

Forward models and image reconstruction in ultrasound
modulated optical tomography

Samuel Powell, BEng, MSc

Department of Medical Physics and Bioengineering, UCL

A thesis submitted for the degree of Doctor of Philosophy

Declaration

I, Samuel Powell, 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.

Abstract

Ultrasound modulated optical tomography is a hybrid imaging modality with numerous potential clinical applications. In this work we develop, validate, and assess the accuracy of a number of forward models which describe the effect. We subsequently derive an inversion procedure which reconstructs images of the optical absorption coefficient in a turbid media from measurements of the optical field autocorrelation function made on the boundary.

We begin with the development of a reference forward model which is accelerated by execution on the parallel architecture of modern graphics processing units. The model is validated against analytical and numerical results from the literature. The acceleration of the model results in improvements in performance of between one and two orders of magnitude when compared with a standard implementation.

Whilst accurate, the reference model is not suited for use in an image reconstruction algorithm. As such, a number of alternative models are derived based upon spherical harmonic expansions of a correlation transport equation. In particular, a third order simplified spherical harmonic approximation is developed which provides a peak improvement in accuracy of over 50% relative to standard diffusion-style approximations. Each of the models is solved by the finite-element method, and validated against the reference Monte-Carlo simulation.

Finally, we develop an image reconstruction procedure for ultrasound modulated optical tomography which employs a finite-element implementation of the proposed forward models. As part of the derivation we investigate the form of the correlation measurement density functions which describe the sensitivity of the technique to perturbations in the optical parameters of the medium. We demonstrate the ability to reconstruct an image of the optical absorption coefficient in a turbid medium from noisy measurements of the field autocorrelation function made on the boundary of the domain. Reconstructions employing data corrupted by 1% Gaussian noise achieve accuracy of circa 80% in the region of peak optical sensitivity, and maintain spatial resolution equivalent to the dimensions of the focused acoustic field probing the domain. Away from the sensitive region of the optical axis, regularisation forces the imaging resolution towards that of conventional diffuse-optical tomography.

Publications

The following publications have resulted from this work.

Peer-reviewed journal articles

- S. Powell, and T. S. Leung. 'Linear reconstruction of absorption perturbations in coherent ultrasound modulated optical tomography'. *Journal of Biomedical Optics* 18(12) 2013: 126020-1-9.
- S. Powell, and T. S. Leung. 'Highly parallel Monte-Carlo simulations of the acousto-optic effect in heterogeneous turbid media'. *Journal of Biomedical Optics* 17 (4) 2012: 045002-1-11.
- T. S. Leung, and S. Powell. 'Fast Monte Carlo simulations of ultrasound-modulated light using a graphics processing unit'. *Journal of Biomedical Optics* 15(5) 2010: 055007-1-7.

Conference proceedings

- S. Powell, and T. S. Leung. 'Image reconstruction techniques for ultrasound-modulated optical tomography'. *SPIE BiOS*, 2013: 857819-1-10.
- S. Powell, and T. S. Leung. 'Acousto-optic monitoring of high-intensity focused ultrasound lesion formation with fibre-coupled autocorrelation detection'. *Proceedings of Meetings on Acoustics* 15 2013: 075004-1-7.
- S. Gunadi, S. Powell, C. E. Elwell, and T. S. Leung. 'Optimization of the acousto-optic signal detection in cylindrical geometry'. Proc. SPIE, 2010, 7564, 756431.

Conference posters

- S. Powell, and T. S. Leung. 'Linear image reconstruction in coherent ultrasound-modulated optical tomography'. *Applied Inverse Problems*, Daejon, South Korea, May 2013.

Acknowledgements

In my supervisors, Terence Leung and Jem Hebden, I have found an effective combination of unerring enthusiasm and support, and constructive, direct criticism. I extend my thanks to each.

Many of my colleagues at UCL have provided significant help as I have completed this work. In particular I am thoroughly indebted to Ben Cox who has always found the time to help me. I offer my thanks also to Simon Arridge.

It would be futile for me to attempt to quantify the individual contributions of my friends. The following people have been sources of support and encouragement to me over the course of this work: Thomas Allan, Annie Bryant, Ian Butterworth, Ruxandra Comanaru, Luke Dunne, Nick Everdell, Alejandro Giacometti, Adam Gibson, Nicholas Hammond, Roman Hoculi, Fiona Holland, Dan Kay, Teedah Sarahtoon, Kirsten Sutcliffe, and Brad Treeby. To my closest friends, Mark Narraway, Michael Babester, and Alex Sutton: you all know what you have done.

I could not have completed this work without the support of my mother, Valerie Taylor, her husband Roger Seifert, my father Nicholas Powell, my aunt and uncle Jude and Nic Rea, my partner Emma Malone, and Betsy.

Samuel Powell, August 2013

Table of contents

Abstract	3
Publications	4
Acknowledgements	5
Table of contents	6
List of figures	9
List of tables	13
1 Introduction	14
1.1 Biomedical optical imaging	14
1.1.1 Photoacoustic tomography	15
1.1.2 Ultrasound modulated optical tomography	16
1.2 Overview and contribution	18
2 Background	20
2.1 Overview	20
2.2 Tissue optics	20
2.2.1 Absorption and scattering	20
2.2.2 Radiative transport	21
2.3 A concise review of ultrasound modulated optical tomography	22
2.3.1 Detection and insonification strategies	22
2.3.2 Modulation, modelling, and inversion	24
2.4 Mechanisms of coherent acousto-optic modulation	27
2.4.1 Displacement of optical scatterers	27
2.4.2 Perturbation of the index of refraction	29
2.4.3 Combining the mechanisms	30
2.5 Models of ultrasound modulated optical tomography	30
2.5.1 UOT models based upon diffusing wave spectroscopy	30

2.5.2	UOT models based upon correlation transport	33
2.6	Autocorrelation ultrasound modulated optical tomography	33
2.6.1	Experimental technique	33
2.6.2	Observables and related quantities	34
3	A highly parallel Monte-Carlo model of correlation transport	36
3.1	Introduction	36
3.1.1	Overview	37
3.2	Theory	38
3.2.1	Monte-Carlo, radiative and correlation transport	38
3.2.2	An approximate adjoint method	42
3.3	Implementation	44
3.3.1	GPU Parallel Programming	44
3.3.2	Architecture	45
3.3.3	The Simulation Domain	49
3.3.4	Geometry Representation	51
3.3.5	Ancillary Data	53
3.3.6	Algorithm	54
3.3.7	Absorption and Roulette	58
3.3.8	Random Number Generator	59
3.3.9	Scored quantities and post-processing	60
3.4	Results	60
3.4.1	Validation: optical	60
3.4.2	Validation: acousto-optical	68
3.4.3	Performance	72
3.4.4	Sensitivity Mapping	76
3.5	Conclusions	78
4	Approximations of correlation transport by spherical harmonic expansions	80
4.1	Introduction	80
4.2	Theory	81
4.2.1	Correlation transport	81
4.2.2	Correlation transport under acoustic modulation	82
4.2.3	Correlation diffusion approximation	85
4.2.4	Correlation diffusion with pre-averaged phase	88
4.2.5	Simplified spherical harmonic approximation	91
4.2.6	Finite element implementation	96
4.3	Results	97
4.3.1	Simplified spherical harmonic approximation	97
4.4	Conclusions	104

5	Linear image reconstruction in ultrasound modulated optical tomography	105
5.1	Introduction	105
5.2	Theory	107
5.2.1	Autocorrelation measurements	107
5.2.2	Linear image reconstruction	108
5.2.3	Forward model	109
5.2.4	Correlation measurement density functions and the Jacobian	110
5.2.5	Regularisation	112
5.3	Results	112
5.3.1	Correlation measurement density functions	113
5.3.2	Image reconstruction: noise and regularisation	120
5.4	Conclusions	129
6	Future work and summary	133
6.1	Future work	133
6.1.1	Forward modelling	133
6.1.2	The inverse problem	135
6.1.3	Towards clinical application	136
6.2	Summary	138
	Bibliography	140

List of figures

2.1	Time varying acoustic displacement of the scatterers of a turbid medium perturbing the optical path between scattering sites.	28
2.2	Time varying index of refraction over the length of an optical path due to perturbation of the index of refraction by an applied acoustic field.	29
2.3	Form of the electric field autocorrelation function and power spectral density in a diffusing wave spectroscopy and acousto-optic measurement.	32
3.1	Architecture of the simulation framework demonstrating the input datasets, organisation of the simulation program code, and relationship to the MATLAB environment	46
3.2	Architecture of the simulation program illustrating the parallel execution of many thousands of individual propagation kernels, each responsible for a single photon packet.	48
3.3	Comparison of the depth resolved unitless fluence in a semi-infinite medium as computed by the simulation program against that computed by CUDAM-CML, in addition to an exponential fit according to diffusion theory.	63
3.4	Plot of the radially resolved reflectance from a multi-layered slab, as calculated by the simulation program, <i>MCML</i> , and a diffusion approximation.	64
3.5	Mesh employed in the simulation of figures 3.6 and 3.7 for the series of concentric shells.	65
3.6	Plot of flux on the outer surface of the concentric spheres vs. angle from the illumination source: case 1	66
3.7	Plot of flux on the outer surface of the concentric spheres vs. angle from the illumination source: case 2	66
3.8	Temporally resolved reflectance (upper curve) and transmittance (lower curve) for a slab of optical properties given in the text, resolved at a radial distance $r = 30\text{mm}$ from the source.	67
3.9	Temporally resolved total transmittance for a slab of optical properties given in the text with relative refractive index of $n = 1.4$ (upper curve) and $n = 1$ (lower curve).	68

3.10	Acousto-optic modulation depth calculated by the simulation program vs. an analytical solution for changes in the optical absorption coefficient μ_a	69
3.11	Acousto-optic modulation depth calculated by the simulation program vs. an analytical solution for changes in the acoustic frequency f_a	70
3.12	Acousto-optic modulation depth calculated by the simulation program vs. an analytical solution for changes in the acoustic displacement amplitude A	70
3.13	Acousto-optic modulation depth calculated by the simulation program vs. an analytical solution for changes in the scattering coefficient μ_s	71
3.14	Modulation depth in a slab geometry insonified by a column of ultrasound	72
3.15	Modulation depth along the reflection and transmission axes of a slab geometry insonified by a column of ultrasound	73
3.16	Spatial sensitivity maps (-% change from reference value). Inward arrows depict the source and outward arrows the detector location. The dotted ellipse illustrates the focus of the ultrasound region.	77
4.1	Zeroth harmonic correlation fluence ($\log_{10} G_0(\mathbf{r})$), as calculated by the reference Monte-Carlo simulation.	98
4.2	Zeroth harmonic correlation fluence ($\log_{10} G_0(\mathbf{r})$), as calculated by the standard correlation diffusion approximation, equivalent to first order simplified spherical harmonic method.	98
4.3	Zeroth harmonic correlation fluence ($\log_{10} G_0(\mathbf{r})$), as calculated by the third order simplified spherical harmonic (SP_3) approximation to the correlation transport equation.	99
4.4	First harmonic correlation fluence ($\log_{10} G_1(\mathbf{r})$), as calculated by the reference Monte-Carlo simulation.	100
4.5	First harmonic correlation fluence ($\log_{10} G_1(\mathbf{r})$), as calculated by the standard correlation diffusion approximation, equivalent to first order simplified spherical harmonic method.	100
4.6	First harmonic correlation fluence ($\log_{10} G_1(\mathbf{r})$), as calculated by the third order simplified spherical harmonic (SP_3) approximation to the correlation transport equation.	101
4.7	Internal first harmonic modulation depth distribution, $MD(\mathbf{r}) [-]$, as calculated by the reference Monte-Carlo simulation.	101
4.8	Internal first harmonic modulation depth distribution, $MD(\mathbf{r}) [-]$, as calculated by the standard correlation diffusion approximation, equivalent to first order simplified spherical harmonic method. Note that the colour scale is limited to that of the reference solution.	102
4.9	Internal first harmonic modulation depth distribution, $MD(\mathbf{r}) [-]$, as calculated by the third order simplified spherical harmonic (SP_3) approximation to the correlation transport equation.	102

4.10	Percentage error modulation depth between P_1 correlation diffusion approximation and reference Monte-Carlo solution.	103
4.11	Percentage error modulation depth between SP_3 approximation and reference Monte-Carlo solution.	104
5.1	Illustration of the form of the measurement in autocorrelation UOT.	107
5.2	UOT absorption sensitivity in reflection mode at lag $\tau = 0$, with a 4cm source-detector spacing.	114
5.3	UOT absorption sensitivity in reflection mode for an AC measurement, with a 4cm source-detector spacing: ultrasound located close to detector.	115
5.4	UOT absorption sensitivity in reflection mode for an AC measurement, with a 4cm source-detector spacing: ultrasound in region of sensitivity.	115
5.5	UOT absorption sensitivity in reflection mode for an AC measurement, with a 4cm source-detector spacing: ultrasound on edge of sensitivity region.	116
5.6	UOT absorption sensitivity in reflection mode for a modulation depth measurement, with a 4cm source-detector spacing: ultrasound on edge of sensitivity region.	116
5.7	UOT absorption sensitivity in reflection mode at lag $\tau = 0$, with a 4cm source-detector spacing: ultrasound on edge of sensitivity region.	117
5.8	UOT absorption sensitivity in reflection mode at lag $\tau = 0$, with a 4cm source-detector spacing: ultrasound on edge of sensitivity region.	118
5.9	UOT absorption sensitivity in reflection mode for a modulation depth measurement over thirty-six ultrasound scan locations, with a 4cm source-detector spacing.	119
5.10	UOT absorption sensitivity in transmission mode at lag $\tau = 0$, with a 5cm source-detector spacing.	120
5.11	UOT absorption sensitivity in transmission mode for an AC measurement, with a 5cm source-detector spacing: ultrasound on edge of sensitivity region.	121
5.12	UOT absorption sensitivity in transmission mode at lag $\tau = 0$, with a 5cm source-detector spacing: ultrasound on edge of sensitivity region.	121
5.13	UOT absorption sensitivity in transmission mode for a modulation depth measurement over thirty-six ultrasound scan locations, with a 5cm source-detector spacing.	122
5.14	Absorption perturbation for homogeneous background linear reconstruction. .	123
5.15	Linear reconstruction of absorption perturbation on an homogeneous background using the rank-one zero-lag Jacobian.	124
5.16	Linear reconstruction of absorption perturbation on an homogeneous background using the UOT AC Jacobian.	125
5.17	Linear reconstruction of absorption perturbation on an homogeneous background using the UOT AC Jacobian with 0.01% Gaussian noise.	126

5.18	Linear reconstruction of absorption perturbation on an homogeneous background using the UOT modulation depth Jacobian.	127
5.19	Linear reconstruction of absorption perturbation on an homogeneous background using the UOT modulation depth Jacobian with 0.1% Gaussian noise. . .	127
5.20	Linear reconstruction of absorption perturbation on an homogeneous background using the UOT modulation depth Jacobian with 1% Gaussian noise. . .	128
5.21	Direct mapping approach to reconstruction of an absorption perturbation using a modulation depth measurement in reflection mode without noise.	129
5.22	Linear reconstruction of absorption perturbation on an homogeneous background using the UOT modulation depth Jacobian in transmission mode without noise.	130
5.23	Linear reconstruction of absorption perturbation on an homogeneous background using the UOT modulation depth Jacobian in transmission mode with 0.1% Gaussian noise.	131
5.24	Linear reconstruction of absorption perturbation on an homogeneous background using the UOT modulation depth Jacobian in transmission mode with 1% Gaussian noise.	131

List of tables

3.1	Comparison of the diffuse reflectance and total transmittance of a thin slab with matched boundary predicted by the simulation program against a diffusion model and alternative Monte-Carlo simulation codes.	62
3.2	Optical properties of the multi-layered slab geometry employed in the simulation of figure 3.4	63
3.3	Optical properties of the concentric shells [mm^{-1}] employed in the simulation of figures 3.6 and 3.7	65
3.4	Execution speed (photon packets/ms) of alternative acousto-optic simulation programs and configurations.	74

Chapter 1

Introduction

The field of biomedical imaging encompasses the study of techniques designed to produce images of clinically or scientifically relevant representations of anatomy and biological processes. The information presented in such images is broad, ranging from the delineation of anatomical structures to the concentrations of molecular biomarkers, and a commensurate number of imaging techniques have been designed to capture the requisite data.

The subject of this work is one such technique—ultrasound-modulated optical tomography (UOT, also known as UMOT/USMOT). This hybrid imaging technique aims to combine the optical contrast of biological tissues at near-infrared wavelengths with the spatial resolution of focused ultrasound fields in order to provide highly localised measurements of the optical properties of biological tissues.

In this chapter we will begin by placing UOT in context with alternative optical imaging techniques, before considering in greater depth the challenges faced in its development. Finally, we provide an overview of this work, and its contribution to the field.

1.1 Biomedical optical imaging

Optical imaging techniques have numerous attractive qualities. Unlike modalities such as X-ray computed tomography (CT) [60], or positron emission tomography (PET)[45], the radiation employed is non-ionising. Furthermore, the small size and low cost of the requisite hardware (contrasted with techniques such as magnetic resonance imaging) suggest that bedside monitoring could be commonplace in the near-future: the pulse oximeter, which relies on diffuse optical technology, is already ubiquitous in clinical environments.

Biomedical techniques designed to produce images with optical contrast at significant depth (as opposed to superficial techniques such as standard microscopy or optical coherence tomography [121]) typically operate in the near-infrared region of the electromagnetic spectrum (circa 600-1100nm). This range of wavelengths, often denoted the ‘optical window’, is chosen for two principle reasons:

1. Biological tissues are relatively transparent in the optical window, and thus light probing the medium can penetrate deeper than is possible using surrounding regions of the spectrum. At wavelengths shorter than those in the optical window, the absorption of oxygenated and deoxygenated haemoglobin present in blood quickly attenuates incident radiation. At longer wavelengths, the same effect occurs due to the increasing absorption of water.
2. When multiple wavelengths of illumination are employed, the data may be combined via knowledge of the absorption spectra of various chromophores (such as the aforementioned haemoglobin) to infer parameters such as oxygen saturation. In turn, such measures may indicate underlying pathologies, e.g., increased vascularisation due to tumour angiogenesis, or serve as an indicator of the progress of treatments which themselves alter the optical properties of tissues, e.g., thermal necrosis by therapeutic high-intensity focused ultrasound (HIFU).

These properties have been extensively exploited in the field of near-infrared spectroscopy (NIRS) [58]. A principal application is in determining the state of blood oxygenation in the brain, which is related to the local oxygen demand, and in turn, functional activation.

Whilst light in the optical window is capable of penetrating deep into tissue, optical scattering causes the input energy to quickly diffuse throughout the tissue, significantly limiting the spatial resolution of such techniques. The desire to localise changes in the concentration of various chromophores, such as haemoglobin, motivated the development of diffuse optical tomography (DOT) [10, 24, 47]. In its most basic form, DOT involves the collection of data from multiple detectors placed around the periphery of an object which are subject to illumination from multiple continuous wave (CW) sources. A model-based image reconstruction procedure is then employed to reconstruct an image of the absorption coefficient within the medium. The use of frequency modulated or pulsed sources permits the recovery not only of the absorption coefficient, but also of the scattering coefficient of the medium [11].

Despite various advances, the spatial resolution in DOT is still limited to around 10mm [47]. The desire to improve the spatial resolution of DOT, whilst maintaining optical contrast, has recently led to the development of two hybrid imaging modalities: ultrasound-modulated optical tomography, the subject of this work, and photoacoustic tomography (PAT) [18].

1.1.1 Photoacoustic tomography

In photoacoustic tomography the medium to be imaged is typically subject to pulsed wide-field illumination from an optical source operating within the optical window. A laser is typically chosen as the optical source, owing to its ability to generate high peak powers and generate short pulses (though the use of high power light emitting diodes has been investi-

gated [5]). The input illumination is absorbed by chromophores in the tissue causing a small local rise in temperature. The resulting thermoelastic expansion generates a broadband acoustic wave which propagates through the medium until it is detected upon the boundary.

By back projecting the detected acoustic field to its origin, an estimate is made of the initial photoacoustic pressure distribution. Since the initial pressure distribution is dependent upon local absorption, so to is the resulting image.

The penetration depth which can be achieved in PAT is primarily dependent upon the attenuation of the input illumination. Since optical absorption is also the source of its contrast, the technique can achieve its best results in the case of highly absorbing targets embedded in weakly absorbing media: depths of up to 40mm have been reported *in vivo* in human breast tissue [66]. Without the use of focused optical excitation, the spatial resolution of PAT is dependent upon the frequency content of the detected acoustic field, the higher frequencies of which are preferentially absorbed by biological tissues. In practice, sub-millimetre resolution can be achieved for penetration depths of around a centimetre [18].

Whilst the ‘raw’ photoacoustic image obtained from acoustic back projection is principally dependent upon absorption, the actual relationship between the initial pressure distribution and the optical coefficients of the medium is more complicated. The initial pressure distribution is in fact proportional the product of the absorption coefficient and the local optical energy density. As such, quantitative photoacoustic imaging requires the use of a model-based inversion which considers the optical transport problem [33]. Such an approach is also of great importance in permitting spectroscopic measurements, since the optical energy density will change with wavelength such that ‘raw’ images can not be considered on a like-for-like basis. The recovery of the scattering coefficient in PAT is challenging, since it is only related to the measurement through the local optical energy density [34].

1.1.2 Ultrasound modulated optical tomography

Ultrasound-modulated optical tomography is, as its name suggests, more closely related to DOT than PAT, in terms of its practical implementation. In UOT coherent optical sources and detectors are placed around the medium and an ultrasound field is introduced which interacts with the optical field. We refer to the mechanisms of interaction between the acoustic and optical field as acousto-optic (AO) modulation. Note that many authors use terms such as acousto-optic imaging (AOI) and acousto-optic tomography (AOT) to refer to UOT, but we shall reserve the use of terms involving ‘acousto-optic’ to refer specifically to the modulation processes.

We will explore the physical basis of acousto-optic modulation further in chapter 2, but for now we state that in coherent acousto-optics the optical field undergoes phase modulation, such that the light which has passed through the acoustic field has frequency components shifted by multiples of the acoustic frequency. Since we have considerable spatial and temporal control over the ultrasound field—it can be focused into most biological me-

dia, and its propagation speed is such that time gating and electronic modulation can be achieved with readily available instrumentation—we may define the region from which the modulated, or ‘tagged’ optical field is generated.

A ‘raw’ UOT image (a direct mapping approach) assigns the intensity of the detected modulated light to a particular region of an image based upon the assumed location of the ultrasound distribution. Our measured data is proportional to the product of the local acoustic intensity and the optical sensitivity (that is the probability of light travelling from a given source, to a given detector, via a particular point). Since the optical sensitivity of the medium is dependent upon the absorption and scattering coefficients, so too is our image.

The spatial resolution which can be achieved in UOT is entirely defined by the size of the insonified region, though smaller regions of acousto-optic interaction result in lower signal to noise ratios in measurement. At shallow depths a resolution of $100\mu\text{m}$ has been demonstrated [115], but resolutions in the millimetre range are more common at depths of a few centimetres. The imaging depth in UOT is dependent upon optical attenuation, since light must reach the point at which it is modulated, before returning to the detector point for measurement. Penetration depths beyond 90mm have been demonstrated *ex-vivo* in a medium of optical properties similar to human breast tissue [67], but it is more typical to experience penetration depths of the same order of magnitude as in DOT (a few centimetres) [38].

Two questions regarding the design of a UOT system follow immediately from our introduction:

- How do we best detect the small ($\sim 10^6\text{Hz}$) shifts in frequency from the background light ($\sim 10^{15}\text{Hz}$)? This question has achieved significant attention from numerous researchers over the past twenty years, and we will review the relevant contributions in chapter 2.
- As is the case in PAT, our measured data and associated raw images do not quantitatively represent the underlying optical coefficients we wish to recover. The reconstruction of the optical properties from the detected data constitutes the inverse problem in acousto-optics, and has received considerably less attention in the literature than the experimental aspects of the problem.

A practical implementation of UOT is likely to offer lower spatial resolution than that which can be achieved in PAT. Nonetheless, UOT provides an order of magnitude improvement in spatial resolution over DOT.

Unlike PAT, UOT has similar relative sensitivities to absorption and scattering as DOT. Whilst knowledge of the local optical absorption coefficient is of significant clinical value, the scattering coefficient also conveys information regarding anatomical structure [24] and functional responses [50]. Moreover, since UOT does not require absorption for contrast (as is the case in PAT), there is the potential to image scattering structures in low-absorption backgrounds to substantially great depth. As a first step towards quantitative imaging, we will

concentrate in this work on the recovery of the absorption coefficient, principally because its clinical relevance is better understood than that of scattering, but also because this forms the basis for the future development of spectroscopic techniques.

1.2 Overview and contribution

In this work we will develop theory and methods which address the forward and inverse problems in acousto-optics. We will primarily concentrate on autocorrelation UOT, the advantages and disadvantages of which are discussed in the theoretical review of chapter 2. The reasons for this choice are twofold: first, the technique is ‘close’ to the underlying physics, allowing direct access to the power spectral density of the detected light; second, this detection technique is under active development within our research group at the time of writing, and thus the results of this work have direct application in ongoing empirical studies.

In chapter 2 we will begin with a review of the development in the field, before providing the theoretical basis from which we will develop our work.

We will begin the investigation-proper in chapter 3 where we will develop a reference forward model based upon a Monte-Carlo solution to the radiative transport equation. We will take advantage of recent developments in the highly parallel architecture of graphics processing units (GPUs) to accelerate the model, and employ a mesh based optical domain to maximise the flexibility of the simulation. This work has implications beyond UOT, demonstrating techniques for efficient simulation in diffuse optics. The work developed in this chapter was published by the author in 2012 [105], following earlier developments in 2010 [73]. The tool developed during this chapter was exploited in work by others in our research group [52], [72].

Whilst accurate and optimised, the MC model of chapter 3 is too slow to be used as the basis for an image reconstruction (inversion) technique; rather, it serves as a useful reference against which alternative models can be validated. Sakadžić and Wang previously published the results of a diffusion-like model which can be solved quickly by deterministic numerical methods [118]. In chapter 4 we provide a previously unpublished derivation for their model, and we do so with a level of generality which allows it to be extended to consider higher order effects not present in the original version. We use this model of the AO interactions in a modified light model which employs the simplified spherical harmonic method, and demonstrate significant improvements in the model over the original with respect to our reference Monte-Carlo simulation.

In chapter 5 we employ the models of chapter 4 to develop an image reconstruction procedure for autocorrelation UOT. We develop discrete-form sensitivity relations which describe the sensitivity of the various measurement types to perturbations in the acoustic absorption in the medium. These sensitivity relations are used as part of a linear reconstruction process which produces images of the optical absorption of a medium from a single optical source-detector pair, and a raster scan of the acoustic field. We examine the qualities of different

measurement types, and the effects of noise on the degree of regularisation required to produce a useful solution. This work is formulated in the style of diffuse optical tomography, and naturally lends itself to a practical implementation. The work of this chapter was published by the author in 2013 [106], following earlier reports at conference [106].

This work closes in chapter 6 with a summary of the work, and an extensive discussion of how the outcomes of this investigation can be developed, improved, and generalised to alternative detection mechanisms and insonification protocols.

Chapter 2

Background

2.1 Overview

In this chapter we will review the theory and literature pertinent to our investigation, develop the theoretical framework of the acousto-optic interactions which underpin this work, and formalise the quantities and measures in the particular experimental configuration targeted in this work.

We begin with a brief overview of tissue optics, introducing the relevant physical quantities and pertinent models of light transport in tissue. In section 2.3.1 we will review the literature of the field, considering separately the experimental and theoretical advances which have taken place in the field. In section 2.4 we will analyse the nature of the acousto-optic modulation in UOT. In section 2.5 we will consider the models which have been proposed which employ the mechanisms of acousto-optic modulation to describe UOT. Finally, in section 2.6, we will describe the specific experimental methods targeted in this thesis, the observables and related quantities.

2.2 Tissue optics

This brief introduction to tissue optics will primarily serve to formally introduce the physical quantities and assumptions made throughout this work, and the most relevant model of light in tissue: radiative transport. For a more complete review and analysis the reader should refer to texts such as [56, 85, 94, 137].

2.2.1 Absorption and scattering

Exogenous light propagating in typical biological tissue undergoes two primary processes: absorption, and scattering.

Absorption occurs when the atoms which constitute biological tissue are excited by some incident radiation, and this energy is partially or fully lost by vibrational relaxation, rather

than by a re-radiation of the incident field. Different atoms existing in varying molecular combinations each have specific wavelengths at which they can be excited, and hence absorption is a wavelength dependent phenomenon, giving rise to our perception of colour. One of the most pertinent *chromophores* in biomedical optical imaging is haemoglobin, which exhibits different spectra in its oxygenated and de-oxygenated states, and provides great insight owing to its functional relevance in physiology. We quantify absorption in biological tissues by the absorption coefficient μ_a [mm^{-1}]. The absorption coefficient will be considered in greater depth in chapter 3.

Scattering is the process by which a number of atoms or molecules collectively excited by the incident radiation re-radiate the field with a different wave-front as that of the input excitation, owing to their size and geometrical arrangement. Scattering is also wavelength dependent, exhibiting a monotonic dependence with wavelength over the ranges typically employed in near infra-red optical imaging [127], such as the techniques described in this work. We quantify the amount of scattering by the scattering coefficient μ_s [mm^{-1}]. This quantity describes the efficiency of the scattering process in tissue, but does not describe the nature of the scattering. It is thus common to also employ a *phase function* which describes the probability density function of the angle at which light is scattered. The implications of the phase function will be elucidated in chapter 3, and this will be made use of further in chapter 5. An important related property is the mean-cosine of the scattering angle, $g = \langle \cos \theta \rangle$. This describes the relationship between how much light is scattered in the incident direction, and how much is back-scattered. In most tissues $g \approx 0.9$, which means that incident radiation, scattered to a degree indicated by μ_s , is mostly scattered in the direction of incidence. This definition gives rise to the transport scattering coefficient $\mu'_s = \mu_s(1 - g)$ [m^{-1}] which describes the amount of scattering in a medium as if each body scattered the incident light in an isotropic fashion. Accordingly, the transport mean free path $l_{\text{tr}} = \mu'^{-1}_s$ [m] describes the distance that light travels (on average) before its direction has been randomised by scattering processes.

2.2.2 Radiative transport

A rigorous and accurate description of light transport in tissue could be provided by solution (by various numerical methods) of Maxwell's equations. However the disordered, random nature of biological tissues would require such a solution to be averaged over all the realisable states of organisation of the underlying tissue.

The most relevant approximation for the purposes of the optical imaging is the radiative transport equation (RTE). The RTE operates under the assumption of weak-scattering, the high-frequency limit of optics, where it is assumed that the effects of interference can be neglected and different components of the optical field can be summed in power. The RTE for continuous wave illumination is

$$[\hat{\mathbf{s}} \cdot \nabla + \mu_a + \mu_s] L(\mathbf{r}, \hat{\mathbf{s}}) = \mu_s \int p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') L(\mathbf{r}, \hat{\mathbf{s}}') d\hat{\mathbf{s}}', \quad (2.1)$$

where $L(\mathbf{r}, \hat{\mathbf{s}})$ [$\text{W m}^{-2} \text{ sr}^{-1}$] is the specific intensity, $p(\hat{\mathbf{s}}, \hat{\mathbf{s}}')$ is the phase function described previously, and the integral takes place over all solid angle. The RTE is an energy balance expression for the flow of energy at \mathbf{r} in direction $\hat{\mathbf{s}}$: the first term is the net change of energy, and this is equated to the losses due to absorption and out-scattering, and in-scattering from all directions $\hat{\mathbf{s}}'$ to $\hat{\mathbf{s}}$.

We will discuss the theoretical development of radiative transfer, and the nature of specific intensity, or radiance, in chapter 3.

2.3 A concise review of ultrasound modulated optical tomography

We will review the field of ultrasound modulated optical tomography in two parts which reflect those aspects of the investigation first raised in our introduction. First, we shall consider the experimental advances which have taken place, specifically considering the alternative detection mechanisms and insonification strategies which have been described in the literature. Second, we shall consider the theoretical developments in the field, to which we refer to work pertaining to the physical mechanisms of modulation, modelling, and image reconstruction. Further details can be found in the recent reviews of Reskink et al. [110], Elson et al. [38], and Morgan et al. [101].

2.3.1 Detection and insonification strategies

Marks et al. [91] first detected the acoustic modulation of a coherent optical field in a turbid medium by band-pass filtering the signal from a single photo-multiplier tube (PMT) at the centre frequency of acoustic pulses. Wang et al. followed a similar approach [133] but employed a focused continuous wave transducer which was scanned through the medium to provide an image of an embedded absorber. Wang and Ku [136] developed their method by employing a frequency swept technique which encoded the location of the optical signal along the axis of the ultrasound field. At a similar time to Wang's initial publication, Leutz and Maret [74] employed a digital correlator to indirectly sample the power spectrum of light detected by a PMT in a photon-counting configuration, in addition to using a Fabry-Pérot interferometer for direct analysis of the spectrum of the detected field. Later, Sakadzic and Wang would use the Fabry-Pérot technique to produce high resolution ($< 100 \mu\text{m}$) images of absorbers embedded in turbid media [115].

Each of the methods demonstrated to this point employed single detector systems. Such methods are both sensitive and allow the evolution of a speckle grain, or *coherence area*, to be followed in real-time — this is important in tissue as decorrelation of the field due to Brownian motion (as studied directly in DWS) would otherwise serve to average out the response. The downside of these methods is that the signal to noise (SNR) is typically optimised by detecting only a single coherence area, from which only a small optical flux (relative to the total exiting flux) can be obtained.

The introduction of parallel detection mechanisms, capable of sampling the flux from multiple coherence areas simultaneously, began with the work of Lévéque et al. [77]. In this work the authors demonstrated an improvement in SNR by detecting multiple mutually incoherent speckle grains on the individual pixels of a charge coupled device (CCD) array. Since the integration period of such a device is far greater than the period of the ultrasonic modulation, a method was derived which determines the modulated component by separately integrating the collected flux over four phases of the acoustic cycle, averaged over multiple cycles. This *lock-in* method has the disadvantage of being less immune to the decorrelation which occurs in living tissue. Continued experimental investigation ensued employing parallel detection techniques. Yao et al. [138] incorporated the ideas of Wang and Ku's earlier work by coupling a frequency swept modulation with parallel lock-in detection. Lévéque [78] demonstrated three-dimensional imaging of an absorbing inclusion in a biological sample media. More recently, Ruan et al. extended the technique to detect modulation from the first overtone generated by non-linear propagation of the acoustic field in the medium [112], and also to employ pulsed ultrasound [111].

Li et al. [80] proposed an alternative parallel CCD method in which the variance of the speckle field integrated over time is measured. The premise of this *speckle-contrast* technique is that under acoustic modulation the individual speckle grains oscillate in their intensity; thus, as the degree of modulation can be derived by considering how spatially 'smooth' the field becomes during averaging. The integration period in this method is shorter than most of the lock-in methods, and hence speckle decorrelation is better tolerated.

Parallel detection methods were further advanced by the integration of interferometric techniques. Gross et al. [51] proposed mixing the speckle pattern from the turbid medium with a reference beam shifted by the acoustic frequency in a form of *digital holography*. A principle advantage of this method is that the desired signal from the sample is amplified by heterodyne mixing with the reference arm: the amplitude of the reference arm can be increased to the point that the technique becomes shot-noise limited. Since a CCD is used for detection, the fast oscillations of the unmodulated light and the shifted reference beam are averaged by the detector. By tilting the angle of the reference beam with respect to the CCD the different spatial frequencies of the speckled field from the sample are spread out in k -space (the two-dimensional Fourier transform of the detected field). The static interference fringes between the shifted reference beam and the desired modulated field occupy a portion of k -space dependent upon the angle of the two beams. The decorrelation noise of the speckle field is primarily limited to the low spatial-frequencies and can thus be separated from the desired signal. The shot noise of the reference beam generates a DC offset across k -space which can be calculated and removed from the desired signal. In practice, the reference beam is shifted by a slightly greater amount from the acoustic frequency in order that a four-phase lock-in detection method can be employed. This technique was adapted by Ruan et al. [113] to detect modulation due to trains of ultrasound pulses, and extract the modulation due to higher acoustic harmonics.

An alternative holographic technique which has generated significant interest is the use of photorefractive crystals. In this method a photorefractive crystal is illuminated by a signal beam from the insonified sample, and a reference (local oscillator) beam. In the approach initially proposed by Murray et al. [102] the local oscillator is at the same frequency as that illuminating the sample. In this case a hologram of the interference pattern between the unmodulated light from the sample and the reference beam is written into the crystal via the Pockels effect. The modulated component of the sample beam varies at a frequency beyond the response time of the crystal, and does not form a component of the hologram. Once formed, a part of the sample beam is diffracted by the holographic grating in the crystal to become phase matched with the reference beam and the resulting coherent wavefront can be integrated over a single detector. This approach allows inference of the modulated component of the light from the reduction in the unmodulated intensity. Axial resolution can be achieved by the use of pulsed ultrasound. In an alternative approach suggested by Ramaz et al. [109] the reference beam is up-shifted to the frequency of ultrasonic insonification such that the static interference pattern in the crystal is a result of only the modulated light emanating from the sample. In this method light from the reference beam is diffracted from the resulting hologram to form a wavefront coherent with the modulated light. The two signals interfere on a detector and the heterodyne signal is extracted. In practice a method of modulation is applied to the ultrasound to distinguish the signal temporally from background noise. Despite the use of a single detector, the optical étendue remains large, defined by the area of the crystal employed in the experiment. The principal limitation of photorefractive techniques is the response time of the crystal: if the crystal responds to changes in the wavefront too slowly, the effects of decorrelation will be recorded as if they were a result of the modulated field, absent of decorrelation.

