11.4 hours, with each iteration requiring between 1.37 minutes on average.

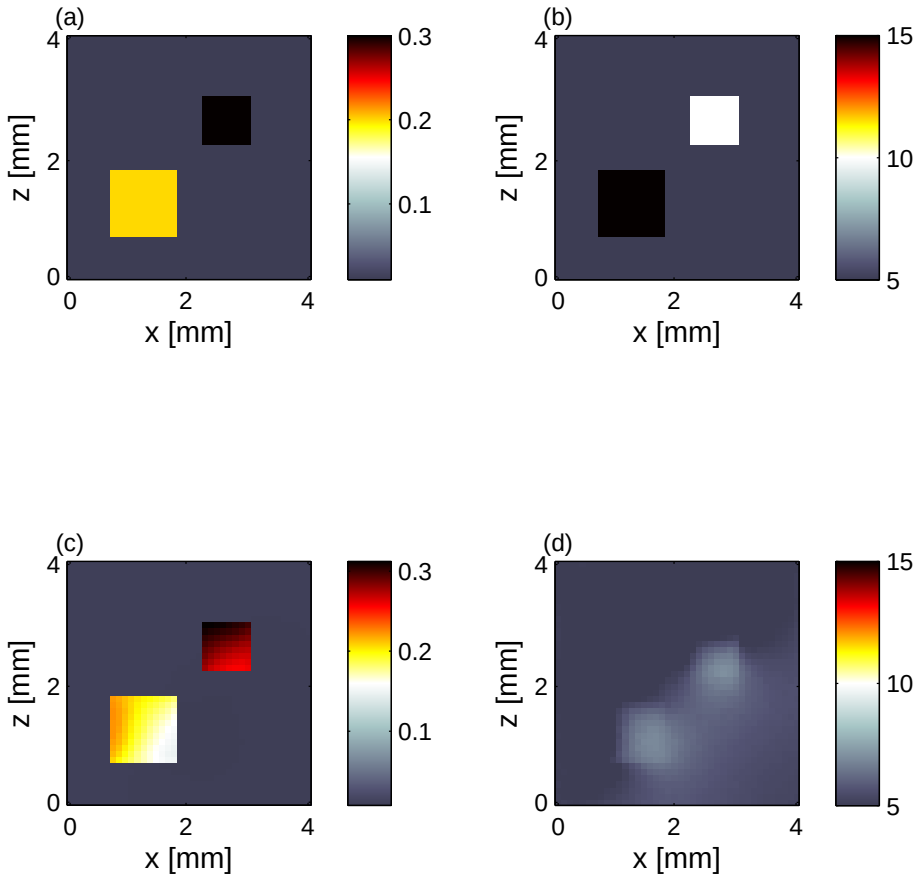

FIGURE 6.8: (a) True absorption coefficient with a background absorption coefficient of 0.01mm^{-1} and inclusions where the absorption coefficient is equal to 0.2mm^{-1} and 0.3mm^{-1}; (b) True scattering coefficient with a background scattering coefficient of 5mm^{-1} and inclusions where the scattering coefficient is equal to 10mm^{-1} and 15mm^{-1}; (c) Estimate of absorption coefficient after 300 iterations, recovered using simultaneous inversion through GD optimisation using 10^7 photons; (c) Estimate of scattering coefficient after 300 iterations, recovered using simultaneous inversion through GD optimisation using 10^7 photons.

Comparing Fig. 6.8(a) with (c), it can be seen that the absorption coefficient is reconstructed accurately far from the sources at z=0mm and x=4mm but many pixels have significant errors in the regions of the inclusions that are nearest the corner of the domain at x=4mm, z=0mm. The estimate of the scattering coefficient has significant errors; this

is highlighted in the plots in Fig. 6.9 which show the average of the estimate of $\mu_{\rm a}$ and $\mu_{\rm s}$ plotted as a function of their true values in the two inclusions, with the maximum and minimum values in the inclusion indicated by the error bars.

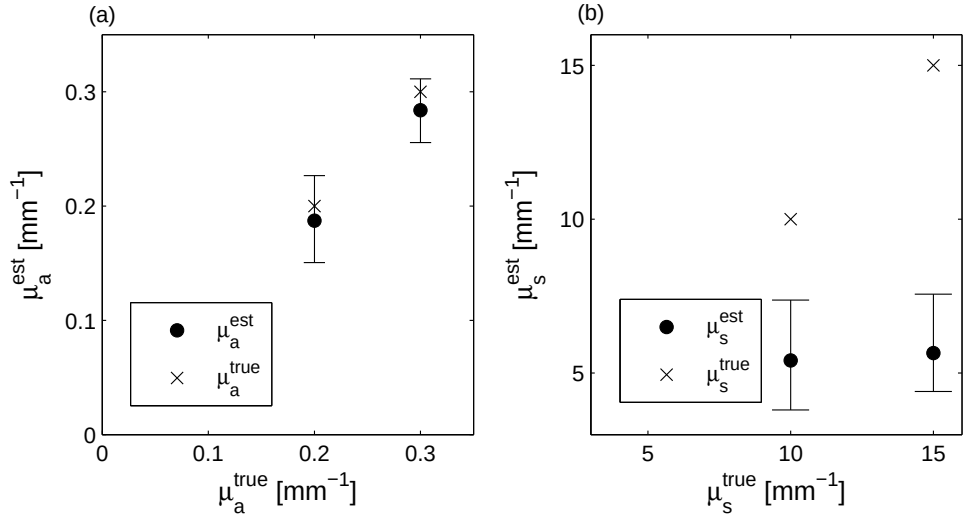

FIGURE 6.9: (a) Error bar plots of average estimated absorption coefficient (recovered using 300 iterations of GD algorithm using 10^7 photons which reconstructed $\mu_{\rm a}$ and $\mu_{\rm s}$ simultaneously) against true absorption coefficient in inclusions in Fig. 6.8(a); (b) Error bar plots of average estimated scattering coefficient (recovered using 300 iterations of GD algorithm using 10^7 photons which reconstructed $\mu_{\rm a}$ and $\mu_{\rm s}$ simultaneously) against true scattering coefficient in inclusions in Fig. 6.8(b).

While the average estimate in the inclusion is quite near the true value for the absorption coefficient, the error is significant for the scattering coefficient in both inclusions. This leads to the following questions:

(1) why is the absorption coefficient reconstructed with a relatively high degree of accuracy?

(2) can the accuracy of the absorption coefficient be improved beyond the estimate shown in Fig. 6.8(c)?

Question (1) can be answered in part with reference to the inversion performed in [98] which attempted to reconstruct simultaneously the absorption coefficient in every pixel and the value of $\mu_{\rm s}$ averaged over all pixels. This showed that the absorption coefficient could be reconstructed quite accurately everywhere in the domain, provided

the mean estimate of the scattering over the domain was fairly accurate. This may be the reason why the estimate of μ_{a} in Fig. 6.8(c) is quite accurate; the average scattering coefficient in the domain, plotted above in Fig. 6.6 was near the true value of 5.96mm^{-1} at every iteration in the optimisation, and the value at the final iteration was 5.17mm^{-1}, coresponding to a 13% relative error compared with the true μ_{s}. It was shown in [98] that the sensitivity of the error in the absorption coefficient was shown to be greatest in the presence of noise in the data and a single line source illumination; plots of the average error in μ_{a}^{est} versus that in μ_{s}^{est} showed that an average relative error of 13% in the scattering ceofficient yielded between an 8-11% average relative error in the absorption. This is greater than the 6.3% relative error in the estimate of μ_{a} in Fig. 6.8(c) when the average error in μ_{s}^{est} was 13%. This leads onto question (2); by estimating the mean of the scattering coefficient in all pixels rather than in each pixel, can the estimate of the absorption coefficient be improved?

The inversion was carried out again, this time only recovering μ_{a} in each pixel and the mean value of μ_{s}. This was achieved by computing the gradient in each pixel , $j \in [1, J]$, (with sphering applied), $\left(\frac{\partial \epsilon}{\partial \hat{\mu}_s}\right)^{(i,j)}$, and then averaging over the gradient before forming the update to the scattering coefficient given by $\alpha^{(i)} \left(-\frac{1}{J}\sum_{j=1}^{J}\left(\frac{\partial \epsilon}{\partial \hat{\mu}_s}\right)^{(i,j)}\right)$. The update to absorption was calculated in the usual way. Again, the foward and adjoint RMC simulations were run with 10^7 photons and 10 Fourier harmonics were recorded. Fig. 6.10 shows error bar plots like those in Fig. 6.9, but the plot on the right-hand side shows the accuracy of the average scattering estimate (which has no spread unlike in Fig. 6.9(b)).

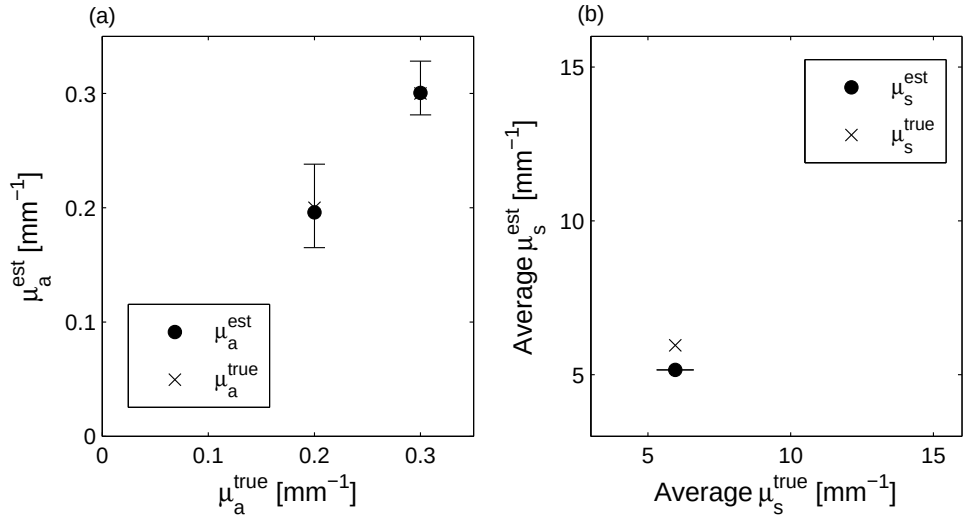

FIGURE 6.10: (a) Error bar plots of average estimated absorption coefficient (recovered using 300 iterations of GD algorithm using 10^7 photons which reconstructed μ_{a} and average value of μ_{s} simultaneously) against true absorption coefficient in inclusions in Fig. 6.8(a); (b) Error bar plots of average estimated scattering coefficient (recovered using 300 iterations of GD algorithm using 10^7 photons which reconstructed μ_{a} and μ_{s} simultaneously) against average of the true scattering coefficient in Fig. 6.8(b).

The plots above show the improved accuracy in the μ_{a} estimate in the two inclusions with the relative error averaged over the whole domain being 5.1% compared with 6.3% when the full distribution of the scattering was estimated. This is a slightly curious result because the average scattering coefficient was estimated to be 5.16mm^{-1} which is slightly lower than the average of the distribution in Fig. 6.8(d). This may be due to the fact that the spread of the scattering coefficient when the full μ_{a} and μ_{s} distributions were being recovered simultaneously was negatively impacting the accuracy of the μ_{a} estimate. It is therefore possible that a statistic other than the mean, e.g. the most frequently occuring μ_{s} in the distribution, may be a better choice but this is not investigated here.

Other approaches may be able to improve the estimate in μ_{a}. Plotting the log of the error functional from the case where μ_{a} and μ_{s} were reconstructed simultaneously at every pixel in the domain in red in Fig. 6.11 shows that a large change in the error functional is seen between the zeroth and 50$^{\text{th}}$ iteration, with an almost linear region between the 50$^{\text{th}}$ and 300$^{\text{th}}$ iteration. This is consistent with observations made in [98], which show that the initial large change in $\epsilon^{(i)}$ is due to partial convergence of the absorption coefficient; however, after $\sim$50 iterations, the estimate of μ_{a} cannot converge without improvement

in the accuracy of μ_s. This is supported by the fact that performing the same inversion using an initial guess $\mu_s^{(0)} = \mu_s^{true}$, the estimate of the absorption coefficient converges to its true value with a <6% error averaged over all pixels within 50 iterations. Under the assumption $\mu_s^{(0)} = \mu_s^{true}$, the relative error in μ_a decreased to 2.5% by the 100th iteration, which is comparable to the case investigated in Section 6.3 where only μ_a was estimated using 10^8 photons in the forward and adjoint RMC simulations. In addition, it is possible to estimate how many iterations would be required for convergence of both μ_a and μ_s given the non-zero slope in the linear region of Fig. 6.11, which is estimated to be $\sim 4.03 \times 10^{-4}$. In order to determine the minimum expected value of the error functional, the model was run using 10^7 photons for each illumination using the true absorption and scattering coefficients, and the least-squares error between this data and H_ξ^{meas} was calculated. The log of this value was -7.22, although this is of course subject to some uncertainty due to variance in H_ξ^{meas} and $H_\xi(\mu_a^{true}, \mu_s^{true})$. Assuming the slope in the linear region remains constant (which it is unlikely to due to a reduction in magnitude of the gradient as the true μ_a and μ_s values are approached) more than 6000 iterations would be required. At ~1.37 minutes per iteration, the inversion would require nearly 90 hours of run-time, which is not possible due to limited access to Legion for such extended times.