Another recent technique which can utilise a single detector is the spectral-hole burning (SHB) method, as demonstrated by Li et al. [81]. Unperturbed, the doped SHB crystal absorbs the incident wavelengths of interest. By pumping the crystal with a reference beam at the acoustically modulated frequency of interest, the appropriate energy levels of the doped crystal are saturated. At this time, and for tens of milliseconds following, the crystal is transparent to the same wavelengths. Whilst this ‘spectral-hole’ exists, the signal beam can be passed through the crystal and the unmodulated light will be absorbed, whilst the relevant signal passes to the detector. The chief disadvantage of this approach is that the crystal must be cryogenically cooled.

2.3.2 Modulation, modelling, and inversion

Leutz and Maret’s previously described experimental work [74] derived the power spectrum of the light from a UOT experiment by taking the Fourier transform of the autocorrelation function calculated by a digital correlator. As part of their work, the authors developed a model of the AO effect in turbid media insonified by a plane wave. This model was based

upon the approach of diffusing wave spectroscopy (DWS) [89, 103]. Such models calculate the expected autocorrelation in an experiment based upon the modulation of optical path-lengths in a medium. In addition to the Brownian motion of DWS, Leutz and Maret incorporated the movement of the optical scatterers in the medium, assuming them to oscillate in harmony with the applied acoustic field.

Kempe et al. [62] extended the analysis of Leutz and Maret to consider a narrow beam of ultrasound and compared this model to experimental data. Mahan et al. [88] proposed the acoustic tagging process to be that of Brillouin scattering, and proceeded to develop an expression of the intensity of the modulated signal and the signal-to-noise ratio in the case of an ultrasound field focused into a small volume.

Theoretical understanding of UOT was advanced by Wang [135] who jointly modelled the two coherent mechanisms of interaction previously investigated: the acoustically induced displacement of optical scatterers and the change in the index of refraction due to compression and rarefaction. Theoretical research continued with the extension of Wang's model to the case of anisotropically scattering media [114]. Each of these works assumed an optical transport mean free path far greater than the acoustic wavelength such that successive scattering events in the medium could be considered uncorrelated. A theory accounting for strong correlations in the interaction at successive scattering sites was presented later [116].

Blonigen et al. [22] investigated the nature of the phase shifts applied to ultrasound modulated light in the context of a photorefractive crystal detection regime.

The DWS style of the analytical models presented thus far were limited to geometries for which an analytic form of the optical path length from source to receiver was known. Further developments by Sakadžić and Wang [117] moved away from this approach with the development of a correlation transfer equation (CTE). The CTE provides a treatment of the acoustic continuous-wave (CW) UOT problem in the style of a transport equation, thus capable of modelling arbitrary optical geometries. The authors provide a rigorous derivation of the model based upon multiple scattering theory (as opposed to transport theory [56]) where the time-domain effects of acoustic modulation are integrated into the mutual coherence function which is ultimately transformed into the more familiar quantity of specific intensity. Owing to the lack of analytical solutions to the transport equation, the authors demonstrate its application with a Monte-Carlo model. In chapter 3, we take a heuristic approach to the development of a similar MC model which seeks to retain a clear distinction between the purely optical, and the acousto-optic aspects of the simulation. The two models will be seen to be equivalent, though much like distinction between analytical and heuristic derivations of the standard radiative transport equation, the former provides a more clear physical basis to the field variable, and the approximations required in the derivation. In later work, the authors reformulate their CTE to account for pulsed ultrasound pulse [119].

In addition to their derivation of the CTE, Sakadžić and Wang present a diffusion approximation for the CW UOT problem [118]. This work is of great importance since the diffusion style formulation can be solved by standard numerical means, such as finite element,

or finite difference methods. The computation of such solutions is relatively inexpensive, when contrasted with MC based approximations to the underlying CTE. These advantages come at the cost of the applicability of the model: various assumptions and approximations are made in the derivation which limit its applications to higher frequency ultrasound fields and moderate ultrasound pressures. The limitations in the acousto-optic aspect of the model are in addition to those introduced due to the diffusion approximation, principally that the model is only valid in regions where the optical scattering coefficient is much greater than the absorption coefficient. Despite its reduced flexibility, the flexibility and moderate computational requirements of the diffusion approximation provides the best potential for use as part of an iterative model-based inversion. For this reason we provide a full derivation of the model in chapter 4 where we also extend the technique with a higher order light transport model. Later, in chapter 5, we employ the model as part of a linear image reconstruction procedure for UOT.

Bal and Schotland [16] developed a model of incoherent acousto-optic modulation in which both the index of refraction and the number density of the scattering and absorbing particles in a given medium are modulated by a standing acoustic wave. The resulting spatially varying absorption coefficient, scattering coefficient, and index of refraction, were incorporated into a diffusion model of light transport. For a given source and detector, the authors propose measuring the Fourier transform of the forward model in the variables of the acoustic wave-number and phase; the absorption and scattering are recovered via the inversion of this transform. This approach to UOT has not yet been demonstrated experimentally (excepting a particular example involving fluorescent biological media) and will not be considered further in this work.

More recently, Bratchenia et al. [28] demonstrated the recovery of the optical absorption coefficient in a turbid medium via a model-based inversion. In their work the UOT signal was described in the frequency domain, and the acoustically-driven modulation of the coherent input light was linearised in a diffusion-style model, similar to that presented by Allmaras and Bangerth [6]. A three-dimensional recovery of the optical absorption coefficient was performed using an iterative non-linear optimisation employing the Levenberg-Marquardt algorithm. In chapters 4 and 5 of this work we take an alternative approach to the forward and inverse problem in UOT: we employ a time-domain model of the UOT signal which naturally provides many of the oft-employed measurement types in UOT, such as the modulation depth. Our forward model considers the non-linearity of the acoustically driven decorrelation of the optical field, which allows us to examine in greater depth the form of the sensitivity functions which arise in the UOT experiment. In an approach common in DOT, we linearise our forward model in the *optical* parameters; this leads to a one-step difference data reconstruction which resolves an absorption perturbation from an assumed known background.

Huynh et al. [55] modelled the detected acousto-optic signal as a temporal convolution of the acoustic and optical intensity profiles as a pulse propagates along the axis of the ultra-

sound transducer. Discretised, the resultant matrix operators were inverted by a maximum likelihood algorithm. Despite the simplicity of the model, good agreement was found with experimental results.

2.4 Mechanisms of coherent acousto-optic modulation

In the review of section 2.3.2 we described two potential coherent mechanisms for the interaction between light and sound discussed in the literature.

- The displacement of optical scatterers from their rest positions by the propagating acoustic field altering the optical paths of diffuse light through the medium [74]. This mechanism is the basis of diffusing wave spectroscopy (where Brownian motion, rather than the acoustic field, generates a displacement of the scatterers), and has been used in this manner to analyse various parameters of optical scatterers within turbid media.
- The perturbation of the local refractive index of the medium through the strain introduced by the propagating acoustic field altering the optical path lengths of diffuse light through the medium [135]. This technique can be applied to both non-turbid and turbid media, and is used in numerous optical instruments such as the acousto-optic modulator (Bragg cell).

2.4.1 Displacement of optical scatterers

As an acoustic field propagates through a turbid medium the endogenous optical scatterers of the medium will be displaced from their rest positions and thus the optical path length between the two scatterers will be altered. To describe this mechanism of modulation we consider an optical path between two scatterers in the medium which are located at \mathbf{r}_a and \mathbf{r}_b , as depicted in figure 2.1.

The phase accrued over the unperturbed optical path is given by

$$\phi_{a,b} = k_0 n_0 |\mathbf{r}_b - \mathbf{r}_a|, \quad (2.2)$$

where k_0 is the optical wave number *in vacuo*, and n_0 the background index of refraction. The particle displacement of a time-varying monochromatic acoustic field is defined locally at \mathbf{r} as

$$A(\mathbf{r}, t) = A_0(\mathbf{r}) \cos(\omega_a t + \phi(\mathbf{r})), \quad (2.3)$$

where $A_0(\mathbf{r})$ and $\phi(\mathbf{r})$ are the local acoustic displacement amplitude and phase offset at \mathbf{r} , and ω_a the acoustic angular frequency. The displacement ξ of an optical scatterer at \mathbf{r} is assumed to follow the background displacement of the medium, accordingly,

$$\xi(\mathbf{r}, t) = \hat{\mathbf{k}}_a(\mathbf{r}) A_0(\mathbf{r}) \cos(\omega_a t + \phi(\mathbf{r})) \quad (2.4)$$

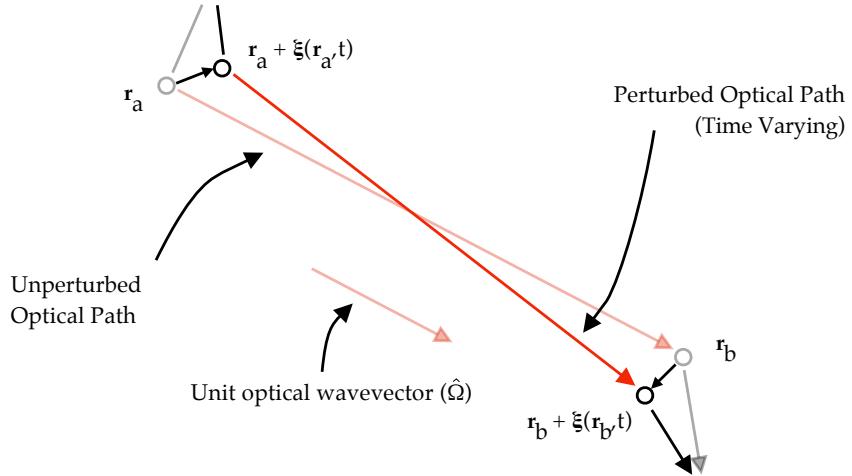


Figure 2.1: Illustration of the time varying acoustic displacement of the scatterers of a turbid medium perturbing the optical path between scattering sites.

where $\hat{\mathbf{k}}_a(\mathbf{r}) = \mathbf{k}_a(\mathbf{r})/|\mathbf{k}_a(\mathbf{r})|$ is the unit acoustic propagation direction and \mathbf{k}_a the local acoustic wave vector at \mathbf{r} . For atypical scatterers, it might be necessary to introduce a phase offset and modified amplitude of oscillation relative to the background medium, though this is not expected to be required for endogenous biological scatterers under moderate acoustic amplitudes and within the frequency range of medical ultrasound systems [117]. We can define the incremental phase of the perturbed path due to the displacement of optical scatterers as

$$\Delta\phi_{d_{a,b}}(t) = k_0 n_0 \hat{\mathbf{k}} \cdot (\xi(\mathbf{r}_b, t) - \xi(\mathbf{r}_a, t)), \quad (2.5)$$

where $\hat{\mathbf{k}}$ is the unit optical propagation direction for the path between the two unperturbed scattering sites. The total time-varying phase accrued over the perturbed path is

$$\phi_{d_{a,b}}(t) = k_0 n_0 \left[|\mathbf{r}_b - \mathbf{r}_a| + \hat{\mathbf{k}} \cdot (\xi(\mathbf{r}_b, t) - \xi(\mathbf{r}_a, t)) \right]. \quad (2.6)$$

Rather than considering the phase due to displacement of optical scatterers as a quantity defined over a path, an alternative approach is to consider the extra phase increment as a function of a particular scattering site. Assume the optical path of figure 2.1 now extends to a third scattering site at \mathbf{r}_c with unit optical propagation direction $\hat{\mathbf{k}}_{b,c}$. The phase increment over the path from \mathbf{r}_a to \mathbf{r}_c will be the summation of the two individual increments,

$$\Delta\phi_{a,c} = k_0 n_0 \left(\hat{\mathbf{k}} \cdot (\xi(\mathbf{r}_b, t) - \xi(\mathbf{r}_a, t)) + \hat{\mathbf{k}}_{b,c} \cdot (\xi(\mathbf{r}_c, t) - \xi(\mathbf{r}_b, t)) \right). \quad (2.7)$$

If the scatterers at \mathbf{r}_a and \mathbf{r}_c are stationary (or the incremental phase due to their oscillation has already been considered in some series expansion),

$$\Delta\phi_{a,c} = k_0 n_0 \left(\hat{\mathbf{k}} \cdot \xi(\mathbf{r}_b, t) - \hat{\mathbf{k}}_{b,c} \cdot \xi(\mathbf{r}_b, t) \right). \quad (2.8)$$

Denoting the scatterer at position \mathbf{r}_b the j^{th} scatterer, the incoming optical propagation direction $\hat{\mathbf{k}}_j = \hat{\mathbf{k}}$ and the outgoing propagation direction $\hat{\mathbf{k}}_{j+1} = \hat{\mathbf{k}}_{b,c}$,

$$\Delta\phi_j = -k_0 n_0 (\hat{\mathbf{k}}_{j+1} - \hat{\mathbf{k}}_j) \cdot \xi(\mathbf{r}_b, t). \quad (2.9)$$

2.4.2 Perturbation of the index of refraction

As an acoustic field propagates through a medium a mechanical strain is developed as regions are locally compressed and rarefied. This strain modifies the local permittivity in the medium, and thus the index of refraction, as depicted in figure 2.2. The magnitude of this interaction is described by the adiabatic piezo-optic coefficient dn/dp of the medium, which is $1.466 \times 10^{-10} \text{ m}^2 \text{ N}^{-1}$ in water [108]. We follow the literature [135] in defining the quantity

$$\eta = \frac{dn}{dp} \rho v_a^2, \quad (2.10)$$

where ρ is the density of the insonified medium and v_a the speed of sound. This permits a convenient form of the resulting variation of the index of refraction when expressed in terms of acoustic particle displacement. Noting the relationship between acoustic particle displacement and pressure in an harmonic plane wave ($\xi = p/(\rho v_a \omega_a)$) allows the variation in the index of refraction from the background medium due to the acoustic field of equation 2.3 to be written as

$$\Delta n(\mathbf{r}, t) = n_0 \eta k_a A_0(\mathbf{r}) \cos(\omega t + \phi(\mathbf{r})). \quad (2.11)$$

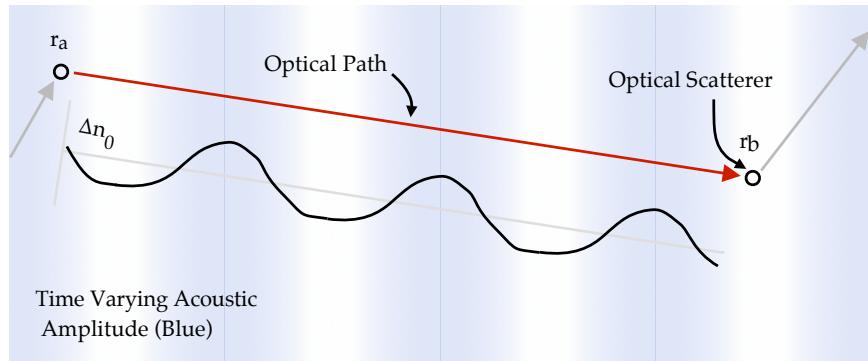


Figure 2.2: Illustration of the time varying index of refraction over the length of an optical path due to perturbation of the index of refraction by an applied acoustic field.

Let us again consider an arbitrary optical path between two scatterers in a medium at \mathbf{r}_a and \mathbf{r}_b . The phase accrued over the unperturbed optical path is given by equation 2.2. The incremental phase accrued over the path due to the spatially varying index of refraction is given by

$$\Delta\phi_{n_{a,b}} = k_0 n_0 \eta k_a \int_{\mathbf{r}_a}^{\mathbf{r}_b} A_0(\mathbf{r}) \cos(\omega t + \phi(\mathbf{r})) d\mathbf{r}. \quad (2.12)$$

The total time-varying phase accrued over the path with perturbed index of refraction is

$$\phi_{n_{a,b}} = k_0 n_0 \int_{\mathbf{r}_a}^{\mathbf{r}_b} [1 + \eta k_a A_0(\mathbf{r}) \cos(\omega t + \phi(\mathbf{r}))] d\mathbf{r}. \quad (2.13)$$

2.4.3 Combining the mechanisms

In theory, the integration of the perturbed index of refraction of section 2.4.2 should be performed over the limits defined by displaced optical scatters of section 2.4.1. A complete expression for the incremental phase over a path between two scattering sites at \mathbf{r}_a and \mathbf{r}_b would then be

$$\Delta\phi_{a,b} = k_0 n_0 \left[\hat{\mathbf{k}} \cdot (\xi(\mathbf{r}_b, t) - \xi(\mathbf{r}_a, t)) + \eta k_a \int_{\mathbf{r}_a + \xi(\mathbf{r}_a, t)}^{\mathbf{r}_b + \xi(\mathbf{r}_b, t)} A_0(\mathbf{r}) \cos(\omega t + \phi(\mathbf{r})) d\mathbf{r} \right], \quad (2.14)$$

and the complete phase increment over the path given by

$$\phi_{a,b} = k_0 n_0 |\mathbf{r}_b - \mathbf{r}_a| + \Delta\phi_{a,b}. \quad (2.15)$$

In practice, the additional phase in the limits of the integration is often neglected from analytical formulations as $|\xi(\mathbf{r}, t)|$ is considered to be much smaller than a typical free path $|\mathbf{r}_b - \mathbf{r}_a|$ at the frequencies and pressures of interest [118]. In this case, the two mechanisms and their associated phase increments are considered independently.

2.5 Models of ultrasound modulated optical tomography

Two primary analytical models of UOT were discussed in the earlier review. The first is based upon an extension of the methods of diffusing wave spectroscopy. The second is based upon a reformulation of the radiative transfer equation. In this chapter we introduce both methods, their limitations and capabilities. We provide only an overview of the latter method here, since the theory is developed further in the body of this work (particularly in chapters 3 and 4).

2.5.1 UOT models based upon diffusing wave spectroscopy

Diffusing wave spectroscopy investigates the nature of a turbid media by analysing the dynamic properties of the multiple scattering of diffuse light passing through the medium as the scatterers undergo Brownian motion. This has a clear corollary with the first mechanism of interaction described in this chapter, and for this reason, many of the methods of analysis applied to this field were subsequently applied to consider the acoustically driven displacement of a turbid medium's scattering objects.

The measure used in the DWS approach, both for empirical measurement and as the output of the analytical formulations, is the un-normalised electric field autocorrelation func-

tion¹ [103],

$$G(\tau) = \langle E(t)E^*(t + \tau) \rangle_t \quad (2.16)$$

where t is time, τ is lag, $E(t)$ is the electric field of the scattered light, $E^*(t)$ is the conjugate of the electric field, and the angle brackets represent averaging over time. The electric field at a given point in space is the summation of all the possible scattering paths taken by the light through the medium from a coherent source. For DWS we expect that as the scattering particles move randomly under Brownian motion, the correlation of the electric field will decrease according to the magnitude of the mean-square particle displacement, which is in turn related to the temperature, size, and diffusivity of the particles [89]. If we neglect Brownian motion and instead consider the autocorrelation due to the application of a harmonic ultrasound field, we might now expect to find an autocorrelation function which itself is harmonic; the acoustic field, via the mechanisms of interaction, modifies the individual scattering paths to a varying extent before returning to their initial values, and so on. In this case, we might find it convenient to consider the power spectral density of the detected electrical field. By the Wiener-Khinchin theorem we can find the optical power spectral density by taking the Fourier transform of the electric field autocorrelation function,

$$S(\omega) = \int_{-\infty}^{\infty} G(\tau) \exp(-j\omega\tau) d\tau. \quad (2.17)$$

The forms of the autocorrelation functions and power spectrum densities for both DWS and acousto-optics are illustrated in figure 2.3.

To develop an expression for the autocorrelation function which considers both mechanisms of interaction under insonification by an acoustic field, we begin by assuming the ratio of the optical transport mean free path and the optical wavelength to be much greater than unity ($l_{tr}/\lambda_0 \gg 1$). As such, we imply there is no correlation between different optical paths: $\langle E_{s=s_1}(t)E_{s \neq s_1}^*(t + \tau) \rangle_t = 0$ where $E_s(t)$ is the fraction of the incident electric field scattered along a path of length s . This assumption of weak scattering allows us to write the autocorrelation function as an integration over all possible path lengths,

$$G(\tau) = \langle E(t)E^*(t + \tau) \rangle_t = \int_s p(s) \langle E_s(t)E_s^*(t + \tau) \rangle_t ds, \quad (2.18)$$

where $p(s)$ is the probability density function of s . In doing so, we also assume that the acoustic amplitude is not so great as to cause variation in $p(s)$ over an acoustic period. Taking the scattering anisotropy $g = 0$ such that $l_{tr} = l$, there will be $N + 1$ paths over which we must accrue phase increments for a path involving $N \approx s/l_{tr}$ scattering events. The resultant

¹In the literature, this often denoted G_1 to distinguish it from the intensity autocorrelation function G_2 . We will neglect the subscript in this work.

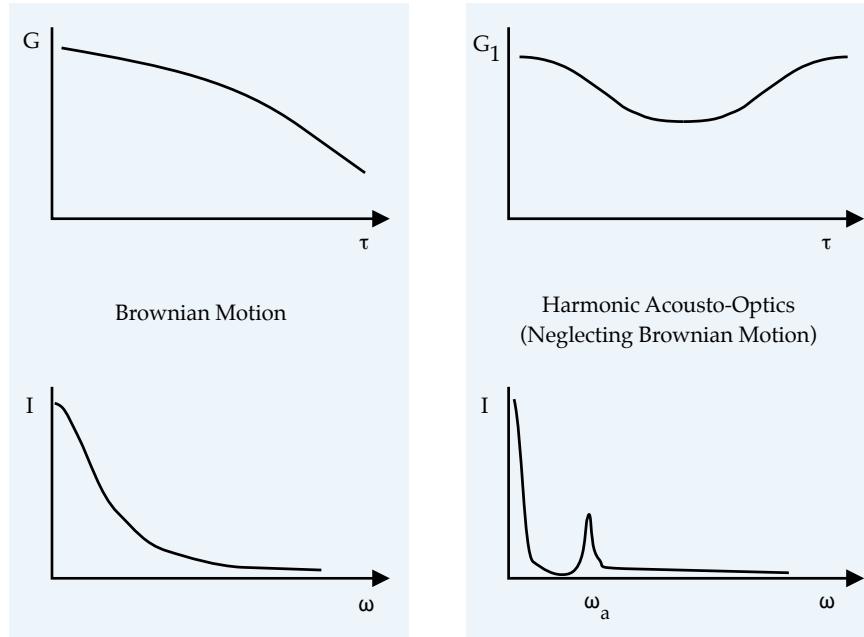


Figure 2.3: Illustration of the form of the electric field autocorrelation function and power spectral density expected from a diffusing wave spectroscopy, and acousto-optic measurement.

field autocorrelation function can be written

$$\langle E_s(t)E_s^*(t+\tau) \rangle_t = \left\langle \exp \left[-j \left(\sum_{j=1}^{N+1} (\Delta\phi_{n,j}(t+\tau) - \Delta\phi_{n,j}(t)) + \sum_{j=1}^N (\Delta\phi_{d,j}(t+\tau) - \Delta\phi_{d,j}(t)) \right) \right] \right\rangle_t, \quad (2.19)$$

where $\Delta\phi_{n,j}$ is the phase increment due to the perturbation of the index of refraction over the path leading to the j^{th} scatterer, as defined in equation 2.12, and $\Delta\phi_{d,j}$ is the phase increment due to displacement of the j^{th} optical scatterer, as defined in equation 2.9. Completing the formulation of the the analytic theory now requires two steps.

- First, the terms for the phase increments must be averaged over all time, path lengths and propagation angles. In Wang's original formulation [135] this algebraic manipulation was performed by assuming an acoustic plane wave under conditions of isotropic scattering in an homogeneous medium, with weak correlations between scattering phase increments ($k_a l_{\text{tr}} \gg 1$). Later work developed the algebraic manipulation to include anisotropic scattering [114] (in the phase increments) and strong correlations between scattering phase increments [116].

- Second, the probability distribution of the potential scattering path lengths ($p(s)$) for a particular geometry and optical source-detector configuration must be sought, in order to permit the integration over s : the diffusion approximation has been used for this purpose in each of the studies of DWS. Solutions have been limited to studies involving slab geometries under planar illumination.

In summary, the analytical solutions to the acousto-optic effect based upon diffusing wave spectroscopy are only capable of producing solutions for simple geometries for which the probability density function of the optical path from source to receiver are known. In all cases, assumptions of weak scattering and a planar acoustic field are employed. Despite their limitations, these solutions do provide insight into the relative importance of the modulation mechanism and their relationship to the acoustic and optical parameters of the problem.

2.5.2 UOT models based upon correlation transport

An alternative formulation of acousto-optic theory has been presented in the literature which aims to avoid the limitations of the DWS approach in permitting the consideration of arbitrary heterogeneous optical properties and ultrasound fields [118]. This formulation applies the methods (and approximations) of radiative transport [97] to develop a correlation transfer equation; this approach has previously been applied to media with scatterers undergoing Brownian motion [2, 35].

As was discussed in our introduction to tissue optics, the primary limitation of transport style models is that it is difficult to find solutions for anything but the simplest of geometries. As with radiative transport, a typical approach involves statistically sampling the equation in a statistical Monte-Carlo model.

We will consider transport style models of UOT (and the approximation of their solution by Monte-Carlo methods) in chapter 3, and further approximations which permit deterministic numerical solution in chapter 4.

2.6 Autocorrelation ultrasound modulated optical tomography

In this work we will concentrate on development of models and reconstruction techniques applied to autocorrelation ultrasound modulated optical tomography, a technique originally demonstrated experimentally by Leutz and Maret [74], and subsequently by others [27], [79].

2.6.1 Experimental technique

The scattering medium is illuminated by a coherent electromagnetic wave which can be point-like, in the case of direct fibre coupling, or over an extended area if the beam is expanded accordingly. The medium is insonified by a focused time-harmonic acoustic field which can

be scanned through the medium in the case of imaging, or fixed in the case of point-like sensing.

A speckled wavefront emanates from the sample which is modulated by the acousto-optic interaction. A single coherence area is selected and detected by an appropriate detector. Owing to the small signal levels available from a single coherence area, this is typically a PMT or avalanche photodiode (APD) operating in photon counting mode: that is to say that individual pulses are output from the detector as each photon is recorded.

The output of the detector can be processed directly by a dedicated digital autocorrelator, or alternatively this can be achieved in software by recording the arrival times of detected photons.

2.6.2 Observables and related quantities

The principal measurement type for autocorrelation UOT is the electric field autocorrelation function,

$$G^\Gamma(\tau) = \langle E^\Gamma(t)E^\Gamma(t + \tau) \rangle_t. \quad (2.20)$$

The superscript Γ denotes that this is a measurement of the modulation depth flux across the boundary, rather than a measurement of the modulation depth fluence at a point within the medium. In practice, this measurement is derived from the intensity autocorrelation function via the Siegert relationship.

Numerous measures derived from the field autocorrelation function have been employed in the literature. Some, such as the *modulation depth*, have been defined differently by various authors. One reason for this ambiguity is that depending upon the context the autocorrelation function may be analysed in the time (lag) domain, or by its Fourier transform in the frequency domain. In this work we will work in both domains, and we now outline the terminology we employ in each case, and how this relates to measures used by other authors. In all cases, we will assume that the autocorrelation function has been post-processed to remove the exponential decay due to Brownian motion evident in real measurements.

Time domain

The field autocorrelation function can be analysed directly in the time domain. A useful quality of the time domain measurements is that an evaluation of the un-normalised function at $\tau = 0$ gives the intensity that would be measured in a standard optical experiment. In the literature of UOT this is often referred to as the ‘DC’ component, though care should be taken since this since this quantity differs from the DC component of the power spectral density (see the following section).

$$G_{\text{DC}}^{(\Gamma)}(\mathbf{r}) = G^{(\Gamma)}(0, \mathbf{r}) = L^{(\Gamma)}(\mathbf{r}). \quad (2.21)$$

In the case that the superscript Γ is present, indicating a measurement of flux accross the boundary, the spatial location \mathbf{r} represents the location of a point detector. In its absence, the measurement is of the correlation fluence within the medium at the specified point.

Less ambiguous is that which is often referred to as the 'AC' component,

$$G_{\text{AC}}^{(\Gamma)}(\mathbf{r}) = G^{(\Gamma)}(0, \mathbf{r}) - G^{(\Gamma)}(T/2, \mathbf{r}), \quad (2.22)$$

which measures the total fluctuation of the autocorrelation function as a result of all spectral components in the signal. This measure is typically applied to the autocorrelation function *after* removal of any exponential decay due to Brownian motion.

One of the most common measurements is the modulation depth. We follow the common definition,

$$\text{MD}^{(\Gamma)}(\mathbf{r}) = \frac{G_{\text{AC}}^{(\Gamma)}(\mathbf{r})}{G_{\text{DC}}^{(\Gamma)}(\mathbf{r})} = \frac{G(0, \mathbf{r}) - G(T/2, \mathbf{r})}{G(0, \mathbf{r})} = 1 - g\left(\frac{1}{2}T, \mathbf{r}\right), \quad (2.23)$$

where T is the acoustic period, and $g(\tau, \mathbf{r})$ is the *normalised* field autocorrelation function.

Frequency domain

By the Wiener-Khinchin theorem, the Fourier transform of the autocorrelation function gives the power-spectral density (PSD) of the signal, $S^{(\Gamma)}(\omega, \mathbf{r})$. Since $E(t)$ is real and wide-sense stationary, the spectral-density is even, real and non-negative. As shall be seen in this work, it is often convenient to work in the frequency domain. Numerous authors have redefined the modulation depth as being the ratio of the PSD at some harmonic of the acoustic frequency to the component at DC, $\omega = 0$. We shall make clear the use of the frequency domain MD measurement at the N^{th} harmonic by referring to the *N^{th} harmonic modulation depth*

$$\text{MD}_N^{(\Gamma)}(\mathbf{r}) = \frac{S(N\omega_a, \mathbf{r})}{S(0, \mathbf{r})}, \quad (2.24)$$

where the superscript takes the same meaning as the previous case. The first harmonic modulation depth is a good approximation for the modulation depth when the modulation is small, $\text{MD} \approx \text{MD}_1$, $\text{MD} \ll 1$.

The total intensity may be found in the frequency domain by integrating over all frequency,

$$L^{(\Gamma)}(\mathbf{r}) = G^{(\Gamma)}(0, \mathbf{r}) = 2 \int_0^\infty S^{(\Gamma)}(\omega, \mathbf{r}) d\omega, \quad (2.25)$$

and hence the time domain modulation depth, which considers all acoustic harmonics, can be calculated by

$$\text{MD}^{(\Gamma)}(\mathbf{r}) = \frac{\int_{\omega_a}^\infty S^{(\Gamma)}(\omega, \mathbf{r}) d\omega}{\int_0^\infty S^{(\Gamma)}(\omega, \mathbf{r}) d\omega}. \quad (2.26)$$

This is rarely used in the literature, since the measurement is more easily acquired from the time-domain, and frequency domain measures are typically employed when the models only consider a finite number of the acoustic harmonics.

Chapter 3

A highly parallel Monte-Carlo model of correlation transport

3.1 Introduction

The problem of light-transport in highly scattering media is well described by the radiative transfer equation (RTE) under conditions which are typically satisfied by biological tissues [56]. Unfortunately, analytical solutions to the RTE are limited to simple geometries.

Approximate numerical solutions to the RTE can, in principle, be sought via deterministic numerical approaches such as the finite element method [131]. Such approaches require that the simulation domain be discretised to form a linear system of equations which can be solved by various direct or iterative techniques. In three dimensions, the RTE is an integro-differential equation which solves for the five-dimensional quantity of specific intensity (defined over space and the *direction of propagation*), this requisite discretisation leads to prohibitively large linear systems in the case of three-dimensional simulations.

An alternative approach is the application of the Monte-Carlo (MC) method in which the desired observables are estimated by a stochastic model [107, 120, 134]. An estimation converges towards its expected value as we run a large number of trials of the stochastic model, each of which samples some state of the system under investigation. In the case of light transport as described by the RTE, each trial consists of the propagation of an amount of energy launched into the medium and propagated according to statistical properties of the medium which themselves can be derived from the electromagnetic wave equation [85]. Whilst extremely flexible, the MC technique is computationally intensive; numerous variance reduction techniques have been employed to improve efficiency, with some researchers employing ‘white’ MC models whereby absorption processes are applied to the sampled trials of a single MC model in post-processing [3, 49, 64, 95]. Improvements in simulation speed have also been achieved by the application of aggressive, platform-specific, optimisations [40, 59].

In recent years, considerable advances have been made in executing MC simulations of

light transport on the highly parallel architecture of graphics processing units (GPUs). The first demonstration by Alerstam et al. [4] considered the problem of light transport in an homogeneous semi-infinite medium; increases in speed over a CPU implementation of up to three orders of magnitude were achieved. Fang and Boas [41] subsequently developed a general purpose code to perform MC simulations in heterogeneous voxelised geometries. Under specific circumstances speed increases over a CPU implementation of up to $\times 350$ were noted. Such improvements in performance have been realised because of a natural mapping of the ‘embarrassingly parallel’ MC algorithm to the GPU hardware architecture.

Monte-Carlo models can also be applied to the correlation transport problem in acousto-optics. Throughout the initial theoretical investigations of the acousto-optic technique, a number of CPU-based MC models were demonstrated to support the analytical solutions under development [114, 116–119, 135]. We recently demonstrated the acceleration of a MC AO simulation [73], achieving performance improvements of up to $\times 100$ over a CPU implementation equivalent to that presented in ref. [135]. We also demonstrated that our GPU-based AO simulations involved an order of magnitude of extra computational-effort over a pure optical simulation.

Whilst our previous implementation of an AO MC simulation [73] was limited to homogeneous semi-infinite slab geometries with plane wave ultrasound, here we will develop a code capable of simulating the AO effect in optically heterogeneous domains, with arbitrary (monochromatic) acoustic field distributions. Previous CPU-based AO MC simulations reported in the literature [118, 119] have implemented a similar level of flexibility by the use of a voxelised simulation domain for the spatial distribution of the optical and acoustic parameters. In this chapter we will demonstrate the implementation of optical heterogeneity via a mesh-based system which avoids the deleterious effects of a voxelised geometry [21], and exploit the computation power of the GPU to avoid approximations in the treatment of the AO phase shifts demonstrated in previous work [118].

Our aim is the development of a tool that has the potential to provide significant guidance in the application and optimisation of the AO technique as both an imaging and a sensing modality for a broad range of tissue geometries and detection mechanisms. A robust and flexible model is also required for the validation of approximate methods such as those described in chapter 4. In the future, a fast Monte-Carlo model may also serve as the basis of image reconstruction techniques for ultrasound-modulated optical tomography, though the computational burden would likely require implementation on a GPU-equipped cluster.

3.1.1 Overview

We begin in section 3.2 with an overview of the relevant theory of correlation transport, and an approximate adjoint technique employed for point-source, point-detector simulations. The implementation of the simulation program and its post-processing techniques are detailed in section 3.3, following a brief overview of the pertinent considerations of parallel

programming for GPU architectures. In section 3.4 we demonstrate the validation of the simulation program, consider its performance relative to other simulation codes, and demonstrate one application in the calculation of explicit spatial sensitivity maps corresponding to experimental studies recently undertaken.

3.2 Theory

3.2.1 Monte-Carlo, radiative and correlation transport

In the context of radiative transport, it is the task of the MC technique to repeatedly propagate fractions of the total input radiance through the medium along different paths which are generated by consideration of the statistical properties of the medium. A good technical overview of the relationship between the RTE and the various MC methods seen in the literature is provided by Sassaroli and Martelli [120]. We will describe the exact nature of our implementation in section 3.3.

It is unfortunate that in the literature each of these fractions of energy is often described as a ‘photon’. Contrary to the very nature of the photon [98] (in the quantum electrodynamical sense) these descriptions ascribe said fractions of energy a scalar value and a location in space. Whilst Ishimaru discussed the link between radiance (or specific intensity) and the average Poynting vector in his 1975 book [57], it is not until more recently that Mischenko provided a complete derivation of the RTE from Maxwell’s Equations [97], and continued to discuss in depth the relationship between the RTE and the physical quantities [99] underlying the concept of radiance. It is perhaps the intangible nature of the radiance which the RTE propagates which encourages this improper use: if the link between the quantities in the RTE (and thus a MC simulation) and an observable were physically obvious, no doubt a different nomenclature would have been chosen.

In the context of MC, we will in this work employ the term ‘photon packet’ to refer to some fraction of the radiance propagating in a medium. For our purposes, we will think of this photon packet as representing the propagation of a quasi plane-wave (of limited spatial extent) which carries energy through the medium along a particular path [137]. We may describe locations at which this ray interacts with a statistically sampled point in the medium. In this sense, some fraction of radiance which propagated along a path, j , can be ascribed a value corresponding to the underlying strength of the electric field of the quasi-plane wave,

$$L_j(\mathbf{r}, \hat{\mathbf{s}}) = \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} E_j(t) E_j^*(t) dt \quad (3.1)$$

where E_j is the complex electric field strength associated with the quasi-plane which constitutes $L_j(\mathbf{r}, \hat{\mathbf{s}})$, and the radiance is not dependent upon time since the system is wide-sense stationary. As the simulation determines the path defining j , the associated fraction of radiance is attenuated according to the Beer-Lambert law [120]. Upon detection the total radiance is found by summing over all fractions, $L(\mathbf{r}, \hat{\mathbf{s}}) = \sum_j L_j(\mathbf{r}, \hat{\mathbf{s}})$.

In the case of acousto-optics, we propagate samples of *correlation* radiance $G_j(\mathbf{r}, \hat{\mathbf{s}}, t, \tau)$ through the domain. By the definition of the un-normalised field autocorrelation function,

$$G(t, \tau) = \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} E(t) E^*(t - \tau) dt, \quad (3.2)$$

we describe the correlation radiance,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, t, \tau) = \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} E_j(t) E_j^*(t - \tau) dt. \quad (3.3)$$

To proceed, we must consider how $G_j(\mathbf{r}, \hat{\mathbf{s}}, t, \tau)$ varies in time along its path. In chapter 2 we explored the two mechanisms by which we consider coherent light to be modulated as it traverses a path between two scattering sites. In the MC procedure each input photon packet (quasi-plane wave) will undergo $N - 1$ sequential scattering events at \mathbf{r}_k , for $1 \leq k \leq N - 1$, $k \in \mathbb{Z}$ before it leaves the turbid medium. If we assume that the optical source does not move with the acoustic field, $\xi_s(\mathbf{r}_0, t) = 0$, we can rewrite the phase increment due to the displacement of scatterers (equation 2.9),

$$\phi_d(\mathbf{r}, t) = -n_0 (\mathbf{k}_{0,k+1} - \mathbf{k}_{0,k}) \cdot \xi_s(\mathbf{r}, t), \quad (3.4)$$

where $\mathbf{k}_{0,k}$ and $\mathbf{k}_{0,k+1}$ represent the optical wave-vectors prior to and following incidence with the scatterer at \mathbf{r}_k , respectively. Now, over the path j involving N free paths and $N - 1$ scattering events we may write the complex phase perturbation,

$$\phi_j(t) = \sum_{k=1}^{N-1} \phi_d(\mathbf{r}_k, t) + \sum_{k=1}^N \phi_n(\mathbf{r}_{k-1}, \mathbf{r}_k, t), \quad (3.5)$$

where j is a particular scattering event, \mathbf{r}_0 is the input location, and $\mathbf{r}_N = \mathbf{r}$ is the location where the path is detected. Since ϕ_d and ϕ_n are expressed in complex analytic form, and the acoustic field is monochromatic, the resultant summation leads to a complex phase perturbation with a particular magnitude and phase offset, which oscillates harmonically at the acoustic frequency. Taking the magnitude $m_j = |\phi_j(0)|$ and phase angle $\varphi_j = \angle \phi_j(0)$ of this perturbation we may express the time varying phase increment as [22],

$$\phi_j(t) = \text{Im} [m_j \exp(j[\omega_a t - \varphi_j])]. \quad (3.6)$$

As such, we can describe E_j in equation 3.3,

$$E_j(t) = a_j \exp(j[\omega_0 t - k l_j + \phi_j(t)]) \quad (3.7)$$

where a_j is some amplitude which has been attenuated during propagation through the scattering medium, l_j is the length of path j , and ω_0 is the optical angular frequency. Inserting our expression for the insonified partial field into our expression of the correlation radiance,

equation 3.3.

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, t, \tau) = \lim_{T \rightarrow +\infty} \frac{1}{T} \int_{-\frac{T}{2}}^{\frac{T}{2}} a_j \exp(j(\omega_o t - kl_j + m \sin(\omega_a t - \varphi_j))) \times a_j \exp(-j(\omega_o(t - \tau) - kl_j + m \sin(\omega_a(t - \tau) - \varphi_j))) dt. \quad (3.8)$$

Since the system is wide-sense stationary and harmonic over an acoustic period, we may reduce the averaging to consider only a single acoustic cycle. Rewriting with $T = 2\pi/\omega_a$,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, t, \tau) = \frac{1}{T} \int_0^T a_j^2 \exp(j\omega_o \tau + jm_j \sin(\omega_a t - \varphi_j) + jm_j \sin(\varphi_j + \omega_a(t - \tau))) dt \quad (3.9)$$

We apply the Jacobi-Anger expansion [1] to remove the sine terms in the exponent,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, t, \tau) = \frac{1}{T} \int_0^T a_j^2 \exp(j\omega_o \tau) \sum_{n=-\infty}^{\infty} J_n(m_j) \exp(jn\omega_a t) \exp(-jn\varphi_j) \times \sum_{k=-\infty}^{\infty} J_k(m_j) \exp(-jk\omega_a t) \exp(jk(\omega_a \tau + \varphi_j)) dt, \quad (3.10)$$

Where J_n is the n^{th} order Bessel function of the first kind. To simplify the integrand we define,

$$A_j(n) = J_n(m_j) \exp(-jn\varphi_j) \quad (3.11)$$

and

$$B_j(k) = J_k(m_j) \exp(jk(\omega_a \tau + \varphi_j)) \quad (3.12)$$

such that,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, t, \tau) = a_j^2 \exp(j\omega_o \tau) \times \frac{1}{T} \int_0^T \sum_{n=-\infty}^{\infty} A_j(n) \exp(jn\omega_a t) \sum_{k=-\infty}^{\infty} B_j(k) \exp(-jk\omega_a t) dt \quad (3.13)$$

We expand the integrand for a term $n = n'$ and $k = k'$,

$$A_j(n') \exp(jn'\omega_a t) B_j(k') \exp(-jk'\omega_a t) = A_j(n') B_j(k') \exp(jn'\omega_a t - jk'\omega_a t) \quad (3.14)$$

and consider the integral with respect to time,

$$A_j(n') B_j(k') \frac{1}{T} \int_0^T \exp(jn'\omega_a t - jk'\omega_a t) dt \quad (3.15)$$

by orthogonality,

$$\frac{1}{T} \int_0^T \exp(jn'\omega_a t - jk'\omega_a t) dt = \begin{cases} 1 & \text{if } k' = n' \\ 0 & \text{otherwise} \end{cases} \quad (3.16)$$

such that we can rewrite equation 3.13,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, t, \tau) = \sum_{n=-\infty}^{\infty} a_j^2 \exp(j\omega_o \tau) A_j(n) B_j(n) \frac{1}{T} \int_0^T \exp(jn\omega_a t - jn\omega_a t) dt. \quad (3.17)$$

Completing the integration, and noting that the system resultant system is independent of time,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, \tau) = \sum_{n=-\infty}^{\infty} a_j^2 \exp(j\omega_o \tau) A_j(n) B_j(n), \quad (3.18)$$

reintegrating the terms $A_j(n)$ and $B_j(n)$,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, \tau) = \sum_{n=-\infty}^{\infty} a_j^2 \exp(j\omega_o \tau) J_n(m_j) \exp(-jn\varphi_j) J_n(m_j) \exp(jn(\omega_a \tau + \varphi_j)), \quad (3.19)$$

and simplifying,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, \tau) = \sum_{n=-\infty}^{\infty} a_j^2 J_n^2(m_j) \exp(j\omega_o \tau) \exp(jn\omega_a \tau). \quad (3.20)$$

We can now again apply the identity $J_{-n}(z) = (-1)^n J_n(z)$ [1], after expanding the summation for $n = 0$,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, \tau) = a_j^2 \exp(j\omega_o \tau) \times \left(J_0^2(m_j) + \sum_{n=1}^{\infty} J_n^2(m_j) \exp(jn\omega_a \tau) + \sum_{n=1}^{\infty} J_n^2(m_j) \exp(-jn\omega_a \tau) \right). \quad (3.21)$$

Now, since $\exp(jx) + \exp(-jx) = 2 \cos(x)$,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, \tau) = a_j^2 \exp(j\omega_o \tau) \left(J_0^2(m_j) + 2 \sum_{n=1}^{\infty} J_n^2(m_j) \cos(n\omega_a \tau) \right). \quad (3.22)$$

In this form, we can see that the correlation radiance contains spectral content at the input optical frequency, with side-bands at integer multiples of $\pm\omega_a$. The phase of each of the element is lost in the autocorrelation function such that the final summation over j sums the power of each field component. In practice we typically neglect the component at the optical frequency since it is not recovered in the transformation from the experimentally recorded intensity autocorrelation (as discussed in chapter 2). Thus,

$$G_j(\mathbf{r}, \hat{\mathbf{s}}, \tau) = a_j^2 \left[J_0^2(m_j) + 2 \sum_{n=1}^{\infty} J_n^2(m_j) \cos(n\omega_a \tau) \right]. \quad (3.23)$$

There are therefore three tasks to be performed in the proposed MC simulation. First, we must statistically generate the random paths, j , along which we will propagate our ‘photon packets’. Second, we must attenuate this photon packets according to the Beer-Lamber law to determine their associated intensity a^2 . Third, we must accumulate the phase increments

which constitute m_j according to 3.5. In this formulation, the first two tasks are identical to the MC in the standard RTE. The accumulation of phase increments precludes the possibility of a ‘white’ MC approach, since the calculations require knowledge of the complete photon-path, rather than the partial path lengths in each region of differing absorption coefficient.

3.2.2 An approximate adjoint method

MC models of light transport are unable to directly simulate point-source, point-detector (PS/PD) regimes which commonly occur in studies involving the use of coherent light, since the random walk which begins at the source location is highly unlikely to impinge directly upon the detection point. In some cases it may be sufficient to integrate the detected flux over a region surrounding the detector location, assuming the detector accepts diffuse light incident from any direction. An alternative approach is to employ the reciprocity of light transport to define the *adjoint* transport process [86].

We will consider the adjoint method with greater rigour in chapter 5 when we employ this procedure to calculate the sensitivity functions required to perform image reconstruction. For now, we consider only the practicalities of an adjoint based approach to performing PS/PD simulations. To begin, we consider the standard optical case. To determine the exiting flux across a boundary for a given detector configuration, L_d^Γ , we perform the following procedure.

1. The MC forward model would be executed to give the radiance distribution throughout the domain for the actual source location. Once normalised we might think of this as a radiance probability density function which describes the probability of the transport of light from the source location to a given point in the domain.
2. The MC forward model would be executed again with the source term replaced by the *adjoint source*. This source has the location and acceptance profile of the *detector* in the simulated experiment, and its intensity is scaled by the boundary conditions of the medium. Once again, following normalisation, we can (given the reciprocity of the problem) consider the resultant adjoint radiance distribution to represent the probability density of radiance at any point in the domain being detected by the actual detector.
3. The product of the two probability functions gives the probability density of radiance generated by the source reaching a given position in the domain, with a particular direction, reaching the detector. By integrating the result over the domain, and over all solid angle, we arrive at the total (normalised) flux which would be detected in the actual experiment (this is a form of the Chapman-Kolmogorov equation[122]). Neglecting the normalisation we may write this,

$$L_d^\Gamma = \int_{4\pi} \int_{\Omega} L(\mathbf{r}, \hat{\mathbf{s}}) L^*(\mathbf{r}, \hat{\mathbf{s}}) d\mathbf{r} d\hat{\mathbf{s}} \quad (3.24)$$

where $L(\mathbf{r}, \hat{\mathbf{s}})$ is the forward radiance due to the source, $L^*(\mathbf{r}, \hat{\mathbf{s}})$ is the adjoint radiance due to the adjoint detector, and integration is performed over the entire domain (Ω), and over all solid angle (4π).

To perform this procedure requires that $L(\mathbf{r}, \hat{\mathbf{s}})$ be recorded throughout the domain. However, typical MC codes, including that under development in this work, usually record the fluence $\Phi(\mathbf{r}) = \int_{4\pi} L(\mathbf{r}, \hat{\mathbf{s}}) d\hat{\mathbf{s}}$ throughout the domain. In this work we will make the following approximation,

$$\int_{4\pi} \int_{\Omega} L(\mathbf{r}, \hat{\mathbf{s}}) L^*(\mathbf{r}, \hat{\mathbf{s}}) d\mathbf{r} d\hat{\mathbf{s}} \approx \int_{\Omega} \left[\int_{4\pi} L(\mathbf{r}, \hat{\mathbf{s}}) d\hat{\mathbf{s}} \right] \left[\int_{4\pi} L^*(\mathbf{r}, \hat{\mathbf{s}}) d\hat{\mathbf{s}} \right] d\mathbf{r}. \quad (3.25)$$

This expression would become an equality in the case that radiance is perfectly isotropic. The error in this approximation is worst near sources and boundaries where the radiance distribution becomes more directionally dependent. Application of this approximation reduces the accuracy of our solution in a similar manner to that which is achieved with a diffusion style model. However, there are certain scenarios which cannot be simulated by diffusion style approaches. We will return to consider this more fully in chapters 4 and 5.

Until this point we have been describing the transport of radiance, rather than correlation. We now expand the idea to the case of AO where energy will be distributed not only at the input optical frequency, but at side-bands shifted by multiples of the acoustic frequency. Consider a pencil beam incident normal to the boundary of the simulation domain at some point \mathbf{r}_s illuminating the system with a given irradiance, $L_s^\Gamma(\mathbf{r}_s, \hat{\mathbf{s}})$. We denote the probability of optical power at some point \mathbf{r}_a in the k^{th} side-band reaching some point \mathbf{r}_b in the n^{th} side-band as $P_k^n(\mathbf{r}_b, \mathbf{r}_a)$. The probability distribution of optical power at each side-band throughout the domain resulting from the source at \mathbf{r}_s is given by the unit-less normalised fluence rate $P_0^n(\mathbf{r}, \mathbf{r}_s) = \Phi^n(\mathbf{r}) / L_s^\Gamma(\mathbf{r}_s, \hat{\mathbf{s}})$, where $\Phi^n(\mathbf{r})$ is the fluence rate at a point \mathbf{r} in the medium in the n^{th} side-band. In replacing our detector with an *adjoint* source of equivalent directivity at \mathbf{r}_d , we may derive another set of optical power probability distributions $P_0^n(\mathbf{r}, \mathbf{r}_d)$. By our approximation,

$$P(\mathbf{r}_d, \mathbf{r}_s) \approx \int_{\Omega} P(\mathbf{r}_d, \mathbf{r}) P(\mathbf{r}, \mathbf{r}_s) d\mathbf{r}. \quad (3.26)$$

In the case of transitions to and from different side-bands we must sum over the relevant permutations:

$$P_0^n(\mathbf{r}_d, \mathbf{r}_s) \approx \int_{\Omega} \sum_{k=0}^{\infty} P_k^n(\mathbf{r}_d, \mathbf{r}) P_0^k(\mathbf{r}, \mathbf{r}_s) d\mathbf{r}. \quad (3.27)$$

We assume that contributions to the $(n+1)^{\text{th}}$ side-band are negligible and thus express the detected exitance at the zeroth side-band,

$$L_{d,0}^\Gamma(\mathbf{r}_d) \approx L_s^\Gamma(\mathbf{r}_s) \int_{\Omega} P_0^0(\mathbf{r}_d, \mathbf{r}) P_0^0(\mathbf{r}, \mathbf{r}_s) d\mathbf{r}, \quad (3.28)$$

where by reciprocity $P_0^0(\mathbf{r}_d, \mathbf{r}) = P_0^0(\mathbf{r}, \mathbf{r}_d)$. The detected exitance at the first side-band is given by,

$$L_{d,1}^\Gamma(\mathbf{r}_d) \approx L_s^\Gamma(\mathbf{r}_s) \int_{\Omega} P_0^1(\mathbf{r}_d, \mathbf{r}) P_0^0(\mathbf{r}, \mathbf{r}_s) + P_1^1(\mathbf{r}_d, \mathbf{r}) P_0^1(\mathbf{r}, \mathbf{r}_s) d\mathbf{r}. \quad (3.29)$$

We find $P_0^1(\mathbf{r}_d, \mathbf{r})$ by optical reciprocity, and make the approximation $P_1^1(\mathbf{r}_d, \mathbf{r}) \approx P_0^0(\mathbf{r}_d, \mathbf{r})$. Thus, the first harmonic modulation depth, MD_1^Γ detected at a point on the boundary is given by,

$$MD_1^\Gamma \approx \frac{\int_{\Omega} P_0^1(\mathbf{r}_d, \mathbf{r}) P_0^0(\mathbf{r}, \mathbf{r}_s) + P_0^0(\mathbf{r}_d, \mathbf{r}) P_0^1(\mathbf{r}, \mathbf{r}_s) d\mathbf{r}}{\int_{\Omega} P_0^0(\mathbf{r}_d, \mathbf{r}) P_0^0(\mathbf{r}, \mathbf{r}_s) d\mathbf{r}}. \quad (3.30)$$

Given the approximations made in developing this expression, it is important that any simulation which employs this approach is demonstrated to tend to the correct result, as may be found by integrating the output over a larger region. We will demonstrate the validity and use of the expression in section 3.4.

3.3 Implementation

3.3.1 GPU Parallel Programming

The design of an efficient GPU accelerated parallel algorithm such as the current simulation requires a different approach to traditional serial or parallel CPU algorithms. In the latter case the efficiency of an algorithm is primarily a function of the number of instructions to be performed for each execution of the algorithm; efficiency is increased by pre-processing data and storing this in memory for retrieval by the algorithm. On a GPU the converse is true, memory access is expensive and computation is cheap, as such the ratio of arithmetic to memory access should be maximised - even to the extent that it may be preferable to re-evaluate intermediate values of a large calculation rather than store them.

The second crucial aspect of writing efficient GPU accelerated code is preventing the divergence of multiple threads due to control flow instructions. This is epitomised by the *single instruction, multiple thread* (SIMT) paradigm. Each of the many streaming multiprocessors (SMs) contained within a GPU contains a single instruction scheduler which passes the required computational instruction to multiple scalar processors (SPs) which each process a single thread. If, through a control flow instruction, any of the SPs within a single SM are executing different instruction to another, execution of those threads residing on that SM must be executed in serial until the threads re-converge.

Since each simulated photon packet has no dependence on the state of another, such algorithms are inherently parallelisable. The large datasets required to store an acoustic field provide challenges for a GPU-based system, as does inherent divergence caused by only a subset of photons currently under simulation transitioning between regions of particular optical properties.

The present simulation is to be implemented in the nVidia CUDA architecture. Whilst tied to a specific platform, the tool-chain provided by nVidia is (at the time of writing) considerably more advanced than other offerings, such as OpenCL. An extensive discussion of best practices when programming in CUDA is provided by nVidia. Our implementation follows three broad strategies.

1. Exposing the parallelism of the algorithm: the problem at hand is inherently parallel as each photon packet is propagated independently of every other. In our algorithm cases will occur when individual photon packets require different treatments. For example some may be reflected or refracted at an optical boundary, whereas others may be incident upon a scattering object. It is important to minimise this divergence since each execution branch will be serialised by the hardware.
2. Use of special purpose memory: whilst such devices are capable of a very high arithmetic instruction throughput, this comes at the expense of the caching of memory. Various special purpose memory spaces exist on the device which are optimised to store particular types of data.
3. Use of *intrinsic* arithmetic functions: CUDA exposes a number of functions which map directly to the hardware level; such instructions are considerably faster than the standard C library implementations, but often have reduced accuracy. Judicious use of these functions (particularly the trigonometrical and transcendental replacements) can considerably increase performance.

3.3.2 Architecture

The simulation framework is built around an executable program (the simulation program) which is configured by variables contained within a set of MATLAB data files: this is depicted in figure 3.1. The requisite datasets can each be built manually according to the specifications provided in section 3.3.3, though a number of scripts have been developed by the author to perform basic tasks such as generating simple acoustic fields and simulation geometries. The simulation program itself consists of two primary components.

1. The core photon-propagation and phase accrual algorithm (the kernel) is written in the C language using the nVidia CUDA language extensions to permit execution on a suitable GPU. The kernel and its support functions are wrapped in C++ functions to provide an external interface (the kernel wrapper). The kernel and its wrapper must be compiled using the nVidia supplied nvcc compiler and link against a number of nVidia libraries.
2. Data input and output, kernel configuration and management is performed by a piece of code written in the C++ language (the driver). The driver links against the kernel

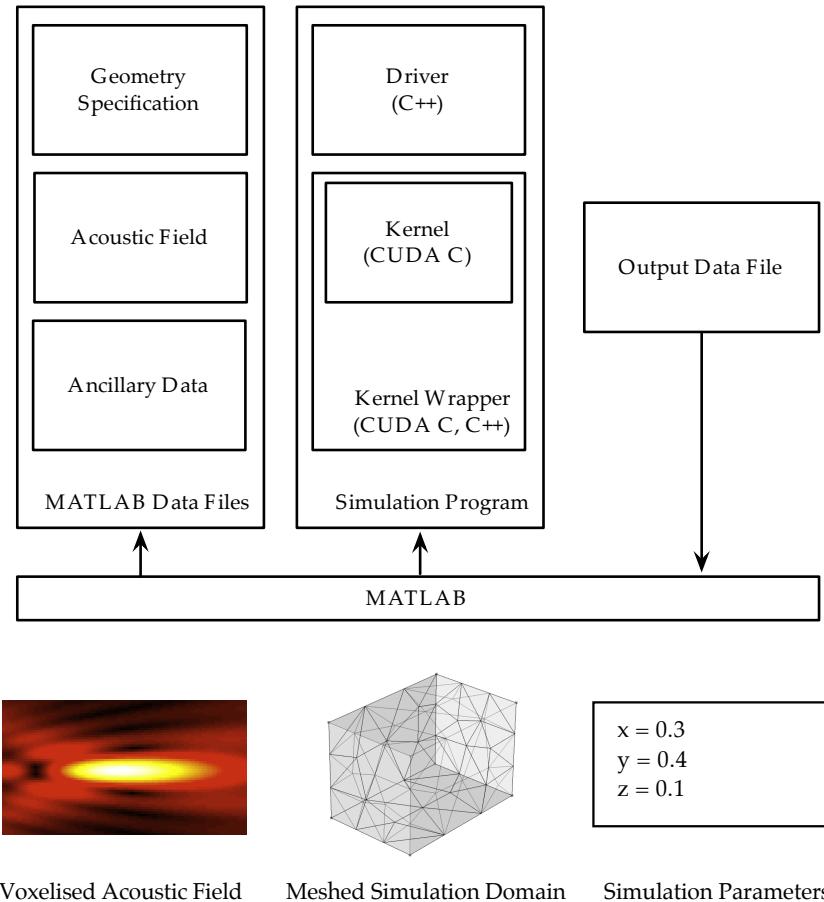


Figure 3.1: Architecture of the simulation framework demonstrating the input datasets, organisation of the simulation program code, and relationship to the MATLAB environment

wrapper, in addition to libraries provided by Mathworks which permit the reading of MATLAB data files. The driver can be compiled by a standard C++ compiler.

Upon execution the simulation kernel is initialised by the simulation driver. After all required data has been uploaded to the GPU, a pointer to a memory allocation for photon state information is passed to the kernel, and from this each parallel thread can record the photons packets which are detected. Once called, each thread of the kernel will initialise a photon and propagate it until it either impinges upon a detector surface, is completely absorbed by the medium, or leaves the simulation domain. In the first case the photon state is written to the output buffer and a new photon is launched, in the latter cases the photon is terminated and a new photon is launched. If the required number of photons for the current

run has been launched or recorded, the thread exits and control returns to the driver. The driver then writes the data to disc before re-executing the simulation kernel repeatedly, until such time that the desired number of photons have been simulated. This process is depicted in figure 3.2

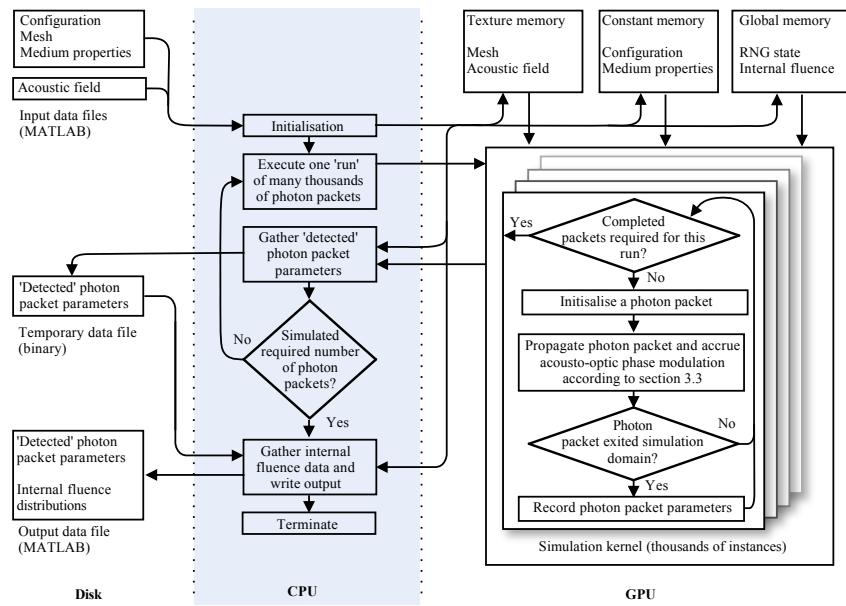


Figure 3.2: Architecture of the simulation program illustrating the parallel execution of many thousands of individual propagation kernels, each responsible for a single photon packet.

3.3.3 The Simulation Domain

The simulation domain consists of a description of the geometry of the system, the optical properties of the different regions of the geometry of the system, the applied acoustic field and a number of physical constants. A representation of these physical systems must be developed which considers the importance of choosing the correct memory architecture for a GPU accelerated application.

Acoustic Field

Of the three sets of data required by the simulation program, the largest is the acoustic dataset. As described in section 3.2, simulation of the acousto-optic modulation process requires knowledge of the following acoustic parameters at each point in the medium,

A_0 the local acoustic displacement amplitude,

$\hat{\mathbf{k}}_a$ the local acoustic propagation direction, and,

ϕ , the local acoustic phase offset, relative to an arbitrary point,

in addition to various constants such as the acoustic wave-number k_a and material dependent properties such as the speed of sound in the medium, v_a .

If it assumed that a single precision number is required for each variable (three for the propagation direction) then 20 bytes of storage are required for each point in the acoustic field. Storing six points per acoustic wavelength (to provide a good approximation to the field shape, in addition the prevention of aliasing), assuming the speed of sound in water $v_a = 1480 \text{ m s}^{-1}$ and a typical ultrasound frequency $f_a = 2 \text{ MHz}$, 42875 elements per cubic millimetre must be stored, occupying 832kB of memory. Given the aim of this work is to consider the use of acousto-optics and its ability to non-invasively target physiologically relevant targets, it is likely that domains of several cubic centimetres will be required; the memory requirements for such fields will quickly become prohibitively large.

Previous CPU based simulations [118] also represent the acoustic data in this voxelised manner. In this work the spacing of grid elements was optimised to provide greater resolution in the axis of the focus of the ultrasound field, and trading this against memory usage off-axis. Even if enough memory is available on the device to store the field, we must still consider that memory access on the device is relatively expensive, more so than in the CPU case. One might suggest calculating the acoustic propagation direction $\hat{\mathbf{k}}_a$ at runtime by taking the derivative of the phase of the local acoustic field $\nabla\phi(\mathbf{r})$, however a numerical differentiation will require sampling the phase at multiple points surrounding that of interest, incurring the significant overhead of device memory access.

The GPU architecture offers a convenient solution: texture memory. Texture memory resides in standard device memory; the process of binding an array to a texture provides an

alternate access route to the specified region of device memory which conveys numerous advantages.

- Texture memory is spatially cached such that reads from the same or similar locations will be fulfilled from the cache, rather than by re-accessing device memory, this improves access time for other texture reads which must retrieve the data from device memory.
- A single read from texture memory can return a compound variable of up to 128 bytes (four single precision floating point numbers or integers).
- Texture memory supports various addressing modes which allow the acoustic field automatically wrap around itself or clamp to the nearest value; this can be selected individually in each dimension allowing the efficient representation of acoustic fields with linear symmetry.
- The texture hardware can perform tri-linear interpolation upon the reading of a specific point. This allows precision to be maintained as the grid is made more coarse.

The texture caching procedure is an opaque feature of the hardware, and this can cause difficulty in optimisation. The cache miss rate can, however, be determined using a profiling tool provided by nVidia.

If the spatially varying acoustic data can be represented in only four floating point numbers, a single texture read could be used to retrieve the required acoustic data at a specific point in the medium with the added advantage of automatic tri-linear interpolation. To achieve this, we scale the acoustic propagation direction (a unit vector) with the local acoustic displacement amplitude to form a compound variable $A_0 \hat{\mathbf{k}}_a$ whose magnitude can be taken at runtime, and the propagation vector restored by normalisation.

Under the assumption that the acoustic field does not exceed the available memory on the device, we now have a relatively efficient way of storing and retrieving the data at runtime. The simulation driver accepts the acoustic data as four three-dimensional MATLAB arrays, indexed according to their spatial location in a Cartesian manner. The first is the local phase offset, and the remaining three are the x , y and z components of the amplitude scaled propagation direction. Noting MATLAB's one-based indexing, the array elements located at [1] [1] [1] represents the acoustic field within the simulation domain located at the point $\mathbf{r} = (0, 0, 0)$. To map the entries of the acoustic field arrays to the simulation domain the complete extent of the acoustic field is recorded in a separate variable.

Upon reading the MATLAB arrays from the data file the simulation driver creates a compound array of four-element floating point numbers which store each of the relevant variables for every point in the three-dimensional array. During this process, the input arrays are permuted from their column-major format (as used by MATLAB) to the C-style row-major format. This permutation is performed such that addressing of the array within the kernel remains intuitive. The simulation driver uploads the permuted compound array to the device's

memory and binds a texture to the appropriate location for later access by the kernel. During simulation the kernel indexes the acoustic field texture using a normalised access method; a texture request is made by submitting a three-element floating point number with each element in the range $[0, 1]$.

3.3.4 Geometry Representation

In many Monte-Carlo simulation programs which support heterogeneous optical properties, including GPU accelerated programs [41] and those which consider acousto-optics [118], the heterogeneity is provided by discretising the domain in the same manner as has been applied to form the acoustic field representation — a voxelised grid (a piecewise constant basis). As a photon packet takes a step through multiple volume elements (voxels), each must be examined to ascertain if there has been a change in optical properties. The approach has the benefit of being conceptually straightforward. As discussed briefly in section 3.3.1 this representation is sub-optimal for a GPU implementation due to the overhead in accessing different memory locations. Furthermore, a photon may move through a variable number of voxels on each step, this implies the use of a `for` or `while` loop in the kernel code which propagates photons: this introduces divergent execution paths and thus serialisation of execution. A different approach is taken here whereby the simulation domain is represented by a tetrahedral mesh. This representation is typically encountered in applications of, e.g., the finite element method, but has also been applied to a CPU based Monte-Carlo photon propagation [40]. By employing this method, photons move around within volume elements which subdivide the simulation geometry. On each photon step, we must check for intersection with any of the four boundaries (facets) of the current element. Whilst we must still interrogate a dataset to determine the location of these facets, the amount of data to be retrieved, and the nature of the processing to be performed, is now fixed across each thread, thus avoiding issues of divergence and potentially reducing the amount of data which must be retrieved from memory¹.

In addition to being a more suitable description of the simulation geometry for technical reasons, a meshed approach allows the boundaries between two optically dissimilar regions to be represented with arbitrary accuracy, rather than as a stepped voxelisation. Some of the advantages of this approach are discussed in [90]. Two tables of data are required to define the geometry of the simulation domain, though two more are required for this particular application. The tables are now considered by their internal names.

nodes The node list records the three-dimensional Cartesian co-ordinate of every node, or vertex, in the mesh. These nodes are used by the element matrix to form the individual tetrahedral elements from which the geometry is constructed.

¹This is of course dependent upon the complexity of the models; homogeneous cuboid geometries would be far more efficiently represented as voxelised structures, however the research aims of this study require the flexibility to consider more physiologically relevant, complex structures.

elements The element table defines each tetrahedral element by four numbers which define the co-ordinates of each vertex of the element by indexing the node list.

neighbours The neighbour list is a requirement of this application. As photons are propagated through the simulation domain they will pass from one element to the next via one of the boundaries (facets) of each tetrahedron. To determine into which element they may transition, we access the element neighbour list with the current element index; this returns the indices of the four elements neighbouring the current element (or zero for the domain boundary). By convention element opposite node j of element i is located in the neighbour list at index $[i, j]$.

properties The properties list is a requirement of this application. Each element in the element table has an entry in the properties list which defines an index to the medium of which the element is composed.

The primary mesh tables (elements and nodes lists) may be generated by various means. For the most simple geometries such as plane slabs, the nodes and element lists can be built by hand. More complex optically homogeneous geometries can be built using triangulation tools such as the MATLAB `delaunay` function. Open source tools such as `iso2mesh` [39] can build meshes from voxelised data, such as MRI scans (potentially pre-segmented to delineate different optical regions). Alternatively, a number of finite element method tools have integrated geometry definition and meshing tools from which a suitable mesh can be exported. In the case that the acoustic field for simulation is derived from a finite-element analysis, a mesh will be readily available. The neighbours table can be generated using the MATLAB `neighbors` function. Note that this function returns NaN entries for external facets, these must be replaced with zeroes for use in the simulation program. The properties list can be generated by hand. It is likely that this list will already exist if the dataset was derived from a finite-element analysis or from pre-segmented data.

Much like the acoustic datasets, each of these tables is stored in the device as a texture, though in this case a one-dimensional texture accessed by conventional array indexing, rather than the normalised case employed for the three-dimensional acoustics texture. The benefits of texture memory discussed previously are also conveyed to the access of the geometry tables.

The simulation driver reads the dataset from the MATLAB data files, transfers these arrays to the device and binds them to a relevant texture. Just prior to their transfer to the device, each of the arrays is padded with an empty row at their beginning; this is to permit element zero to represent that which is external to the specified simulation domain. If the datasets have been constructed in MATLAB's one-based indexing regime, this requires no attention from the user. Otherwise, it may be necessary to increment each entry in the element and neighbours tables by one.

3.3.5 Ancillary Data

In addition to the acoustic field and geometry data, three more sets of data are required to complete a simulation. These comprise of the medium specifications (primarily the optical properties) as indexed by the geometry tables, simulation physical constants (such as the ultrasound frequency, optical wavelength) and simulation control constants.

Different mediums are specified in a data structure containing the following constants, all of which are specified at the simulated optical wavelength.

mua is the absorption coefficient, μ_a [m^{-1}].

mus is the scattering coefficient, μ_s [m^{-1}].

g is the anisotropy factor, g , of the medium [-].

n0 is the index of refraction of the medium [-].

eta is related to the acousto-optic coefficient, and was defined in equation 2.10. [m^{-1}].

The simulation driver retrieves the medium entries from a MATLAB array named **medium** in the supplied data files. An entry must be provided for every medium referenced in the element properties table according to MATLAB's one based-indexing, the simulation driver automatically pads the array with a defunct zero entry. This dataset is uploaded to the constant memory area of the graphics card which provides cacheing for such read only data.

The simulation program requires the following constants to be presented.

k0 is the wavenumber of the optical source, k_0 in units of rad m^{-1} .

ka is the acoustic wavenumber, k_a in units of rad m^{-1} .

init_photon_location is a three-component variable specifying the initial location, in units of the meter, from which photon packets are launched into the simulation domain.

init_photon_direction is a three-component unit magnitude variable specifying the initial direction which the photon packet takes when launched into the simulation domain.

init_photon_weight is the notional weight of a photon packet and can be set to an arbitrary value, typically unity.

init_element is an index to the element into which the photon is launched into the simulation domain. This should be determined from the initial photon packet location either by hand, or by use of a function such as MATLAB's **tsearchn** function.

`init_face` is the facet from which the photon packet is launched into this simulation domain. Setting this variable disables an intersection test on this face during the first run of the photon packet propagation code. If the photon packet is injected at a point within a domain element, this variable can be set greater than three to disable the test.

`n0_external` is the refractive index external to the simulation domain.

`detector_index` is the medium index which represents the detector. Photon packets entering an element of the specified medium will be detected and terminated as they pass into this region.

The simulation driver retrieves the simulation constant variables from a single element MATLAB array consisting of a structure with the documented field names. This dataset is also uploaded to the constant memory area of the device. Finally, the simulation program requires a number of simulation constants which dictate how execution continues. These include the number of photons to simulate, and the output filename.

3.3.6 Algorithm

Given a representation of the simulation domain the algorithms which propagate photon packets through the simulation domain, accrue their phase, and control the execution of the simulation can now be described.

Many of the fundamental calculations performed are of a standard form for MC light transport codes. Whilst there are significant changes to the ordering of the algorithm, and the way in which the geometry is represented, Wang's description [134] of his *de-facto* standard code, MCML serves as a reference for each part of the photon propagation routine.

As described in section 3.3.2, following the upload of required data from the host to the device, the simulation driver repeatedly calls the simulation kernel to request a certain number of photon packets to be propagated². On each call, tens of thousands of individual threads perform the process described at a high-level in algorithm 1; the important processes are subsequently described in greater depth.

Photon Packet Step Size

With knowledge of the interaction coefficient of the current medium, $\mu_t = \mu_a + \mu_s$, the photon step size is calculated,

$$s = \frac{-\log(1 - \xi)}{\mu_t} \quad (3.31)$$

where ξ is a uniformly distributed random number distributed over the interval $[0, 1)$.

²The entire simulation task could be completed in a single call to the kernel, however this would require a memory buffer large enough to store the results of every photon packet to be recorded. Instead, a smaller buffer is used, and this is written out to disk by the simulation driver between each kernel invocation.

Algorithm 1: The Simulation Kernel Algorithm

Data: photons requested

```

1  get photon index from global photon counter;
2  while photon index < photons requested do
3      initialise photon packet to initial state;
4      repeat
5          sample the photon packet step size;
6          if photon packet intersects a facet then
7              set step size to take photon packet to boundary;
8              if photon reflects then
9                  set next direction according to reflection;
10             else
11                 set next direction according to refraction;
12                 update element and face;
13                 if element is a detector then
14                     flag detection;
15                 end
16             end
17         else
18             set next direction by scattering;
19             absorb photon packet weight into medium ;
20         end
21         accumulate phase due to displaced scatterer ;
22         accumulate phase due to modulated index of refraction ;
23         update location according to step size;
24         update direction according to next direction;
25         perform photon roulette (3.3.7);
26     until the photon packet has zero weight, detection is flagged, or the photon packet
27     has exited the simulation domain;
28     if detected then
29         write photon location, weight and phase to buffer;
30         increment global photon counter;
31     end
32 end

```

Geometry Intersection

Given the current location, direction and step size of a photon packet, we must now test for interaction with the simulation geometry, e.g., intersection of any of the facets of the current element.

By retrieving the four nodes of the current element, we construct each of the four facets of the tetrahedron by iterating over choices of the three nodes $(2, 3, 4)$, $(1, 3, 4)$, $(1, 2, 3)$, and $(1, 2, 3)$. We test for intersection with each facet, ignore the result of any test corresponding to the face over which the photon packet previous travelled.

The intersection of a line-segment with a triangle is a well understood geometric operation. The most basic form of this test involves testing for intersection of the line with the plane defined by the three points of the triangle, before determining if the intersection point occurs within length of the current segment. Alternative accelerated algorithms are available in the literature pertaining to computer graphics, one of the more popular being that of Möller and Trumbore[100]. However, when implemented without infinite precision, such tests require epsilon based algorithms which test results of operations not against zero, but against a small range around zero. As such, the algorithm is no longer robust. This can often be tolerated in graphics rendering application, however in this application, the consequences of false-positives or negatives is that a photon will end an iteration of the algorithm at a position which does not correspond to its current element. All subsequent intersection tests will fail, and the photon will propagate indefinitely.