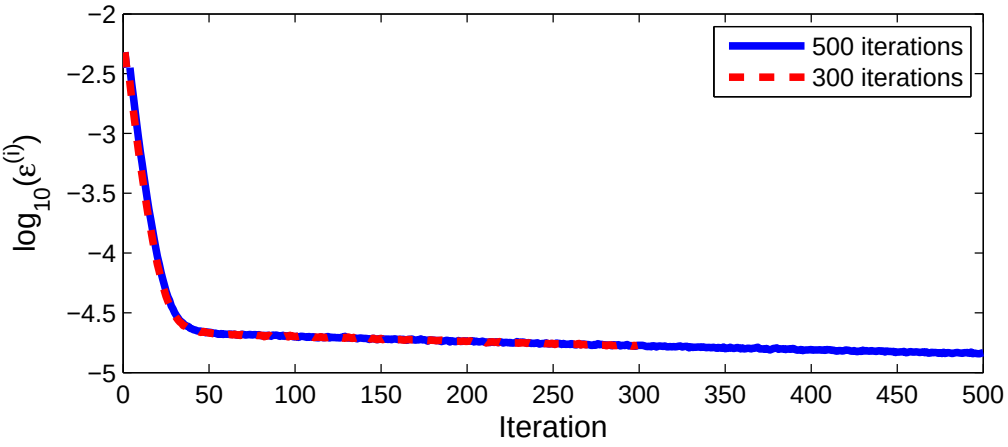

FIGURE 6.11: Plot of error functional from simultaneous reconstruction of μ_a and μ_s using GD optimisation and 10^7 photons. Plots for optimisations run for 300 and 500 iterations.

Running the inversion for 500 iterations helps illustrate the above point that iterating

for longer is likely to result in the 'absolute' minimum being attained, subject to the variance in the data and the forward and adjoint simulations of course. The minor decrease in the error functional over the additional 200 iterations in Fig. 6.11, is due to an improvement of the estimate of the scattering coefficient as negligible change was observed in the absorption coefficient, reflected by the error bar plots in Fig. 6.12 which shows the new range of $\mu_{\rm a}$ and $\mu_{\rm s}$ estimates in the inclusions overlayed with the estimates after 300 iterations, ploted using the asterisk. Thus, provided the SNR in the gradients remains sufficient, the true values of $\mu_{\rm a}$ and $\mu_{\rm s}$ could be estimated to within the level of uncertainty in the modelled absorbed energy density.

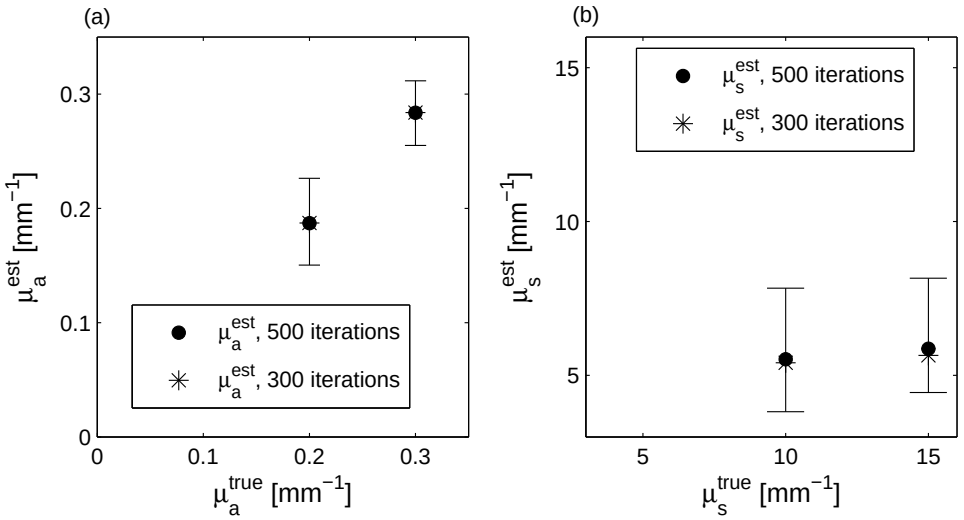

FIGURE 6.12: (a) Error bar plots comparing average estimated absorption coefficient recovered using 500 iterations of GD algorithm using 10^7 photons which reconstructed $\mu_{\rm a}$ and $\mu_{\rm s}$ simultaneously against the absorption coefficient estimated using the same optimisation after 300 iterations. Average estimate of $\mu_{\rm a}$ plotted against true $\mu_{\rm a}$ in inclusions shown in Fig. 6.8(a); (b) Error bar plots comparing average estimated scattering coefficient recovered using 500 iterations of GD algorithm using 10^7 photons which reconstructed $\mu_{\rm a}$ and $\mu_{\rm s}$ simultaneously against the scattering coefficient estimated using the same optimisation after 300 iterations. Average estimate of $\mu_{\rm s}$ plotted against true $\mu_{\rm s}$ in inclusions shown in Fig. 6.8(b).

This is an exciting and promising result because the inversions performed in this chapter, particularly those in which $\mu_{\rm a}$ and $\mu_{\rm s}$ were recovered simultaneously, were not carried out in optimal conditions. In practice, the data may be corrupted with artefacts from acoustic reconstruction and greater levels of noise. However, existing model-based inversion approaches using the FE models of the DA [190], RTE [107] or delta-Eddington

approximation [191] have achieved levels of accuracy in estimates of μ_{a} and μ_{s} that is comparable to estimates presented above (although in the literature μ_{s} was sometimes smooth spatially allowing more accurate reconstruction). The key challenge needing to be overcome with a RMC-based inversion is the presence of noise in any of the quantities estimated using MC; firstly, noise in the functional gradients means that the search direction of the linesearch is not always 'well chosen'. This has the consequence that a descent direction is not guaranteed. In isolation, this can be rectified using second order information in an l-BFGS optimisation as well as through a rigorous linesearch, but noise in $\epsilon^{(i)}$ results in an inefficient linesearch in inversions for μ_{a} where SNR is sufficient for the Wolfe conditions to be satisfied, or a completely ineffective linesearch when recovering μ_{s}. A backtracking linesearch with a well-chosen maximum and minimum step length allowed the optimisation to progress, but was found to be very unstable as reflected in the plot of the error functional at each iteration in Fig. 6.5(c).

6.6 Summary

This chapter investigated the possibility of estimating distributions of the absorption and scattering coefficients from simulated PAT images in 2D using a gradient-based optimisation approach. The reconstructions were peformed under the assumptions that the PAT image was reconstructed perfectly and that the Grüneisen parameter was known. A key question that needed to be tackled first was which optimisation algorithm would be used. The l-BFGS algorithm was believed intially to be the best candidate given its superlinear convergence and self-correcting properties. Nevertheless, it was found that noise in the functional gradients, which was propagated into the Hessian approximation, meant that these properties did not apply. The superior performance of the gradient descent approach with noisy gradients meant that it was used throughout the chapter when estimating μ_{a} and/or μ_{s}. The absorption coefficient could be estimated very accurately when 10^8 photons was used in the forward and adjoint RMC simulations, yielding an average relative error in the estimate after 39 iterations of 1.9%, even with the data containing $\sim$1% noise from MC simulation. However, the much reduced sensitivity of the error functional to the scattering coefficient meant that noise in ϵ prevented progress of

the optimisation and required the implementation of a backtracking linesearch with carefully chosen maximum and minimum step lengths. The modifications to the optimisation used to recover μ_{a} yielded reasonably accurate reconstructions of the scattering coefficient, which led to the question of whether μ_{a} and μ_{s} could be estimated simultaneously with additional data. A second image was simulated with an illumination along another boundary, and it was found that the absorption could be estimated with a 6% average relative uncertainty compared with its true value; it is believed that the relatively high level of accuracy is due to the fact that the average value of the scattering coefficient was estimated reasonably accurately, with a 13% relative uncertainty compared with its true value. The optimisation terminated after the maximum number of permitted iterations was reached, but it was observed that the error functional had not converged by this point, suggesting the estimtates could be further improved with many more iterations of the optimisation.

6.7 Future work

There are a number of features of the RMC-based inversions above that would be worth exploring further in future. These largely fall into two categories: managing the noise in the optimisation or reducing noise in the forward and adjoint RMC simulations.

Upward steps in the optimisation were sometimes required in order for the optimisation to progress, however this would ordinarily result in the termination condition being satisfied, therefore resulting in premature termination of the optimisation. Schemes to mitigate against premature termination have been presented in the literature through nonmonotone linesearch algorithms [192, 193]. One such algorithm [193] uses the mean value of the error functional over several step lengths to test modified forms of the Wolfe conditions. This approach may be able to better manage noise the linesearch which was found to be such a significant obstacle in Sections 6.4 and 6.5. Applying a penalty term to the error functional, either through Tikhonov or TV regularisation, may also help manage noise in the inversion by penalising updates with high spatial frequencies from noise in the functional gradients. Another way of managing the noise is

through resampling of the gradient. As the estimate of the functional gradient may be an inaccurate sample of the true gradient, it may be worth resampling the gradient which would involve re-running the forward and adjoint RMC simulations. As each sample can be run at lower computational cost, the true value can be inferred using a sufficiently large number of inexpensive samples, as is performed in the MCMC approach [110]. Running a great number of simulations with a lower number of photons may however still result in excessive run-times and this does not solve the problem caused by noise in the error functional when performing the linesearch (although a similar approach could be applied in resolving this issue). It may therefore be worthwhile to just perform the gradient calculation using simulations with a greater value of N_p, which would require a more efficient model; suitable approaches to accelerating RMC are discussed below.

The number of photons required in a forward or adjoint simulation to reduce noise levels in $H^{(i)}$, $\nabla\epsilon^{(i)}$ or $\epsilon^{(i)}$ to a level where it no longer impacts progress in the invesion may be several orders of magnitude greater than the N_p used in the optimisations above, particularly in 3D. To avoid computation times increasing (significantly) beyond the $\sim$10-13 hours currently required for simultaneous estimation of μ_a and μ_s, the MC model would have to be implemented for execution on massively parallel computing architectures. MCX [132] and MMC [133] are both 3D Monte Carlo models written in highly optimised CUDA code. MCX can simulate 10^{12} photons in a 5mm cubic domain with realistic tissue properties in approximately 20 minutes. However, there can be easily more than 100 runs of the model in an inversion, leading to over 30 hours for an optimisation. In order to maintain the run-time for an optimisation at around 10 hours, parallelisation to multiple GPUs would be required. Such computing facilities exist, but implementation is then complicated by the fact that the data from multiple compute nodes, each with their own GPU(s), would have to be collected on a common, host node, thus introducing further implementational challenges. When inverting large 3D datasets, consideration would have to be given to bottlenecks introduced through transferring large quantities of data across the interconnect used to pass data between nodes. Extending RMC to such an implementation is possible but would require substantial computing expertise.

Chapter 7

Conclusion and future work

Conclusion

There have been two principal objectives of this thesis; firstly, it was of interest to demonstrate the conditions in which a linear inversion, i.e. an inversion that assumes wavelength-independence of the fluence, can yield accurate oxygenation estimates and the conditions in which it fails. This is valuable because the linear inversion is frequently used in the literature without an appreciation of its shortcomings. The second objective of this thesis was to develop a novel method of performing inversions where the fluence is accounted for using an accurate, scalable and potentially highly efficient model of light transport. The outcomes and contributions of this thesis are reviewed below.

- *Using a linear inversion, the accuracy with which blood oxygenation can be estimated is limited and depends strongly on the illumination wavelengths and depth of the structure being imaged.*

  Under the assumption that PAT images can be acquired without noise and that the initial acoustic pressure can be perfectly reconstructed, the accuracy of a linear inversion used to estimate blood oxygenation from such images relies on two distinct features. Firstly, the spectra of the chromophores being estimated – in this

case, oxy- and deoxyhaemoglobin – must have sufficient linear independence, otherwise the inversion becomes ill-posed. Secondly, the images must exhibit minimal spectral colouring as the distortion of the fluence's spectrum results in nonlinearity between the absorption coefficient and the image. Chapter 3 demonstrated that, for blood-filled tissues, spectral colouring and ill-conditioning of the matrix of molar absorption coefficients, α, are prevalent for combinations of wavelengths from certain regions of the NIR-visible spectrum. The effect of spectral colouring is most severe at wavelengths in the 500-600nm range where the absorption of blood is high and undergoes large changes from one wavelength to another, while ill-conditioning of α is significant only at wavelengths where the absorption spectrum of oxy- and deoxy-haemoglobin both undergo little change. Thus, while the latter occurs for a small region of the spectrum between 840nm and 920nm, spectral colouring impacts the accuracy of a linear inversion over a large region of the spectrum. Moreover, the 500-600nm range offers significantly higher SNR compared with wavelengths $>$600nm due to the greater magnitude of the absorption coefficient of blood in this spectral range. This poses an additional issue and one that was demonstrated to markedly limit the applicability of a linear inversion in practice; in the presence of noise in the acquisition of acoustic time-series, the change in pressure from one wavelength to another will fall below the noise equivalent pressure of the detection system, meaning the inversion will become ill-posed or wholly inaccurate. It was shown that for a simple phantom consisting of a blood-filled tube submerged to a depth of 1mm in an turbid medium, the inversion will yield sO_2 estimates with greater than 50% absolute error at most wavelengths across the NIR-visible spectrum even in the presence of generously low noise levels. This was observed for a range of physiological sO_2 values. The exception to this trend was that certain wavelength pairs in the 600-1000nm range yielded sO_2 estimates of reasonable accuracy to a depth of a few mm. However, as the specific wavelength sets that yield an acceptable level of accuracy vary with tissue composition and geometry, applying this in any real imaging situation is likely to provide oxygenation estimates of misleading accuracy because the ground truth sO_2 is not known. Thus, without an optimal imaging scenario (e.g. low noise,

limited depth, minimal acoustic artefacts, high blood oxygenation), modelling of the fluence would be required to obtain sO_2 estimates that are of a useful level of accuracy.