To avoid this problem, we use an alternative approach based upon the triple scalar product [63]. Here we rely only upon the sign, rather than the value, of various arithmetic operations. This approach is mathematically equivalent to the use of Plücker co-ordinates, (as described by Fang [40]), but avoids the associated complexity of analysis and notation.

The signed volume of a tetrahedron defined by four vertices \mathbf{p}_0 – \mathbf{p}_3 is calculated by[63],

$$V(\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3) = \frac{1}{6} [(\mathbf{p}_1 - \mathbf{p}_0) \times (\mathbf{p}_2 - \mathbf{p}_0) \cdot (\mathbf{p}_3 - \mathbf{p}_0)] \quad (3.32)$$

where \cdot is a dot-product and \times the cross product operator. To determine if the ray defined by the current photon packet direction intersects the current tetrahedral element we begin by retrieving the four vertices of the element from the geometry cache. We test for intersection with each of the four facets of the tetrahedron by iterating over the four choices of the three nodes. We ignore the result of a test against the facet, if any, over which the photon packet previously reflected or refracted.

For each of the four facets we calculate three signed volumes,

1. $V(\mathbf{r}_a, \mathbf{r}_b, \mathbf{p}_1, \mathbf{p}_2)$,
2. $V(\mathbf{r}_a, \mathbf{r}_b, \mathbf{p}_2, \mathbf{p}_0)$,
3. $V(\mathbf{r}_a, \mathbf{r}_b, \mathbf{p}_0, \mathbf{p}_1)$.

In the case that the sign of each of these volumes is equal, the ray intersects the facet. To determine if the photon packet will intersect the facet given its current step size we calculate the two signed volumes $V_a = V(\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{r}_a)$ and $V_b = V(\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2, \mathbf{r}_b)$; in the case that the signs of these volumes are opposite, and that $V_a > 0$, intersection occurs at the parametric distance $V_a/(V_a + V_b)$. If any interactions occur the smallest intersection distance is used to scale the current photon packet step size.

If an intersection occurs, the photon packet step size is updated to take it to the boundary of the element. We then perform a series of tests to determine if the photon packet reflects, refracts or simply transmits into the next medium.

1. If the index of refraction on either side of the boundary is matched, the photon packet transmits into the next element with the same propagation direction.
2. Otherwise, we calculate the angle of incidence of the photon packet α_i using the normal to the plane defined by the three points of the face we have intersected. With knowledge of the index of refraction in the incident medium n_i and the medium into which we may transmit, n_t , we next check for total internal reflection which occurs when $n_t^2 < n_i^2(1 - \cos^2(\alpha_i))$. In this case, we reflect the ray in the plane defined by the face,

$$\hat{\mathbf{s}} \leftarrow \hat{\mathbf{s}} - 2\hat{\mathbf{n}} \cos(\alpha_i) \quad (3.33)$$

where $\hat{\mathbf{n}}$ is the unit normal to the face.

3. If total internal reflection does not occur, we must calculate the internal reflectance function,

$$R(\alpha_i) = \frac{1}{2} \left[\frac{\sin^2(\alpha_i - \alpha_t)}{\sin^2(\alpha_i + \alpha_t)} + \frac{\tan^2(\alpha_i - \alpha_t)}{\tan^2(\alpha_i + \alpha_t)} \right] \quad (3.34)$$

where the transmission angle α_t is found by Snell's law ($n_i \sin(\alpha_i) = n_t \sin(\alpha_t)$). We now generate a uniformly distributed random ξ in $[0, 1]$. If $\xi \leq R(\alpha_i)$ then the photon packet is reflected back into the current element, and its next direction is calculated according to the equation for total internal reflection. Otherwise, if $\xi > R(\alpha_i)$, the ray refracts into the next medium and its new direction is calculated,

$$\hat{\mathbf{s}} \leftarrow \frac{n_i}{n_t} \hat{\mathbf{s}} - \mathbf{n} \left(\frac{n_i}{n_t} \cos(\alpha_i) + \sqrt{1 - \sin^2(\alpha_t)} \right) \quad (3.35)$$

Scattering

If the photon packet does not intersect a facet of the current tetrahedron, it is considered to have interacted with the medium. In this case the photon packet first undergoes scattering. In this application the photon packet is scattered by the Henyey-Greenstein phase function

[54]. Expressed as a probability density function,

$$p(\cos(\theta)) = \frac{1 - g^2}{2(1 + g^2 - 2g \cos(\theta))^{3/2}} \quad (3.36)$$

For anisotropic scattering where $g \neq 0$ the resultant deflection cosine, $\cos(\theta)$ is found by [134],

$$\cos(\theta)|_{g \neq 0} = \frac{1}{2g} \left[1 + g^2 - \left(\frac{1 - g^2}{1 - g + 2g\xi} \right)^2 \right] \quad (3.37)$$

For isotropic scattering with $g = 0$,

$$\cos(\theta)|_{g=0} = 2\xi - 1 \quad (3.38)$$

where ξ is a uniformly distributed random number distributed over the interval $(0, 1)$. The uniformly distributed azimuthal angle is then found.

$$\psi = 2\pi\xi \quad (3.39)$$

The direction cosines of the next direction of the photon packet are updated according to the sampled values of the scattering deflection and azimuth.

3.3.7 Absorption and Roulette

Following a scattering event, some proportion of the photon packet weight is absorbed into the medium. The photon packet weight is updated [134],

$$w \leftarrow w - w \frac{\mu_a}{\mu_t} \quad (3.40)$$

where w is the photon packet weight.

In order to reduce the number of photons packets with a very low weight (and thus little statistical significance) the variance reduction technique of *roulette* is employed. Photon packets with a weight less than a certain threshold are assigned a probability m of surviving the roulette procedure. A random number ξ is generated in $[0, 1]$ and if $(\xi > 1/m)$ the photon packet weight is reduced to zero, for termination on the next iteration. In order to conserve energy, a surviving photon packet has its weight updated $w \leftarrow mw$.

Index of refraction phase accrual

The simulation offers two methods to calculate the index of refraction phase increment based upon equation 2.6.

In the first formulation we approximate the equation with a numerical integration over the path of the current photon packet step. We divide the photon packet path length into N subintervals of equal widths $h = |\mathbf{r}_b - \mathbf{r}_a|/N$ and employ the composite midpoint rule,

$$\phi_n(\mathbf{r}_a, \mathbf{r}_b, t) \approx k_0 n_0 \eta h \sum_{k=0}^{N-1} P_0(\mathbf{r}_k) \exp [j(\phi_a(\mathbf{r}_k) - \omega_a t)] \quad (3.41)$$

where $\mathbf{r}_k = \mathbf{r}_a + (k + 0.5)(\mathbf{r}_b - \mathbf{r}_a)/n$. The number of subintervals is defined at runtime as a particular number of points per acoustic wavelength and thus differs for each photon packet under propagation; the divergence of this approach leads to sub-optimal performance. Furthermore, a large number of reads from the acoustic-dataset are required to minimise the error in the numerical integral. This approach is still viable owing to the performance benefits of the parallel implementation and provides excellent accuracy irrespective of the nature of the underlying acoustic field.

An alternative approach to the calculation of the index of refraction phase increments is to assume that over the length of single photon packet step the acoustic field can be represented by a plane wave of given pressure amplitude, propagation direction and phase offset as defined at a point \mathbf{r}_c ,

$$p_c(\mathbf{r}, t) = P_0(\mathbf{r}_c) \exp(j[\phi_a(\mathbf{r}_c) + \mathbf{k}_a(\mathbf{r}_c) \cdot (\mathbf{r} - \mathbf{r}_c) - \omega_a t]). \quad (3.42)$$

This expression is substituted into equation 2.6 and the integration completed over the straight line path of the photon packet step from \mathbf{r}_a to \mathbf{r}_b ,

$$\begin{aligned} \phi_n(\mathbf{r}_a, \mathbf{r}_b, t) \approx & \frac{k_0 n_0 \eta P_0(\mathbf{r}_c) \exp(j(\phi_a(\mathbf{r}_c) - \mathbf{k}_a(\mathbf{r}_c) \cdot \mathbf{r}_c - \omega_a t))}{j \mathbf{k}_a(\mathbf{r}_c) \cdot (\mathbf{r}_b - \mathbf{r}_a)} \\ & \times [\exp(j(\mathbf{k}_a(\mathbf{r}_c) \cdot \mathbf{r}_b)) - \exp(j(\mathbf{k}_a(\mathbf{r}_c) \cdot \mathbf{r}_a))]. \end{aligned} \quad (3.43)$$

In the case that the acoustic wavevector is a slowly changing function of space, this approximation can offer significant performance improvements over the numerical integration when the photon packet step size is large compared to an acoustic wavelength. This is achieved due to reduced memory access, and its non-divergent execution.

Irrespective of the integration method, the real and imaginary parts of each expression are evaluated and stored individually by expansion of the complex exponential.

3.3.8 Random Number Generator

On numerous occasions within the photon propagation algorithm, a uniformly distributed random number must be generated. As a parallel algorithm, this requires special treatment. A standard random number generator (RNG) could be used for each thread (photon propagation kernel), though if it were seeded with the same value, each thread would perform the same calculations. One approach might be to seed the RNG with different values for each thread, but in this case one would have to ensure that each RNG output sequence did not overlap during simulation; to ensure this were true, an RNG with a large period would be required, thus implying a large amount of state information (memory usage). This application uses an implementation by Alerstam [4] of Marsaglia's multiply-with-carry RNG [92]. As detailed in the references, the algorithm requires only eight bytes of state information and has a period of greater than 2^{60} .

3.3.9 Scored quantities and post-processing

The simulation code stores two sets of data for later analysis: the internal fluence deposition and the parameters of photon packets which exit the simulation domain.

The internal fluence rate is recorded at the zeroth and first optical harmonics at the nodes of each element within the domain. Upon an absorption event the contribution to the total fluence rate is found according to equation 3.40, with a proportion $J_0^2(m_s)$ distributed to the zeroth harmonic and $2J_1^2(m_s)$ to the first. The fluence at the absorption event is distributed to the nodes of the mesh according to linear weightings of the distance of the photon packet to each node. Following simulation the fluence rates are normalised by the total input optical power, and the Voronoi volume of each node in the mesh. Consequently, the recorded fluence is approximated on the mesh in a linear basis. Various standard MATLAB functions can be used to investigate the dataset as defined on the nodes of the mesh. Alternatively, the data can be interpolated to a regular grid for further analysis.

Upon leaving the simulation domain the location, weight, absolute path length, and phase modulation terms (m_s, φ_a of equation 3.6) of each photon packet are recorded in a buffer residing in the device memory of the GPU, subsequently being transferred to the host and saved to disk. The autocorrelation function of each photon packet may be calculated using equation 3.23.

3.4 Results

The results of this investigation are presented in three subsections.

1. The algorithm is validated by comparing the results of both optical and acousto-optical simulations against analytical and numerical results available in the literature.
2. We consider the performance of the code by comparing it with existing mesh-based optical simulation codes, and our previous GPU implementation.
3. We demonstrate the use of the code in building explicit perturbation sensitivity maps for an acousto-optic experiment, the results of which are published elsewhere.

As described earlier, one of the core uses of the Monte-Carlo algorithm is as a tool for validation of other approximate numerical models of the acousto-optic effect; and this tool will be used for that purpose in chapter 4.

3.4.1 Validation: optical

Validation of the light transport algorithm is performed by two techniques. In the first case we compare the output of the simulation code with approximate analytical solutions for a variety of geometries. The analytic solutions here are based upon the diffusion approximation which is valid for distances greater than one scattering mean free path ($l_s = 1/\mu'_s$, where μ'_s

is the reduced scattering coefficient) from the source. In addition to differences near sources, discrepancies between the analytical solutions and the Monte-Carlo results will also occur when boundary conditions are not matched. Specifically, the simulation code operates in a domain of limited spatial extent, whereas many of the presented solutions are valid only for semi-infinite domains. These effects will be noted where present. In the second case, the simulation program has been cross-validated against alternative Monte-Carlo simulations to demonstrate agreement near the source.

The choice of optical scenarios chosen for validation is motivated by the validation techniques applied to other Monte-Carlo simulation codes, the availability of analytical solutions, and an understanding of the elements of the algorithm tested by differing scenarios. The latter consideration is the most important, and motivates the following scenarios.

1. Domains involving large numbers of small and/or poor quality elements tests the resilience of the program's intersection codes.
2. Multiple subdomains of differing refractive indices ensure reflection and refraction are properly handled.
3. Multiple subdomains of differing optical properties test the registration of the subdomains with the tetrahedral mesh.

In all cases, the simulation code must predict the proper path length and weight of each photon packet (related to the scattering and absorption in the medium respectively) in order to adequately predict the quantities of interest. In the following sections, a variety of increasingly complicated scenarios (in terms of geometry and heterogeneity of optical properties) are demonstrated.

Transmittance and Diffuse Reflectance of a Thin Slab

Following the procedure outlined in [134] the diffuse reflectance and total transmittance of a slab of a scattering medium of thickness $d = 0.02\text{cm}$ with optical properties $\mu_a = 10\text{cm}^{-1}$, $\mu_s = 90\text{cm}^{-1}$, and anisotropy factor $g = 0.75$ are calculated in response to a coherent pencil beam perpendicularly incident on the lower surface. Ten simulations of 50,000 photon packets are performed, and the average values of the aforementioned parameters are calculated alongside standard error.

In the simulation program the simulation domain is of limited but suitable spatial extent in the $x-y$ plane — e.g., 0.5m. The tetrahedral geometry is constructed using a point cloud consisting of $10 \times 10 \times 2$ points in the x , y , and z directions, respectively. The reflectance and transmittance are calculated by summing the weight of all of the photon packets escaping the simulation domain on the appropriate surfaces. Table 3.1 demonstrates the results of the process against other models, as provided in [134].

The simulation program demonstrates a 0.01% error³ in the diffuse reflectance and a

³All errors presented with respect to the data of van de Hulst

Source	R_d Average	R_d Error	T_t Average	T_t Error
van de Hulst, 1980,	0.09739		0.66096	
Code under development	0.09740	0.00025	0.66070	0.00033
MCML	0.09734	0.00035	0.66096	0.00020
Prahl <i>et al.</i> , 1989	0.09711	0.00033	0.66159	0.00049

Table 3.1: Comparison of the diffuse reflectance and total transmittance of a thin slab with matched boundary predicted by the simulation program against a diffusion model and alternative Monte-Carlo simulation codes.

–0.39% error in the total transmittance. The standard error lays within the bounds defined by the other simulation programs. The output of the program has been tested to be insensitive to the complexity of the tetrahedral geometry (e.g., increases in the number of elements in any axis). Increasing the number of averages to 50 generates an $R_d = 0.0973$ (–0.092%) and $T_t = 0.6609$ (–0.0091%), with consistent standard errors. This validation experiment supports the correct operation of light transport routines of the simulation program.

Depth Resolved Internal Fluence of a Thin Slab (Ad-Hoc Voxelisation)

Once again following the procedure of [134], the internal fluence is recorded as a function of depth through a semi-infinite slab of a scattering medium with optical properties $\mu_a = 0.1\text{cm}^{-1}$, $\mu_s = 100\text{cm}^{-1}$, and anisotropy factor $g = 0.9$. The photon packet absorption deposition is recorded in the simulation program using one dimensional array of 200 hundred elements spanning a 10mm depth of the medium, and converted to fluence through division by the homogeneous absorption coefficient. Diffusion theory suggests that at a depth in the medium of greater than the effective transport coefficient,

$$\mu_{\text{eff}} = \frac{1}{\sqrt{3\mu_a(\mu_s + \mu_s(1-g))}} \approx 5.7\text{mm} \quad (3.44)$$

fluence will decay exponentially according to,

$$\phi(z) = \phi_0 \exp(-z/D) \quad (3.45)$$

where ϕ_0 is the incident irradiance, and z the depth within the medium.

Figure 3.3 compares the output of the simulation program with that of CUDAMCML [4], which independently validates against other Monte-Carlo codes, including MCML. An exponential fit according to the presented theory is also overlaid on the graph, using the value of μ_{eff} previously calculated. The exponential fit to diffusion theory is found to be accurate at distances considerably more shallow than the penetration depth of the semi-infinite slab. As highlighted by Wang, the curves intersect at around a single transport mean free path into the medium ($1/(\mu_a + \mu_s') \approx 1\text{mm}$). Beyond this distance, the two Monte-Carlo codes and the exponential fit are in excellent agreement.

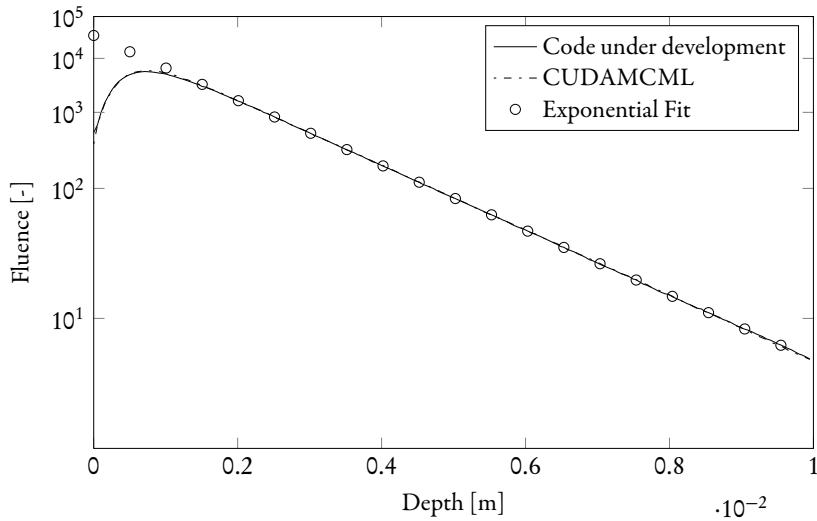


Figure 3.3: Comparison of the depth resolved unitless fluence in a semi-infinite medium as computed by the simulation program against that computed by CUDAMCML, in addition to an exponential fit according to diffusion theory.

Layer	n	μ'_s [1/cm]	μ_a [1/cm]	Depth [cm]
1	1	12	0.1	0.5
2	1	11	0.3	0.2
3	1	13	0.2	0.2

Table 3.2: Optical properties of the multi-layered slab geometry employed in the simulation of figure 3.4

Reflectance of a Multi-Layer Slab

In more recent work, Liemert and Kienle [83, 84] derive a diffusion approximation for an N -layered cylindrical slab with isotropic scattering. Software is provided by the author to numerically solve this equation for an arbitrary number of layers, with differing optical properties. We employ this software, and another Monte-Carlo model, to validate the simulation code in the case of multi-layered slab geometry with index-matched layers. A multi-layered slab is conceived with the optical properties given in table 3.2.

The N -layered diffusion approximation is configured to calculate the reflectance for a cylinder of radius 100mm to minimise any boundary effects when comparing the result with a semi-infinite multi-layered slab. The program is configured to report the reflectance over a 0.1mm grid from the center of the cylinder extending outwards over 300 grid points.

The simulation program was configured by meshing a multi-layered slab of width and depth 10cm, and appropriate height according to table 3.2. A total of 10^7 photons were

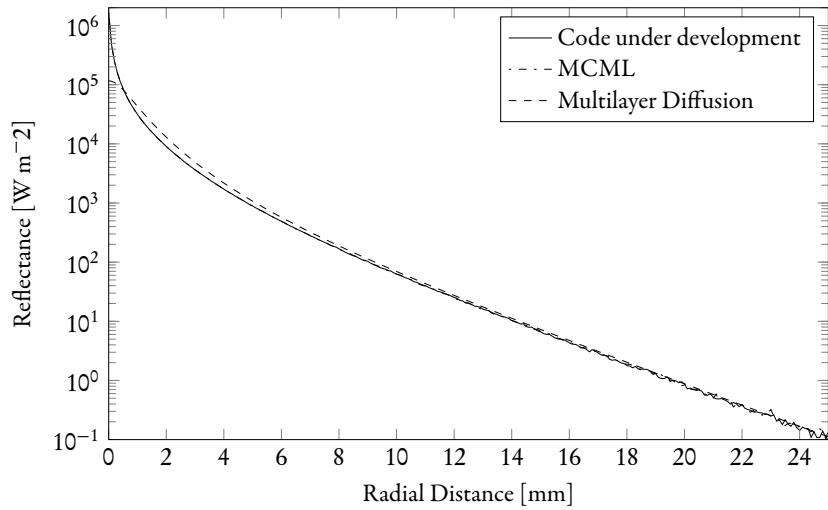


Figure 3.4: Plot of the radially resolved reflectance from a multi-layered slab, as calculated by the simulation program, *MCML*, and a diffusion approximation.

launched at the centre of the slab, and the program recorded all photons exiting the slab. To minimise variance, the reflectance of the simulation program output was calculated by integrating the output flux around thin concentric rings with radii centred on the desired point. This was performed in post-processing, where appropriate scaling was performed to provide reflectance in units of [W m^{-2}]. For further validation close to the source, the MCML program was configured to simulate the same optical configuration in a semi-infinite multi-layer slab. The same number of photons were used as with the simulation program. MCML automatically reports the reflectance in its output file, this was scaled to the same units as produced by the simulation program. Figure 3.4 illustrates the results of the simulation procedure. Both Monte-Carlo models agree excellently over the entire radial range considered. Beyond circa 16mm, statistical noise is evident in these curves owing to the reduced number of photon packets binned at larger distances from the source. The diffusion approximation would not be expected to be accurate close to the source location, and this is indeed noted. As the radial distance increases, the approximation tends towards the Monte-Carlo curves. This result supports the correct operation of the simulation program.

Reflectance of a System of Concentric Shells

Sikora et. al. present a solution to the diffusion equation for the case of a series of concentric shells with heterogeneous optical properties [126]. To demonstrate the functioning of the simulation program with a more complex geometry, we here reproduce figures four and five of their work. These graphs demonstrate the intensity on the surface of the outer shell against the angle from the source of illumination for two different configurations of optical

Shell	Case 1 μ_a	Case 1 μ'_s	Case 2 μ_a	Case 2 μ'_s
Outer	0.005	1	0.005	1
Middle	0.005	1	0.005	0.75
Inner	0.01	1.5	0.01	1.5

Table 3.3: Optical properties of the concentric shells [mm^{-1}] employed in the simulation of figures 3.6 and 3.7

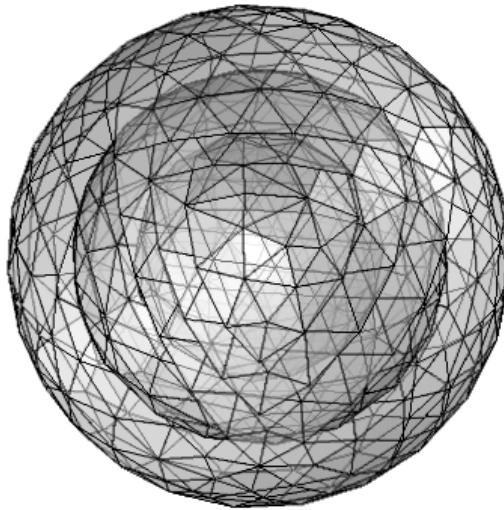


Figure 3.5: Mesh employed in the simulation of figures 3.6 and 3.7 for the series of concentric shells.

properties, as given in table 3.3. The mesh for the concentric shells, of radii 20, 15, and 10mm was built using COMSOL Multiphysics, and is displayed in figure 3.5. The simulations were performed with 10^7 photons launched inwards, normal to the external surface of the sphere. The analytical solutions computed using a Mathematica script provided by an author of the cited work. Variance in the Monte-Carlo results was reduced by integrating the output over bands around the sphere at each of the calculation points before appropriate scaling to match the output of the analytical solution. The results of the two cases are shown in figures 3.6 and 3.7.

The amplitude in each case has been normalised to the value of the analytical solution at an angle of $\pi/2$ radians from the source location. In both cases it can be seen that excellent agreement is found at an angle from the source exceeding $\pi/4$ radians. Whilst the diffusion approximation is not expected to provide complete agreement with the radiative transfer equation (as approximated by the simulation program) near the source, this is exacerbated in this case as the source in the analytical solution is in fact a Gaussian beam, rather than a

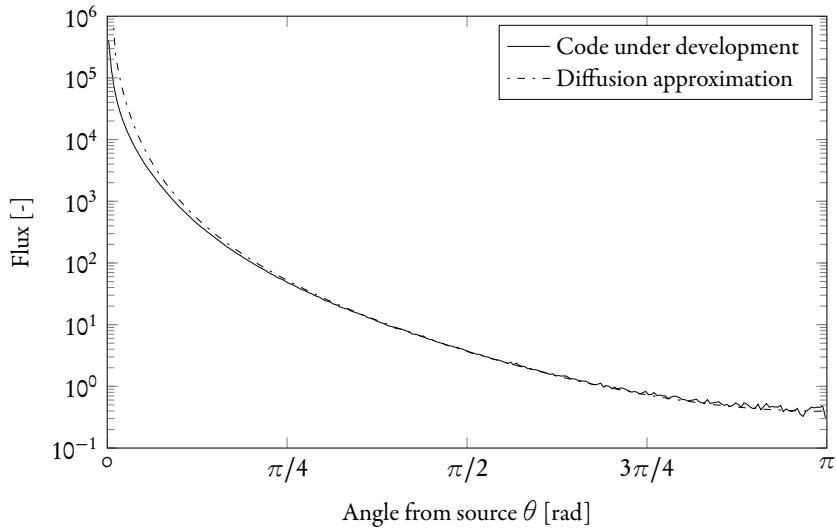


Figure 3.6: Plot of flux on the outer surface of the concentric spheres vs. angle from the illumination source for the parameters of case one of table 3.3.

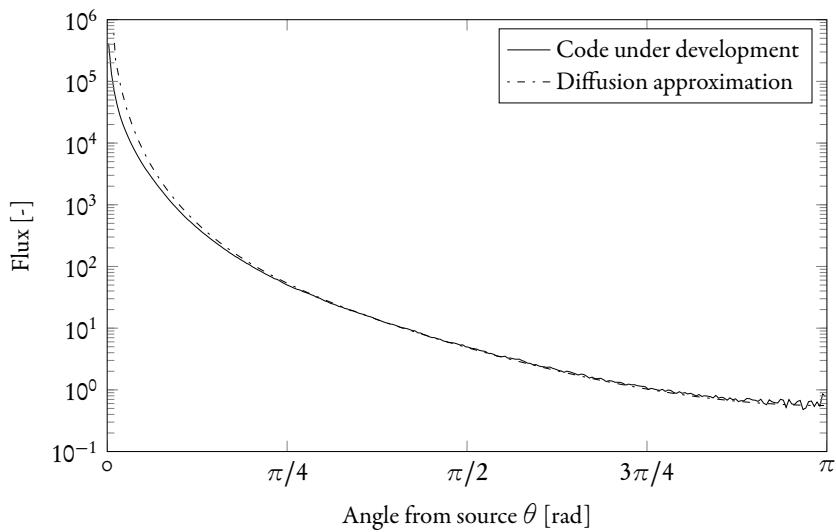


Figure 3.7: Plot of flux on the outer surface of the concentric spheres vs. angle from the illumination source for the parameters of case two of table 3.3.

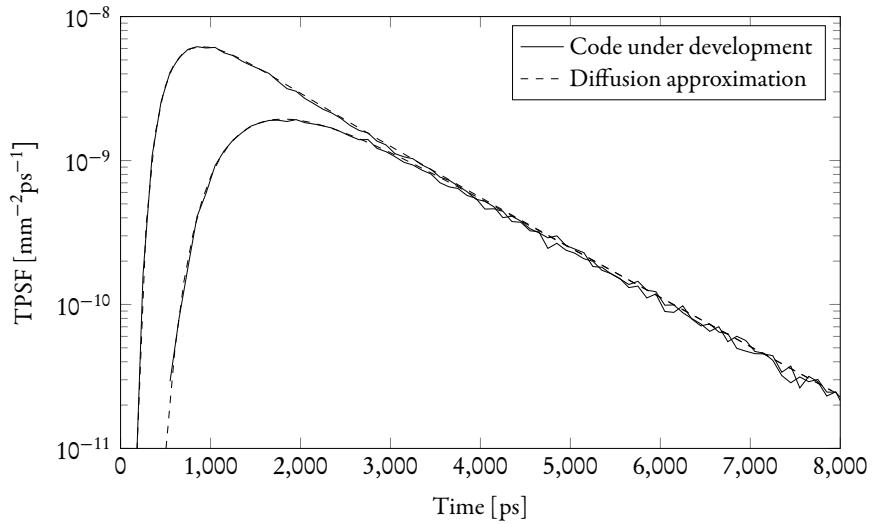


Figure 3.8: Temporally resolved reflectance (upper curve) and transmittance (lower curve) for a slab of optical properties given in the text, resolved at a radial distance $r = 30\text{mm}$ from the source.

pencil beam. The agreement in the nature and differences in the curves at large angles from the differing source provides reasonable confidence in the proper operation of light transport algorithms and geometry handling of the simulation program.

Temporal Point Spread Functions for Semi-Infinite Slabs

Contini et. al. provide solutions for various measures of the optical response of semi-infinite slabs under the diffusion approximation [32]. These solutions include the temporally and spatially resolved reflectance and transmittance of the slabs to instantaneous and continuous-wave illumination from a coherent pencil-beam. To demonstrate that the correct optical path lengths are being taken by the simulated photon packets, and that the correct absorption is being applied, we reproduce two figures from a subsequent paper in which Martelli et. al. compare their analytical results with Monte-Carlo models [93]. In the first case, corresponding to figure two of [93], we examine the time-resolved reflectance and transmittance of a slab of thickness 40mm, $\mu'_s = 0.75\text{mm}^{-1}$, refractive index $n_2 = 1.4$ and relative refractive index to the background medium of $n = 1.33$. The values are resolved at a radial distance $\rho = 30\text{mm}$ from the source position. The results are presented in figure 3.8. Each curve demonstrates excellent agreement with the analytical solution, as would be expected in a region far from the source. Statistical noise is seen in the Monte-Carlo data owing to the limited number of photons resolved at this radial distance.

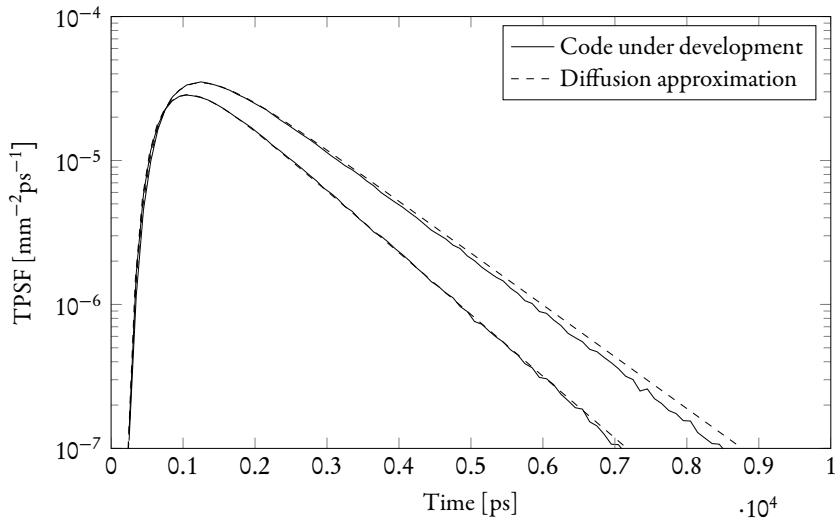


Figure 3.9: Temporally resolved total transmittance for a slab of optical properties given in the text with relative refractive index of $n = 1.4$ (upper curve) and $n = 1$ (lower curve).

In the second case, corresponding to figure three of [93], we consider the total time-resolved transmittance for the same slab, though this time with $\mu'_s = 0.5\text{mm}^{-1}$, refractive index $n_2 = 1.4$ and relative refractive indices of $n = 1.4$ and $n = 1$. The results are presented in figure 3.9. In this case, excellent agreement is found for the case of matched refractive index, though in the case of $n = 1.4$ the simulation appears to under-report the fluence at large time intervals. We note that the same discrepancy is also seen in the original work, and as such we consider this to be a consequence of the diffusion approximation rather than an error in the Monte-Carlo code. The figures presented in this section support the correct operation of the simulation program.

3.4.2 Validation: acousto-optical

Validating the phase accrual algorithm presents a greater challenge than the optical case since fewer analytical formulations are available. Furthermore, those which are presented in the literature typically employ considerable approximations, or are limited to simple geometries. Two sets of analytical results are employed in this work.

Plane Wave Acousto-Optics

Sakadžić and Wang present an analytical solution for the modulation depth developed during coherent illumination of a semi-infinite slab insonified by plane-wave ultrasound [114]. As part of this work, a set of figures is presented comparing their analytical solution with a Monte-Carlo model which they have developed. These figures compare the expected mod-

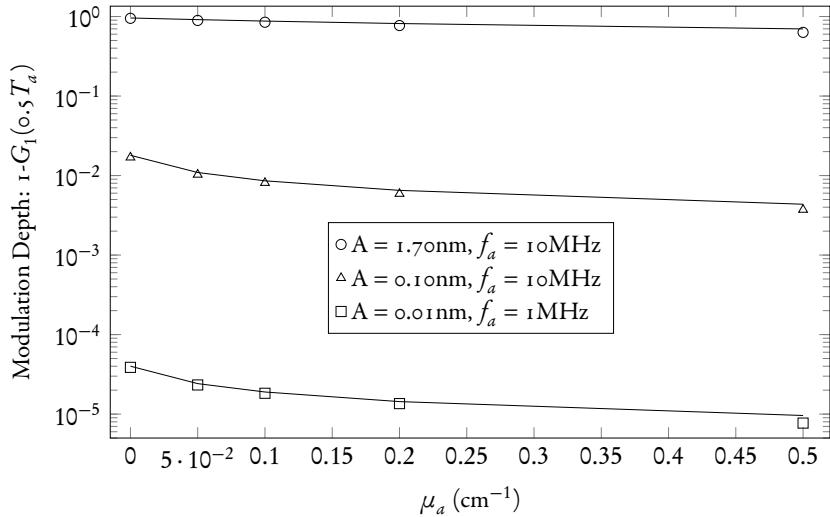


Figure 3.10: Acousto-optic modulation depth calculated by the simulation program vs. an analytical solution for changes in the optical absorption coefficient μ_a

ulation depth against the optical absorption coefficient, acoustic frequency, acoustic amplitude, and optical scattering coefficient. The baseline optical properties of the slab are $v_a = 1480$ [m s^{-1}], $\eta = 0.3211$, $n_0 = 1.33$, $L = 2$ [cm], $\mu_s = 20$ [cm^{-1}], $\mu_a = 0$ and $g = 0$. Figures 3.10 through 3.13 compare the analytical solution with the results calculated by the simulation program. Given the considerable approximations made in the analytical development, excellent agreement is seen between the analytical solution and the output of the simulation program. In figure 3.11 we observe that the modulation depth increases with acoustic frequency due to increased acousto-optic tagging efficiency. In typical biological media we would expect increasing acoustic attenuation with frequency to act contrary to this trend: this is not modelled in the present example.

Partially Insonified Medium

An alternative analytical formulation is derived by Sakadžić and Wang based upon application of the diffusion approximation to a correlation transfer equation [118]. In this work the correlation diffusion equation is solved by a finite difference technique for a semi-infinite slab geometry containing a cylinder of plane-wave ultrasound. We will explore such models further in chapter 4.

Here we recreate this simulation in a cuboid mesh of width $x = 20$ mm. The depth and height of the mesh ($y = 100$ mm, $z = 40$ mm) are chosen to minimise boundary effects. A cylinder of ultrasound is created of radius 3.175mm. 10^7 photon packets were launched into the domain at $x = 0$, $y = 10$ mm and allowed to propagate until they were completely

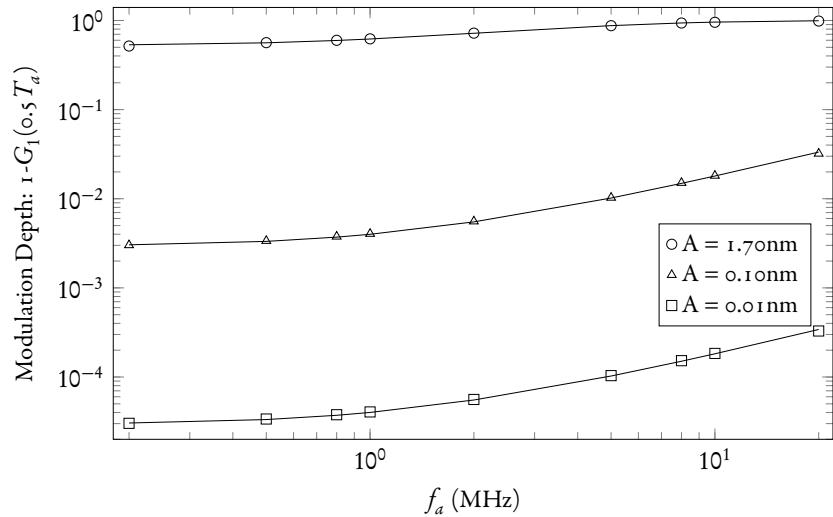


Figure 3.11: Acousto-optic modulation depth calculated by the simulation program vs. an analytical solution for changes in the acoustic frequency f_a

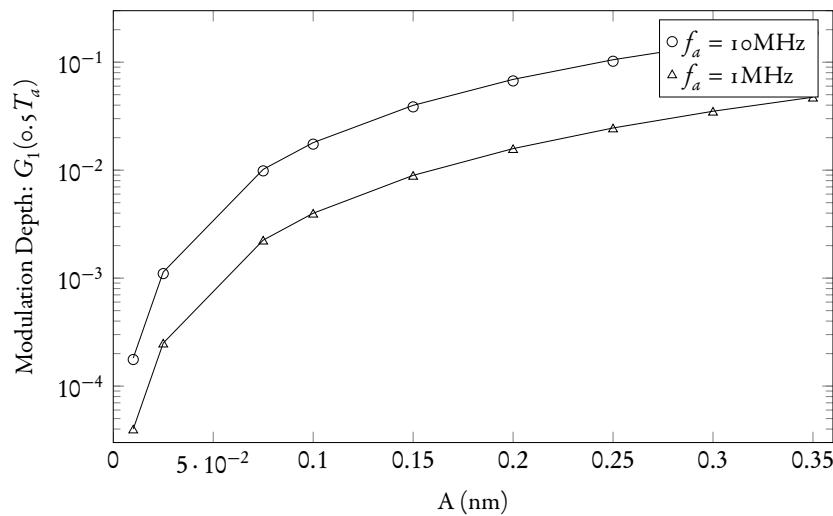


Figure 3.12: Acousto-optic modulation depth calculated by the simulation program vs. an analytical solution for changes in the acoustic displacement amplitude A

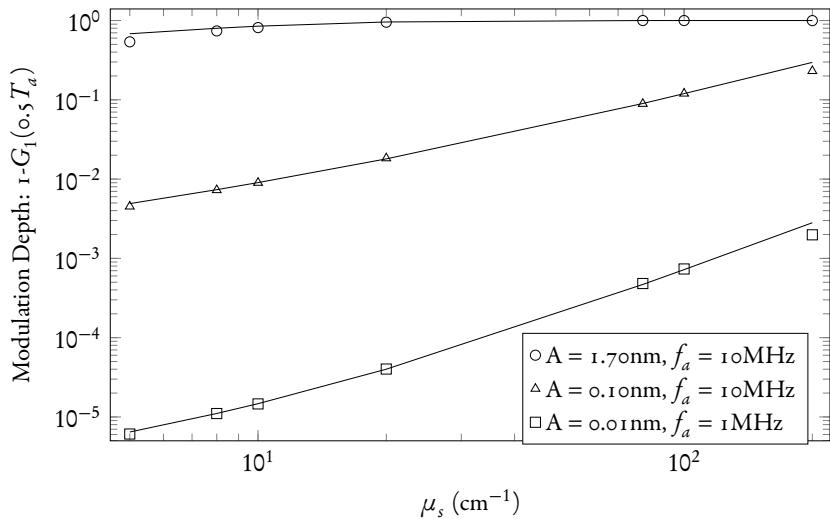


Figure 3.13: Acousto-optic modulation depth calculated by the simulation program vs. an analytical solution for changes in the scattering coefficient μ_s

absorbed into the medium, or exited the domain via one of the boundaries. The simulation program recorded the internal fluence distribution on each node throughout the medium at the input optical frequency, $\Phi_0(\mathbf{r})$, and at the first shifted side-band, $\Phi_1(\mathbf{r})$. The internal first side-band modulation depth was calculated as $MD_1 = \Phi_1(\mathbf{r})/\Phi_0(\mathbf{r})$ prior to interpolation to a regular grid using standard MATLAB functions.

Figure 3.14 plots the internal first side-band modulation depth MD_1 inside the slab in the plane of the source. Figure 3.15 compares the MC simulation to the analytical solution along the transmission and reflection planes of the slab.

Both figures demonstrate a minimum in the recorded MD_1 close to the source region where the large unmodulated intensity is dominant. In figure 3.15 we note considerable disagreement between the MC and analytical models in the same region. This discrepancy occurs because the analytical model, employing the diffusion approximation, requires the approximation of a collimated source incident on the medium with a point source internal to the medium: that is to say that we expect the analytical solution to be invalid in these regions.

Close to the ultrasound location the optical field is principally diffuse. Given that the ultrasound column has constant intensity over its extent, modulated light radiates outwards from this ‘virtual source’ in a relatively uniform manner (contrasted with the exponentially decaying unmodulated source incident on the boundary). The MD_1 , as the ratio of the aforementioned quantities, is thus found to be maximised in those areas closest to the virtual source where the unmodulated flux is smallest: this effect is manifest as a shadow of high MD_1 emanating from the ultrasound column, directed away from the unmodulated source position.

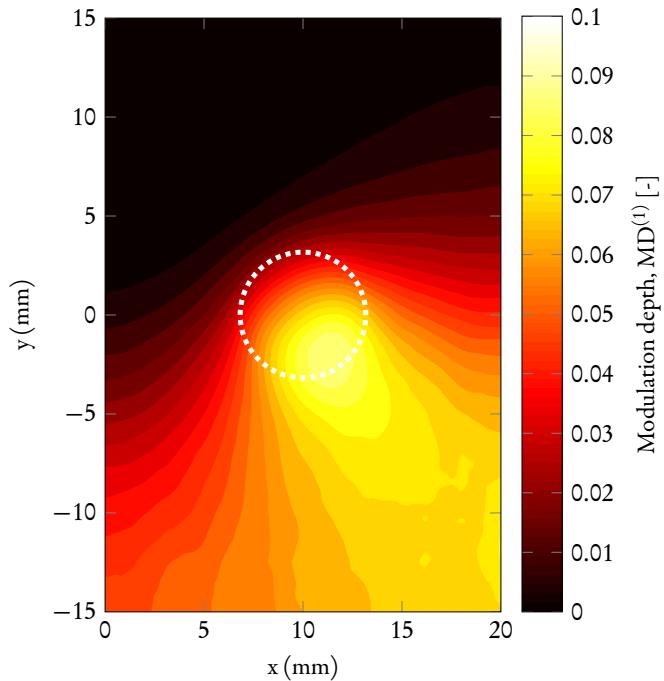


Figure 3.14: First harmonic internal modulation depth for a cut through the slab in the plane of the source. The source is incident normal to the slab surface at $x = 0\text{mm}$, $y = 10\text{mm}$. The the column of ultrasound is centred at $x = 10\text{mm}$, $y = 0\text{mm}$, as indicated by the white circle.

Overall, excellent agreement is found over the region in which the diffusion approximation is valid. The MC results agree excellently with those presented in the original work [118].

In each of the simulations performed for the validation procedure the explicit acoustic integration technique was employed since this is exact in the case of plane-wave insonification. Use of the numerical integration method with six points per wavelength produced results which differed negligibly for this particular set of acoustic and optical properties.

3.4.3 Performance

To illustrate the performance of the simulation code, we compare its optical performance against a recent CPU mesh-based MC model, MMCM [40]. We execute the simulation of the inhomogeneous sphere in a scattering slab using the same mesh (MMCM Mesh 2) as used by Fang, who reports a speed of 5.6 photon packets per millisecond (pp/ms) for a single thread, scaling to 21.2 pp/ms when executed over four threads. Our GPU implementation processes 39.8 pp/ms whilst executing on a single nVidia GTX 250 graphics card.

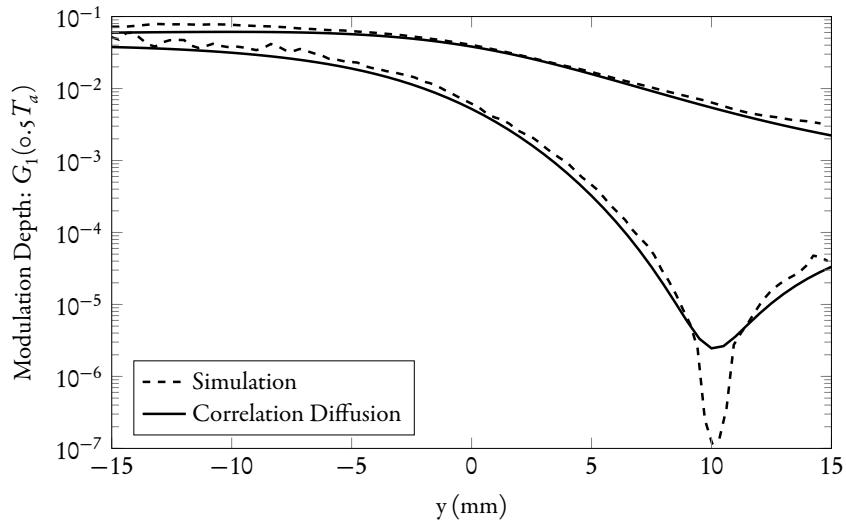


Figure 3.15: First harmonic modulation depth evaluated along the reflection ($x = 0$ mm, lower curve) and transmission ($x = 20$ mm, upper curve) surfaces of the slab in the plane of the source as calculated by the simulation program and the correlation diffusion equation. Excellent agreement is seen with the diffusion equation away from the source location in the reflection surface ($x = 0, y = 10$ mm).

The GPU implementation executes at almost eight times the speed of the single core CPU implementation, or almost twice that of the four-core parallel implementation of Fang. The GPU implementation may be executed on a server incorporating multiple GPUs, where performance will scale linearly according to the number of devices available. Whilst calculation of the AO phase shifts was disabled at run-time, the overhead of variable definitions at compile time still has considerable performance implications owing to the number of threads which may execute simultaneously. Removal of the AO aspects of the simulation code should be expected to significantly improve performance if only an optical simulation is required.

Very little data has been published with which to compare the performance of the code in the AO case. We recently presented the results of a GPU based slab geometry plane-wave AO code and compared these to a CPU implementation [73] of the same. We now compare the performance of the current simulation code to that of both the CPU and GPU codes demonstrated previously.

A semi-infinite slab of depth 2cm (represented in the current simulation code as a slab of lateral extent 10cm) with optical properties $n_0 = 1.33$, $\mu_a = 0\text{cm}^{-1}$, $g = 0$, $\eta = 0.321$ and varying μ_s is illuminated on one surface and the MD_1 determined given insonification by an plane-wave acoustic field of displacement amplitude $A = 1.7\text{nm}$ and varying frequency.

μ_s (cm $^{-1}$)	Execution Speed (photon packets/ms)					
	1MHz			10MHz		
	CPU[73]	GPU Orig.[73]	GPU (Num. Int.)	GPU (Explicit)	GPU (Num. Int.)	GPU (Explicit)
5	9.43	1449.27 ($\times 154$)	232.15 ($\times 25$)	251.10 ($\times 27$)	71.85 ($\times 8$)	294.43 ($\times 31$)
10	4.60	1030.93 ($\times 224$)	166.17 ($\times 36$)	193.24 ($\times 42$)	62.08 ($\times 13$)	191.86 ($\times 42$)
20	2.25	568.18 ($\times 253$)	106.22 ($\times 47$)	109.80 ($\times 49$)	49.51 ($\times 22$)	109.24 ($\times 49$)
50	0.89	253.44 ($\times 285$)	44.54 ($\times 50$)	46.77 ($\times 53$)	30.38 ($\times 34$)	46.19 ($\times 52$)
100	0.44	87.57 ($\times 199$)	20.83 ($\times 47$)	20.64 ($\times 47$)	16.38 ($\times 37$)	19.52 ($\times 44$)
200	0.22	29.18 ($\times 133$)	8.71 ($\times 40$)	7.96 ($\times 36$)	8.29 ($\times 38$)	8.18 ($\times 37$)

Table 3.4: Execution speed (photon packets/ms) of alternative acousto-optic simulation programs and configurations. Numbers in parenthesis demonstrate the speed-up relative to the CPU implementation.

Table 3.4 details the speed of the simulations in numbers of photon packets per millisecond. The previously reported simulation figures are given in the column labeled GPU Orig., the speed of the current simulation is described for two frequencies and for both the numerical integration of phase of equation 3.41, and the explicit integration of equation 3.43.

To understand the resultant performance data, we must consider the limiting factors in each of the simulation algorithms. The performance of the original CPU and GPU codes is primarily a function of the scattering coefficient, μ_s , since this dictates the number of calculations which will be performed for a given photon packet's simulation (as discussed further in our earlier work [73]).

The new simulation code retains the same performance dependency whilst introducing two further factors which may affect performance.

1. If the average size of the mesh describing the simulation domain becomes smaller than the average step size of a photon packet (i.e. at low values of μ_s and with a dense mesh), photon packet steps will routinely be truncated and geometry intersection tests will begin to dominate the processing workload. In this region of operation, the performance will be dependent upon the dimensions of the mesh.
2. When operating in the numerical integration mode a large photon packet step size relative to the acoustic wavelength will result in a multiple memory reads and sets of arithmetic operations to determine the value of the integral (i.e. $N > 1$ in equation 3.41). In this region, performance will depend upon the relative sizes of the average step size of a photon packet and the acoustic wavelength. This limitation does not affect the AO explicit integration mode, which currently performs only one calculation per photon packet step.

The aforementioned qualities of the performance of the new simulation code are evident in the performance data of table 3.4. Under insonification by ultrasound at a frequency of 1MHz (when the acoustic wavelength is long compared a mean free path) geometry intersection tests (and their associated memory access) dominate the execution time and the new simulation code performs similarly under both integration regimes.

The increased flexibility and associated complexity of the new simulation decreases the performance of the new code relative to the previously reported GPU implementation. The new code achieves a performance varying between 1/8th and 1/3th of the previously reported GPU code; as μ_s increases the speeds of the two GPU codes approach one another, since fewer geometrical intersection tests are processed in the new code relative to the total computational workload.

Under 10MHz insonification the explicit integration mode of the simulation maintains similar execution speed to that of 1MHz insonification, though the numerical integration shows a significant relative degradation in performance at low values of μ_s .

In all but one of the examples investigated here the new simulation code is at least an order of magnitude faster than the CPU implementation; though as discussed performance

could be degraded by introducing a finer mesh, or in the case of the numerical integration mode, a higher frequency of insonification.

3.4.4 Sensitivity Mapping

In recent work, Gunadi and Leung [52] demonstrated the measurement of spatial sensitivity maps for the AO sensing technique. In this work an optical absorber of $\mu_a = 42\text{cm}^{-1}$, $\mu_s = 12\text{cm}^{-1}$, and dimensions $5 \times 5 \times 12\text{mm}$ was moved around a tank of intralipid of negligible absorption and equal scattering to that of the absorber. In the first case a transmission mode geometry was investigated, with an optical source and detector separated by 35mm. The absorber was moved around the central region in 1mm steps over a 25×25 grid, in each case the detected optical intensity (in the absence of ultrasound) and the AO MD₁ (in the presence of ultrasound) were recorded. A reference measurement was taken in the absence of an absorber such that the sensitivity can be represented as $J = (X_{\text{ref}} - X)/X_{\text{ref}} \times 100\%$, where X may be optical intensity or AO MD₁, and X_{ref} is the corresponding reference measurement.

We thus recreate the experiment in simulation though we reduce the absorber size to a $1 \times 1 \times 12\text{mm}$ cube to achieve a higher spatial resolution. The linear acoustic field is generated by distributing monopoles over the surface of a spherical cap of equal dimension to the transducer used in practice, and summing their contributions in the domain. The amplitude is scaled to match that recorded by Gunadi and Leung. The mesh and absorptive inclusion is built on the fly using the iso2mesh toolbox. Following simulation, the approximate adjoint method described in section 3.2.2 is used to calculate the detected MD⁽¹⁾. Figure 3.16a and 3.16b depict the optical and AO spatial sensitivity recovered by 1250 independent MC simulations, each of two million photon packets.

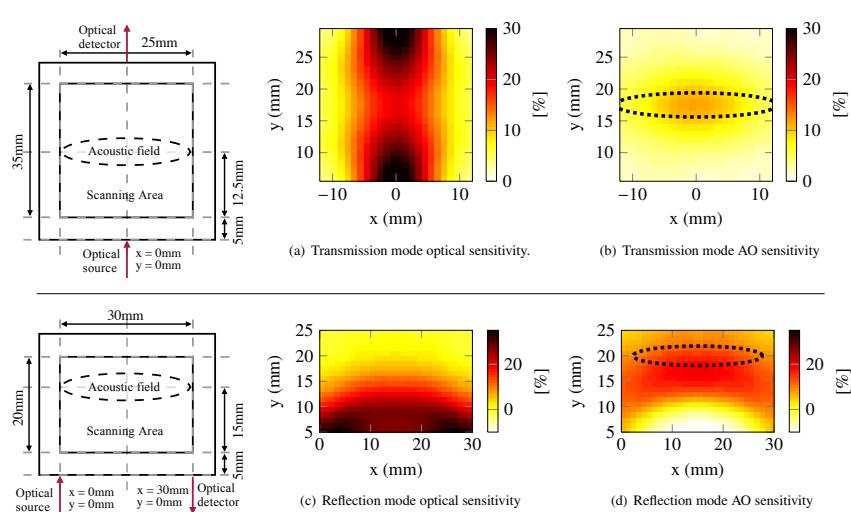


Figure 3.16: Spatial sensitivity maps (-% change from reference value). Inward arrows depict the source and outward arrows the detector location. The dotted ellipse illustrates the focus of the ultrasound region.

Whilst the smaller absorptive inclusion leads to differences in the absolute sensitivity demonstrated by measurement and simulation, the form of the sensitivity distribution retains the same qualities as described by the experimental observations of [52]. As would be expected, the optical sensitivity remains concentrated around the source and detector locations. The AO sensitivity demonstrates localisation to the acoustic field and, more significantly, a very low level of sensitivity close to the source and detector.

The second case studied in [52] is a reflection mode geometry with source-detector spacing of 30mm. The ultrasound field is located centrally at a depth of 20mm from the input and detection surface. The same optical absorber was moved around a scanning region of 30×20 mm, located centrally and starting at a depth of 5.5mm from the input surface (owing to the dimensions of the experimental absorber). This experiment was recreated in simulation; figure 3.16c and 3.16d depict the optical and acoustic-optic spatial sensitivity recovered by 1800 independent MC simulations of two million photon packets.

Once again, the absolute sensitivities of the simulation differ slightly to that of the experimental work owing to the differing absorber dimensions. The optical sensitivity has the expected ‘banana’ shape between the source and detector, with sensitivity rapidly declining at distances greater than around 10mm. The AO sensitivity is concentrated at a depth of around 17mm from the input/output plane; the peak relative sensitivity is not within the ultrasound focus region as the lack of light at these depths outweighs the higher acoustic pressures. Close to the source/detector plane a region of negative AO sensitivity is found. Negative sensitivity suggests that an absorptive perturbation has the effect of increasing the measurement quantity. This can occur when the perturbation reduces the amount ‘unmodulated’ light reaching the measurement position by an amount proportionally greater than it reduces light which has been modulated into the first side-band. The form of the simulated sensitivities shows excellent agreement with the experimental work of Gunadi and Leung. We will revisit such sensitivity maps in chapter 5 where they will be employed in an image reconstruction problem.

3.5 Conclusions

We have demonstrated the implementation of a highly parallel AO MC simulation code. The optical model has been extensively validated against various analytical and numerical solutions, and the AO model against an appropriate analytical example from the literature.

The performance of the simulation code has been compared against CPU and GPU implementations of pre-existing optical and AO simulation codes. Whilst our previous GPU implementation of an AO simulation code provides better performance, this is at the cost of significantly reduced flexibility in the form of the domain. The flexibility and performance of the code have also been demonstrated in the generation of spatial sensitivity maps which support published experimental work[52].

At the time of writing, the program developed here is the first GPU accelerated, mesh-based, Monte-Carlo simulation code which has been demonstrated in the literature. Beyond its use as a fast forward model in AO, various aspects of its implementation may be valuable to the wider diffuse optics community. The removal of the AO calculations from the code base is expected to considerably improve performance such that it may be appropriate for use where voxelised simulations are currently used, despite their disadvantages [21].

The validity of the model is limited by the two components of the algorithm: the light transport model, and the phase accrual expressions. The former component, derived from the scalar RTE, is valid for the vast majority of conceivable experimental scenarios where the weak-scattering approximation is valid; excepting measurements affected by interference phenomena such as coherent back-scattering [56]. The latter component was derived here in considerable generality, though we have made the assumption that the optical paths through the medium are not themselves altered by the ultrasonically induced perturbations in the index of refraction. Sakadžić states this assumption to be reasonable for pressure amplitudes experienced in biomedical applications at moderate ultrasound frequencies $f_a < 10^6 \text{ Hz}$ [117].

The primary limitation of the current model is that it is designed purely for monochromatic acoustic excitation. The implementation of a polychromatic acoustic field distribution would permit the simulation of modulated continuous wave inputs, or pulsed ultrasound, as employed to gain spatial selectivity in the axis of the ultrasound. This could potentially be implemented in the style of Sakadžić's voxelised CPU simulation [119].

It would also be of great value to extend the simulation to directly record the correlation radiance in the meshed domain. This would permit the use of an unapproximated adjoint method to calculate the result of PS/PD simulations, and more importantly, allow the model to be used as part of an iterative inversion procedure to reconstruct unknown optical parameters from acousto-optic measurements. We will examine this idea in greater depth in chapter 5 where we employ a correlation diffusion approximation to achieve this goal.

Chapter 4

Approximations of correlation transport by spherical harmonic expansions

4.1 Introduction

Forward models based upon truncated spherical harmonic (P_N) expansions of the radiative transport equation are ubiquitous in diffuse optical tomography [12]. The diffusion approximation, found from the application of further simplifying assumptions to the P_1 approximation, has enjoyed particular attention. One attraction of the latter is that the resultant parabolic (in the time dependent case) or elliptic (in the continuous-wave case) equations are amenable to solution by fast and rigorous deterministic numerical methods.

The finite element method in particular has been used extensively for the numerical solution of the diffusion approximation [8, 15, 123, 124]. Owing to the tractability of the finite element, the system matrices which define the discretised problem are of significant value in forming sensitivity relations required for various image reconstruction techniques. Such matters are discussed further in chapter 5.

If an improved approximation is desired one option is to employ higher order P_N approximations [129]. In three dimensions, this techniques quickly becomes too computationally demanding to be used as the forward model in, for example, an iterative reconstruction algorithm, since the number of equations grows as $(N + 1)^2$. Alternative methods such as discrete ordinates (the S_N approximation) also suffer from similar computational issues, with the number of equations growing as $N(N + 2)$.

An alternative approximation, originally applied to neutron transport in nuclear reactors, is that of the *simplified* spherical harmonic approximations (SP_N) [46]. As will be seen, this approximation results in a system of parabolic/elliptic equations which are coupled by their coefficients. These approximations were originally derived from the planar-form of the transport equation before the gradient operators were extrapolated to their three-dimensional equivalents. Whilst providing improved results over the standard diffusion equation (if be-

low that of the full higher order spherical harmonic methods) the lack of a rigorous mathematical foundation prevented their widespread adoption in other fields such as the modelling of light transport in diffuse optical tomography. Larsen et al. subsequently demonstrate that the SP_N approximations can be derived by an asymptotic approach which expands the transport equation in a Neumann series [68–71]. More recently, Klose [65] applied the technique to diffuse optical imaging.

A diffusion-style approach has also been applied to the problem of UOT. In work by Sakadžić and Wang [118], a simplified correlation transport equation was developed, and the authors stated a diffusion-like approximation.

In this chapter we provide a full derivation for the diffusion approximation presented by Sakadžić and Wang, which has until now not been published. We extend this work by performing the derivation in a general way which permits the consideration of a higher order approximation of the phase increments which describe the underlying correlation transport process. We also take another approach to simplifying the underlying transport equation which permits us to form an SP_N approximation for UOT. This requires further approximations in the treatment of the acousto-optic phase increments which drive modulation of the optical field, but should provide a more accurate model of the optical energy distribution in a given domain. Each of these models are solved by the finite-element method, and the results compared with our reference Monte-Carlo simulation, the development of which was described in section 3.

4.2 Theory

4.2.1 Correlation transport

Ackerson et al. [2] originally proposed a correlation transport equation for use in the field of diffusing wave spectroscopy. Dougherty et al. [35] later provided a more rigorous derivation based upon multiple scattering (analytic) theory with moving scatterers, building on the work of Ishimaru and Hong [57] and others [87, 128]. Dougherty’s integral formulation reads,

$$G(\mathbf{r}_b, \hat{\mathbf{s}}, \tau) = G_0(\mathbf{r}_b, \hat{\mathbf{s}}, \tau) + \mu_s \int_{\mathbf{r}_0}^{\mathbf{r}_b} \int_{4\pi} p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') G(|\mathbf{r}_b - \mathbf{r}_a|, \hat{\mathbf{s}}, \tau) g^1(\hat{\mathbf{s}}, \hat{\mathbf{s}}', \tau) \exp(-\mu_t |\mathbf{r}_b - \mathbf{r}_a|) d|\mathbf{r}_b - \mathbf{r}_a| d\hat{\mathbf{s}}'. \quad (4.1)$$

This equation states that the radiance at a particular point in space \mathbf{r}_b with direction $\hat{\mathbf{s}}$ is given by the incident direct field (G_0) plus the diffuse contribution scattered from directions $\hat{\mathbf{s}}'$ into $\hat{\mathbf{s}}$ along all the points \mathbf{r}_a along some line starting at a distant \mathbf{r}_0 to the point of evaluation \mathbf{r}_b . The equation differs from the radiative transport equation by the addition of the term $g^1(\hat{\mathbf{s}}, \hat{\mathbf{s}}', \tau)$; the *single scattering correlation function* which describes the decorrelation of the optical field which occurs when partial waves undergo a scattering process from

moving particles. We might think of this temporal decorrelation in the frequency domain as the broadening of the line-width of a coherent source by a Doppler shift imposed by the moving scatterers. An expression for g^1 under Brownian motion of the scatterers was presented in some of the original investigations of DWS as described by Maret and Wolf [89], and Pine et al. [103],

$$g^1(\hat{\mathbf{s}}, \hat{\mathbf{s}}', \tau) = \exp \left[-\frac{1}{3} k_0^2 6 D_B \tau (1 - \hat{\mathbf{s}} \cdot \hat{\mathbf{s}}') \right], \quad (4.2)$$

where D_B is the Brownian diffusion component and $6 D_B \tau = \langle \Delta r^2(\tau) \rangle$ where $\langle \Delta r^2(\tau) \rangle$ is the mean-square displacement of the scattering particles. As a function of $\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}'$ this term can be managed in the same manner as the phase function under expansion into a series of spherical harmonics. Specifically, when $\tau \ll (2 D_B k_0^2)^{-1}$ the term can be Taylor expanded,

$$g^1(\hat{\mathbf{s}}, \hat{\mathbf{s}}', \tau) = 1 - 2 D_B k_0^2 \tau + 2 D_B k_0^2 \tau \frac{4\pi}{3} \sum_{m=-1}^1 Y_{1,m}^*(\hat{\mathbf{s}}') Y_{1,m}(\hat{\mathbf{s}}). \quad (4.3)$$

By this method, a diffusion style equation for correlation transport was derived by a number of authors; see Boas [23] for a full derivation.

4.2.2 Correlation transport under acoustic modulation

In the case of a medium under harmonic insonification, Sakadžić and Wang citeSakadzic:2006tx derived a correlation transport equation amenable to solution by P_N approximation methods. The ‘locality’ required for such an approximation is achieved by various approximations and assumptions which we will explore here. The approximated integral form of the equation reads,

$$\begin{aligned} G(\mathbf{r}_b, \hat{\mathbf{s}}, t, \tau) &= G_0(\mathbf{r}_b, \hat{\mathbf{s}}, t, \tau) \\ &+ \mu_s \int_{r_0}^{r_b} \int_{4\pi} p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') G(\mathbf{r}_a, \hat{\mathbf{s}}', t, \tau) \exp(-\mu_t |\mathbf{r}_b - \mathbf{r}_a|) \Psi(\mathbf{r}_b, \mathbf{r}_a, \hat{\mathbf{s}}, t, \tau) d|\mathbf{r}_b - \mathbf{r}_a| d\hat{\mathbf{s}}'. \end{aligned} \quad (4.4)$$

In this equation the function $\Psi(\mathbf{r}_b, \mathbf{r}_a, \hat{\mathbf{s}}, t, \tau)$ describes the temporal decorrelation of the field due to transport in an insonified medium between two scattering points located at \mathbf{r}_a and \mathbf{r}_b following scattering into $\hat{\mathbf{s}}$ from direction $\hat{\mathbf{s}}'$ at \mathbf{r}_a . The function is given by

$$\Psi(\mathbf{r}_b, \mathbf{r}_a, \hat{\mathbf{s}}, t, \tau) = \exp(j\Delta\varphi), \quad (4.5)$$

and

$$\begin{aligned} \Delta\varphi &= 2\Lambda \sin\left(\frac{1}{2}\omega_a \tau\right) \sin\left(\mathbf{k}_a \cdot \frac{\mathbf{r}_b - \mathbf{r}_a}{2}\right) \\ &\times \left[S_a(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a) \sin\left(\omega_a t + \frac{1}{2}\omega_a \tau - \mathbf{k}_a \cdot \frac{\mathbf{r}_a + \mathbf{r}_b}{2} + \phi - \phi_a\right) \right. \\ &\quad \left. - \frac{\eta}{(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)} \sin\left(\omega_a t + \frac{1}{2}\omega_a \tau - \mathbf{k}_a \cdot \frac{\mathbf{r}_a + \mathbf{r}_b}{2} + \phi\right) \right], \quad (4.6) \end{aligned}$$

where $\Lambda = 2k_0 n_0 P_0 / (k_a \rho v_a^2)$, k_0 is the optical wavenumber *in vacuo*, n_0 is the refractive index of the medium, P_0 is the local pressure amplitude of the acoustic field, ρ the density of the medium, v_a the speed of sound, η was defined in section 2.4.2, S_a and ϕ_a are the relative scattered displacement amplitude and phase offset, respectively, $\hat{\mathbf{s}}$ is a unit vector representing the direction of propagation of the partial wave, and $\hat{\mathbf{s}}_a$ is the local acoustic propagation direction. As is evident from their coefficients, the first term describes the increment of phase due to the displacement of the optical scatterers, and the second describes the phase increment due to modulation of the refractive index.

Approximating the phase increment

To develop this correlation transport equation into a form suitable for expansion in spherical harmonics, we will follow the approach of Sakadžić and Wang in averaging the phase increment over space, and over free paths. The N^{th} order approximation is given by [118]

$$\Psi(\mathbf{r}_b, \mathbf{r}_a, \hat{\mathbf{s}}, t, \tau) \approx 1 - \Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_N = 1 - \frac{|\mathbf{r}_b - \mathbf{r}_a|}{l_s} \sum_{n=1}^N (-j^n) \frac{\langle \Delta\varphi^n \rangle_{\mathbf{r}_c, l}}{n!}, \quad (4.7)$$

where $l_s = \mu_t^{-1}$ is the optical mean free path, and $\langle \Delta\varphi^n \rangle_{\mathbf{r}_c, l}$ is the average of the phase increment over space ($\mathbf{r}_c = (\mathbf{r}_a + \mathbf{r}_b)/2$) and free paths ($l = |\mathbf{r}_b - \mathbf{r}_a|$). This approximation limits the acoustic pressures over which the model is valid. Sakadžić and Wang demonstrated the development of the second order phase perturbation, but we shall proceed with generality to permit, e.g, higher ultrasound pressures to be tolerated, by the inclusion of a greater number of terms in the phase increment.

To find $\langle \Delta\varphi^n \rangle_{\mathbf{r}_c, l}$ we first average $\Delta\varphi^N$ over the centre-of-gravity co-ordinate $\mathbf{r}_c = (\mathbf{r}_a + \mathbf{r}_b)/2$ in a volume $V_0 \sim l_s^3$,

$$\langle \Delta\varphi^n \rangle_{\mathbf{r}_c} = \frac{1}{V_0} \int_{V_0} \Delta\varphi^n d^3\mathbf{r}. \quad (4.8)$$

The dependent variable appears in $\Delta\varphi$ (equation 4.6) only as the argument of a sine function. By de Moivre's identity, and employing the binomial theorem, we note that

$$\lim_{x_0 \rightarrow \infty} \frac{1}{x_0} \int_0^{x_0} \sin^n(x) dx = \begin{cases} 0, & \text{for } n \text{ odd,} \\ \frac{n!}{2^n ((n/2)!)^2}, & \text{for } n \text{ even.} \end{cases} \quad (4.9)$$

We now make the assumption that the transport mean free path is much greater than the acoustic wavelength, $k_a l_{\text{tr}} \gg 1$. This approximation might be thought of as the high-frequency limit of acousto-optic modulation, permitting us to treat the individual phase increments accumulated between scattering events as uncorrelated (at least in the case of isotropic scattering).

Assuming a typical transport mean free path of 1mm^{-1} , and a speed of sound $v_a = 1480\text{ms}^{-1}$, this approximation requires that we employ ultrasound frequencies in excess of

2MHz (which gives a value of $(k_a l_{tr}) \approx 8.5$). Experimentally it is desirable to use lower frequencies of ultrasound to minimise the increasing effects of acoustic attenuation; this model may be unsuitable in such scenarios. In a practical application, one may wish to minimise the size of the ultrasonic transducer— this can be achieved without increasing the volume of the focal region (and hence decreasing the spatial resolution) by the use of higher acoustic frequencies. Throughout the development of UMOT, increasingly sensitive detection mechanisms have been designed which can ameliorate the reduction in detected modulation experienced at greater depth due to the application of higher acoustic frequencies. We thus consider the restrictions imposed by this approximation to be acceptable.

Within this regime we may use the relationship in equation 4.9 to perform the integral. The resulting averages are zero for odd n , and for even n ,

$$\langle \Delta\varphi^n \rangle_{r_c} = \frac{n!}{((n/2)!)^2} \frac{\Lambda^n}{(\hat{s} \cdot \hat{s}_a)^n} \sin^n \left(\frac{1}{2} \omega_a \tau \right) \sin^n \left(\mathbf{k}_a \cdot \frac{\mathbf{r}_b - \mathbf{r}_a}{2} \right) \times \left[\eta^2 - 2S_a \eta \cos(\phi_a) (\hat{s} \cdot \hat{s}_a)^2 + S_a^2 (\hat{s} \cdot \hat{s}_a)^4 \right]^{n/2}. \quad (4.10)$$

As highlighted by Sakadžić, the assumption $k_a l_{tr} \gg 1$ has, via the averaging process, removed the time dependence of the equation, thus rendering it wide-sense stationary. We proceed to calculate $\langle \Delta\varphi^n \rangle_{r_c, l}$ by finding the expected value of the expression given the probability distribution of a mean free path,

$$\langle \Delta\varphi^n \rangle_{r_c, l} = \frac{1}{l_s} \int_0^\infty \exp \left(-\frac{l}{l_s} \right) \langle \Delta\varphi^n \rangle_{r_c} dl, \quad (4.11)$$

which is performed by noting that

$$\sin^n \left(\mathbf{k}_a \cdot \frac{\mathbf{r}_b - \mathbf{r}_a}{2} \right) = \sin^n \left(\frac{k_a l (\hat{s} \cdot \hat{s}_a)}{2} \right). \quad (4.12)$$

Completing the integration we have, for even $N > 0$,

$$\langle \Delta\varphi^n \rangle_{r_c, l} = \left[\prod_{m=1}^{N/2} \frac{(2m-1)^2}{1 + (mk_a l_{tr} (\hat{s} \cdot \hat{s}_a))^2} \right] \Lambda^n (k_a l_{tr})^n \sin^n \left(\frac{\omega_a \tau}{2} \right) \times \left[\eta^2 - 2S_a \eta \cos(\phi_a) (\hat{s} \cdot \hat{s}_a)^2 + S_a^2 (\hat{s} \cdot \hat{s}_a)^4 \right]^{N/2}. \quad (4.13)$$

The differential form

Following the averaging of the phase-increments, the (optically steady-state) correlation transport equation is no longer dependent on time, but only lag τ . The phase-averaged integral form of the correlation transport equation reads

$$G(\mathbf{r}_b, \hat{s}, \tau) = G_0(\mathbf{r}_b, \hat{s}, \tau) + \int_{r_0}^{r_b} \int_{4\pi} \mu_s p(\hat{s}, \hat{s}') G(\mathbf{r}_a, \hat{s}', \tau) \exp(-\mu_t |\mathbf{r}_b - \mathbf{r}_a|) [1 - \Psi(\hat{s}, \hat{s}_a, \tau)_N] d|\mathbf{r}_b - \mathbf{r}_a| d\hat{s}', \quad (4.14)$$

where the N^{th} order approximation to the acousto-optic coefficients $\Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_N$ was given in equation 4.7. We convert this to the more familiar differential form by application of $\hat{\mathbf{s}} \cdot \nabla$ to the integral equation [56],

$$[\hat{\mathbf{s}} \cdot \nabla + \mu_t +] G(\mathbf{r}, \hat{\mathbf{s}}, \tau) = \mu_s \int_{4\pi} p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') G(\mathbf{r}, \hat{\mathbf{s}}', \tau) [1 - \Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_N] d\hat{\mathbf{s}}' + S(\mathbf{r}, \hat{\mathbf{s}}, \tau), \quad (4.15)$$

where $S(\mathbf{r}, \hat{\mathbf{s}}, \tau)$ is a source of correlation radiance.

4.2.3 Correlation diffusion approximation

We will now approximate the correlation transport equation of equation 4.15 to form a correlation diffusion approximation following the style of Ishimaru [56]. A necessary condition of the derivation of the correlation transport equation was that the phase increments accrued during successive free paths are uncorrelated; this requires the use of an isotropic phase function,

$$p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') = \frac{1}{4\pi}. \quad (4.16)$$

We approximate the correlation radiance,

$$G(\mathbf{r}, \hat{\mathbf{s}}, \tau) = \left[\frac{1}{4\pi} G(\mathbf{r}, \tau) + \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, \tau) \cdot \hat{\mathbf{s}} \right] \quad (4.17)$$

where

$$G(\mathbf{r}, \tau) = \int_{4\pi} G(\mathbf{r}, \hat{\mathbf{s}}, \tau) d\hat{\mathbf{s}} \quad (4.18)$$

is the correlation density, and

$$\mathbf{F}(\mathbf{r}, \tau) = F(\mathbf{r}, \tau) \hat{\mathbf{s}}_f = \int_{4\pi} G(\mathbf{r}, \hat{\mathbf{s}}, \tau) \hat{\mathbf{s}} d\hat{\mathbf{s}} \quad (4.19)$$

is the correlation flux.