- *The Monte Carlo method for light transport can be readily extended to simulate the radiance which can be stored compactly using frequency-domain angular discretisation.*

  Despite Monte Carlo models of light transport often being considered the gold standard model of the RTE in biomedical optics, they have found limited application in QPAT. This is in part due to the computational load associated with simulating millions, if not billions, of photons to obtain a sufficiently low-variance fluence estimate, but it is also due to the fact that without both forward and adjoint simulations of the radiance, accurate functional gradients with respect to the absorption and scattering coefficients are extremely computationally intensive to compute. Since the Monte Carlo method already models the directional portion of the RTE, access to the radiance using the MC method is straightforward because directional information of the field is carried by the photon packets. Thus, simply a method of measuring and storing the angular part of the field is required, which was demonstrated in 2D using a Fourier basis. Storing the field in the frequency-domain reduces the computational load of angular discretisation compared with the discrete ordinate method. The models developed offer a parallelisable approach to simulating the radiance, which was validated in 2D against analytic solutions for homogeneous media. The models can be readily applied to the optical inverse problem in QPAT through calculation of functional gradients with respect to the absorption and scattering coefficients.

- *Forward and adjoint simulations using the RMC algorithm allow accurate and tractable calculation of functional gradients with respect to μ_a and μ_s.*

  Chapter 5 evaluated the accuracy of functional gradients computed using the 2D forward and adjoint RMC models compared with gradients calculated via the finite difference approach using the MC forward model. A major obstacle uncovered

in this process is the sensitivity of the finite difference estimate of the functional gradient to noise in fluence estimates obtained using RMC. The finite difference method requires that the perturbation is sufficiently small for convergence of the gradient; however, convergence could not be explicitly tested due to the high levels of noise in the gradient estimates. When the perturbation was large (15% of the background μ_{s}) and 10^9 photons were simulated, validation of the adjoint-assisted and finite difference functional gradients was not possible. Reducing the perturbation would require simulation of significantly more photons, particularly in 3D due to the inherently higher variance in 3D MC models (this was not however investigated in this thesis). The adjoint-assisted approach to gradient calculation is much more robust to noise compared with the finite difference method meaning that the adjoint method could be used to estimate distributions of μ_{a} and μ_{s} in QPAT, which had not been achieved using a MC model before.

- *Estimation of the distribution of μ_{a} and μ_{s}, independently and simultaneously, is possible to a practically useful degree of accuracy using a gradient-based optimisation and functional gradients obtained using forward and adjoint RMC simulations.*
  Using the methods presented in Chapter 5 for gradient calculation, inversions for the absorption and scattering coefficients could be attempted. A significant challenge encountered in solving the inverse problem was the suitability of the optimisation technique. As the number of photons simulated in the forward and adjoint Monte Carlo simulations will determine the uncertainty in scored quantities like fluence estimates and Fourier harmonics, this uncertainty is propagated into the estimates of the functional gradients and the error functional. The impact of this is two-fold; firstly, the search direction is not optimal and, secondly, noise in the error functional will reduce the chance of a step length satisfying the Wolfe conditions in the linesearch. The former led to the gradient descent algorithm out-performing the l-BFGS approach in the recovery of μ_{a} as the SNR in the functional gradient and error functional was sufficient when inverting for the absorption coefficient meaning the inversion could progress efficiently. The inversion yielded an estimate of the absorption coefficient whose accuracy was to within approximately the same

noise level on average over the domain as in the measured data. However, when estimating the scattering coefficient or both μ_{a} and μ_{s} simultaneously, SNR in the gradients and error functional became problematic and the optimisation had to be forced to take steps, otherwise it would not have progressed. This approach produced relatively accurate reconstructions for the absorption coefficient, but convergence of the scattering coefficient was comparatively much worse (although this is in part due to the diffusive nature of scattering, i.e. this problem is not restricted to RMC-based inversions). Since it is the absorption coefficient that is of more value from a clinical perspective, it can be argued the the accuracy of the scattering must just be sufficient to provide an accurate estimate of the absorption coefficient. Preliminary inversions in Chapter 6 revealed that recovery of the average value of μ_{s} over the domain may allow the inversion to yield estimates of the absorption coefficient that are of a clinically useful level of accuracy.

The highly limited accuracy that can be obtained using inversions that do not account for the fluence was apparent in Chapter 3. There are, however, many practical and implementational challenges standing in the way of the methods proposed in Chapters 4–6 being useful in a clinical or pre-clinical setting. First, the inversions were attempted in a near-ideal scenario in which it was assumed that the Grüneisen parameter was known and that the PAT image was perfectly reconstructed. As discussed in Chapter 1, these assumptions rarely apply in practice, but approaches to remedy these issues are already in existence. For example, the Grüneisen parameter can be divided out of the image using by taking a ratio of images at two different wavelengths [98] or it can be recovered using additional data [104], and limited-vew artefacts can be reduced using a pair of orthogonal detectors [194]. The implementational challenges associated with massive parallelisation of RMC cited at the end of Chapter 6 would need to be tackled in order for a RMC-based inversion approach to be viable. The primary reason for this is that although the implementation of radiance Monte Carlo in 3D is quite similar to that in 2D, several orders of magnitude more photons would be required to minimise noise levels in the inversions, which would only be possible via quite substantial parallelisation. Nevertheless, the hardware and expertise is available for such advancements meaning Monte

Carlo methods may become a more commonly used tool in modelling light transport in QPAT.

Future work

In its current state, the field of QPAT is faced with many challenges. This thesis attempts to answer two key questions; (a) under what conditions is it necessary to model the fluence to obtain reasonably accurate quantitative information (sO_2 , chromophore concentrations, etc.)? And (b) how can the computational challenges associated with modelling the fluence in a model-based inversion scheme be overcome? Although the conclusions highlighted in the previous section provide a significant contribution to the field through answering the above questions, the answers are not complete and further research is desirable. This section aims to underline the aspects of further research that are expected to yield results that advance the field of QPAT.

Chapter 3 demonstrated that there are a few circumstances in which a linear inversion will yield accurate results, namely when wavelengths are 'well-chosen', the object of interest is shallow and the noise in the image is low. These parameters are of course interrelated; additional questions that are therefore worth attempting to answer is what level of SNR would be necessary for well-chosen wavelengths not to be corrupted by noise, and how do optimal wavelengths change as the target structure changes with depth or if the constituent chromophores change (for example through the introduction of a CSF layer in the brain)? This systematic analysis could be extended to investigate how the accuracy of a linear inversion changes when an artery is in close proximity to a vein, which has a differing absorption spectrum.

It was also found in Chapter 3 that if the sO_2 estimates obtained using a series of two-wavelength linear inversions using wavelengths from across the NIR/visisble spectrum are plotted as a histogram, the mean of the histogram appears to be near the true oxygenation of the target. This superficial investigation requires a more careful study to understand why the average of many two-wavelength sO_2 estimates tends to the average vessel oxygenation and whether this observation applies more generally to different

tissue structures and constititutions. A first approach would be to perform additional simulations, ideally using a 3D transport model such as Monte Carlo, to investigate how the shape of the histrogram of sO_2 estimates changes with target depth and in the presence of more vasculature. This could then be extended to tube phantom or *in vitro* studies, provided the true chromophore concentration ratio can be ascertained via an independent measurement.

A more in-depth understanding of the circumstances in which a linear inversion is likely to yield reliablly accurate sO_2 estimates would also help identify when a model-based inversion is required. Chapter 4 presented a Monte Carlo model of radiance which was applied within a gradient-based inversion scheme to estimate the absorption and scattering coefficients. Two limitations of the model are that it is 2D and has undergone minimal parallelisation. This had the consequence that inversions in Chapter 6 could only be carried out on 2D problems and computation times when inverting for the scattering coefficient were in the 11-14 hour range. A 3D GPU-accelerated implementation of the algorithm, using a spherical harmonic angular basis, would mean that 3D problems could be solved efficiently and, potentially, improve the estimate of the scattering, given lower variance gradient estimates could be obtained.

Writing efficient GPU code, e.g. in NVIDIA CUDA, is challenging; for this reason, writing a 3D spherical harmonic RMC from scratch may not be a worthwhile endeavour, especially with the availability of open-source codes such as MCX [132]. In order to modify this code, the spherical harmonic coefficients would have to be calculated for each photon packet in parallel, which would result in slower photon propagation times. Nevertheless, this would most likely be the most straightforward means of arriving at a GPU-accelerated 3D radiance MC model.

While these areas of research remain key to advancing the field of QPAT, they are by no means the only challenges. Chapter 2 highlighted a number of other obstacles standing in the way of obtaining accurate quantitative information using PAT, including the absorption-scattering nonuniqueness problem and the Grüneisen parameter being unknown. One strategy for overcoming ill-posedness due to an unknown Grüneisen parameter involves using a ratio of PAT images as the data [104]. This is unattractive because

ratio data contains more noise; however, it is possible to reconstruct the Grüneisen parameter simultaneously with the absorption and scattering coefficient with additional data [104]. To reduce instability when inverting for these three parameters simultaneously, Pulkkinen et al. have shown that spectral constraints on the scattering using a Mie scattering relationship and the inclusion of spatial priors that promote structure in the parameter estimates yields improved accuracy of the reconstruction [195]. Constraints introduced through spectral classification via a mixture of Gaussians model have also been shown to improve the accuracy of absorption and scatering estimates [196], and this method could be extended to reconstruct the Grüneisen parameter simultaneously. These approaches can be applied in conjunction with the radiance Monte Carlo forward and adjoint models proposed in this thesis. While parallelisation of the algorithm remains a priority, the use of multiwavelength data and constraining the wavelength-dependent scattering coefficient could potentially help reduce ill-posedness and yield better estimate of μ_{s}, improving upon results obtained in Chapter 6.

There is a growing body of methods used to solve inverse problems that have not yet been applied in QPAT. Much of the focus in QPAT has been on model-fitting approaches, but the scale of the problem, the spectral case in particular, and the computational demands of the forward and inverse models are typically so large that these techniques are impractical. Methods applied in mature fields like image segmentation may provide worthwhile avenues for research in QPAT; techniques such as machine learning, Markov random fields and atlas-based methods may be able to exploit the spatial and spectral information in PAT images to extract quantitative measures.

Bibliography

[1] Paul Suetens. *Fundamentals of Medical Imaging*. Cambridge University Press, Cambridge, UK, 2004. URL http://dx.doi.org/10.1017/CBO9780511596803.

[2] M Symms, H R Jäger, K Schmierer, and T a Yousry. A review of structural magnetic resonance neuroimaging. *Journal of neurology, neurosurgery, and psychiatry*, 75(9):1235–44, September 2004. ISSN 0022-3050. doi: 10.1136/jnnp.2003.032714. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1739217&tool=pmcentrez&rendertype=abstract.

[3] William Swindell and Stephen E D Webb. X-ray transmission computed tomography. In Stephen E D Webb, editor, *The Physics of Medical Imaging*, chapter 4, pages 98–127. Adam Hilger, Bristol, 1st edition, 1988.

[4] P N T Wells. Ultrasound imaging. *Physics in medicine and biology*, 51(13):R83–98, July 2006. ISSN 0031-9155. doi: 10.1088/0031-9155/51/13/R06. URL http://www.ncbi.nlm.nih.gov/pubmed/16790922.

[5] Alexander Graham Bell. Upon the production and reproduction of sound by light. *Journal of the Society of Telepraph Engineers*, 9(30):404–426, 1880.

[6] S J Davies, C Edward, G S Taylor, and S B Palmers. Laser-generated ultrasound: its properties, mechanisms and multifarious applications. *Journal of Physics D: Applied Physics*, 26:329–348, 1993.

[7] Markus W. Sigrist. Laser generation of acoustic waves in liquids and gases. *Journal of Applied Physics*, 60(7):R83, 1986. ISSN 00218979. doi: 10.1063/1.337089. URL http://link.aip.org/link/JAPIAU/v60/i7/pR83/s1&Agg=doi.

[8] RA Kruger, P Liu, and CR Appledorn. Photoacoustic ultrasound (PAUS)âĂŤreconstruction tomography. *Medical Physics*, 22(10):1605–1609, 1995. URL http://link.aip.org/link/?MPHYA6/22/1605/1.

[9] Steven L Jacques. Optical properties of biological tissues: a review. *Physics in Medicine and Biology*, 58(14):5007–5008, July 2013. ISSN 0031-9155. doi: 10.1088/0031-9155/58/14/5007. URL http://stacks.iop.org/0031-9155/58/i=14/a=5007?key=crossref.c531ac37cfa1f77bbc0d33a8e92de8c9.

[10] Konstantin Maslov, Hao F Zhang, Song Hu, and Lihong V Wang. Optical-resolution photoacoustic microscopy for in vivo imaging of single capillaries. *Optics Letters*, 33(9):929–31, May 2008. ISSN 0146-9592. URL http://www.ncbi.nlm.nih.gov/pubmed/18451942.

[11] Edward Zhang, Jan Laufer, and Paul Beard. Backward-mode multiwavelength photoacoustic scanner using a planar Fabry-Perot polymer film ultrasound sensor for high-resolution three-dimensional imaging of biological tissues. *Applied Optics*, 47(4):561–577, March 2008. ISSN 0003-6935. URL http://www.ncbi.nlm.nih.gov/pubmed/18239717.

[12] PC Beard and TN Mills. Miniature optical fibre ultrasonic hydrophone using a Fabry-Perot polymer film interferometer. *Electronics Letters*, 33(9):801 – 803, 1997. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=590258.

[13] GJ Diebold, T Sun, and MI Khan. Photoacoustic monopole radiation in one, two, and three dimensions. *Physical Review Letters*, 67(24):3384–3387, 1991. URL http://link.aps.org/doi/10.1103/PhysRevLett.67.3384.

[14] GJ Diebold and PJ Westervelt. The photoacoustic effect generated by a spherical droplet in a fluid. *The Journal of the Acoustical Society of America*, 84:2245–2251, 1988. URL http://link.aip.org/link/?JASMAN/84/2245/1.

[15] G. J. Diebold and T. Sun. Properties of Photoacoustic Waves in One, Two and Three Dimensions. *Acustica*, 80:339–351, 1994.

[16] M Khan. The photoacoustic effect generated by laser irradiation of an isotropic solid cylinder. *Ultrasonics*, 34(1):19–24, March 1996. ISSN 0041624X. doi: 10.1016/0041-624X(95)00128-P. URL http://linkinghub.elsevier.com/retrieve/pii/0041624X9500128P.

[17] Yae-Lin Sheu and Pai-Chi Li. Simulations of thermally induced photoacoustic wave propagation using a pseudospectral time-domain method. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 56(5):1104–12, May 2009. ISSN 1525-8955. doi: 10.1109/TUFFC.2009.1144. URL http://www.ncbi.nlm.nih.gov/pubmed/19473928.

[18] Bradley E Treeby and B T Cox. k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields. *Journal of Biomedical Optics*, 15(2):021314, 2010. ISSN 1560-2281. doi: 10.1117/1.3360308. URL http://www.ncbi.nlm.nih.gov/pubmed/20459236.

[19] Robert A. Kruger, Richard B. Lam, Daniel R. Reinecke, Stephen P. Del Rio, and Ryan P. Doyle. Photoacoustic angiography of the breast. *Medical Physics*, 37(11):6096, 2010. ISSN 00942405. doi: 10.1118/1.3497677. URL http://link.aip.org/link/MPHYA6/v37/i11/p6096/s1&Agg=doi.

[20] Xueding Wang, Yongjiang Pang, Geng Ku, Xueyi Xie, George Stoica, and Lihong V Wang. Noninvasive laser-induced photoacoustic tomography for structural and functional in vivo imaging of the brain. *Nature Biotechnology*, 21(7):803–806, July 2003. ISSN 1087-0156. doi: 10.1038/nbt839. URL http://www.ncbi.nlm.nih.gov/pubmed/12808463.

[21] Konstantin Maslov, Gheorghe Stoica, and Lihong V Wang. In vivo dark-field reflection-mode photoacoustic microscopy. *Optics Letters*, 30(6):625–7, March 2005. ISSN 0146-9592. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2644744&tool=pmcentrez&rendertype=abstract.

[22] Hao F Zhang, Konstantin Maslov, George Stoica, and Lihong V Wang. Functional photoacoustic microscopy for high-resolution and noninvasive in vivo imaging.

Nature biotechnology, 24(7):848–51, July 2006. ISSN 1087-0156. doi: 10.1038/nbt1220. URL http://www.ncbi.nlm.nih.gov/pubmed/16823374.

[23] Z. Xie, S. Jiao, H.F. Zhang, and C.A. Puliafito. Laser-scanning optical-resolution photoacoustic microscopy. *Optics Letters*, 34(12):1771–1773, 2009. ISSN 1539-4794. URL http://www.opticsinfobase.org/abstract.cfm?URI=ol-34-12-1771.

[24] KP Köstli, Martin Frenz, HP Weber, G Paltauf, and H. Optoacoustic tomography: time-gated measurement of pressure distributions and image reconstruction. *Applied*, 2001. URL http://www.opticsinfobase.org/abstract.cfm?&id=64951.

[25] Leonid a Kunyansky. A series solution and a fast algorithm for the inversion of the spherical mean Radon transform. *Inverse Problems*, 23(6):S11–S20, December 2007. ISSN 0266-5611. doi: 10.1088/0266-5611/23/6/S02. URL http://stacks.iop.org/0266-5611/23/i=6/a=S02?key=crossref.8d8e7d1fb7715c832fe23432cc59c22b.

[26] Kun Wang and Mark A Anastasio. A simple Fourier transform-based reconstruction formula for photoacoustic computed tomography with a circular or spherical measurement geometry. *Physics in Medicine & Biology*, 57:493–499, 2012.

[27] Leonid a Kunyansky. Explicit inversion formulae for the spherical mean Radon transform. *Inverse Problems*, 23(1):373–383, February 2007. ISSN 0266-5611. doi: 10.1088/0266-5611/23/1/021. URL http://stacks.iop.org/0266-5611/23/i=1/a=021?key=crossref.17fcb756c4dac79c8ed5f29d96abbc59.

[28] Minghua Xu and Lihong V. Wang. Universal back-projection algorithm for photoacoustic computed tomography. *Physical Review E*, 71(1):1–7, January 2005. ISSN 1539-3755. doi: 10.1103/PhysRevE.71.016706. URL http://link.aps.org/doi/10.1103/PhysRevE.71.016706.

[29] Minghua Xu and Lihong V Wang. Time-domain reconstruction for thermoacoustic tomography in a spherical geometry. *IEEE Transactions on Medical Imaging*, 21

(7):814–22, July 2002. ISSN 0278-0062. doi: 10.1109/TMI.2002.801176. URL http://www.ncbi.nlm.nih.gov/pubmed/12374318.

[30] Leonid Kunyansky. Fast reconstruction algorithms for the thermoacoustic tomography in certain domains with cylindrical or spherical symmetries. *Inverse Problems and Imaging*, 6(1):111–131, February 2012. ISSN 1930-8337. doi: 10.3934/ipi.2012.6.111. URL http://www.aimsciences.org/journals/displayArticlesnew.jsp?paperID=7012.

[31] Yuan Xu, Dazi Feng, and Lihong V Wang. Exact frequency-domain reconstruction for thermoacoustic tomography–I: Planar geometry. *IEEE transactions on medical imaging*, 21(7):823–8, July 2002. ISSN 0278-0062. URL http://www.ncbi.nlm.nih.gov/pubmed/12374319.

[32] David Finch, Sarah K. Patch, and Rakesh. Determining a Function from Its Mean Values Over a Family of Spheres. *SIAM Journal on Mathematical Analysis*, 35(5):1213–1240, 2004. ISSN 00361410. doi: 10.1137/S0036141002417814. URL http://link.aip.org/link/SJMAAH/v35/i5/p1213/s1&Agg=doi.

[33] Yuan Xu and Lihong Wang. Time Reversal and Its Application to Tomography with Diffracting Sources. *Physical Review Letters*, 92(3):3–6, January 2004. ISSN 0031-9007. doi: 10.1103/PhysRevLett.92.033902. URL http://link.aps.org/doi/10.1103/PhysRevLett.92.033902.

[34] B T Cox, S Kara, S R Arridge, and P C Beard. k-space propagation models for acoustically heterogeneous media: application to biomedical photoacoustics. *The Journal of the Acoustical Society of America*, 121(6):3453–64, June 2007. ISSN 1520-8524. doi: 10.1121/1.2717409. URL http://www.ncbi.nlm.nih.gov/pubmed/17552697.

[35] Bradley E Treeby and B T Cox. A k-space Green's function solution for acoustic initial value problems in homogeneous media with power law absorption. *The Journal of the Acoustical Society of America*, 129(6):3652, June 2011. ISSN 1520-8524. doi: 10.1121/1.3583537. URL http://www.ncbi.nlm.nih.gov/pubmed/21682390.

[36] Bradley E Treeby, Edward Z Zhang, and B T Cox. Photoacoustic tomography in absorbing acoustic media using time reversal. *Inverse Problems*, 26(11):115003, November 2010. ISSN 0266-5611. doi: 10.1088/0266-5611/26/11/115003. URL http://stacks.iop.org/0266-5611/26/i=11/a=115003http://stacks.iop.org/0266-5611/26/i=11/a=115003?key=crossref.6b88d3cc9ced08fc3be5de2bc14bd8ee.

[37] Yulia Hristova, Peter Kuchment, and Linh Nguyen. Reconstruction and time reversal in thermoacoustic tomography in acoustically homogeneous and inhomogeneous media. *Inverse Problems*, 24(5):055006, October 2008. ISSN 0266-5611. doi: 10.1088/0266-5611/24/5/055006. URL http://stacks.iop.org/0266-5611/24/i=5/a=055006?key=crossref.560c30dd771f35dc518aa2a27ffd80cc.

[38] Christian Lutzweiler and Daniel Razansky. Optoacoustic Imaging and Tomography: Reconstruction Approaches and Outstanding Challenges in Image Performance and Quantification. *Sensors*, 13(6):7345–7384, June 2013. ISSN 1424-8220. doi: 10.3390/s130607345. URL http://www.mdpi.com/1424-8220/13/6/7345/.

[39] Amir Rosenthal, Daniel Razansky, and Vasilis Ntziachristos. Fast semi-analytical model-based acoustic inversion for quantitative optoacoustic tomography. *IEEE transactions on medical imaging*, 29(6):1275–85, June 2010. ISSN 1558-0062. doi: 10.1109/TMI.2010.2044584. URL http://www.ncbi.nlm.nih.gov/pubmed/20304725.

[40] Kun Wang, Sergey a Ermilov, Richard Su, Hans-Peter Brecht, Alexander a Oraevsky, and Mark a Anastasio. An imaging model incorporating ultrasonic transducer properties for three-dimensional optoacoustic tomography. *IEEE transactions on medical imaging*, 30(2):203–14, February 2011. ISSN 1558-254X. doi: 10.1109/TMI.2010.2072514. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3033994&tool=pmcentrez&rendertype=abstract.

[41] Mark a Anastasio, Jin Zhang, Xiaochuan Pan, Yu Zou, Geng Ku, and Lihong V Wang. Half-time image reconstruction in thermoacoustic tomography. *IEEE*

transactions on medical imaging, 24(2):199–210, February 2005. ISSN 0278-0062. URL http://www.ncbi.nlm.nih.gov/pubmed/15707246.

[42] X Luís Deán-Ben, Andreas Buehler, Vasilis Ntziachristos, and Daniel Razansky. Accurate model-based reconstruction algorithm for three-dimensional optoacoustic tomography. *IEEE Transactions on Medical Imaging*, 31(10):1922–8, October 2012. ISSN 1558-254X. doi: 10.1109/TMI.2012.2208471. URL http://www.ncbi.nlm.nih.gov/pubmed/23033065.

[43] B. E. Treeby, T. K. Varslot, E. Z. Zhang, J. G. Laufer, and P. C. Beard. Automatic sound speed selection in photoacoustic image reconstruction using an autofocus approach. *Journal of Biomedical Optics*, 16(9):090501, 2011. ISSN 10833668. doi: 10.1117/1.3619139. URL http://link.aip.org/link/JBOPFO/v16/i9/p090501/s1&Agg=doi.

[44] Jithin Jose, Rene G H Willemink, Wiendelt Steenbergen, C H Slump, Ton G van Leeuwen, and Srirang Manohar. Speed-of-sound compensated photoacoustic tomography for accurate imaging. *Medical Physics*, 39(12):7262–71, December 2012. ISSN 0094-2405. doi: 10.1118/1.4764911. URL http://www.ncbi.nlm.nih.gov/pubmed/23231277.