Conservation of energy

To develop the approximation we begin by forming an energy balance equation by integrating equation 4.15 over all 4π of solid angle, having introduced the isotropic phase function,

$$\begin{aligned} & \int_{4\pi} [\hat{\mathbf{s}} \cdot \nabla + \mu_t] G(\mathbf{r}, \hat{\mathbf{s}}, \tau) d\hat{\mathbf{s}} \\ &= \frac{\mu_s}{4\pi} \int_{4\pi} \left[1 - \Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_N \int_{4\pi} G(\mathbf{r}, \hat{\mathbf{s}}', \tau) d\hat{\mathbf{s}}' + S(\mathbf{r}, \hat{\mathbf{s}}, \tau) \right] d\hat{\mathbf{s}}. \end{aligned} \quad (4.20)$$

We use the previous definitions of correlation density and correlation flux to complete the integration of the left hand side of equation 4.20, and the integral over $\hat{\mathbf{s}}'$ involving $G(\mathbf{r}, \hat{\mathbf{s}}', \tau)$ on the right hand side. Rearranging, we get

$$\nabla \cdot \mathbf{F}(\mathbf{r}, \tau) + \mu_a G(\mathbf{r}, \tau) = -\frac{\mu_s}{4\pi} G(\mathbf{r}, \tau) \int_{4\pi} [\Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_N] d\hat{\mathbf{s}} + S_0(\mathbf{r}, \tau), \quad (4.21)$$

where $S_0(\mathbf{r}, \tau) = \int_{4\pi} S(\mathbf{r}, \hat{\mathbf{s}}, \tau) d\hat{\mathbf{s}}$ is an isotropic source of correlation fluence. To complete the energy balance expression we must integrate the phase increment; we will demonstrate this procedure to the second order here. From equation 4.13 with $N = 2$,

$$\frac{1}{4\pi} \int_{4\pi} \Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_2 d\hat{\mathbf{s}} = \frac{\Lambda^2}{8\pi} \sin^2\left(\frac{\omega_a \tau}{2}\right) \int_{4\pi} \frac{(k_a l_{tr})^2}{(k_a l_{tr})^2(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)^2 + 1} \\ \times \left[\eta^2 - 2\eta S_a \cos(\phi_a)(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)^2 + S_a^2(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)^4 \right] d\hat{\mathbf{s}}. \quad (4.22)$$

We proceed by integrating the terms in η which describes the phase increments due to index of refraction variation,

$$\frac{\eta^2}{4\pi} \int_{4\pi} \frac{(k_a l_{tr})^2}{(k_a l_{tr})^2(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)^2 + 1} d\hat{\mathbf{s}} = \eta^2 (k_a l_{tr}) \tan^{-1}(k_a l_{tr}), \quad (4.23)$$

the terms in S_a^2 ,

$$\frac{S_a^2}{4\pi} \int_{4\pi} \frac{(k_a l_{tr})^2(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)^4}{(k_a l_{tr})^2(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)^2 + 1} d\hat{\mathbf{s}} = S_a^2 \left[\frac{1}{3} - \frac{1}{(k_a l_{tr})^2} + \frac{\arctan(k_a l_{tr})}{(k_a l_{tr})^3} \right], \quad (4.24)$$

and finally the cross terms,

$$\frac{2S_a \eta \cos \phi_a}{4\pi} \int_{4\pi} \frac{(k_a l_{tr})^2(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)^2}{(k_a l_{tr})^2(\hat{\mathbf{s}} \cdot \hat{\mathbf{s}}_a)^2 + 1} d\hat{\mathbf{s}} = 2S_a \eta \cos \phi_a \left[1 - \frac{\arctan(k_a l_{tr})}{(k_a l_{tr})} \right]. \quad (4.25)$$

Thus,

$$\frac{1}{4\pi} \int_{4\pi} \Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_2 d\hat{\mathbf{s}} = h(\mathbf{r}, \tau)_2 = \frac{\Lambda^2}{2} \left[\eta^2 (k_a l_{tr}) \tan^{-1}(k_a l_{tr}) \right. \\ \left. + S_a^2 \left[\frac{1}{3} - \frac{1}{(k_a l_{tr})^2} + \frac{\arctan(k_a l_{tr})}{(k_a l_{tr})^3} \right] + S_a \eta \cos \phi_a \left[1 - \frac{\arctan(k_a l_{tr})}{(k_a l_{tr})} \right] \right], \quad (4.26)$$

which we substitute into equation 4.21,

$$\nabla \cdot \mathbf{F}(\mathbf{r}, \tau) + [\mu_a + \mu_s h(\mathbf{r}, \tau)_2] G(\mathbf{r}, \tau) = S_0(\mathbf{r}, \tau). \quad (4.27)$$

In this form the nature of this balance equation is revealed. The divergence of the correlation flux ($\nabla \cdot \mathbf{F}$) represents the net gain or loss of correlation flux for an infinitesimal point in the medium. Our derivation shows this to be equal to any sources at that point, less losses due to absorption ($\mu_a G(\mathbf{r}, \tau)$), and decorrelation ($\mu_s h(\mathbf{r}, \tau)_2$).

Fick's Law

The second element of the approximation requires that we form an equivalent of Fick's law for the correlation transport process. In an analogue of its optical counterpart, this expression will describe the diffusion of correlation from regions of high correlation density to those of

low correlation density by properly defining the correlation flux vector with respect to the diffusion coefficient. To form the expression we multiply equation 4.15 by \hat{s} and integrate over all 4π steradians. After inserting the isotropic phase function,

$$\begin{aligned} & \int_{4\pi} [\hat{s} \cdot \nabla + \mu_t] G(\mathbf{r}, \hat{s}, \tau) \hat{s} d\hat{s} \\ &= \frac{\mu_s}{4\pi} \int_{4\pi} \left[\int_{4\pi} G(\mathbf{r}, \hat{s}', \tau) [1 - \Psi(\hat{s}, \hat{s}_a, \tau)_N] d\hat{s}' + S(\mathbf{r}, \hat{s}) \right] \hat{s} d\hat{s}. \quad (4.28) \end{aligned}$$

For arbitrary vector \mathbf{A} [56],

$$\int_{4\pi} \hat{s} [\hat{s} \cdot \nabla(\mathbf{A} \cdot \hat{s})] d\hat{s} = 0 \quad (4.29)$$

and,

$$\int_{4\pi} \hat{s} (\mathbf{A} \cdot \hat{s}) d\hat{s} = \frac{4\pi}{3} \mathbf{A} \quad (4.30)$$

such that we may complete the integral involving the term in $\hat{s} \cdot \nabla$. The term in μ_t can be integrated given our definition of correlation flux. Accordingly,

$$\frac{1}{3} \nabla G(\mathbf{r}, \tau) + \mu_t \mathbf{F}(\mathbf{r}, \tau) = \mathbf{S}_1(\mathbf{r}, \tau). \quad (4.31)$$

where $\mathbf{S}_1(\mathbf{r}, \tau) = \int_{4\pi} S(\mathbf{r}, \hat{s}, \tau) \hat{s} d\hat{s}$ is a source of correlation flux, and the term in $G(\mathbf{r}, \hat{s}', \tau)$ has reduced to zero under integration over \hat{s} .

Forming the diffusion approximation

Having derived the energy balance equation, and an equivalent of Fick's law, we may proceed to combine the resulting equations to form the diffusion approximation. Rearranging equation 4.31 in terms of the flux vector \mathbf{F} ,

$$\mathbf{F}(\mathbf{r}, \tau) = -D [\nabla G(\mathbf{r}, \tau) - 3S_1(\mathbf{r}, \tau)], \quad (4.32)$$

where we define the isotropic diffusion coefficient,

$$D = \frac{1}{3(\mu_a + \mu_s)}. \quad (4.33)$$

We complete our diffusion approximation by inserting this expression of the flux vector into the energy balance expression of equation 4.27 and assuming only isotropic sources such that $S_1(\mathbf{r}, \hat{s}) = 0$,

$$[-\nabla \cdot D \nabla + \mu_a + \mu_s b_2(\mathbf{r}, \tau)] G(\mathbf{r}, \tau) = S_0(\mathbf{r}, \tau). \quad (4.34)$$

By analogy with light transport we equip the diffusion approximation with a Robin (or impedance) boundary condition,

$$G(\mathbf{r}, \tau) + 2A\hat{\mathbf{n}} \cdot D \nabla G(\mathbf{r}, \tau) = 0, \quad (4.35)$$

where $\hat{\mathbf{n}}$ is the outward unit normal to the boundary, the term $A = (1+R)/(1-R)$ accounts for the potential mismatch of refractive indices at a boundary, and R is a reflection parameter which can be obtained empirically, or by Snell's laws [61]. In the case of a matched refractive index $A = 1$ and the Robin boundary condition provides a straightforward expression for the outward correlation flux on the boundary (our measurement) $G^\Gamma(\mathbf{r}, \tau) = G(\mathbf{r}, \tau)/2$, $\mathbf{r} \in \partial\Omega$.

To arrive at the expression stated by Sakadžić and Wang [118] we recall that our derivation made the assumption that $(k_a l_{\text{tr}}) \gg 1$. As such, terms in $b_2(\mathbf{r}, \tau)$ which contain $(k_a l_{\text{tr}})$ raised to lower than some power may be discounted without significant error. All terms with $(k_a l_{\text{tr}})$ to a power of zero or below were removed from the expression presented in the original work.

4.2.4 Correlation diffusion with pre-averaged phase

In the context of the standard diffusion approximation to light transport in scattering media, it is understood that the spatial distribution of light remains approximately constant under variation of the scattering coefficient μ_s and scattering anisotropy g , under the condition that the *transport* scattering coefficient $\mu'_s = \mu_s(1-g)$ remains constant. Sakadžić and Wang also demonstrated an approximate similarity relation for acousto-optic modulation [114]: the phase-increments in an anisotropically scattering medium can be approximated by replacing μ_s by μ'_s in equation 4.47. In this case, the average phase increment is calculated with a maximum error of circa 10% for values of the product $k_a l \approx 2$.

In equation 4.7 an expression was provided for the approximate phase perturbation over a given path length, averaged over both the spatial co-ordinate and the mean-free path. We might attempt to employ this similarity relation formally by replacing the scattering mean-free path with the transport mean free path and repeating the derivation. If we were to take this approach, we must consider that the transport mean free path $l_{\text{tr}} = l_s/(1-g)$ is defined in terms of the scattering anisotropy $g = \int_{4\pi} p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') \hat{\mathbf{s}} \cdot \hat{\mathbf{s}}' d\hat{\mathbf{s}}$ and thus a product of the phase function and acousto-optic phase increment must be integrated. Both the phase function and the acousto-optic phase increment individually define some closed surface with symmetry, but the angle of symmetry in the phase function (around the incident direction), and the phase increment (around the ultrasound propagation direction), are different, and hence the standard approaches taken to form the diffusion approximation can no longer be employed.

Even if the integrals were straightforward, this value of this approach is debatable since the phase perturbation expression is only approximate in the case of anisotropic scattering, and the presence of the phase function is only to scale the phase perturbation according to the mean cosine of the scattering, angle g . By performing the averaging procedure of section 4.2.2 we have already decoupled the accumulation of phase increments from the light transport model.

Angular pre-averaging

An alternative approach is to further average the phase increments over all solid angle prior to forming our diffusion approximation. That is to say that we will redefine the approximation of equation 4.7,

$$\begin{aligned}\Psi(\mathbf{r}_b, \mathbf{r}_a, \hat{\mathbf{s}}, t, \tau) &\approx 1 - (1-g) \frac{1}{4\pi} \int_{4\pi} \Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_N d\hat{\mathbf{s}} \\ &= 1 - (1-g) \frac{|\mathbf{r}_b - \mathbf{r}_a|}{l_s} \sum_{n=1}^N (-j)^n \frac{1}{4\pi} \int_{4\pi} \frac{\langle \Delta\varphi^n \rangle_{\mathbf{r}_c, l}}{n!} d\hat{\mathbf{s}},\end{aligned}\quad (4.36)$$

Whilst this process further reduces the accuracy of our approximation of the phase increments, this decoupling permits the derivation of higher order (in the spherical harmonic sense) light transport models, and the incorporation of standard phase function models. In the context of diffusing wave spectroscopy, this procedure was demonstrated by Ackerson, where he referred to the procedure as pre-averaging [2]. Under this modified approximation the transport equation of equation 4.15 becomes

$$[\hat{\mathbf{s}} \cdot \nabla + \mu_t] G(\mathbf{r}, \hat{\mathbf{s}}, \tau) = \mu_s [1 - (1-g)b(\mathbf{r}, \tau)_N] \int_{4\pi} p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') G(\mathbf{r}, \hat{\mathbf{s}}', \tau) d\hat{\mathbf{s}}' + S(\mathbf{r}, \hat{\mathbf{s}}, \tau),\quad (4.37)$$

where

$$b(\mathbf{r}, \tau)_N = \frac{1}{4\pi} \int_{4\pi} \Psi(\hat{\mathbf{s}}, \hat{\mathbf{s}}_a, \tau)_N d\hat{\mathbf{s}}\quad (4.38)$$

was calculated to the second order ($N = 2$) during our derivation of the conservation of energy expression (equations 4.22 through 4.25) in the preceding section.

Correlation diffusion with pre-averaged phase increments

Having averaged the angular dependence of the acousto-optic phase term, we may now proceed to re-derive the correlation diffusion approximation with an anisotropic phase function. Here we choose the oft-employed Henyey-Greenstein phase function [54] to describe the probability density function of the change in scattering angle between incoming and outgoing correlation radiance. This derivation is performed in the same way as the isotropic case. Under the modified approximation, the conservation of energy expression becomes

$$\nabla \cdot \mathbf{F}(\mathbf{r}, \tau) + [\mu_a + \mu'_s b(\mathbf{r}, \tau)_N] G(\mathbf{r}, \tau) = S_0(\mathbf{r}, \tau),\quad (4.39)$$

where the $(1-g)$ term causes replacement of μ_s with $\mu'_s = \mu_s(1-g)$. There are somewhat greater differences in the derivation of Fick's law. We recommence the derivation from equation 4.28 which now reads

$$\begin{aligned}\int_{4\pi} [\hat{\mathbf{s}} \cdot \nabla + \mu_t] G(\mathbf{r}, \hat{\mathbf{s}}, \tau) \hat{\mathbf{s}} d\hat{\mathbf{s}} \\ = \mu_s [1 - (1-g)b(\mathbf{r}, \tau)_N] \int_{4\pi} \left[\int_{4\pi} p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') G(\mathbf{r}, \hat{\mathbf{s}}', \tau) d\hat{\mathbf{s}}' + S(\mathbf{r}, \hat{\mathbf{s}}) \right] \hat{\mathbf{s}} d\hat{\mathbf{s}}.\end{aligned}\quad (4.40)$$

We insert our approximation of the correlation radiance from equation 4.17 before performing the integral on the left hand side, and of the source term, as per the isotropic case,

$$\begin{aligned} \frac{1}{3} \nabla G(\mathbf{r}, \tau) + \mu_t \mathbf{F}(\mathbf{r}, \tau) &= \mu_s [1 - (1 - g)h(\mathbf{r}, \tau)_N] \\ &\times \int_{4\pi} \left[\int_{4\pi} p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') \left[\frac{1}{4\pi} G(\mathbf{r}, \tau) + \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, \tau) \cdot \hat{\mathbf{s}}' \right] d\hat{\mathbf{s}}' \right] \hat{\mathbf{s}} d\hat{\mathbf{s}} + S_1(\mathbf{r}, \tau). \end{aligned} \quad (4.41)$$

We first consider the integral over $\hat{\mathbf{s}}'$: the first term can be completed according to our definition of correlation fluence; the second term is completed by recognising that due to the phase function, only the component of the flux vector along $\hat{\mathbf{s}}$ provides a contribution to the integral. For a diagrammatic explanation, see [94]. We align our co-ordinate system to the incident scattering direction to simplify the resulting integral,

$$\begin{aligned} \frac{1}{3} \nabla G(\mathbf{r}, \tau) + \mu_t \mathbf{F}(\mathbf{r}, \tau) &= \mu_s [1 - (1 - g)h(\mathbf{r}, \tau)_N] \\ &\times \int_{4\pi} \left[\frac{1}{4\pi} G(\mathbf{r}, \tau) + \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, \tau) \cdot \hat{\mathbf{s}} \int_{4\pi} p(\theta) \cos(\theta) d\hat{\mathbf{s}} \right] \hat{\mathbf{s}} d\hat{\mathbf{s}} + S_1(\mathbf{r}, \tau), \end{aligned} \quad (4.42)$$

where θ is the angle between $\hat{\mathbf{s}}$ and $\hat{\mathbf{s}}'$. We previously described the scattering anisotropy factor g as being defined as the mean cosine of the scattering angle, hence,

$$\begin{aligned} \frac{1}{3} \nabla G(\mathbf{r}, \tau) + \mu_t \mathbf{F}(\mathbf{r}, \tau) &= \mu_s [1 - (1 - g)h(\mathbf{r}, \tau)_N] \\ &\times \int_{4\pi} \left[\frac{1}{4\pi} G(\mathbf{r}, \tau) + \frac{3g}{4\pi} \mathbf{F}(\mathbf{r}, \tau) \cdot \hat{\mathbf{s}} \right] \hat{\mathbf{s}} d\hat{\mathbf{s}} + S_1(\mathbf{r}, \tau). \end{aligned} \quad (4.43)$$

The integral of $G(\mathbf{r}, \tau)$ goes to zero under the integration in product with $\hat{\mathbf{s}}$, and using equation 4.30 we can complete the integration of the flux vector,

$$\frac{1}{3} \nabla G(\mathbf{r}, \tau) + \mu_t \mathbf{F}(\mathbf{r}, \tau) = \mu_s g [1 - (1 - g)h(\mathbf{r}, \tau)_N] \mathbf{F}(\mathbf{r}, \tau) + S_1(\mathbf{r}, \tau). \quad (4.44)$$

We may now proceed to form the diffusion approximation as per the anisotropic case of section 4.2.3. Rearranging equation 4.44 in terms of the flux vector, recalling that $\mu'_s = \mu_s(1 - g)$,

$$\mathbf{F}(\mathbf{r}, \tau) = -D [\nabla G(\mathbf{r}, \tau) - 3S_1(\mathbf{r}, \tau)], \quad (4.45)$$

where the anisotropic diffusion coefficient is now

$$D = \frac{1}{3 [\mu_a + \mu'_s [1 + g h(\mathbf{r}, \tau)_N]]}. \quad (4.46)$$

In the case of isotropic scattering with $g = 0$ this expression collapses to the previously derived diffusion coefficient. Inserting the definition of the flux vector into the anisotropic energy balance expression, and assuming isotropic sources, we arrive at the final form of the diffusion approximation,

$$[-\nabla \cdot D \nabla + \mu_a + \mu'_s h_N(\mathbf{r}, \tau)] G(\mathbf{r}, \tau) = S_0(\mathbf{r}, \tau). \quad (4.47)$$

In practice, since $h(\mathbf{r}, \tau)_N \ll 1$, the standard diffusion coefficients of $D = 3\mu_{tr}^{-1}$ or $D = 3\mu_s'^{-1}$ may be employed with a negligible error.

4.2.5 Simplified spherical harmonic approximation

Much work performed in the context of nuclear reactors has come to be used in the modelling of diffuse optics. This is particularly true in Monte-Carlo applications, but diffusion theory, and extensions thereof, are of common interest. One of the tools to have been developed in this context was the *simplified* spherical harmonic method first explored by Gelbard [46]. Here we will apply this method to the correlation transport problem, employing the angularly pre-averaged expression developed in the preceding section.

When we wish to consider problems in optical tomography close to sources, or boundaries, significant errors arise in the diffusion approximation. One approach to provide improved results is to rely upon a higher order spherical harmonic (P_N) approximation. Unfortunately the number of equations to be solved grows with the square of the order of the spherical harmonic approximation. When such systems are subsequently solved using numerical techniques, the size of the discretised problem grows similarly. Attempts have been made to mitigate this problem via hybrid and variable-order schemes [129, 132]. In the original derivation of the simplified spherical harmonic technique (SP_N), the *one-dimensional* transport problem is approximated by the P_N method to a given order, and the derivatives of the resulting equations are replaced by their multi-dimensional counterparts in a somewhat *ad-hoc* manner. By this approach, the set of equations to be solved scales linearly with the order of the approximation. Whilst this method was demonstrated to provide improved accuracy over the standard diffusion approach, it did not find significant application in diffuse optics when contrasted with the standard diffusion approximation. Arguably, this was due to the intuitive replacement of the one-dimensional derivatives with their multi-dimensional counterparts, a process which lacks the rigour of the full P_N approximations.

More recently the SP_N approach was solved by an asymptotic derivation in both the diffusion limit [69–71] and more generally in the context of a transport approximation [43, 68, 104]. This approach yielded the same set of equations as the previous ad-hoc method, thus validating the approach. The older method still retains some advantages, notably in the simplicity by which suitable boundary conditions can be found.

The first formal application of the SP_N method to light transport in tissue was provided by Klose and Larsen [65]. In this work, Klose and Larsen replace the neutron scattering function with the Henyey-Greenstein model, suitable for use in diffuse optics. The authors derive the simplified spherical harmonic approximation to the seventh order according to the original ad-hoc method, including the partially reflective boundary conditions required for diffuse optics. The resultant model is found to provide a significant improvement in accuracy compared to the diffusion approximation, and requires only 2.5 times the computational effort: the authors contrast this with a discrete ordinates (S_N) [131] solution which provides

similar results but at the cost of three-orders of magnitude increased computational effort over the diffusion model.

Subsequent to the work of Klose, Chu and Dehghani have demonstrated demonstrated the reconstruction of diffuse optical tomography images using the SP_N model [30].

The plane-parallel correlation transport equation

The plane-parallel geometry under consideration is an infinite slab in which optical properties are constant in the x and y directions such that the correlation radiance $G(\mathbf{r}, \hat{\mathbf{s}}, \tau)$ is dependent only upon θ and z . We assume that a collimated source of correlation radiance is incident perpendicularly to the surface of the slab. We denote $\mu = \cos \theta$ and thus the differential solid angle element $d\hat{\mathbf{s}} = \sin \theta d\theta d\phi = -d\mu d\phi$. Rewriting the correlation transport equation of equation 4.37 in this geometry,

$$\begin{aligned} & \left[\mu \frac{d}{dz} + \mu_t \right] G(z, \mu, \phi, \tau) \\ &= \mu_s [1 - (1-g)b(\mathbf{r}, \tau)] \int_{-1}^1 \int_0^{2\pi} p(\mu, \phi, \mu', \phi') G(z, \mu', \phi', \tau) d\phi' d\mu' + S(z, \mu, \phi, \tau). \end{aligned} \quad (4.48)$$

We proceed by completing the integration of equation 4.48 over ϕ ,

$$\begin{aligned} & \left[\mu \frac{d}{dz} + \mu_t \right] G(z, \mu, \tau) \\ &= \mu_s [1 - (1-g)b(\mathbf{r}, \tau)] \int_{-1}^1 p_0(\mu, \mu') G(z, \mu', \tau) d\mu' + 2\pi S(z, \mu, \tau), \end{aligned} \quad (4.49)$$

where

$$p_0(\mu, \mu') = \int_0^{2\pi} \int_0^{2\pi} p(\mu, \phi, \mu', \phi') d\phi' d\phi. \quad (4.50)$$

The angle between (μ, ϕ) and (μ', ϕ') is denoted γ and

$$\cos \gamma = \cos \theta \cos \theta' + \sin \theta \sin \theta' \cos(\phi - \phi').$$

If the phase function is symmetric about the incident wave, and thus is a function of only γ , it can be expanded in a series of Legendre polynomials,

$$p(\mu, \phi, \mu', \phi') = p(\cos \gamma) = \sum_{n=0}^{\infty} W_n P_n(\cos \gamma). \quad (4.51)$$

The term $P_n(\cos \gamma)$ is expanded by a summation formula [1] and integrated over ϕ and ϕ' ,

$$p_0(\mu, \mu') = \sum_{n=0}^{\infty} W_n P_n(\mu) P_n(\mu'). \quad (4.52)$$

If we employ the Henyey-Greenstein phase function, as per the previous derivation, the weights W_n are given by

$$W_n = \frac{2n+1}{2} g^n. \quad (4.53)$$

The plane-parallel P_N approximation

We define the Legendre moments of the correlation radiance,

$$G_n(z, \tau) = \int_{-1}^1 P_n(\mu) G(z, \mu, \tau) d\mu, \quad (4.54)$$

multiply the azimuthally integrated planar CTE (equation 4.49) by $P_n(\mu)$ and integrate over μ ,

$$\begin{aligned} & \left[\mu \frac{d}{dz} + \mu_t \right] G_n(z, \tau) \\ &= \mu_s [1 - (1-g)b(\mathbf{r}, \tau)] \int_{-1}^1 P_n(\mu) \int_{-1}^1 p_0(\mu, \mu') G(z, \mu', \tau) d\mu' d\mu + \delta_n 4\pi S(z, \mu, \tau), \end{aligned} \quad (4.55)$$

where the Kronecker delta limits contributions from the source to the zeroth Legendre polynomial. Into this expression we substitute the Legendre expansion of the phase function from equation 4.52,

$$\begin{aligned} & \left[\mu \frac{d}{dz} + \mu_t \right] G_n(z, \tau) = \mu_s [1 - (1-g)b(\mathbf{r}, \tau)] \frac{2m+1}{2} g^m \int_{-1}^1 P_n(\mu) P_m(\mu) \\ & \quad \times \int_{-1}^1 P_m(\mu') G(z, \mu', \tau) d\mu' d\mu + \delta_n 4\pi S(z, \mu, \tau). \end{aligned} \quad (4.56)$$

The integral over μ' can be replaced with the expression of the Legendre moments of correlation radiance from 4.54,

$$\begin{aligned} & \left[\mu \frac{d}{dz} + \mu_t + \right] G_n(z, \tau) \\ &= \mu_s [1 - (1-g)b(\mathbf{r}, \tau)] G_m(z, \tau) \frac{2m+1}{2} g^m \int_{-1}^1 P_n(\mu) P_m(\mu) d\mu + \delta_n 4\pi S(z, \mu, \tau). \end{aligned} \quad (4.57)$$

The Legendre polynomials have an orthogonality relationship,

$$\int_{-1}^1 P_m(\mu) P_n(\mu) d\mu = \frac{2}{2n+1} \delta_{m,n} \quad (4.58)$$

which is used to simplify the final integration over μ , and to redefine the indices of the implied summations such that $m = n$ as the Kronecker delta eliminates all cross terms,

$$\left[\mu \frac{d}{dz} + \mu_t \right] G_n(z, \tau) = \mu_s [1 - (1-g)b(\mathbf{r}, \tau)] G_n(z, \tau) g^n + \delta_n 4\pi S(z, \mu, \tau). \quad (4.59)$$

Rearranging, we get

$$\left[\mu \frac{d}{dz} + \mu_t - \mu_s g^n [1 - (1-g)b(\mathbf{r}, \tau)] \right] G_n(z, \tau) = \delta_n 4\pi S(z, \mu, \tau). \quad (4.60)$$

We seek to develop this expression using the recurrence relation for the Legendre polynomials,

$$\mu P_n(\mu) = \frac{n+1}{2n+1} P_{n+1}(\mu) + \frac{n}{2n+1} P_{n-1}(\mu), \quad (4.61)$$

which is applied to the first term of the operator in square brackets in equation 4.60,

$$\begin{aligned} \frac{d}{dz} \left[\frac{n+1}{2n+1} G_{n+1} + \frac{n}{2n+1} G_{n-1} \right] \\ + [\mu_t - \mu_s g^n [1 - (1-g)b(\mathbf{r}, \tau)]] G_n(z, \tau) = \delta_n 4\pi S(z, \mu, \tau). \end{aligned} \quad (4.62)$$

To proceed we convert this system of first order differential equations into a series of second order equations by eliminating the odd order moments of the correlation radiance. Rearranging in terms of $G_n(z, \tau)$ for n odd,

$$G_n(z, \tau) = -\frac{1}{\mu_{c,n}(z, \tau)} \frac{d}{dz} \left[\frac{n+1}{2n+1} G_{n+1} + \frac{n}{2n+1} G_{n-1} \right], \quad (4.63)$$

and inserting this expression into equation 4.62 for the even order moments,

$$\begin{aligned} -\frac{n+1}{2n+1} \frac{d}{dz} \frac{1}{\mu_{c,n+1}} \frac{d}{dz} \left[\frac{n+2}{2n+3} G_{n+2}(z, \tau) + \frac{n+1}{2n+3} G_n(z, \tau) \right] \\ -\frac{n}{2n+1} \frac{d}{dz} \frac{1}{\mu_{c,n-1}} \frac{d}{dz} \left[\frac{n}{2n-1} G_n(z, \tau) + \frac{n-1}{2n-1} G_{n-2}(z, \tau) \right] \\ + \mu_{c,n} G_n(z, \tau) = \delta_n 4\pi S(z, \mu, \tau). \end{aligned} \quad (4.64)$$

In the latter expressions, adopting the nomenclature of Klose [65], the n^{th} order *correlation* absorption coefficient is given by,

$$\mu_{c,n} = \mu_t - \mu_s g^n [1 - (1-g)b(\mathbf{r}, \tau)].$$

It should be noted that this series of equations is still equivalent to the original planar correlation transport equation — no approximation of the underlying transport equation has yet taken place.

The SP_N approximation

We arrive at the three-dimensional SP_N approximations by the replacement of the derivative operators with their multi-dimensional counterparts. Accordingly equation 4.64 becomes,

$$\begin{aligned} -\frac{n+1}{2n+1} \nabla \cdot \frac{1}{\mu_{c,n+1}} \nabla \left[\frac{n+2}{2n+3} G_{n+2}(\mathbf{r}, \tau) + \frac{n+1}{2n+3} G_n(\mathbf{r}, \tau) \right] \\ -\frac{n}{2n+1} \nabla \cdot \frac{1}{\mu_{c,n-1}} \nabla \left[\frac{n}{2n-1} G_n(\mathbf{r}, \tau) + \frac{n-1}{2n-1} G_{n-2}(\mathbf{r}, \tau) \right] \\ + \mu_{c,n} G_n(\mathbf{r}, \tau) = \delta_n 4\pi S(\mathbf{r}, \mu, \tau). \end{aligned} \quad (4.65)$$

We will now find the first two coupled moments in this series which will be used to form the specific approximations of a given order. For $n = 0$ in equation 4.65,

$$-\nabla \cdot \frac{1}{\mu_{c,1}} \nabla \left[\frac{2}{3} G_2(\mathbf{r}, \tau) + \frac{1}{3} G_0(\mathbf{r}, \tau) \right] + \mu_{c,0} G_0(\mathbf{r}, \tau) = 4\pi S(\mathbf{r}, \tau). \quad (4.66)$$

For $n = 2$ in equation 4.65, we obtain

$$\begin{aligned} -\frac{3}{5} \nabla \cdot \frac{1}{\mu_{c,3}} \nabla \left[\frac{4}{7} G_4(\mathbf{r}, \tau) + \frac{3}{7} G_2(\mathbf{r}, \tau) \right] \\ -\frac{2}{5} \nabla \cdot \frac{1}{\mu_{c,1}} \nabla \left[\frac{2}{3} G_2(\mathbf{r}, \tau) + \frac{1}{3} G_0(\mathbf{r}, \tau) \right] + \mu_{c,2} G_2(\mathbf{r}, \tau) = 0. \end{aligned} \quad (4.67)$$

We derive the SP_1 approximation by setting all contributions to the correlation radiance $G_n(\mathbf{r}, \tau) = 0$ for $n > 0$,

$$-\nabla \cdot \frac{1}{3\mu_{c,1}} \nabla G_0(\mathbf{r}, \tau) + \mu_{c,0} G_0(\mathbf{r}, \tau) = 4\pi S(\mathbf{r}, \tau). \quad (4.68)$$

From equation 4.2.5 we find $\mu_{c,0} = \mu_a + \mu'_s h(\mathbf{r}, \tau)_N$, and $\mu_{c,1} = \mu_a + \mu'_s [1 + g h(\mathbf{r}, \tau)]$, ergo the SP_1 approximation is identical to the anisotropic diffusion approximation presented in equation 4.47. In turn, with $g = 0$ (and under the approximations regarding (k_a, l_{tr}) stated earlier) the expression is equivalent to the original result of Sakadžić and Wang [118].

The SP_3 approximation is found by discounting all contributions to the correlation radiance for $n > 3$ which generates the following coupled expressions:

$$-\nabla \cdot \frac{1}{3\mu_{c,1}} [G_0(\mathbf{r}, \tau) + 2G_2(\mathbf{r}, \tau)] + \mu_{c,0} G_0(\mathbf{r}, \tau) = 4\pi S(\mathbf{r}, \tau), \quad (4.69a)$$

$$-\nabla \cdot \frac{9}{35} \nabla G_2(\mathbf{r}, \tau) + \mu_{c,2} G_2(\mathbf{r}, \tau) = \frac{2}{5} [\mu_{c,0} G_0(\mathbf{r}, \tau) - S(\mathbf{r}, \tau)], \quad (4.69b)$$

where the second expression is found by substitution of the first into the term for $G_0(\mathbf{r}, \tau)$ in equation 4.67.

Boundary conditions

One of the advantages of deriving the SP_N equations by extrapolating the three-dimensional form from the planar transport form, is that the boundary conditions can be determined by the same technique: that is to say that the planar boundary conditions are developed and extrapolated to their three dimensional form.

We will not repeat this process here, but the derivation for the case of diffuse optics can be found in [65]. If we assume that our acoustic field is zero on the boundary, then exactly those conditions found in the aforementioned work are valid, otherwise the n^{th} order *correlation* absorption coefficients must be used in place of the standard coefficients.

4.2.6 Finite element implementation

To compare each of the models developed in this chapter with our reference Monte-Carlo implementation, we solve the resultant partial differential equations in a domain of interest. A direct analytical solution is impractical for all but the most simple of geometries and acoustic field configurations, and thus we approach this problem numerically.

In this work, we solve each of the equations by the finite element method. It is beyond the scope of this work to provide a detailed treatment of the method, so only a brief overview will be provided for the case of the isotropic correlation diffusion equation. The anisotropic case is identical, save for modified coefficients. Solution of the equations of the SP_N approximation are slightly more involved, since we must solve for two unknowns, but the principles are the same. For details regarding the solution of the SP_N approximations by the finite element method, see [30]. In any case, the application of the finite element method to the solution of elliptic partial differential equations (or coupled systems thereof) is well detailed in the literature, for particular examples in the case of diffuse optics, see [8, 13, 53].

To solve the diffusion approximation, equation 4.47 is multiplied by a test function which obeys the boundary conditions, and whose zeroth and first derivatives are integrable over the domain. The boundary conditions are incorporated by the subsequent integration by parts. The domain is subdivided into a mesh of non-overlapping elements joined at N vertex nodes. On this mesh we define a set of piecewise linear basis functions such that $u_i(\mathbf{r}_j) = \delta_{ij}$ for $i, j = 1, \dots, N$ where \mathbf{r}_j located at the j^{th} vertex node. We subsequently approximate the solution $G(\mathbf{r}, \tau) \approx \sum_j^N u_j(\mathbf{r}) G_j(\tau)$. Selecting the basis functions in the weak formulation to be the same as the mesh basis allows us to write the resulting linear system of equations:

$$\mathbf{A}(\tau)[x]G(\tau) = \mathbf{q}, \quad (4.70)$$

where \mathbf{A} is the finite element system matrix. We express the parameters of the forward model, $x(\mathbf{r})$, the absorption-like decorrelation function $b(\mathbf{r})$, and the diffusion coefficient D using the same basis functions such that, for example, $\mu_a(\mathbf{r}) = \sum_k u_{\mu,k}(\mathbf{r}) \mu_{a,k}$. Consequently,

$$\begin{aligned} \mathbf{A}(\tau)[x]_{ij} = & \int_{\Omega} \sum_k \left[u_{D,k}(\mathbf{r}) D_k \nabla u_i(\mathbf{r}) \cdot \nabla u_j(\mathbf{r}) \right. \\ & + \left. \left[u_{\mu,k}(\mathbf{r}) \mu_{a,k} + u_{b,k}(\mathbf{r}) b_k(\tau) \right] u_i(\mathbf{r}) u_j(\mathbf{r}) \right] d^n \mathbf{r} \\ & + \frac{1}{2A} \int_{\partial\Omega} u_i(\mathbf{r}) u_j(\mathbf{r}) d^{n-1} \mathbf{r}, \end{aligned} \quad (4.71)$$

and,

$$\mathbf{q}_j = \int_{\Omega} q_j u_j(\mathbf{r}) d^n \mathbf{r}. \quad (4.72)$$

Solution of the matrix equation 4.70 yields the correlation fluence throughout the domain. To determine the measurable exiting flux on the boundary we employ the boundary condition of equation 4.35. In discrete form, with a matched boundary, the exiting flux at

some point \mathbf{r}_d is therefore given by application of a linear operator $\mathbf{D}_j^T = 1/2\mu_j(\mathbf{r}_d)$ to the solution vector $\mathbf{G}(\tau)$.

4.3 Results

4.3.1 Simplified spherical harmonic approximation

To assess the capabilities of the SP_3 correlation approximation developed in this chapter, we will compare the results of a simulation with the original P_1 method proposed by Sakadžić and Wang [118], and the reference Monte-Carlo code developed in chapter 3.

For this study we create a cuboid simulation domain of dimensions $(30, 30, 20)\text{mm}$ with optical properties of $\mu_a = 0.01\text{mm}^{-1}$, $\mu_s = 20\text{mm}^{-1}$, an anisotropy factor of $g = 0.9$, and hence a reduced scattering coefficient of $\mu'_s = 2\text{mm}^{-1}$. Inside this domain a small cuboid region of absorption coefficient $\mu_a = 1\text{mm}^{-1}$, and with dimensions $(2, 2, 12)\text{mm}$ is located with its centre at $(13, 9, 10)\text{mm}$. The medium has density $\rho = 10^3\text{kg m}^{-3}$, and the whole space is assigned a refractive index of $n = 1.33$. A pencil beam of coherent light of wavelength 532nm is incident perpendicular to the slab at $(15, 0, 10)\text{mm}$ travelling in the positive y direction: this is implemented directly in the Monte-Carlo simulations, and in the correlation diffusion simulations this was approximated by a point-source of coherent light at a distance of one transport mean free path under the surface at $(1/\mu'_s)$, as is standard practice in the diffusion approximation. The medium is insonified by a plane-wave ultrasound beam with a Gaussian amplitude profile of full-width half-maximum of 4.7mm , travelling in the positive x axis with its centre located at $(9, 10)\text{mm}$ in the $y-z$ plane. The peak amplitude of the beam is 200kPa .