[45] Dimple Modgil, Mark a. Anastasio, Kun Wang, and Patrick J. LaRiviere. Image reconstruction in photoacoustic tomography with variable speed of sound using a higher order geometrical acoustics approximation. *Proceedings of SPIE*, 7177: 71771A–71771A–8, 2009. doi: 10.1117/12.809001. URL http://link.aip.org/link/PSISDG/v7177/i1/p71771A/s1&Agg=doi.

[46] K. Daoudi, A. Hussain, E. Hondebrink, and W. Steenbergen. Correcting photoacoustic signals for fluence variations using acousto-optic modulation. *Optics Express*, 20(13):14117, June 2012. ISSN 1094-4087. doi: 10.1364/OE.20.014117. URL http://www.opticsinfobase.org/abstract.cfm?URI=oe-20-13-14117.

[47] Lu Yin, Qiang Wang, Qizhi Zhang, and Huabei Jiang. Tomographic imaging of absolute optical absorption coefficient in turbid media using combined photoacoustic and diffusing light measurements. *Optics Letters*, 32(17):2556, 2007. ISSN

0146-9592. doi: 10.1364/OL.32.002556. URL http://www.opticsinfobase.org/abstract.cfm?URI=ol-32-17-2556.

[48] Adam Q. Bauer, Ralph E. Nothdurft, Todd N. Erpelding, Lihong V. Wang, and Joseph P. Culver. Quantitative photoacoustic imaging: correcting for heterogeneous light fluence distributions using diffuse optical tomography. *Journal of Biomedical Optics*, 16(9):096016, September 2011. ISSN 10833668. doi: 10.1117/1.3626212. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3188642&tool=pmcentrez&rendertype=abstracthttp://link.aip.org/link/JBOPFO/v16/i9/p096016/s1&Agg=doi.

[49] BT Cox, JG Laufer, and PC Beard. Quantitative Photoacoustic Image Reconstruction using Fluence Dependent Chromophores. *Biomedical Optics Express*, 1(1): 5830–5835, 2010. URL http://www.opticsinfobase.org/abstract.cfm?uri=boe-1-1-201&seq=0.

[50] X. Luis Dean-Ben, Andre C Stiel, Yuanyuan Jiang, Vasilis Ntziachristos, Gil G Westmeyer, and Daniel Razansky. Light fluence normalization in turbid tissues via temporally unmixed multispectral optoacoustic tomography. *Optics Letters*, 40(20):4691, 2015.

[51] Konstantin Maslov, Mathangi Sivaramakrishnan, Hao F. Zhang, George Stoica, and Lihong V. Wang. Technical considerations in quantitative blood oxygenation measurement using photoacoustic microscopy in vivo. *Proc. of SPIE, Photons Plus Ultrasound: Imaging and Sensing*, 6086:60860R–60860R–11, February 2006. ISSN 0277786X. doi: 10.1117/12.646265. URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?articleid=1273958http://link.aip.org/link/PSISDG/v6086/i1/p60860R/s1&Agg=doi.

[52] Scott A Prahl. *Light Transport in Tissue*. PhD thesis, University of Texas at Austin, January 1988. URL http://www.ncbi.nlm.nih.gov/pubmed/19336841.

[53] Tanja Tarvainen. *Computational Methods for Light Transport in Optical Tomography*. PhD thesis, University of Kuopio, 2006.

[54] H. De Raedt, K. Michielsen, J.S. Kole, and M.T. Figge. Solving the Maxwell equations by the Chebyshev method: a one-step finite-difference time-domain algorithm. *IEEE Transactions on Antennas and Propagation*, 51(11):3155–3160, November 2003. ISSN 0018-926X. doi: 10.1109/TAP.2003.818809. URL http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1243513.

[55] B. T. Cox, J. G. Laufer, and P. C. Beard. The challenges for quantitative photoacoustic imaging. *Proceedings of SPIE*, 7177:717713–717713–9, 2009. doi: 10.1117/12.806788. URL http://link.aip.org/link/PSISDG/v7177/i1/p717713/s1&Agg=doi.

[56] Xosé Luís Deán-Ben, Erwin Bay, and Daniel Razansky. Functional optoacoustic imaging of moving objects using microsecond-delay acquisition of multispectral three-dimensional tomographic data. *Nature Scientific Reports*, 4:5878, January 2014. ISSN 2045-2322. doi: 10.1038/srep05878. URL http://www.ncbi.nlm.nih.gov/pubmed/25073504.

[57] Akira Ishimaru. *Wave Propagation and Scattering in Random Media. Vol 1: Single Scattering and Transport Theory*. Academic Press, 1978.

[58] Michael I Mishchenko, Larry Travis, and Andrew A Lacis. *Scattering, Absorption, and Emission of Light by Small Particles*. Cambridge University Press, Cambridge, 3rd edition, 2002. ISBN 0-521-78252-X. URL http://onlinelibrary.wiley.com/doi/10.1002/cbdv.200490137/abstract.

[59] Jorge Ripoll. Derivation of the scalar radiative transfer equation from energy conservation of Maxwell's equations in the far field. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 28(8):1765–75, August 2011. ISSN 1520-8532. URL http://www.ncbi.nlm.nih.gov/pubmed/21811340.

[60] A. J. Welch and Martin J. C. van Gemert, editors. *Optical-Thermal Response of Laser-Irradiated Tissue*. Springer Berlin Heidelberg, 2nd edition, 1995.

[61] LG Henyey and JL Greenstein. Diffuse radiation in the galaxy. In *Annales d'Astrophysique*, volume 3, pages 70–84, 1940. URL http://adsabs.harvard.edu/full/1940AnAp....3..117H.

[62] JH Joseph and WJ Wiscombe. The delta-Eddington approximation for radiative flux transfer. *Journal of the Atmospheric Sciences*, 33:2452–2459, 1976. URL http://ams.allenpress.com/perlserv/?request=get-abstract&doi=10.1175/1520-0469(1976)033<2452:TDEAFR>2.0.CO;2.

[63] S.R. Arridge. Optical tomography in medical imaging. *Inverse problems*, 15: R41–R93, 1999. URL http://iopscience.iop.org/0266-5611/15/2/022.

[64] P. Surya Mohan, Tanja Tarvainen, Martin Schweiger, Aki Pulkkinen, and Simon R. Arridge. Variable order spherical harmonic expansion scheme for the radiative transport equation using finite elements. *Journal of Computational Physics*, 230 (19):7364–7383, August 2011. ISSN 00219991. doi: 10.1016/j.jcp.2011.06.004. URL http://linkinghub.elsevier.com/retrieve/pii/S0021999111003536.

[65] WJ Wiscombe. The delta-M method: rapid yet accurate radiative flux calculations for strongly asymmetric phase functions. *Journal of the Atmospheric Sciences*, 34:1408–1422, 1977. URL http://journals.ametsoc.org/doi/abs/10.1175/1520-0469(1977)034%3C1408:TDMRYA%3E2.0.CO%3B2.

[66] Andre Liemert and Alwin Kienle. Infinite space GreenâĂŹ s function of the time-dependent radiative transfer equation. *Biomedical Optics Express*, 3(3):543–551, 2012.

[67] André Liemert and Alwin Kienle. Radiative transfer in two-dimensional infinitely extended scattering media. *Journal of Physics A: Mathematical and Theoretical*, 44(50):505206, December 2011. ISSN 1751-8113. doi: 10.1088/1751-8113/44/50/505206. URL http://stacks.iop.org/1751-8121/44/i=50/a=505206?key=crossref.a5becdf9af26a10d10d813161a4eb4f5.

[68] André Liemert and Alwin Kienle. Novel analytical solution for the radiance in an anisotropically scattering medium. *Applied optics*, 54(8):1963–9, March 2015. ISSN 1539-4522. URL http://www.ncbi.nlm.nih.gov/pubmed/25968372.

[69] Fabrizio Martelli, Angelo Sassaroli, Yukio Yamada, and Giovanni Zaccanti. Analytical approximate solutions of the time-domain diffusion equation in layered slabs. *Journal of the Optical Society of America A*, 19(1):71, 2002. ISSN 1084-7529. doi:

10.1364/JOSAA.19.000071. URL http://www.opticsinfobase.org/abstract.cfm?URI=JOSAA-19-1-71.

[70] Simon R Arridge and Jeremy C Hebden. Optical imaging in medicine: II. Modelling and reconstruction. *Physics in Medicine and Biology*, 42(5): 841–853, May 1997. ISSN 0031-9155. doi: 10.1088/0031-9155/42/5/008. URL http://stacks.iop.org/0031-9155/42/i=5/a=008?key=crossref.a26fa5f4e9425e09e5cef5bc6cfd26ca.

[71] André Liemert and Alwin Kienle. Light transport in three-dimensional semi-infinite scattering media. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 29(7):1475–81, July 2012. ISSN 1520-8532. URL http://www.ncbi.nlm.nih.gov/pubmed/22751416.

[72] André Liemert and Alwin Kienle. Green's functions for the two-dimensional radiative transfer equation in bounded media. *Journal of Physics A: Mathematical and Theoretical*, 45(17):175201, May 2012. ISSN 1751-8113. doi: 10.1088/1751-8113/45/17/175201. URL http://stacks.iop.org/1751-8121/45/i=17/a=175201?key=crossref.a2bc108b9fd65fce29c0a6978782be53.

[73] André Liemert and Alwin Kienle. Analytical approach for solving the radiative transfer equation in two-dimensional layered media. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 113(7):559–564, May 2012. ISSN 00224073. doi: 10.1016/j.jqsrt.2012.01.013. URL http://linkinghub.elsevier.com/retrieve/pii/S0022407312000374.

[74] Andre Liemert. Personal Communication, 30-07-2015.

[75] WM Star, JPA Marijnissen, and M.J.C. Gemert. Light dosimetry in optical phantoms and in tissues: I. Multiple flux and transport theory. *Physics in Medicine and Biology*, 33(4):437, 1988. URL http://iopscience.iop.org/0031-9155/33/4/004.

[76] Martin J. C. Van Gemert, Ashley J. Welch, Willem M. Star, Massoud Motamedi, and Wai-Fung Cheong. Tissue optics for a slab geometry in the diffusion approximation. *Lasers in Medical Science*, 2(4):295–302, December 1987. ISSN

0268-8921. doi: 10.1007/BF02594174. URL http://link.springer.com/10.1007/BF02594174.

[77] Bedrettin Subasilar. Analytical approaches to the delta-Eddington model of the radiative transfer through vertically inhomogeneous optical depths. *Applied Mathematical Modelling*, 32(4):514–534, April 2008. ISSN 0307904X. doi: 10.1016/j.apm.2007.02.005. URL http://linkinghub.elsevier.com/retrieve/pii/S0307904X07000261.

[78] Daniele Contini, Fabrizio Martelli, and Giovanni Zaccanti. Photon migration through a turbid slab described by a model based on diffusion approximation. I. Theory. *Applied Optics*, 36(19):4587, July 1997. ISSN 0003-6935. doi: 10.1364/AO.36.004587. URL https://www.osapublishing.org/ao/abstract.cfm?uri=ao-36-19-4587.

[79] AH Hielscher. Comparison of finite-difference transport and diffusion calculations for photon migration in homogeneous and heterogeneous tissues. *Physics in Medicine and Biology*, 43(1285), 1998. URL http://iopscience.iop.org/0031-9155/43/5/017.

[80] M Keijzer, W M Star, and P R Storchi. Optical diffusion in layered media. *Applied optics*, 27(9):1820–4, May 1988. ISSN 0003-6935. URL http://www.ncbi.nlm.nih.gov/pubmed/20531660.

[81] Arnold D. Kim and Miguel Moscoso. Radiative transfer computations for optical beams. *Journal of Computational Physics*, 185(1):50–60, February 2003. ISSN 00219991. doi: 10.1016/S0021-9991(02)00047-5. URL http://linkinghub.elsevier.com/retrieve/pii/S0021999102000475.

[82] OC Zienkiewicz, RL Taylor, OC Zienkiewicz, and RL Taylor. *The finite element method*. McGraw-Hill, Woburn, fifth edition, 1977. ISBN 0750650494. URL http://onlinelibrary.wiley.com/doi/10.1002/cbdv.200490137/abstracthttp://civil.dept.shef.ac.uk/current/module/CIV4135.pdf.

[83] Alexander D. Klose and Edward W. Larsen. Light transport in biological tissue based on the simplified spherical harmonics equations. *Journal of Computational Physics*, 220(1):441–470, December 2006. ISSN 00219991. doi: 10.1016/j.jcp.2006.07.007. URL http://linkinghub.elsevier.com/retrieve/pii/S0021999106003421.

[84] W. Cong, H. Shen, a. Cong, Y. Wang, and G. Wang. Modeling photon propagation in biological tissues using a generalized Delta-Eddington phase function. *Physical Review E*, 76(5):1–5, November 2007. ISSN 1539-3755. doi: 10.1103/PhysRevE.76.051913. URL http://link.aps.org/doi/10.1103/PhysRevE.76.051913.

[85] SR Arridge, M Schweiger, and M. A finite element approach for modeling photon transport in tissue. *Medical Physics*, 20(2):299–309, 1993. URL http://engineering.dartmouth.edu/~Amir_H_Golnabi/ENGG168_files/Papers/arridge_fem.pdf.