In the Monte-Carlo simulation, 3×10^9 photon packets were simulated and the fluences at the zeroth and first harmonics were directly recorded on the mesh. In the finite element simulations, eleven runs of the forward model were executed for both the P_1 and SP_3 approximations at evenly spaced points in lag, τ , over a single acoustic cycle. The un-normalised field autocorrelation coefficient was recorded at each node of the mesh, and this was converted to the fluence values at the zeroth and first harmonic by application of a discrete cosine transform. In both cases the first harmonic modulation depth was then found by taking the ratio of the first to the zeroth harmonic of the correlation fluence. The resultant fields were evaluated in their linear-basis over a regular grid extending throughout the $x-y$ plane, at a height of $z = 10 \times 10^{-3}$. We will begin by inspecting the fluence distributions in each of the individual harmonics.

Figures 4.1 through 4.3 show the correlation fluence distribution at the zeroth harmonic for the Monte-Carlo simulation, the P_1 approximation and the SP_3 approximation, each displayed to the same colour scale¹. Dashed lines indicate the extent of the ultrasound field.

¹The striations visible in the images are a graphical artefact, and whilst of a similar scale to the ultrasound wavelength are not an effect of the insonification.

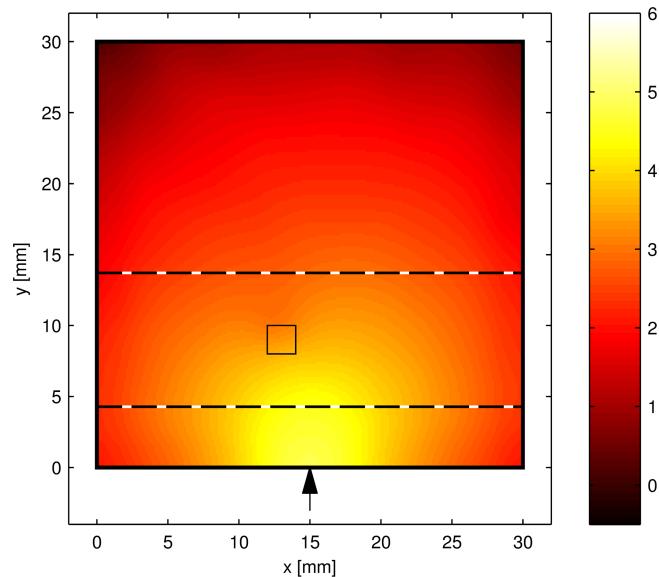


Figure 4.1: Zeroth harmonic correlation fluence ($\log_{10} G_0(\mathbf{r})$), as calculated by the reference Monte-Carlo simulation.

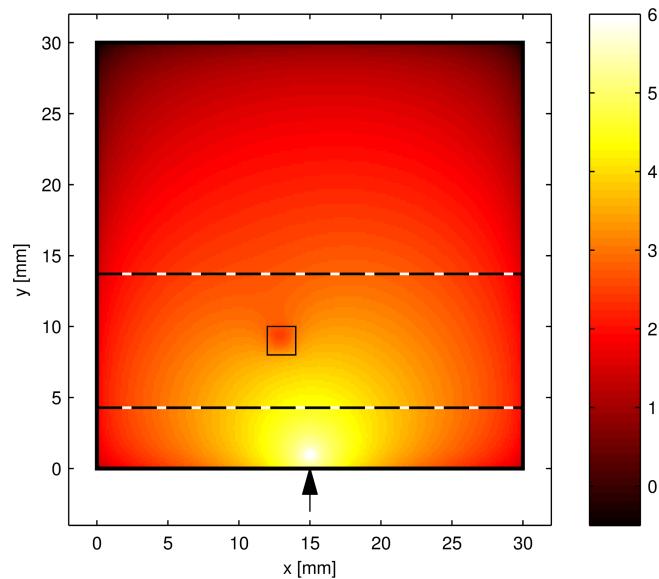


Figure 4.2: Zeroth harmonic correlation fluence ($\log_{10} G_0(\mathbf{r})$), as calculated by the standard correlation diffusion approximation, equivalent to first order simplified spherical harmonic method.

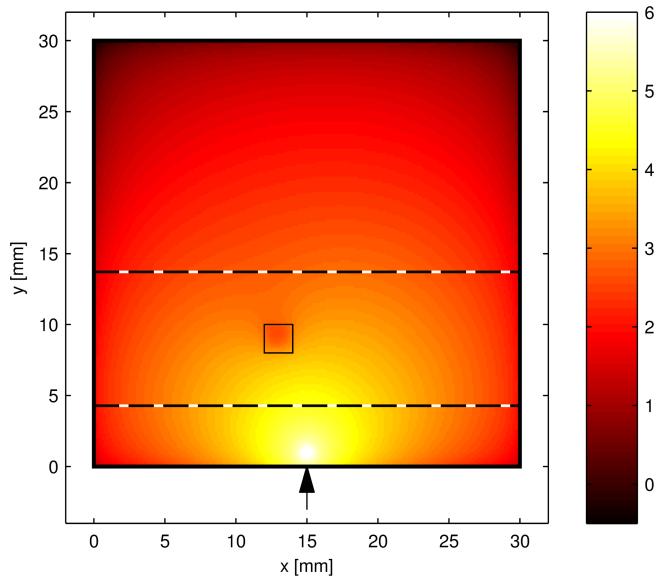


Figure 4.3: Zeroth harmonic correlation fluence ($\log_{10} G_0(\mathbf{r})$), as calculated by the third order simplified spherical harmonic (SP_3) approximation to the correlation transport equation.

The most noticeable difference between the two simulations is seen around the source region. This is unsurprising given the approximation of a collimated source used in each of the spherical harmonic approximations: the point-source in each of figures 4.2 and 4.3 is visible in the diagrams as the region of greatest correlation fluence. The P_1 simulation demonstrates a slightly reduced correlation fluence at areas far from the source, when compared to the Monte-Carlo and SP_3 simulations.

Figures 4.4 through 4.6 show the correlation fluence distribution at the first harmonic for each of the simulations. These figures demonstrate more significant differences than were seen at the zeroth harmonic. It is immediately evident that the SP_1 approximation significantly over-estimates the first harmonic correlation fluence, when compared with both the reference Monte-Carlo simulation, and the SP_3 approximation. This explains the reduced zeroth order correlation fluence in the P_1 approximation: the energy in the fluence in the first harmonic is sourced from the input energy at the zeroth harmonic. Small contour perturbations in the first harmonic correlation fluence are seen in both of the spherical harmonic approximations; a greater reduction is seen in the SP_3 approximation which we expect to provide improved performance at higher absorption values. This reduction of correlation fluence is not seen in the Monte-Carlo result due to the spatial averaging required to smooth the solution.

Figures 4.7 through 4.9 show the modulation depth distribution for each of the simulations. The first figure indicates the reference solution. Close to the source position, the mod-

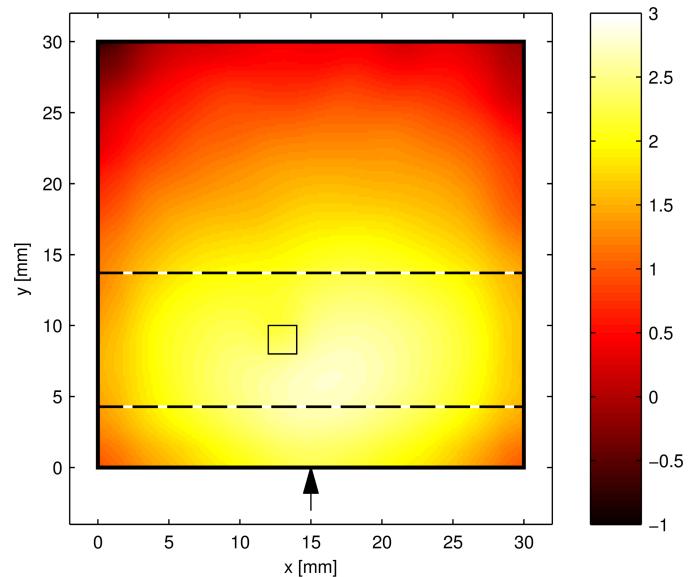


Figure 4.4: First harmonic correlation fluence ($\log_{10} G_1(\mathbf{r})$), as calculated by the reference Monte-Carlo simulation.

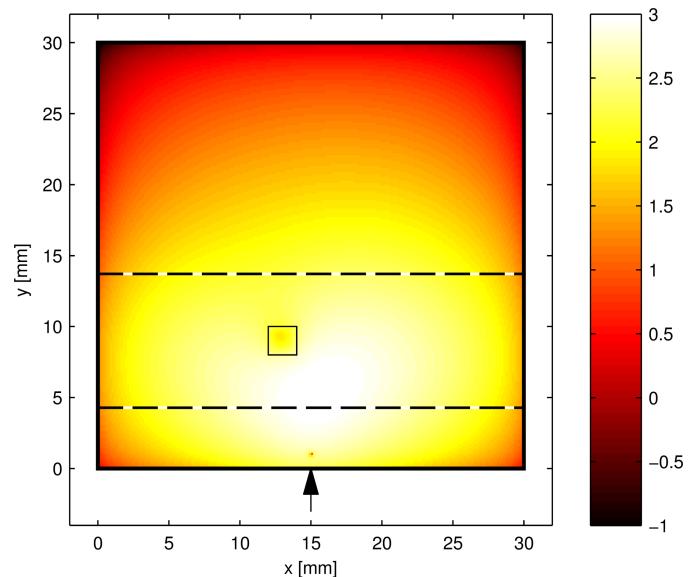


Figure 4.5: First harmonic correlation fluence ($\log_{10} G_1(\mathbf{r})$), as calculated by the standard correlation diffusion approximation, equivalent to first order simplified spherical harmonic method.

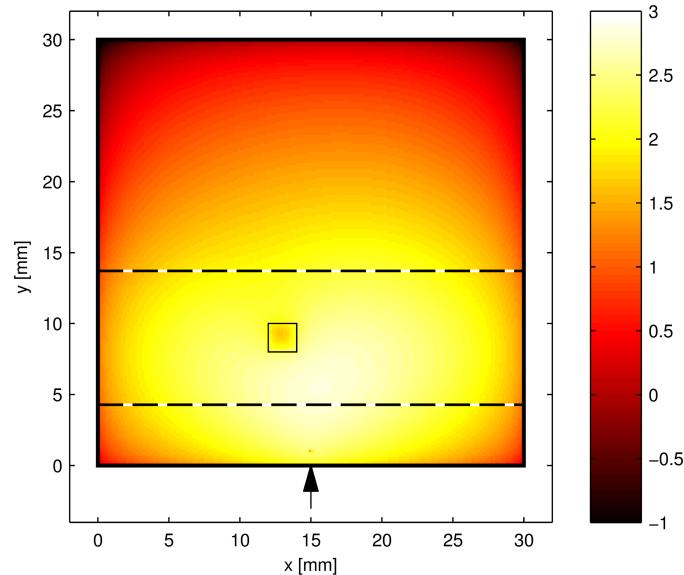


Figure 4.6: First harmonic correlation fluence ($\log_{10} G_1(\mathbf{r})$), as calculated by the third order simplified spheric harmonic (SP_3) approximation to the correlation transport equation.

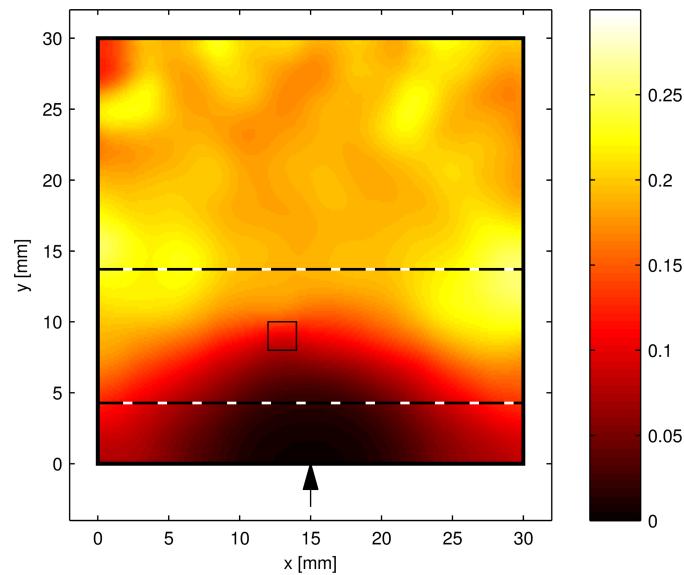


Figure 4.7: Internal first harmonic modulation depth distribution, $MD(\mathbf{r}) [-]$, as calculated by the reference Monte-Carlo simulation.

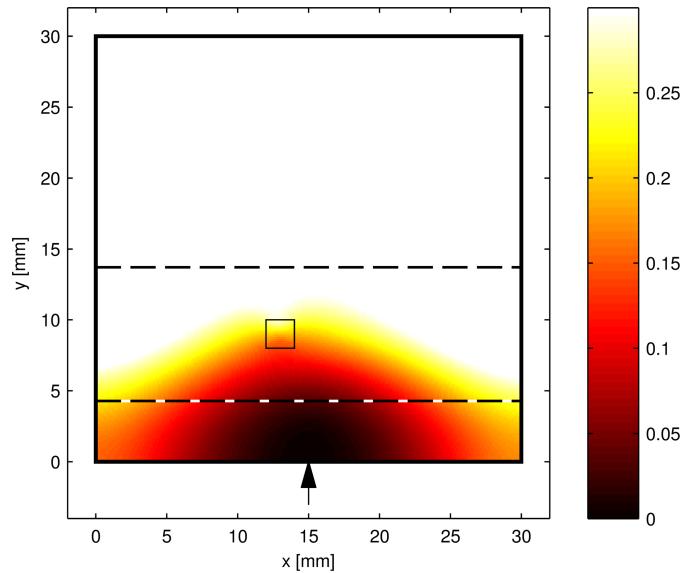


Figure 4.8: Internal first harmonic modulation depth distribution, $MD(\mathbf{r})$ [-], as calculated by the standard correlation diffusion approximation, equivalent to first order simplified spherical harmonic method. Note that the colour scale is limited to that of the reference solution.

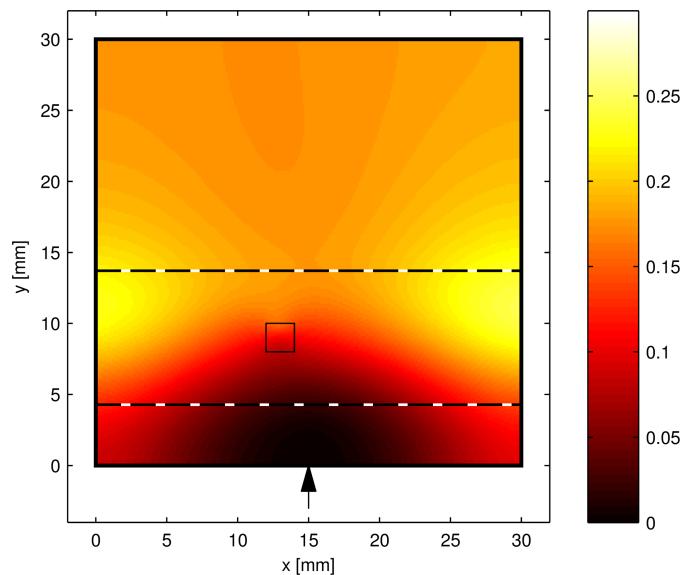


Figure 4.9: Internal first harmonic modulation depth distribution, $MD(\mathbf{r})$ [-], as calculated by the third order simplified spherical harmonic (SP_3) approximation to the correlation transport equation.

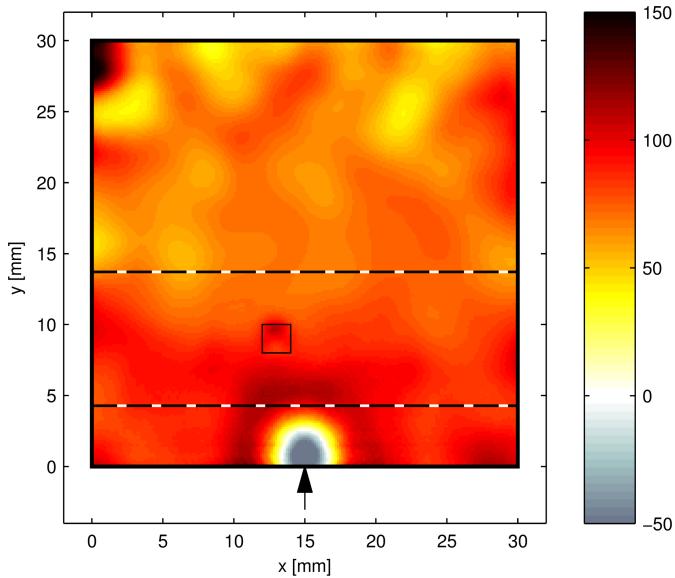


Figure 4.10: Percentage error modulation depth between P_1 correlation diffusion approximation and reference Monte-Carlo solution.

ulation depth is suppressed by the large amount of unmodulated light being directed into the medium. The modulation depth peaks in the directly insonified regions furthest from the source location. Away from both the source and the acoustic insonification the modulation depth stabilises and becomes more uniform — whilst all of the harmonics are being attenuated in these regions, this occurs proportionally such that the modulation depth remains constant. We see good agreement when comparing the SP_3 simulation of figure 4.9 with the reference Monte-Carlo simulation of figure 4.4. The ‘steady-state’ modulation depth of circa 0.2 matches up with the reference solution, as does the peak modulation at around 0.25. In figure 4.8 we find that the original P_1 simulation performs less well: the modulation depth through the domain is over-estimated when compared with the reference solution. Note that the large white region in the figure is due to limiting of the colour scale to the reference solution. We will compare these errors more carefully in the following figures which plot the percentage difference between the finite-element and Monte-Carlo simulations.

In figure 4.10 the over-estimation of the correlation diffusion approximation is evident. An expected error is seen near to the source location, but deep in the medium where the correlation diffusion equation should provide a reasonable approximation, the modulation depth is calculated with approximately an 80% error. Furthermore, we note that an increased error is found close to the region of high absorption. In figure 4.11 the SP_3 approximation is shown to provide a solution much closer to the reference. In the region of insonification a maximum error of approximately 20% is found. Deeper into the medium the modulation depth is under-estimated by this approximation, by around 30%.

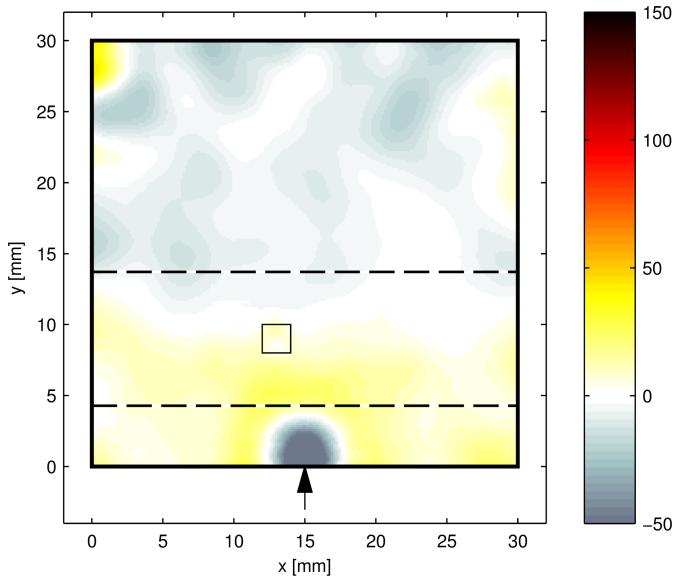


Figure 4.11: Percentage error modulation depth between SP_3 approximation and reference Monte-Carlo solution.

4.4 Conclusions

In this chapter we have provided a full derivation of a correlation diffusion approximation originally proposed by Sakadžić and Wang [118]. This equation has been derived with generality such that it can be used to calculate higher order phase increments.

We have also taken an alternative approach in which the phase increments were averaged over angle. This has permitted the direct use of an anisotropic phase function in a standard diffusion approximation. This approach also permitted approximation by the simplified spherical harmonic method.

Significant errors were found in the anisotropic simulation using the diffusion (or first order simplified spherical harmonic) approximation. The third order (SP_3) approximation demonstrated good agreement with the Monte-Carlo solution, however. These results are of particular value: despite further approximation of the phase increment, the improved model can be solved without significant additional computational effort. This model could easily be extended to a higher order in terms of both the light transport (by inclusion of more terms in the SP_N series), and also of the acousto-optic phase increment (by inclusion of more terms in $h(\mathbf{r}, \tau)_N$).

Chapter 5

Linear image reconstruction in ultrasound modulated optical tomography

5.1 Introduction

The purpose of imaging and sensing techniques such as ultrasound-modulated optical tomography is the production of images or spatially defined quantitative measurements of some parameter of interest. In the field of optical biomedical imaging this is typically an image of the optical absorption coefficient within a diffuse medium at one or more wavelengths. The vastly different spatial resolutions of the acoustic and diffuse optical aspects of ultrasound-modulated optical tomography encourages various approaches to the generation of images or spatially localised measurements depending upon the nature of the optical configuration and ultrasound excitation.

The most basic approach is that of direct mapping; when the measurement inherently possesses sufficient and controllable spatial selectivity a measured datum may be assigned directly to the region of assumed sensitivity in an image. This approach is common in systems where the energy which is used to probe the system travels in straight lines without diffuse reflection from interfaces or multiple scattering by a participating medium. For example, an ultrasound image may be constructed from multiple A-lines, the signals from which are assumed to reflect the change in the acoustic properties of the medium perpendicular to the transducer face. A standard two-dimensional X-ray is another example in which the directivity of the probe is assumed in order that an image may be formed directly (this is common to most optical-imaging systems). In a sensing application the ‘image’ is essentially a single region, but the implicit assumption of spatial resolution remains. An example of a sensing application is the ubiquitous pulse oximeter; this instrument relies upon the differential absorption of two wavelengths of light which propagate diffusely within the tissue. A measurement is made with the implicit (and in this case reasonable) assumption that the light samples a region of tissue that is perfused with a representative sample of blood.

In each of these cases the measurement may be a proxy for the parameter; in the example of a standard two-dimensional X-ray the transformation for the measured data (the integrals of X-ray absorption along paths from the source and detector) to the relevant parameter (such as the presence of bone) is made by the person interpreting the X-ray. Furthermore, processing subsequent to acquisition of the measured data may be applied to transform the measurement into a more useful form; a set of two-dimensional X-rays may be transformed to produce cross sectional images by tomographic reconstruction.

In UOT the sensitivity of the system is defined approximately by the product of the optical sensitivity and acoustic amplitude in the system. Depending upon the application, the spatial resolution afforded by a focussed or time gated ultrasound field is arguably sufficient to delineate regions of biological interest. If we neglect acoustic absorption and assume a constant ultrasound excitation amplitude probes each region in a given experiment, a UOT image produced via direct mapping records the scaled optical sensitivity in the medium (rather than the optical absorption).

In transmission mode configurations of a UOT experiment with a large optical étendue (such as those employing photo-refractive [26, 102, 109], spectral-hole burning [82], CCD [77], digital holography [51], or speckle contrast [80] based detection mechanisms) the optical sensitivity through an homogeneous medium transverse to the optical axis is relatively constant when contrasted with the exponentially varying sensitivity parallel to the axis of light collection. Since optical absorbers embedded in a turbid region reduce the optical sensitivity of the system in their vicinity, a transverse scan of the ultrasound probe can provide useful images of the embedded absorbers. Few demonstrations have been provided of more complicated absorption patterns in which shadowing of the optical sensitivity may distort the image, or of systems where an image is generated by a scan parallel to the optical axis in a high étendue system.

In low étendue systems that employ point-like source and/or detection mechanisms, the optical sensitivity varies significantly throughout the image in all optical configurations and scan directions irrespective of the homogeneity of the medium. In the context of UOT this point was highlighted succinctly by the images produced by Lev and Sfez [75, 76]. In this case, and in that of a high étendue system with significant heterogeneity, the fundamental problem is that the optical sensitivity of the system is not uniform.

In the case of pure diffuse optical tomography (DOT), it is the very non-uniformity of the optical sensitivity which is used to generate images of optical absorption. In DOT it is recognised that the spatial resolution for a pair of point-like optical source and detector is unsuitable for direct image reconstruction, and instead a model-based inversion procedure is employed. In the most general sense, this approach employs a forward model to predict the spatial sensitivity of the measurement and this is used to drive the update of an image based upon the measured data. The DOT problem remains ill-posed, and a-priori information regarding the solution is typically imposed to constrain the reconstruction. A similar approach was taken in UOT by Bratchenia et al. [28].

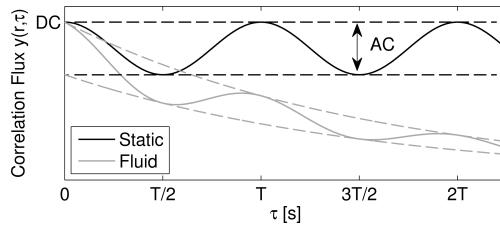


Figure 5.1: Illustration of the form of the measurement in autocorrelation UOT. The DC and AC measurement types are found from the non-decaying autocorrelation flux measured on the boundary of a static medium (black line). A fitting algorithm must be used in the case of measurements in a fluid exhibiting decorrelation to Brownian motion (grey line).

In this work we specifically consider the case of single fibre point-source and point-detector measurements of the field autocorrelation under monochromatic acoustic excitation. In chapters 3 and 4 we examined alternative forward models which describe the acousto-optic effect in ultrasound-modulated optical tomography. In this chapter we will apply the techniques typically employed in DOT to develop a basic and robust model-based image reconstruction approach, that of a linear difference-data reconstruction, to autocorrelation UOT.

5.2 Theory

5.2.1 Autocorrelation measurements

In an autocorrelation based UOT experiment we measure the lag (τ) dependent autocorrelation of the flux exiting the medium at a particular location (\mathbf{r}) on the boundary, $y(\mathbf{r}, \tau)$, $\mathbf{r} \in \delta\Omega$ in response to the insonification of the coherently illuminated medium. In practice, this value must be derived from the intensity autocorrelation function via the Seigert relation [19]. Neglecting an exponential decay in the correlation function due to Brownian motion, the data contain an oscillatory component at the acoustic frequency and its harmonics; this is depicted in figure 5.1. In this work we consider ultrasound pressures of amplitudes small enough to ensure that only the fundamental frequency is significant; more details of the assumptions employed are presented later when we consider the forward model in section 5.2.3.

We consider three measurement types:

1. $y(\mathbf{r})^{\text{DC}} = G_{\text{DC}}^{\Gamma}(\mathbf{r})$ is equal to the continuous wave DOT intensity;
2. $y(\mathbf{r})^{\text{AC}} = G_{\text{AC}}^{\Gamma}(\mathbf{r})$ is the AC correlation flux;
3. $m(\mathbf{r}) = G_{\text{AC}}^{\Gamma}(\mathbf{r})/G_{\text{DC}}^{\Gamma}(\mathbf{r})$ is the modulation depth flux, which is defined as the quotient of the previous two quantities.

5.2.2 Linear image reconstruction

In this work we will employ a linear (difference data) reconstruction technique. Whilst strictly limited to the recovery of qualitative information [25], its simple formulation allows us to investigate many basic aspects of the UOT inversion, including its sensitivity to noise, and the effects of the a-priori information required to circumvent the ill-posedness of the underlying problem.

Consider a turbid medium parameterised by its spatially varying absorption coefficient, $x(\mathbf{r}) = \mu_a(\mathbf{r})$. The measured data $y(\mathbf{r}, \tau)$ are found by application of the non-linear operator F ,

$$y(\mathbf{r}, \tau) = F(\tau)[x(\mathbf{r})]. \quad (5.1)$$

The operator F incorporates the physics of the problem under the given parameterisation, the boundary conditions, a given set of optical sources and detectors, and a particular ultrasound configuration.

In a linear reconstruction we typically consider the change in a measurement $\Delta y = y(\mathbf{r}, \tau) - y_0(\mathbf{r}, \tau)$ brought about by a perturbation in the parameters $\Delta x(\mathbf{r}) = x(\mathbf{r}) - x_0(\mathbf{r})$. Expanding equation 5.1 in a Taylor series around x_0 and dropping the spatial and lag dependent notation momentarily,

$$y = y_0 + F'[x_0](x - x_0) + \frac{1}{2}(x - x_0)^T F''[x_0](x - x_0) + \dots, \quad (5.2)$$

and then discounting the higher order terms and rearranging, we get:

$$\Delta x(\mathbf{r}) \approx (F(\tau)'[x_0])^{-1} \Delta y(\mathbf{r}, \tau). \quad (5.3)$$

F' is the first-order Fréchet derivative of the forward model [12]. In the case of the modulation depth measurement,

$$m(\mathbf{r}) = \frac{y(\mathbf{r})^{\text{AC}}}{y(\mathbf{r})^{\text{DC}}} = 1 - \frac{F(T/2)[x(\mathbf{r})]}{F(0)[x(\mathbf{r})]}. \quad (5.4)$$

Expanding this expression in a series around a baseline x_0 , discounting higher order terms and rearranging,

$$\Delta x(\mathbf{r}) = \left[\frac{y_0(\mathbf{r})_{\text{AC}} F(0)'[x_0] - y_0(\mathbf{r})_{\text{DC}} F(T/2)'[x_0]}{[y_0(\mathbf{r})_{\text{DC}}]^2} \right]^{-1} \Delta m(\mathbf{r}). \quad (5.5)$$

Our task is thus to develop an expression for F' : in this work we will develop the discrete form of this operator which is the Jacobian matrix, \mathbf{J} [14]. The inversion procedure of equation 5.3 (or its equivalent modulation depth formulation) is then given by,

$$\Delta x_k = (\mathbf{J}(\tau)[x_0]_{mk})^{-1} \Delta y_m. \quad (5.6)$$

5.2.3 Forward model

In chapter 4 we provided a full derivation for a diffusion style approximation to a correlation transport equation, originally proposed by Sakadžić and Wang [118]. In this chapter we employ the anisotropic correlation diffusion approximation with second order phase perturbations as our forward model. We restate the pertinent theory here, but refer the reader to the previous chapter for details of the limitations and assumptions made during the derivation. From equation 4.47 our model of the correlation fluence in the domain is given by,

$$[-\nabla \cdot D \nabla + \mu_a + \mu'_s b_2(\mathbf{r}, \tau)] G(\mathbf{r}, \tau) = S_0(\mathbf{r}, \tau). \quad (5.7)$$

where $b_2(\mathbf{r}, \tau)$, defined in chapter 4, is proportional to the square of the applied acoustic field pressure, $G(\mathbf{r}, \tau)$ is the correlation fluence in the medium, S_0 is an isotropic source term, $\mu'_s = \mu_s(1-g)$ is the reduced scattering coefficient, and g is the scattering anisotropy. $D = (3\mu'_s)^{-1}$ is the diffusion coefficient, which we have modified from the original form given in chapter 4 by exploiting the fact that the diffusion approximation is valid only for $\mu_s \gg \mu_a$. This further approximation will slightly simplify the image reconstruction procedure to be outlined in this chapter. For more details on the nature of the diffusion coefficient, a topic which has generated considerable debate, see e.g. [7, 17, 36, 37, 44, 48].

Sources, boundary conditions and detectors

A collimated source of coherent light is approximated by an isotropic point source located at a depth $1/\mu'_s$ below the incident surface. We employ a modified Robin boundary condition,

$$G(\mathbf{r}, \tau) + 2A\hat{\mathbf{n}} \cdot D \nabla G(\mathbf{r}, \tau) = 0, \quad \mathbf{r} \in \delta\Omega \quad (5.8)$$

where $\hat{\mathbf{n}}$ is the outward normal to the boundary at \mathbf{r} , and A depends upon the refractive index mismatch across the boundary [124]. The outgoing correlation flux is given by,

$$y(\mathbf{r}, \tau) = -\mathbf{n} \cdot D \nabla G(\mathbf{r}, \tau), \quad \mathbf{r} \in \delta\Omega. \quad (5.9)$$

In the case of an indexed matched boundary where $A = 1$ equations 5.8 and 5.9 may be combined to give

$$y(\mathbf{r}, \tau) = \frac{1}{2} G(\mathbf{r}, \tau), \quad \mathbf{r} \in \delta\Omega, A = 1. \quad (5.10)$$

Finite-element implementation

We will employ the same finite-element approach to the solution of the forward problem as was detailed in chapter 4. Having derived the weak form of the equations, and described the domain, we arrived in equation 4.70 at an expression for the correlation fluence throughout the domain,

$$\mathbf{A}(\tau)[x] \mathbf{G}(\tau) = \mathbf{q} \quad (5.11)$$

To make a measurement at a point \mathbf{r}_i equation 5.10 is implemented by a linear operator \mathbf{D}^T in which the appropriate elements of the row vector are initialised. $\mathbf{D}_j^T = 1/2u_j(\mathbf{r}_i)$. We may now express a single measurement of equation 5.1 in a discrete form,

$$\mathbf{y}(\tau)_m = \mathbf{F}(\tau)[x]_m = \mathbf{D}_m^T \mathbf{A}(\tau)[x]_m^{-1} \mathbf{q}_m. \quad (5.12)$$

The subscript m indicates that this expression considers a single UOT experiment involving one source location $\mathbf{r}_{j,m}$, one detector location $\mathbf{r}_{i,m}$, and a single ultrasound pressure distribution $p(\mathbf{r}, \tau)_m$. For a single measurement the linear inversion of equation 5.3 may thus be written,

$$\Delta \mathbf{x} \approx \left[(\mathbf{D}_m^T \mathbf{A}(\tau)[x_0]_m^{-1} \mathbf{q}_m)' \right]^{-1} \Delta \mathbf{y}(\tau)_m. \quad (5.13)$$

5.2.4 Correlation measurement density functions and the Jacobian

In equation 5.13, the term $(\mathbf{D}_m^T \mathbf{A}(\tau)[x_0]_m^{-1} \mathbf{q}_m)'$ represents the sensitivity of the measurement to a perturbation in the parameters of the forward model. In the context of diffuse optical tomography, Arridge generalised such sensitivity functions in the framework of *photon-measurement density functions* [9, 13]. This work followed earlier derivations of the measurement sensitivity to perturbations in absorption by Schotland et al. [122], who termed this the *photon hitting density*, and others [29, 42, 125].

We now explore the *correlation* measurement density functions (CMDFs) which arise in correlation based UOT. We assume that both the scattering coefficient μ_s and the acoustic field distribution $p(\mathbf{r}, \tau)$ do not change between measurement of the baseline and the perturbed state in our linear reconstruction. At higher acoustic frequencies where the model is valid (that is to say where $(k_a l_{\text{tr}}) \gg 1$), the term $b_2(\mathbf{r}, \tau)$ is only weakly dependent on the optical properties, and we neglect this dependence in the following analysis. Under these assumptions, and given that the coherent source term is independent of absorption, we take the derivative of the forward model, expanding the derivative of the inverse system matrix [1],

$$(\mathbf{D}_m^T \mathbf{A}(\tau)[x_0]_m^{-1} \mathbf{q}_m)' = -\mathbf{D}_m^T \mathbf{A}_m^{-1} (\mathbf{A}(\tau)[x_0]_m)' \mathbf{A}_m^{-1} \mathbf{q}_m. \quad (5.14)$$

The system is parameterised by only the absorption coefficient of the medium, and we thus take the derivative of the system matrix with respect to each basis coefficient $\mu_{a,k}$,

$$(\mathbf{A}(\tau)[x_0]_m)'_{\mu_{a,k}} = \int_{\Omega} u_{\mu,k}(\mathbf{r}) u_i(\mathbf{r}) u_j(\mathbf{r}) d\mathbf{r}. \quad (5.15)$$

We substitute this result back into equation 5.14 which provides an expression for the k^{th} element of the m^{th} CMDF. By definition, this also defines the k^{th} column of the m^{th} row of the Jacobian which we seek.

$$(\mathbf{J}(\tau)[x_0]_m^T)_k = -\mathbf{D}_m^T \mathbf{A}(\tau)[x_0]_m^{-1} (\mathbf{A}(\tau)[x_0]_m)'_{\mu_{a,k}} \mathbf{A}(\tau)[x_0]_m^{-1} \mathbf{q}_m. \quad (5.16)$$

To transform this expression into a more useful form, we exploit the reciprocity of the correlation diffusion equation. Arridge expresses this reciprocity in the context of the diffusion equation [10],

$$\mathbf{D}_m^T \mathbf{A}[x]^{-1} \mathbf{q}_m = \mathbf{q}_m^T \mathbf{A}[x]^{-1} \mathbf{D}_m. \quad (5.17)$$

This expression states that the correlation flux across a point on the boundary of the medium due to an internal source of correlation fluence is equal to the correlation fluence measured at the same point in the medium due to an *adjoint* source of correlation flux across the same point on the boundary of the medium, when this source is scaled according to the appropriate measurement operator, here \mathbf{D} [9]. In the discrete case this reciprocity manifests itself in the symmetry of the finite element system matrix. To proceed we apply the reciprocity expressed in equation 5.17 to the expression of the Jacobian in equation 5.16,

$$(\mathbf{J}(\tau)[x_0]_m^T)_k = - \left(\mathbf{A}(\tau)[x_0]_m^{-1} \mathbf{q}_m \right)^T \left(\mathbf{A}(\tau)[x_0]_m \right)'_{\mu_{a,k}} \mathbf{A}(\tau)[x_0]_m^{-1} \mathbf{D}_m. \quad (5.18)$$

where we have exploited the symmetry of the system matrix, its inverse and derivatives, and the properties of the transpose operation. Finally we denote the solution to the forward problem $\mathbf{G}(\tau)_{f,m} = \mathbf{A}(\tau)[x_0]_m^{-1} \mathbf{q}_m$, and the solution to the *adjoint* problem (in which the source term is the measurement operator) $\mathbf{G}(\tau)_{a,m} = \mathbf{A}(\tau)[x_0]_m^{-1} \mathbf{D}_m$. Thus,

$$(\mathbf{J}(\tau)[x]_m^T)_k = -\mathbf{G}(\tau)_{f,m}^T (\mathbf{A}(\tau)[x_0]_m)'_{\mu_{a,k}} \mathbf{G}(\tau)_{a,m}. \quad (5.19)$$

Here we see that the Jacobian for a perturbation of absorption in the k^{th} node can be found by taking the inner product of the correlation fluence in the domain due to the (standard) source term, with the product of the basis function derivative term and the correlation fluence in the domain due to the adjoint source. We may thus compute the correlation measurement density function for a given ultrasound configuration by two computations of the forward model.