[86] Caigang Zhu and Quan Liu. Review of Monte Carlo modeling of light transport in tissues. *Journal of Biomedical Optics*, 18(5):50902, May 2013. ISSN 1560-2281. doi: 10.1117/1.JBO.18.5.050902. URL http://www.ncbi.nlm.nih.gov/pubmed/23698318.

[87] Rinat O. Esenaliev. Application of light and ultrasound for medical diagnostics and treatment. *Proceedings of SPIE*, 4707(409):158–164, 2002. ISSN 0277786X. doi: 10.1117/12.475583. URL http://link.aip.org/link/?PSI/4707/158/1&Agg=doi.

[88] Elena V. Savateeva, Alexander Karabutov, Vladimir Solomatin, and Alexander Oraevsky. Optical properties of blood at various levels of oxygenation studied by time resolved detection of laser-induced pressure profiles. *Proceedings of SPIE*, 4618:63–75, 2002. ISSN 0277786X. doi: 10.1117/12.469849. URL http://link.aip.org/link/?PSI/4618/63/1&Agg=doi.

[89] Jan Laufer, Clare Elwell, Dave Delpy, and Paul Beard. In vitro measurements of absolute blood oxygen saturation using pulsed near-infrared photoacoustic spectroscopy: accuracy and resolution. *Physics in Medicine and Biology*, 50(18):

4409–4428, September 2005. ISSN 0031-9155. doi: 10.1088/0031-9155/50/18/011. URL http://www.ncbi.nlm.nih.gov/pubmed/16148401.

[90] Zhongjiang Chen, Sihua Yang, and Da Xing. In vivo detection of hemoglobin oxygen saturation and carboxyhemoglobin saturation with multiwavelength photoacoustic microscopy. *Optics Letters*, 37(16):3414, August 2012. ISSN 0146-9592. doi: 10.1364/OL.37.003414. URL http://www.opticsinfobase.org/abstract.cfm?URI=ol-37-16-3414.

[91] Raul Fainchtein. Local determination of hemoglobin concentration and degree of oxygenation in tissue by pulsed photoacoustic spectroscopy. *Proceedings of SPIE*, 3916:19–33, 2000. ISSN 0277786X. doi: 10.1117/12.386335. URL http://link.aip.org/link/?PSI/3916/19/1&Agg=doi.

[92] Keith M. Stantz. Photoacoustic spectroscopic imaging of intra-tumor heterogeneity and molecular identification. *Proceedings of SPIE*, 6086:608605–608605–12, 2006. ISSN 0277786X. doi: 10.1117/12.645106. URL http://link.aip.org/link/PSISDG/v6086/i1/p608605/s1&Agg=doi.

[93] W Verkruysse, G W Lucassen, J F de Boer, D J Smithies, J S Nelson, and M J van Gemert. Modelling light distributions of homogeneous versus discrete absorbers in light irradiated turbid media. *Physics in medicine and biology*, 42(1):51–65, January 1997. ISSN 0031-9155. URL http://www.ncbi.nlm.nih.gov/pubmed/9015808.

[94] Changhui Li and Lihong V Wang. Photoacoustic tomography and sensing in biomedicine. *Physics in medicine and biology*, 54(19):R59–97, October 2009. ISSN 1361-6560. doi: 10.1088/0031-9155/54/19/R01. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2872141&tool=pmcentrez&rendertype=abstract.

[95] Geoffrey P. Luke, Seung Yun Nam, and Stanislav Y. Emelianov. Optical wavelength selection for improved spectroscopic photoacoustic imaging. *Photoacoustics*, 1(2):36–42, May 2013. ISSN 22135979. doi: 10.1016/j.pacs.2013.08.001. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=

4134901&tool=pmcentrez&rendertype=abstracthttp://linkinghub.elsevier.com/retrieve/pii/S2213597913000177.

[96] Yuriy Y Petrov, Irina Y Petrova, Igor a Patrikeev, Rinat O Esenaliev, and Donald S Prough. Multiwavelength optoacoustic system for noninvasive monitoring of cerebral venous oxygenation: a pilot clinical test in the internal jugular vein. *Optics letters*, 31(12):1827–9, June 2006. ISSN 0146-9592. URL http://www.ncbi.nlm.nih.gov/pubmed/16729084.

[97] Thomas Jetzfellner, Daniel Razansky, Amir Rosenthal, Ralf Schulz, K.-H. Englmeier, and Vasilis Ntziachristos. Performance of iterative optoacoustic tomography with experimental data. *Applied Physics Letters*, 95(1):013703, 2009. ISSN 00036951. doi: 10.1063/1.3167280. URL http://link.aip.org/link/APPLAB/v95/i1/p013703/s1&Agg=doi.

[98] Teedah Saratoon. *Gradient-based methods for quantitative photoacoustic tomography*. PhD thesis, University College London, 2014.

[99] Biswanath Banerjee, Srijeeta Bagchi, Ram Mohan Vasu, and Debasish Roy. Quantitative photoacoustic tomography from boundary pressure measurements: noniterative recovery of optical absorption coefficient from the reconstructed absorbed energy map. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 25(9):2347–56, September 2008. ISSN 1084-7529. URL http://www.ncbi.nlm.nih.gov/pubmed/18758563.

[100] Jorge Nocedal and S Wright. *Numerical optimization*. Springer, New York, New York, USA, 2nd edition, 1999. URL http://books.google.com/books?hl=en&lr=&id=VbHYoSyelFcC&oi=fnd&pg=PR17&dq=Numerical+Optimization&ots=30P7vpA2XO&sig=nG-1aPgoKq9E-isthO10ico9UqI.

[101] Roger J Zemp. Quantitative photoacoustic tomography with multiple optical sources. *Applied Optics*, 49(18):3566–72, June 2010. ISSN 1539-4522. URL http://www.ncbi.nlm.nih.gov/pubmed/20563210.

[102] Lei Yao, Yao Sun, and Huabei Jiang. Quantitative photoacoustic tomography based on the radiative transfer equation. *Optics Letters*, 34(12):1765–7, June 2009. ISSN 0146-9592. URL http://www.ncbi.nlm.nih.gov/pubmed/19529696.

[103] Guillaume Bal and Gunther Uhlmann. Inverse diffusion theory of photoacoustics. *Inverse Problems*, 26(8):085010, August 2010. ISSN 0266-5611. doi: 10.1088/0266-5611/26/8/085010. URL http://stacks.iop.org/0266-5611/26/i=8/a=085010?key=crossref.2e08ec4bcda4c93b6517697febaf5d32.

[104] Guillaume Bal and Kui Ren. Multi-source quantitative photoacoustic tomography in a diffusive regime. *Inverse Problems*, 27(7):075003, July 2011. ISSN 0266-5611. doi: 10.1088/0266-5611/27/7/075003. URL http://stacks.iop.org/0266-5611/27/i=7/a=075003?key=crossref.c8e0bcca40ce59b0652d1e94ebcc035c.

[105] WW Hager and Hongchao Zhang. A survey of nonlinear conjugate gradient methods. *Pacific journal of Optimization*, 2(1), 2006. URL http://www.caam.rice.edu/~yzhang/caam554/pdf/cgsurvey.pdf.

[106] B. T. Cox, S. R. Arridge, and P. C. Beard. Gradient-based quantitative photoacoustic image reconstruction for molecular imaging. *Proceedings of SPIE*, 6437: 64371T–64371T–10, 2007. ISSN 0277786X. doi: 10.1117/12.700031. URL http://link.aip.org/link/PSISDG/v6437/i1/p64371T/s1&Agg=doi.

[107] T Saratoon, T Tarvainen, B T Cox, and S R Arridge. A gradient-based method for quantitative photoacoustic tomography using the radiative transfer equation. *Inverse Problems*, 29(7):075006, July 2013. ISSN 0266-5611. doi: 10.1088/0266-5611/29/7/075006. URL http://stacks.iop.org/0266-5611/29/i=7/a=075006?key=crossref.59b7980a7c745f0fd199925f802ae55f.

[108] Tian Ding, Kui Ren, and Sarah Vallélian. A one-step reconstruction algorithm for quantitative photoacoustic imaging. *Inverse Problems*, 31(9): 095005, September 2015. ISSN 0266-5611. doi: 10.1088/0266-5611/31/9/095005. URL http://stacks.iop.org/0266-5611/31/i=9/a=095005?key=crossref.e19749ea6e370c78709163fb573492fb.

[109] Markus Haltmeier, Lukas Neumann, and Simon Rabanser. Single-stage reconstruction algorithm for quantitative photoacoustic tomography. *Inverse Problems*, 31(6):065005, 2015. ISSN 0266-5611. doi: 10.1088/0266-5611/31/6/065005. URL http://stacks.iop.org/0266-5611/31/i=6/a=065005?key=crossref.252450754d0db87a14f6dc3bfc32ffd1.

[110] Jari Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*. Springer, 2005. ISBN 0387220739. URL http://books.google.com/books?id=kCJGRAAACAAJ&pgis=1.

[111] Tanja Tarvainen, Aki Pulkkinen, Ben T. Cox, Jari P. Kaipio, and Simon R. Arridge. Image reconstruction in quantitative photoacoustic tomography using the radiative transfer equation and the diffusion approximation. In Vasilis Ntziachristos and Charles P. Lin, editors, *Proceedings of SPIE*, volume 8800, page 880006, June 2016. doi: 10.1117/12.2032498. URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2032498.

[112] A Pulkkinen, B T Cox, S R Arridge, J P Kaipio, and T Tarvainen. A Bayesian approach to spectral quantitative photoacoustic tomography. *Inverse Problems*, 30(6):065012, June 2014. ISSN 0266-5611. doi: 10.1088/0266-5611/30/6/065012. URL http://stacks.iop.org/0266-5611/30/i=6/a=065012?key=crossref.86c2e2177db76885efb652adc88c8a51.

[113] Ben Cox, Tanja Tarvainen, and Simon Arridge. Multiple Illumination Quantitative Photoacoustic Tomography using Transport and Diffusion Models. *Contemporary Mathematics*, 559:1–12, 2011.

[114] T Tarvainen, B T Cox, J P Kaipio, and S R Arridge. Reconstructing absorption and scattering distributions in quantitative photoacoustic tomography. *Inverse Problems*, 28(8), August 2012. ISSN 0266-5611. doi: 10.1088/0266-5611/28/8/084009. URL http://stacks.iop.org/0266-5611/28/i=8/a=084009?key=crossref.521af4c637d0a261ac00bf77e6e9f2d7.

[115] B T Cox, S R Arridge, and P C Beard. Estimating chromophore distributions from multiwavelength photoacoustic images. *Journal of the Optical Society of*

America. A, Optics, image science, and vision, 26(2):443–455, February 2009. ISSN 1084-7529. URL http://www.ncbi.nlm.nih.gov/pubmed/19183699.

[116] Jan Laufer, Ben Cox, Edward Zhang, and Paul Beard. Quantitative determination of chromophore concentrations from 2D photoacoustic images using a nonlinear model-based inversion scheme. *Applied Optics*, 49(8):1219–33, March 2010. ISSN 1539-4522. URL http://www.ncbi.nlm.nih.gov/pubmed/20220877.

[117] Guillaume Bal and Kui Ren. On multi-spectral quantitative photoacoustic tomography in diffusive regime. *Inverse Problems*, 28(2):025010, February 2011. ISSN 0266-5611. doi: 10.1088/0266-5611/28/2/025010. URL http://stacks.iop.org/0266-5611/28/i=2/a=025010?key=crossref.f8ef1a6e29e46bf4567d82d5a057b6d1http://stacks.iop.org/0266-5611/27/i=7/a=075003?key=crossref.c8e0bcca40ce59b0652d1e94ebcc035c.

[118] Emma R Malone, Samuel Powell, Ben T. Cox, and Simon R Arridge. Reconstruction-classification method for quantitative photoacoustic tomography. *arXiv*, August 2015. doi: 10.1002/adma.201501142. URL http://arxiv.org/abs/1508.0083.

[119] Marco Gerling, Ying Zhao, Salvatore Nania, K Jessica Norberg, Caroline S Verbeke, Benjamin Englert, Raoul V Kuiper, Asa Bergström, Moustapha Hassan, Albrecht Neesse, J Matthias Löhr, and Rainer L Heuchel. Real-time assessment of tissue hypoxia in vivo with combined photoacoustics and high-frequency ultrasound. *Theranostics*, 4(6):604–13, January 2014. ISSN 1838-7640. doi: 10.7150/thno.7996. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3982131&tool=pmcentrez&rendertype=abstract.

[120] P Vaupel, K Schlenger, C Knoop, and M Höckel. Oxygenation of human tumors: evaluation of tissue oxygen distribution in breast cancers by computerized O2 tension measurements. *Cancer research*, (51):3316–3322, 1991. URL http://cancerres.aacrjournals.org/content/51/12/3316.short.

[121] KK Kwong and JW Belliveau. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National*

Academy of Sciences of the United States of America, 89(June):5675–5679, 1992. URL http://www.pnas.org/content/89/12/5675.short.

[122] A Kleinschmidt, H Obrig, M Requardt, K D Merboldt, U Dirnagl, A Villringer, and J Frahm. Simultaneous recording of cerebral blood oxygenation changes during human brain activation by magnetic resonance imaging and near-infrared spectroscopy. *Journal of Cerebral Blood Flow and Metabolism*, 16(5):817–26, September 1996. ISSN 0271-678X. doi: 10.1097/00004647-199609000-00006. URL http://www.ncbi.nlm.nih.gov/pubmed/8784226.

[123] Lun-De Liao, Meng-Lin Li, Hsin-Yi Lai, Yen-Yu I Shih, Yu-Chun Lo, Siny Tsang, Paul Chang-Po Chao, Chin-Teng Lin, Fu-Shan Jaw, and You-Yin Chen. Imaging brain hemodynamic changes during rat forepaw electrical stimulation using functional photoacoustic microscopy. *NeuroImage*, 52(2):562–70, August 2010. ISSN 1095-9572. doi: 10.1016/j.neuroimage.2010.03.065. URL http://www.ncbi.nlm.nih.gov/pubmed/20362680.

[124] ES Brannon, AJ Merrill, JV Warren, and EA Jr Stead. The cardiac output in patients with chronic anemia as measured by the technique of right atrial catheterization. *Journal of Clinical Investigation*, 24(3):332–336, 1945. URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC435462/.

[125] S Ogawa and TM Lee. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, 87:9868–9872, 1990. URL http://www.pnas.org/content/87/24/9868.short.

[126] Nikos K Logothetis. The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 357(1424): 1003–37, August 2002. ISSN 0962-8436. doi: 10.1098/rstb.2002.1114. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1693017&tool=pmcentrez&rendertype=abstract.

[127] A P Gibson, J C Hebden, and S R Arridge. Recent advances in diffuse optical imaging. *Physics in Medicine and Biology*, 50(4):R1–R43, February 2005. ISSN 0031-9155. doi: 10.1088/0031-9155/50/4/R01. URL http://stacks.iop.org/0031-9155/50/i=4/a=R01?key=crossref.e5794ad4b1ac401d7bf971db7c0e10d6.

[128] A. Villringer, J. Planck, C. Hock, L. Schleinkofer, and U. Dirnagl. Near infrared spectroscopy (NIRS): A new tool to study hemodynamic changes during activation of brain function in human adults. *Neuroscience Letters*, 154(1-2):101–104, May 1993. ISSN 03043940. doi: 10.1016/0304-3940(93)90181-J. URL http://linkinghub.elsevier.com/retrieve/pii/030439409390181J.

[129] John W Severinghaus and Yoshiyuka Honda. History of Blood Gas Analysis. *Journal of Clinical Monitoring*, 3(2):135–138, 1987. URL http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:History+of+Blood+Gas+Analysis#1.

[130] I Y Petrov, Y Petrov, D S Prough, D J Deyo, I Cicenaite, and R O Esenaliev. Optoacoustic monitoring of cerebral venous blood oxygenation through extracerebral blood. *Biomedical Optics Express*, 3(1):125–36, January 2012. ISSN 2156-7085. doi: 10.1364/BOE.3.000125. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3255330&tool=pmcentrez&rendertype=abstract.

[131] Jun Xia, Amos Danielli, Yan Liu, Lidai Lihong V Wang, and Konstantin Maslov. Calibration-free quantification of absolute oxygen saturation based on the dynamics of photoacoustic signals. *Optics letters*, 38(15):2800–3, August 2013. ISSN 1539-4794. URL http://www.ncbi.nlm.nih.gov/pubmed/23903146.

[132] Qianqian Fang and David A Boas. Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units. *Optics Express*, 17(22):20178–90, October 2009. ISSN 1094-4087. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2863034&tool=pmcentrez&rendertype=abstract.

[133] Qianqian Fang. Mesh-based Monte Carlo method using fast ray-tracing in Plücker coordinates. *Biomedical Optics Express*, 1(1):165–75, August 2010. ISSN 2156-7085. doi: 10.1364/BOE.1.000165. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3003331&tool=pmcentrez&rendertype=abstract.

[134] Philip R. Bevington and D. Keith Robinson. Instrumental and statistical uncertainties. In *Data Reduction and Error Analysis for the Physical Sciences*, chapter 3.1, pages 36–39. McGraw-Hill, New York, 3rd edition, 1969.

[135] Jan Laufer, Edward Zhang, Gennadij Raivich, and Paul Beard. Three-dimensional noninvasive imaging of the vasculature in the mouse brain using a high resolution photoacoustic scanner. *Applied Optics*, 48(10):D299–306, April 2009. ISSN 1539-4522. URL http://www.ncbi.nlm.nih.gov/pubmed/19340121.

[136] Robert J Paproski, Andrew Heinmiller, Keith Wachowicz, and Roger J Zemp. Multi-wavelength photoacoustic imaging of inducible tyrosinase reporter gene expression in xenograft tumors. *Scientific Reports*, 4:5329, January 2014. ISSN 2045-2322. doi: 10.1038/srep05329. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4060505&tool=pmcentrez&rendertype=abstract.

[137] Roman Hochuli, Paul C. Beard, and Ben Cox. Effect of wavelength selection on the accuracy of blood oxygen saturation estimates obtained from photoacoustic images. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Proc. of SPIE, Photons Plus Ultrasound: Imaging and Sensing*, volume 9323, pages 93231V–8, March 2015. doi: 10.1117/12.2081429. URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2081429.

[138] Martin Schweiger and Simon Arridge. The Toast++ software suite for forward and inverse modeling in optical tomography. *Journal of biomedical optics*, 19(4): 040801, April 2014. ISSN 1560-2281. doi: 10.1117/1.JBO.19.4.040801. URL http://www.ncbi.nlm.nih.gov/pubmed/24781586.

[139] Geoffrey P Luke and Stanislav Y Emelianov. Optimization of in vivo spectroscopic photoacoustic imaging by smart optical wavelength selection. *Optics Letters*, 39

(7):2214–7, April 2014. ISSN 1539-4794. doi: 10.1364/OL.39.002214. URL http://www.ncbi.nlm.nih.gov/pubmed/24686714.

[140] Eun Sul Lee and Ronald E. Forthofer. Strategies for Variance Estimation. In *Analyzing Complex Servey Data*, volume 22, chapter 4, pages 22–39. SAGE Publications, Thousand Oaks, 2nd edition, 2006.

[141] Junjie Yao and Lihong V. Wang. Photoacoustic microscopy. *Laser & Photonics Reviews*, 7(5):758–778, September 2013. ISSN 18638880. doi: 10.1002/lpor.201200060. URL http://doi.wiley.com/10.1002/lpor.201200060.

[142] Jan Laufer, Peter Johnson, Edward Zhang, Bradley Treeby, Ben Cox, Barbara Pedley, and Paul Beard. In vivo preclinical photoacoustic imaging of tumor vasculature development and therapy. *Journal of Biomedical Optics*, 17(5):056016–1 – 056016–8, May 2012. ISSN 10833668. doi: 10.1117/1.JBO.17.5.056016. URL http://link.aip.org/link/JBOPFO/v17/i5/p056016/s1&Agg=doihttp://www.ncbi.nlm.nih.gov/pubmed/22612139.

[143] Amit P. Jathoul, Jan Laufer, Olumide Ogunlade, Bradley Treeby, Ben Cox, Edward Zhang, Peter Johnson, Arnold R. Pizzey, Brian Philip, Teresa Marafioti, Mark F. Lythgoe, R. Barbara Pedley, Martin a. Pule, and Paul Beard. Deep in vivo photoacoustic imaging of mammalian tissues using a tyrosinase-based genetic reporter. *Nature Photonics*, 9(April):239–246, March 2015. ISSN 1749-4885. doi: 10.1038/nphoton.2015.22. URL http://www.nature.com/doifinder/10.1038/nphoton.2015.22.

[144] Zijian Guo, Song Hu, Christopher P. Favazza, Todd N. Erpelding, Ladislav Jankovic, and Lihong V. Wang. Quantification of optical absorption coefficients from acoustic spectra with photoacoustic tomography. *Proceedings of SPIE*, 7899:78990W–78990W–9, 2011. doi: 10.1117/12.875015. URL http://link.aip.org/link/PSISDG/v7899/i1/p78990W/s1&Agg=doi.

[145] Janggun Jo and Xinmai Yang. Functional photoacoustic imaging to observe regional brain activation induced by cocaine hydrochloride. *Journal of Biomedical Optics*, 16(9):090506, September 2011. ISSN 1560-2281. doi: 10.1117/

1.3626576. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3189975&tool=pmcentrez&rendertype=abstract.

[146] Erich W Stein, Konstantin Maslov, and Lihong V Wang. Noninvasive, in vivo imaging of blood-oxygenation dynamics within the mouse brain using photoacoustic microscopy. *Journal of Biomedical Optics*, 14(2):020502, 2010. ISSN 1083-3668. doi: 10.1117/1.3095799. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2676448&tool=pmcentrez&rendertype=abstract.

[147] Gina F Lungu, Meng-Lin L Li, Xueyi Y Xie, Lihong V H V Wang, and George Stoica. In vivo imaging and characterization of hypoxia-induced neovascularization and tumor invasion. *International Journal of Oncology*, 30(1):45–54, January 2007. ISSN 1019-6439. URL http://www.ncbi.nlm.nih.gov/pubmed/17143511.

[148] Xueding Wang, David L Chamberland, and Guohua Xi. Noninvasive reflection mode photoacoustic imaging through infant skull toward imaging of neonatal brains. *Journal of Neuroscience Methods*, 168(2):412–21, March 2008. ISSN 0165-0270. doi: 10.1016/j.jneumeth.2007.11.007. URL http://www.ncbi.nlm.nih.gov/pubmed/18155298.

[149] Xueding Wang, Xueyi Xie, Geng Ku, Lihong V Wang, and George Stoica. Noninvasive imaging of hemoglobin concentration and oxygenation in the rat brain using high-resolution photoacoustic tomography. *Journal of Biomedical Optics*, 11(2):024015, 2006. ISSN 1083-3668. doi: 10.1117/1.2192804. URL http://www.ncbi.nlm.nih.gov/pubmed/16674205.

[150] Song Hu, Ping Yan, Konstantin Maslov, Jin-Moo Lee, and Lihong V Wang. Intravital imaging of amyloid plaques in a transgenic mouse model using optical-resolution photoacoustic microscopy. *Optics Letters*, 34(24):3899–901, December 2009. ISSN 1539-4794. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2854007&tool=pmcentrez&rendertype=abstract.

[151] Lun-De Liao, Chin-Teng Lin, Yen-Yu I Shih, Timothy Q Duong, Hsin-Yi Lai, Po-Hsun Wang, Robby Wu, Siny Tsang, Jyh-Yeong Chang, Meng-Lin Li, and You-Yin Chen. Transcranial imaging of functional cerebral hemodynamic changes in single blood

vessels using in vivo photoacoustic microscopy. *Journal of Cerebral Blood Flow and Metabolism*, 32(6):938–951, April 2012. ISSN 1559-7016. doi: 10.1038/jcbfm.2012.42. URL http://www.ncbi.nlm.nih.gov/pubmed/22472612.

[152] Song Hu and Lihong V Wang. Neurovascular photoacoustic tomography. *Frontiers in Neuroenergetics*, 2(June):10, January 2010. ISSN 1662-6427. doi: 10.3389/fnene.2010.00010. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2899522&tool=pmcentrez&rendertype=abstract.

[153] David Abookasis, Christopher C Lay, Marlon S Mathews, Mark E Linskey, Ron D Frostig, and Bruce J Tromberg. Imaging cortical absorption, scattering, and hemodynamic response during ischemic stroke using spatially modulated near-infrared illumination. *Journal of Biomedical Optics*, 14(2):024033, 2010. ISSN 1083-3668. doi: 10.1117/1.3116709. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2868516&tool=pmcentrez&rendertype=abstract.

[154] L S Hansen, J E Coggle, J Wells, and M W Charles. The influence of the hair cycle on the thickness of mouse skin. *The Anatomical Record*, 210(4):569–73, December 1984. ISSN 0003-276X. doi: 10.1002/ar.1092100404. URL http://www.ncbi.nlm.nih.gov/pubmed/6524697.

[155] Alexander J Lin, Maya a Koike, Kim N Green, Jae G Kim, Amaan Mazhar, Tyler B Rice, Frank M LaFerla, and Bruce J Tromberg. Spatial frequency domain imaging of intrinsic optical property contrast in a mouse model of Alzheimer's disease. *Annals of Biomedical Engineering*, 39(4):1349–57, April 2011. ISSN 1573-9686. doi: 10.1007/s10439-011-0269-6. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3069335&tool=pmcentrez&rendertype=abstract.

[156] M Firbank, M Hiraoka, M Essenpreis, and D T Delpy. Measurement of the optical properties of the skull in the wavelength range 650-950 nm. *Physics in Medicine & Biology*, 38(4):503–10, April 1993. ISSN 0031-9155. URL http://www.ncbi.nlm.nih.gov/pubmed/8488176http://iopscience.iop.org/0031-9155/38/4/002.

[157] Zijian Guo, Christopher Favazza, Alejandro Garcia-Uribe, and Lihong V. Wang. Quantitative photoacoustic microscopy of optical absorption coefficients from

acoustic spectra in the optical diffusive regime. *Journal of Biomedical Optics*, 17 (6):066011, 2012. ISSN 10833668. doi: 10.1117/1.JBO.17.6.066011. URL http://link.aip.org/link/JBOPFO/v17/i6/p066011/s1&Agg=doi.

[158] Roman Hochuli, Paul C. Beard, and Ben Cox. Accuracy of approximate inversion schemes in quantitative photacoustic imaging. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Proc. of SPIE, Photons Plus Ultrasound: Imaging and Sensing*, volume 8943, pages 89435V–1 – 89435V–10, March 2014. doi: 10.1117/12.2039825. URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2039825.

[159] Christakis Constantinides, Richard Mean, and Ben J Janssen. Effects of isoflurane anesthesia on the cardiovascular function of the C57BL/6 mouse. *ILAR journal/National*, 52, 2011. URL http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3508701/.

[160] Akira Ishimaru. Theory and Application of Wave Propagation and Scattering in Random Media. *IEEE*, 65(7):1030–1061, 1977.

[161] Angelo Sassaroli and Fabrizio Martelli. 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(10):2110–7, October 2012. ISSN 1520-8532. URL http://www.ncbi.nlm.nih.gov/pubmed/23201658.

[162] SA Prahl and M Keijzer. A Monte Carlo model of light propagation in tissue. *Dosimetry of Laser*, I(1989):102–111, 1989. URL http://engineering.dartmouth.edu/~Amir_H_Golnabi/ENGG168_files/Papers/prahl89.pdf.

[163] Erik Alerstam, Stefan Andersson-Engels, and Tomas Svensson. White Monte Carlo for time-resolved photon migration. *Journal of Biomedical Optics*, 13(4):041304, 2008. ISSN 1083-3668. doi: 10.1117/1.2950319. URL http://www.ncbi.nlm.nih.gov/pubmed/19021312.

[164] DE Peplow. Direction cosines and polarization vectors for Monte Carlo photon scattering. *Nucl Sci Eng*, 131:132–136, 1999. URL http://www.ornl.gov/~5pe/p007.pdf.

[165] Jenni Heino, Simon Arridge, Jan Sikora, and Erkki Somersalo. Anisotropic effects in highly scattering media. *Physical Review E*, 68(3):1–8, September 2003. ISSN 1063-651X. doi: 10.1103/PhysRevE.68.031908. URL http://link.aps.org/doi/10.1103/PhysRevE.68.031908.

[166] David Boas, J Culver, J Stott, and A Dunn. Three dimensional Monte Carlo code for photon migration through complex heterogeneous media including the adult human head. *Optics express*, 10(3):159–70, February 2002. ISSN 1094-4087. URL http://www.ncbi.nlm.nih.gov/pubmed/19424345.

[167] Lihong Wang, SL Jacques, and L Zheng. MCML–Monte Carlo modeling of light transport in multi-layered tissues. *Computer methods and programs in biomedicine*, 47:131–146, 1995. URL http://linkinghub.elsevier.com/retrieve/pii/016926079501640F.

[168] Vijitha Periyasamy and Manojit Pramanik. Monte Carlo simulation of light transport in tissue for optimizing light delivery in photoacoustic imaging of the sentinel lymph node. *Journal of biomedical optics*, 18(10):106008, January 2013. ISSN 1560-2281. doi: 10.1117/1.JBO.18.10.106008. URL http://www.ncbi.nlm.nih.gov/pubmed/24108574.

[169] Vijitha Periyasamy and Manojit Pramanik. Monte Carlo simulation of light transport in turbid medium with embedded object–spherical, cylindrical, ellipsoidal, or cuboidal objects embedded within multilayered tissues. *Journal of biomedical optics*, 19(4):045003, April 2014. ISSN 1560-2281. doi: 10.1117/1.JBO.19.4.045003. URL http://www.ncbi.nlm.nih.gov/pubmed/24727908.

[170] T Binzoni, T S Leung, R Giust, D Rüfenacht, and a H Gandjbakhche. Light transport in tissue by 3D Monte Carlo: influence of boundary voxelization. *Computer methods and programs in biomedicine*, 89(1):14–23, January 2008. ISSN 0169-2607. doi: 10.1016/j.cmpb.2007.10.008. URL http://www.ncbi.nlm.nih.gov/pubmed/18045725.

[171] Lihong V Wang and Steven L. Jacques. Monte Carlo Modeling of Light Transport in Multi-Layered Tissues in Standard C. 1992. URL http://spx.arizona.edu/OPTI630/Homework/MonteCarlo/Mcman.pdf.

[172] Ivan T. Lima, Anshul Kalra, and Sherif S. Sherif. Improved importance sampling for Monte Carlo simulation of time-domain optical coherence tomography. *Biomedical Optics Express*, 2(5):1069, April 2011. ISSN 2156-7085. doi: 10.1364/BOE.2.001069. URL http://www.opticsinfobase.org/abstract.cfm?URI=boe-2-5-1069.

[173] Quan Liu and Nirmala Ramanujam. Sequential estimation of optical properties of a two-layered epithelial tissue model from depth-resolved ultraviolet-visible diffuse reflectance spectra. *Applied Optics*, 45(19):4776, 2006. ISSN 0003-6935. doi: 10.1364/AO.45.004776. URL http://www.opticsinfobase.org/abstract.cfm?URI=ao-45-19-4776.

[174] J S Hendricks and T E Booth. MCNP Variance Reduction Overview. In Raymond Alcouffe, Robert Dautray, Arthur Forster, Guy Ledanois, and B. Mercier, editors, *Monte-Carlo Methods and Applications in Neutronics, Photonics and Statistical Physics*, volume 240 of *Lecture Notes in Physics*, pages 83–92. Springer-Verlag, Berlin/Heidelberg, 1985. ISBN 3-540-16070-1. doi: 10.1007/BFb0049029. URL http://www.springerlink.com/index/10.1007/BFb0049029.

[175] Hironobu Iwabuchi. Efficient Monte Carlo Methods for Radiative Transfer Modeling. *Journal of the Atmospheric Sciences*, 63(9):2324–2339, September 2006. ISSN 0022-4928. doi: 10.1175/JAS3755.1. URL http://journals.ametsoc.org/doi/abs/10.1175/JAS3755.1.

[176] DR Kirkby and DT Delpy. Parallel operation of Monte Carlo simulations on a diverse network of computers. *Physics in Medicine & Biology*, 42:1203–1208, 1997. URL http://iopscience.iop.org/0031-9155/42/6/016.

[177] Erik Alerstam, Tomas Svensson, and Stefan Andersson-Engels. Parallel computing with graphics processing units for high-speed Monte Carlo simulation of photon

migration. *Journal of Biomedical Optics*, 13(6):060504, 2015. ISSN 1083-3668. doi: 10.1117/1.3041496. URL http://www.ncbi.nlm.nih.gov/pubmed/19123645.

[178] Erik Alerstam, Tomas Svensson, and Stefan Andersson-engels. CUDAMCML User manual and implementation notes, 2009.

[179] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. Julia: A Fresh Approach to Numerical Computing. *arXiv*, 1411.1607:37, November 2014. URL http://arxiv.org/abs/1411.1607.

[180] Tanja Tarvainen, Marko Vauhkonen, Ville Kolehmainen, Jari P. Kaipio, and Simon R. Arridge. Utilizing the Radiative Transfer Equation in Optical Tomography. *PIERS Online*, 4(6):655–660, 2008. ISSN 1931-7360. doi: 10.2529/PIERS071219142458. URL http://piers.mit.edu/piersonline/piers.php?year=2008&volume=4&number=6&page=655.

[181] Tanja Tarvainen, Marko Vauhkonen, Ville Kolehmainen, and Jari P Kaipio. Hybrid radiative-transfer-diffusion model for optical tomography. *Applied optics*, 44(6): 876–86, February 2005. ISSN 0003-6935. URL http://www.ncbi.nlm.nih.gov/pubmed/15751677.

[182] Lihong Wang and Steven L. Jacques. Hybrid model of Monte Carlo simulation and diffusion theory for light reflectance by turbid media. *Journal of the Optical Society of America A*, 10(8):1746, August 1993. ISSN 1084-7529. doi: 10.1364/JOSAA.10.001746. URL http://www.opticsinfobase.org/abstract.cfm?URI=josaa-10-8-1746.

[183] F Inanc and AF Rohach. A nodal solution of the multigroup neutron transport equation using spherical harmonics. *Annals of Nuclear Energy*, 15(10): 501–509, 1988. URL http://www.sciencedirect.com/science/article/pii/0306454988900667.

[184] A.G. Buchan, A.a. Calloo, M.G. Goffin, S. Dargaville, F. Fang, C.C. Pain, and I.M. Navon. A POD reduced order model for resolving angular direction in

neutron/photon transport problems. *Journal of Computational Physics*, 296:138–157, September 2015. ISSN 00219991. doi: 10.1016/j.jcp.2015.04.043. URL http://linkinghub.elsevier.com/retrieve/pii/S0021999115003009.

[185] Abhinav K Jha, Matthew a Kupinski, Harrison H Barrett, Eric Clarkson, and John H Hartman. Three-dimensional Neumann-series approach to model light transport in nonuniform media. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 29(9):1885–99, September 2012. ISSN 1520-8532. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3963433&tool=pmcentrez&rendertype=abstract.

[186] Y a Ilyushin and V P Budak. Narrow-beam propagation in a two-dimensional scattering medium. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 28(2):76–81, February 2011. ISSN 1520-8532. URL http://www.ncbi.nlm.nih.gov/pubmed/21293513.

[187] WW Hager and Hongchao Zhang. A new conjugate gradient method with guaranteed descent and an efficient line search. *SIAM Journal on Optimization*, 16(1):170–192, 2005. URL http://epubs.siam.org/doi/abs/10.1137/030601880.

[188] a. De Cezaro, F. Travessini De Cezaro, and J. Sejje Suarez. Regularization approaches for quantitative Photoacoustic tomography using the radiative transfer equation. *Journal of Mathematical Analysis and Applications*, 429(1):415–438, September 2015. ISSN 0022247X. doi: 10.1016/j.jmaa.2015.03.079. URL http://linkinghub.elsevier.com/retrieve/pii/S0022247X15003157.

[189] L.I. Rudin and Stanley Osher. Total variation based image restoration with free local constraints. In *IEEE Proceedings of 1st International Conference on Image Processing*, volume 1, pages 31–35, Austin, 1994. IEEE Comput. Soc. Press. ISBN 0-8186-6952-7. doi: 10.1109/ICIP.1994.413269. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=413269http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=413269.

[190] Aki Pulkkinen, Ben T. Cox, Simon R. Arridge, Jari P. Kaipio, and Tanja Tarvainen. Quantitative photoacoustic tomography using illuminations from a single

direction. *Journal of Biomedical Optics*, 20(3):036015, March 2015. ISSN 1083-3668. doi: 10.1117/1.JBO.20.3.036015. URL http://biomedicaloptics.spiedigitallibrary.org/article.aspx?doi=10.1117/1.JBO.20.3.036015.

[191] T. Saratoon, T. Tarvainen, S. R. Arridge, and B. T. Cox. 3D quantitative photoacoustic tomography using the δ-Eddington approximation. *Proc. of SPIE, Photons Plus Ultrasound: Imaging and Sensing*, 8581:85810V–85810V–12, March 2013. doi: 10.1117/12.2004105. URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2004105.

[192] L Grippo, F Lampariello, and S Lucidi. A nonmonotone line search technique for Newton's method. *SIAM Journal on Numerical Analysis*, 23(4):707–716, 1986. URL http://epubs.siam.org/doi/abs/10.1137/0723046.

[193] Hongchao Zhang and William W. Hager. A nonmonotone line search technique and its application to unconstrained optimization. *SIAM Journal on Optimization*, 14(4):1043–1056, January 2004. ISSN 1052-6234. doi: 10.1137/S1052623403428208. URL http://epubs.siam.org/doi/abs/10.1137/S1052623403428208.

[194] Robert Ellwood, Edward Zhang, Paul Beard, and Ben Cox. Photoacoustic imaging using acoustic reflectors to enhance planar arrays. *Journal of Biomedical Optics*, 19(12):126012, 2014. doi: 10.1117/1.JBO.

[195] Aki Pulkkinen, Ben T. Cox, Simon R. Arridge, Jari P. Kaipio, and Tanja Tarvainen. Bayesian parameter estimation in spectral quantitative photoacoustic tomography. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Proceedings of SPIE*, volume 9708, page 97081G, March 2016. doi: 10.1117/12.2205009. URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2205009.

[196] Emma Malone, Ben Cox, and Simon Arridge. Multispectral reconstruction methods for quantitative photoacoustic tomography. In Alexander A. Oraevsky and Lihong V. Wang, editors, *Proceedings of SPIE*, volume 9708, page

970827, March 2016. doi: 10.1117/12.2212440. URL http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.2212440.