The AC and modulation depth CMDFs and Jacobian

The CMDFs and Jacobian for the DC measurement type described in section 5.2.1 are given directly by equation 5.19. For the AC measurement type the appropriate expression is straightforward,

$$(\mathbf{J}^{\text{AC}}[x]_m^T)_k = (\mathbf{J}(0)[x]_m^T)_k - (\mathbf{J}(T/2)[x]_m^T)_k, \quad (5.20)$$

The modulation depth sensitivity functions are given by substituting equation 5.19 into equation 5.5, as follows:

$$(\mathbf{J}^{\text{MD}}[x]_m^T)_k = \frac{y_{0,m}^{\text{AC}} (\mathbf{J}(0)[x]_m^T)_k - y_{0,m}^{\text{DC}} (\mathbf{J}(T/2)[x]_m^T)_k}{[y_{0,m}^{\text{DC}}]^2}. \quad (5.21)$$

5.2.5 Regularisation

Whilst we expect an improvement in the localisation of our UOT sensitivity functions over that of pure DOT, our present model still describes a diffusion process: thus the problem of reconstructing an internal parameter distribution from data measured on the boundary remains ill-posed, and effectively under-determined. The consequence of these considerations is that even if sufficient measurements were available to form a square Jacobian, direct solution of equation 5.6 would fail to produce a stable solution, as noise in the measurements amplifies high-frequency oscillations in the resultant reconstruction.

In order to seek a useful and stable solution, it is conventional to apply some form of regularisation to the problem. One commonly employed approach is to make the assumption that a ‘good’ solution is one which balances the norm of the residual $\|(\mathbf{J}(\tau)[\mathbf{x}_0])\Delta\mathbf{x}_g - \Delta\mathbf{y}\|$ with the norm of the solution $\|\Delta\mathbf{x}_g\|$, or its derivatives. This approach is that of the general form of *Tikhonov* regularisation in which we replace the matrix equation 5.6 with a minimisation problem,

$$\Delta\mathbf{x}_g = \arg \min_{\mathbf{x}} \left(\|\mathbf{J}(\tau)[\mathbf{x}_0]\Delta\mathbf{y} - \Delta\mathbf{x}\|_2^2 + \lambda^2 \|\Delta\mathbf{x}\|_2^2 \right). \quad (5.22)$$

For the under-determined case, this minimisation procedure is equivalent to solving [20],

$$\Delta\mathbf{x}_g = (\mathbf{J}(\tau)[\mathbf{x}_0]^T \mathbf{J}(\tau)[\mathbf{x}_0] + \lambda^2 \mathbf{L})^{-1} \mathbf{J}(\tau)[\mathbf{x}_0]^T \Delta\mathbf{y}. \quad (5.23)$$

In the following section we employ zeroth-order Tikhonov regularisation such that $\mathbf{L} = \mathbf{I}$, where \mathbf{I} is the identity matrix. We select the regularisation parameter λ according to the l-curve method, though the point of maximum curvature was chosen by inspection when the algorithm failed to find the correct point.

5.3 Results

Having derived a technique for the linear reconstruction of images in ultrasound-modulated optical tomography we will now investigate the technique numerically.

1. We will observe the form of the correlation measurement density functions for absorption perturbations. This analysis will involve both reflection and transmission mode measurements under all of the previously described measurement types. Comparison of the UOT measurement types with the continuous DOT case will demonstrate the potential improvement in spatial resolution achieved under acoustic modulation, and direct observation of the UOT sensitivity functions will provide some qualitative appreciation of the nature of UOT sensitivity in correlation based measurements.
2. The reconstruction of a simple absorption perturbation distribution against an homogeneous background will be demonstrated. A single optical source and detector will

be employed such that the effects of regularisation under increasing noise will be evident in the reconstructed image. We will compare the reconstructed results with a direct mapping approach to demonstrate the problems suggested in section 5.1.

5.3.1 Correlation measurement density functions

In section 5.2.4 we derived the correlation measurement density functions for three measurement types: the zero-lag measurement (equivalent to continuous wave diffuse optical tomography), the AC measurement type, and the modulation depth measurement. We will now inspect these sensitivity functions to gain insight into the sensitivity and spatial localisation offered by these measurements.

A two-dimensional square domain of side 50mm is located with its bottom left corner at $(0, -25)$ mm. The domain is assigned a uniform absorption coefficient $\mu_a = 0.01\text{mm}^{-1}$ and reduced scattering coefficient $\mu'_s = 1\text{mm}^{-1}$ with isotropic scattering $g = 0$. The refractive index of the domain is 1.4 and a matched boundary is employed. An ultrasonic field with a Gaussian profile of full width half maximum 2mm and peak amplitude of 0.2MPa is projected through the plane.

Reflection mode

In the first instance we investigate a reflection mode geometry. An isotropic point source is placed at position $(1, 20)$ mm to approximate a collimated source incident perpendicular to the boundary at $(0, 20)$ mm. The diffuse detector is located at $(0, -20)$ and integrates the outgoing correlation flux according to a Gaussian profile with full width half maximum of 0.1mm.

To begin we investigate the sensitivity to absorption in the zero-lag measurement which is equivalent to the continuous wave sensitivity in diffuse optical tomography. Figure 5.2 depicts the negative zero-lag (intensity) sensitivity function for a given measurement, $-\mathbf{J}(0)[x_0]_m^T$. The approximated collimated source is indicated with the inwards arrow at the appropriate location, and the centre of the diffuse detector with the outwards arrow. The figure demonstrates the characteristic ‘banana’ shaped absorption sensitivity profile often seen in the literature of diffuse optical tomography [9, 42, 122]. The scale of the linear plot is truncated in the positive direction; regions close to the source and detector demonstrate extremely high sensitivity to absorption perturbations. This sensitivity function is insensitive to the acoustic field location: each row of the DC Jacobian which corresponds to a particular optical source and detector pair will be identical. Since in this work we consider only one such pair, a reconstruction using this information can only result in a scaled and regularised image of the same form as the sole unique sensitivity function depicted in figure 5.2. To produce a more detailed image would require that many more optical sources and detectors be employed.

We now turn our attention to the absorption sensitivity for the AC measurement type. Figures 5.3, 5.4, and 5.5 demonstrate the absorption sensitivity with the acoustic field scanned

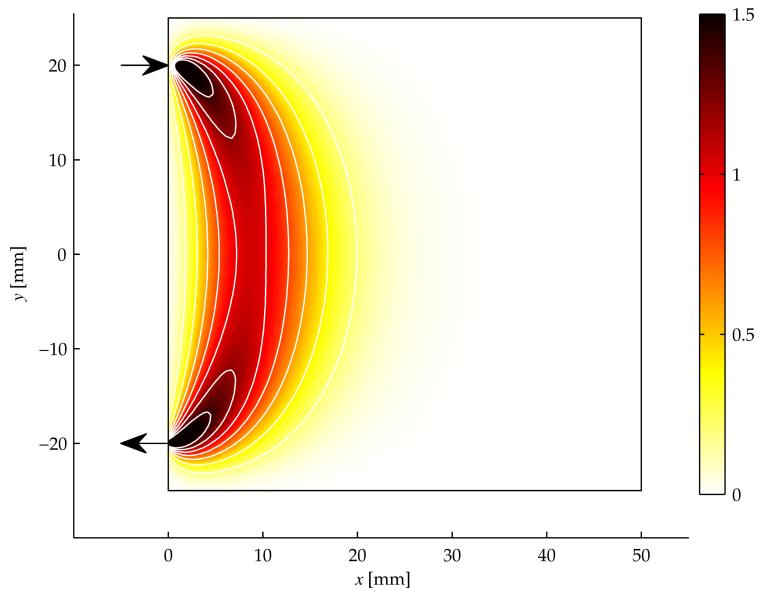


Figure 5.2: $-\mathbf{J}(0)[x_0]_k^T [\text{W m}^{-1}]$: UOT absorption sensitivity in reflection mode at lag $\tau = 0$, with a 40mm source-detector spacing. This sensitivity function is equivalent to that of an absorption photon measurement density function in continuous-wave diffuse optical tomography.

to locations of $(5, -15)$, $(15, -10)$, and $(20, 10)$ mm, respectively. The location of the acoustic field is indicated by the crossed circle. In the first of these images the acoustic scan location is in a region of high optical sensitivity close the source. In these circumstances the form of the sensitivity function does not deviate too greatly from the zero-lag case against which these figures may be compared, though it is evident that a larger region of sensitivity is formed around the acoustic focal region. In the latter two figures the acoustic field is scanned to a location of lower optical sensitivity where the effect of the acoustically driven decorrelation of the field is more evident. In both cases we see a region of increased sensitivity at the focus of the acoustic field which extends to the source and detector locations. As the acoustic field is scanned to regions of decreased optical sensitivity, the AC intensity falls commensurately.

Finally we consider the modulation depth absorption sensitivity functions. In figures 5.6 to 5.8 the ultrasound field is scanned to the same locations as figures 5.3 to 5.5. The first and most noticeable quality of each of the figures is the bipolar nature of the sensitivity functions (note however that the colour scale significantly enhances the contrast of the negative change over that of the positive). Each figure demonstrates that this measurement type is insensitive to perturbations close to the source, detector, and on a smooth path connecting the two. Placing an optical absorber along this path of insensitivity will reduce the amount of correlated and uncorrelated light reaching the source from the detector by an equal amount such

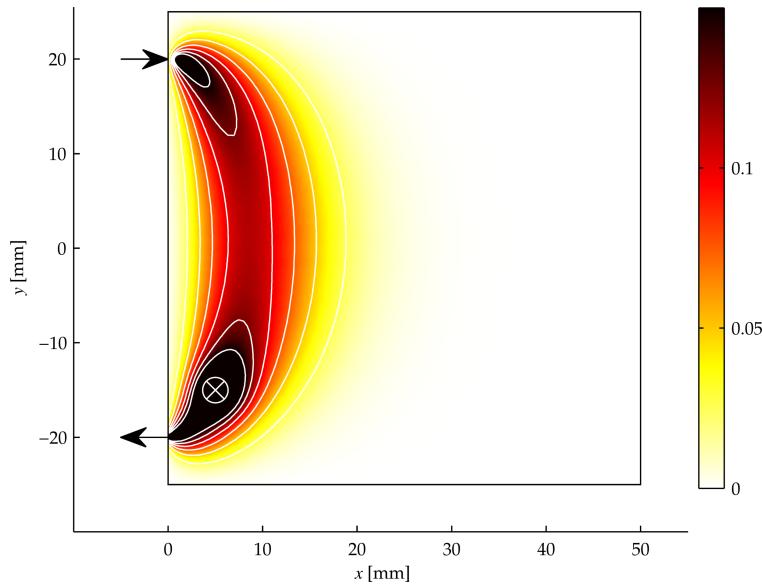


Figure 5.3: $-\mathbf{J}^{\text{AC}}[x_0]_k^T [\text{W m}^{-1}]$: UOT absorption sensitivity in reflection mode for an AC measurement type, with a 4cm source-detector spacing. The ultrasound focus is located at (5, -15)mm, close to the detection position.

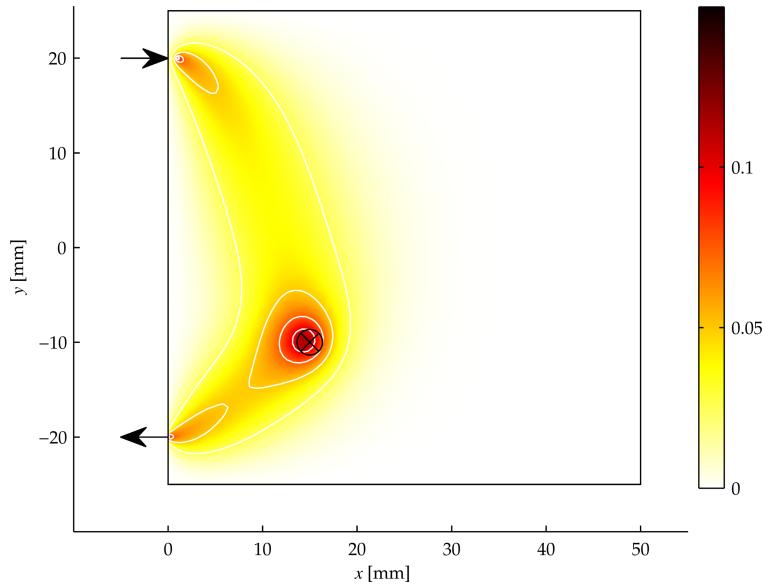


Figure 5.4: $-\mathbf{J}^{\text{AC}}[x_0]_k^T [\text{W m}^{-1}]$: UOT absorption sensitivity in reflection mode for an AC measurement type, with a 4cm source-detector spacing. The ultrasound focus is located at (15, -10)mm in the region of interest which is typically assumed to be probed during a DOT experiment.

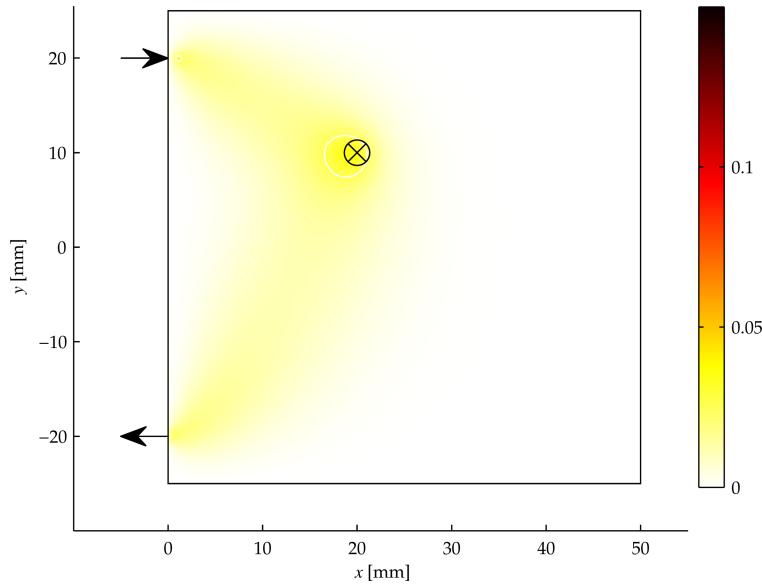


Figure 5.5: $-\mathbf{J}^{\text{AC}}[x_0]_k^T [\text{W m}^{-1}]$: UOT absorption sensitivity in reflection mode for an AC measurement type, with a 4cm source-detector spacing. The ultrasound focus is located at (20,10)mm on the periphery of the sensitive region.

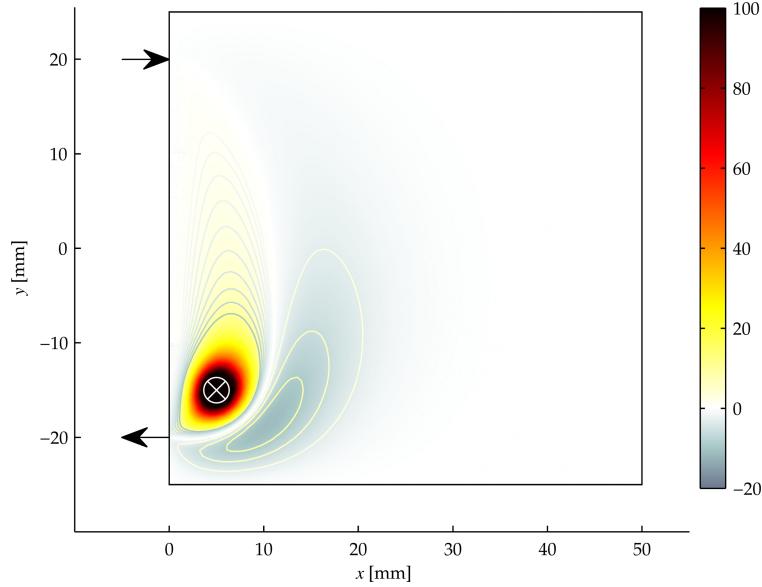


Figure 5.6: $-\mathbf{J}^{\text{MD}}[x_0]_k^T [\text{m}^{-1}]$: UOT absorption sensitivity in reflection mode for a modulation depth measurement, with a 4cm source-detector spacing. The ultrasound focus is located at (5,-15)mm, close to the detection region.

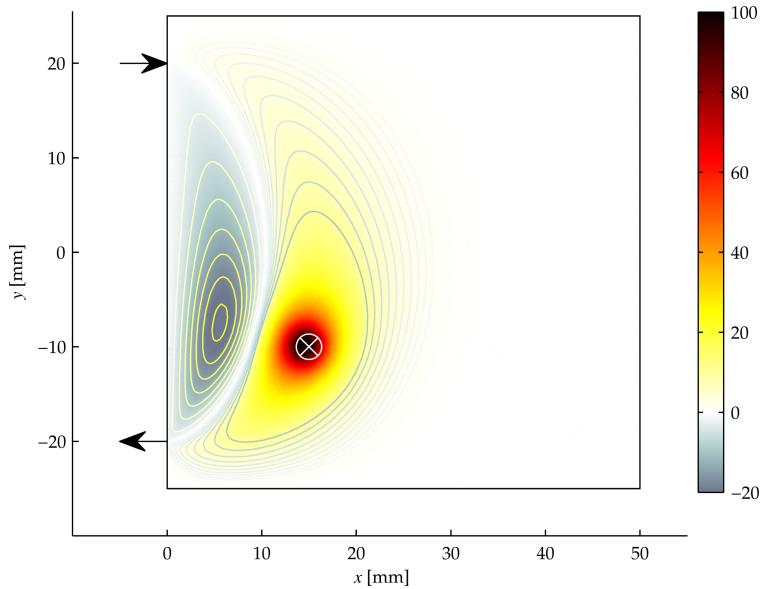


Figure 5.7: $-\mathbf{J}^{\text{MD}}[x_0]_k^T [\text{m}^{-1}]$: UOT absorption sensitivity in reflection mode for a modulation depth measurement, with a 4cm source-detector spacing. The ultrasound focus is located at (15, -10)mm, in the region of interest which is typically assumed to be probed during a DOT experiment.

that whilst the overall light level falls, the modulation depth remains constant. This argument is clearly supported by the analytical form of the sensitivity function given in equation 5.4. The modulation depth sensitivity functions demonstrate high sensitivity in the region of the applied acoustic field. Moving away from the acoustic focus, across the path of zero sensitivity, we encounter a region of negative sensitivity where a perturbation in absorption will attenuate more correlated light, than uncorrelated light.

Comparing figures 5.6 and 5.7 with figure 5.8 we can clearly see that the modulation depth sensitivity reduces as we move further from the region of optical sensitivity. In figure 5.9 we plot the sensitivity functions for an $x - y$ scan of the acoustic field through the medium where the bipolar nature of the sensitivity, and its dependence upon the local optical sensitivity can clearly be seen. Note that the scale in each of the images is fixed. Note in particular that when the acoustic field location is placed directly in the path of greatest optical sensitivity the sensitivity function appears to have the tightest spatial localisation, with two regions of negative sensitivity appearing in the sensitivity function. This may be problematic in a direct mapping approach in the sense that complex absorption perturbations will be averaged over the entire region of sensitivity, but is automatically accounted for in the image reconstruction process.

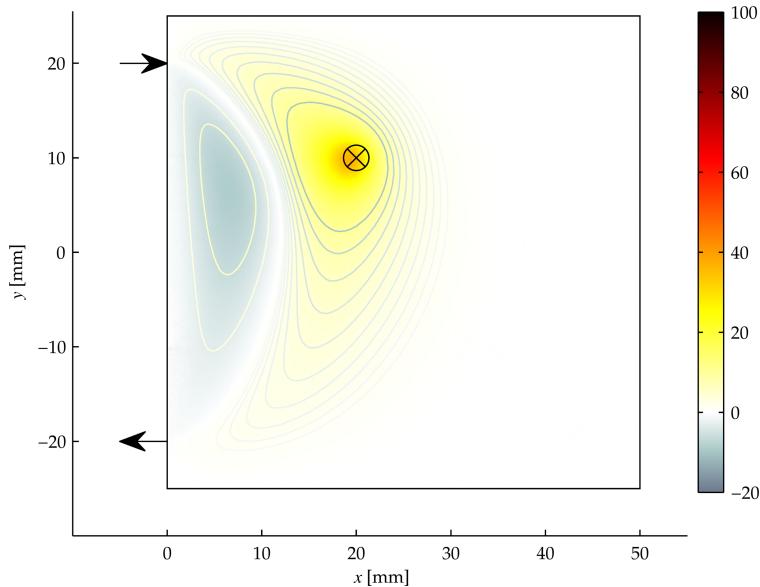


Figure 5.8: $-\mathbf{J}^{\text{MD}}[x_0]_k^T [\text{m}^{-1}]$: UOT absorption sensitivity in reflection mode for a modulation depth measurement. The ultrasound focus is located at $(20, 10)$ mm, on the periphery of the sensitive region.

Transmission mode

We now consider a transmission mode system. An isotropic point source is placed at position $(0.1, 0)$ mm to approximate a collimated source incident perpendicular to the boundary at $(0, 0)$. A diffuse detector of the same configuration as the reflection mode case is located at $(50, 0)$ mm.

We again begin our inspection with the zero-lag case depicted in figure 5.10. The figure depicts a characteristic DOT sensitivity similar to that found in the reflection mode case of figure 5.2. Since the source and detector spacing in this experimental configuration has increased to 50 mm, from 40 mm in reflection mode, the peak sensitivity is slightly reduced.

The transmission mode absorption sensitivity for the correlation measurement made for an AC measurement is displayed in figure 5.11. The results in transmission mode echo those in reflection mode: a region of high sensitivity is found at the location of the acoustic focus which extends outwards to the source and detector location.

We plot the modulation depth sensitivity for the same acoustic field location. Figure 5.12 demonstrates the modulation depth sensitivity for the transmission mode measurement with the acoustic field scanned to a location $(20, 10)$ mm. Once again we find a bipolar function with positive sensitivity in the region of the acoustic focus. Close to the locus of peak optical sensitivity a zero-crossing is present upon which perturbations in absorption will equally affect the modulated and unmodulated components of the detected correlation flux, leading

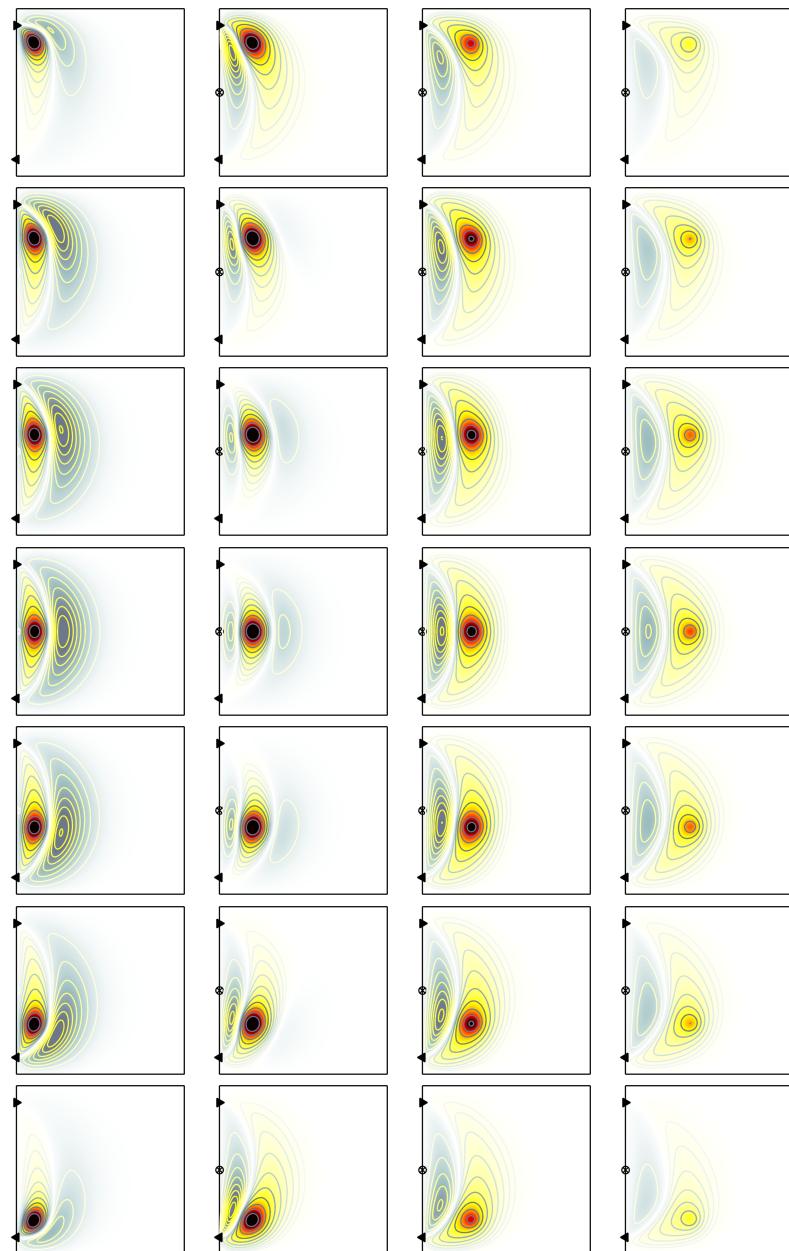


Figure 5.9: UOT absorption sensitivity in reflection mode for a modulation depth measurement over thirty-six of the ultrasound scan locations employed in the reflection mode reconstruction sequence.

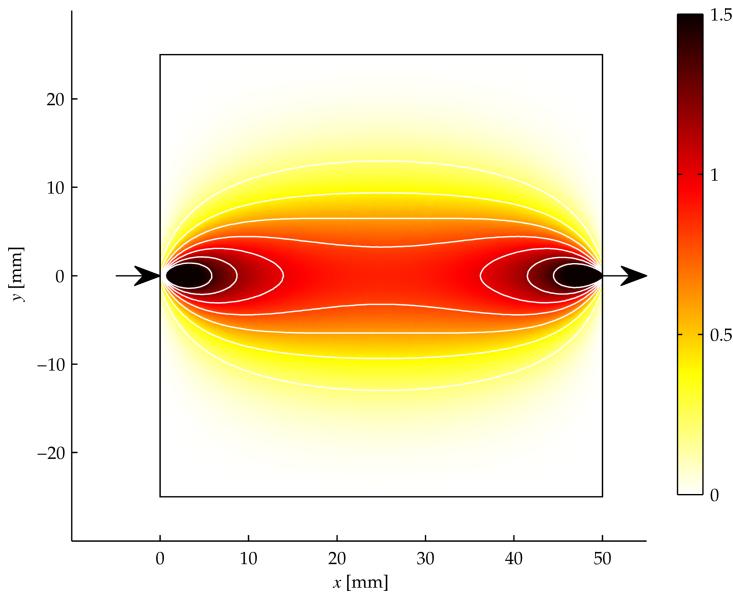


Figure 5.10: $-\mathbf{J}(0)[x_0]_k^T [\text{W m}^{-1}]$: UOT absorption sensitivity in transmission mode at lag $\tau = 0$, with a 5cm source-detector spacing. This sensitivity function is equivalent to that of an absorption photon measurement density function in continuous-wave diffuse optical tomography.

to zero net change in the modulation depth. The peak magnitude of the modulation depth sensitivity is comparable to that of the reflection mode measurement for the same acoustic scan location, since the optical sensitivities at this locations are similar for both measurement modes. In figure 5.13 we demonstrate the modulation depth sensitivity as we scan the acoustic field through the medium, with a constant colour scale.

5.3.2 Image reconstruction: noise and regularisation

Having examined the form of the individual correlation measurement density functions which together form the Jacobian, we will proceed to reconstruct an image in both reflection and transmission modes, under varying degrees of noise. The purpose of this investigation is twofold: first, the procedure will evaluate the functioning of the reconstruction process, second, the regularisation required as a function of noise will be assessed. We will perform the reconstruction on both the AC and modulation depth measurements.

The amount of noise added to the simulated data has been chosen to demonstrate the maximum proportional noise which can be tolerated before the presence of individual absorption features along the path of optical sensitivity can no longer be discerned. It is self-evident that at such levels of noise, not only will a qualitative interpretation of the results be severely degraded, but also the quantitative values returned by the reconstruction procedure.

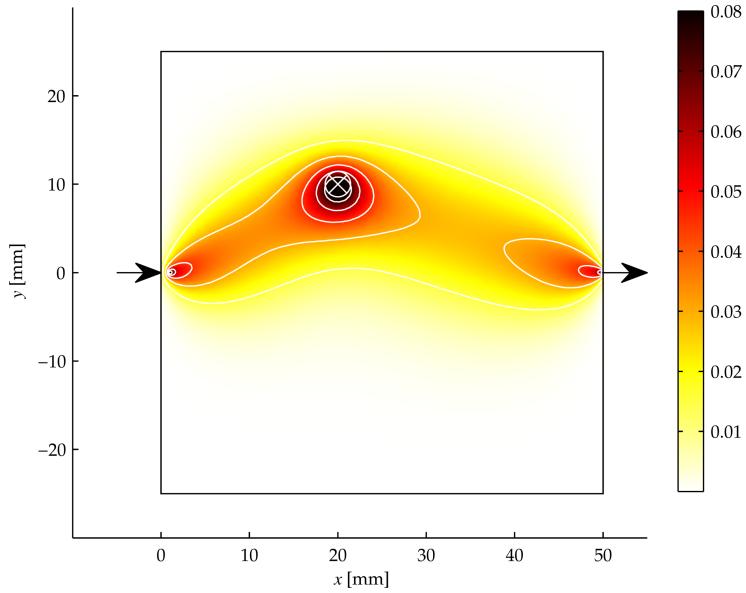


Figure 5.11: $-\mathbf{J}^{\text{AC}}[x_0]_k^T$ [W m⁻¹]: UOT absorption sensitivity in transmission mode for an AC measurement, with a 5cm source-detector spacing. The ultrasound focus is located at (20,10)mm on the periphery of the sensitive region.

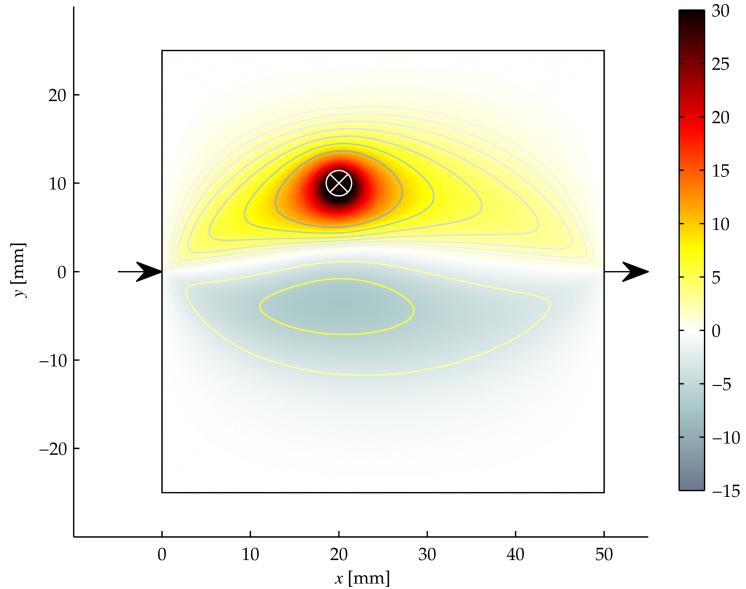


Figure 5.12: $\mathbf{J}^{\text{MD}}[x_0]_k^T$ UOT absorption sensitivity in transmission mode for a modulation depth measurement. The ultrasound focus is located at (20,10)mm on the periphery of the sensitive region.

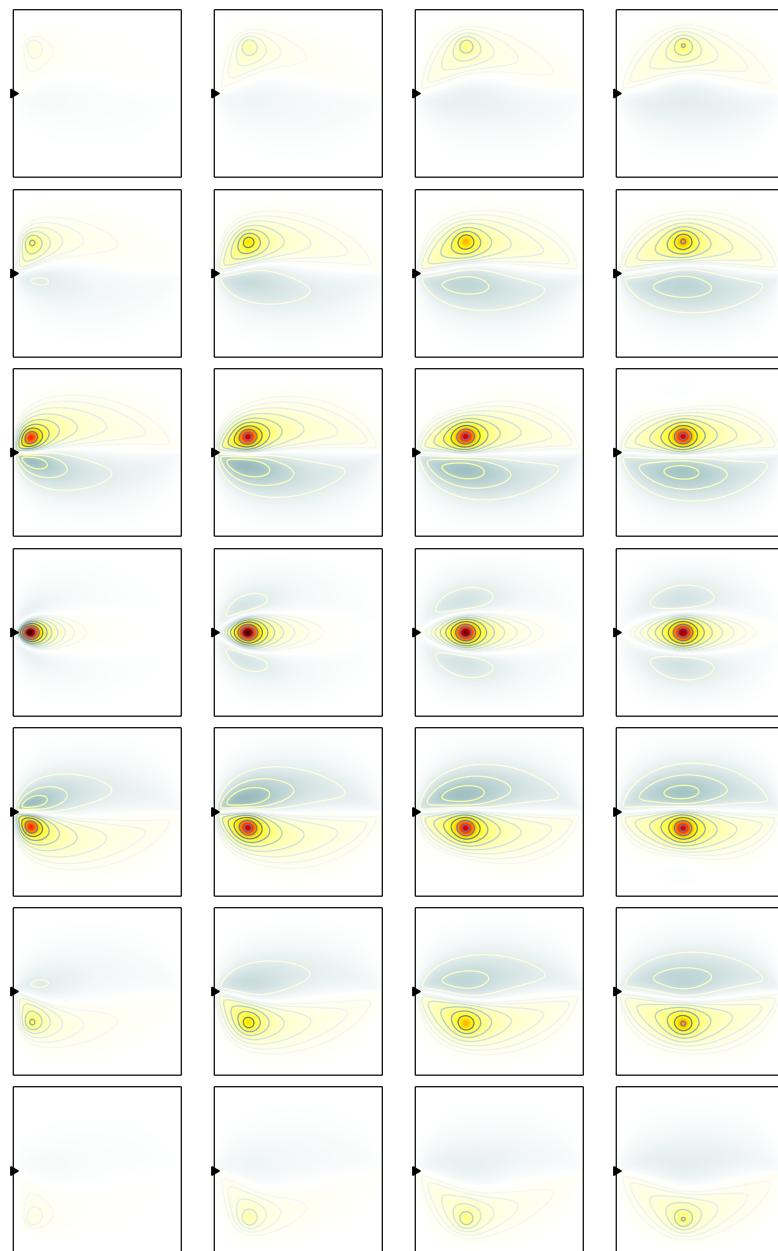


Figure 5.13: UOT absorption sensitivity in transmission mode for a modulation depth measurement over thirty-six ultrasound scan locations, with a 5cm source-detector spacing, and a fixed colour scale.

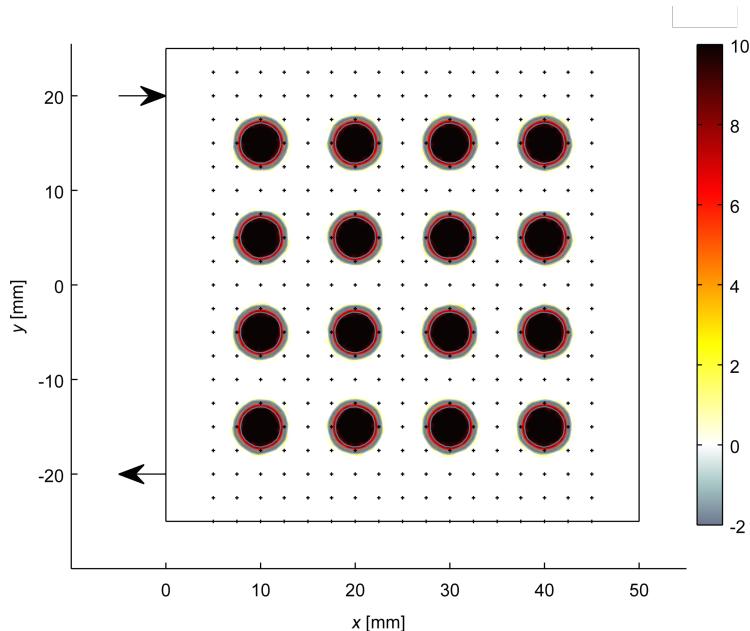


Figure 5.14: Absorption perturbation $\Delta\mu_a$ [m^{-1}] interpolated from the perturbation mesh to the baseline mesh upon which the Jacobian is computed. The baseline mesh does not include features of the perturbed geometry and hence the boundaries are seen to be slightly blurred. Each of the 357 points in the image indicate a single acoustic scan location.

Thus, we shall determine the limit at which our data can no-longer provide physiologically significant information. We will further consider the tolerance of our reconstruction to additive noise in the subsequent discussion.

The reconstruction will be performed on an homogeneous background with a regular set of absorption perturbations such that any aberrations in the reconstructed image will be evident. The background domain is the same as employed in section 5.3.1. The perturbation domain is formed by adding an array of sixteen circular perturbations of radius 2.5mm with a modified absorption coefficient of 0.02mm^{-1} . The perturbed measurements were generated using an independent mesh from the background measurements lest we commit inverse crime of generating measurements by the same model by which they are inverted [31]. To demonstrate an ‘ideal’ reconstruction, the perturbed absorption coefficient was interpolated from its mesh to the baseline mesh upon which the Jacobian is computed, the resultant image is displayed in figure 5.14.

The measurement set was generated by scanning an ultrasonic field with a Gaussian profile of full width half maximum 2mm and peak amplitude of 0.2MPa through the domain from $(5, -25)\text{mm}$ to $(-25, 25)\text{mm}$ at 0.5mm increments in x and y . This resulted in a total of $17 \times 21 = 357$ measurements for both the baseline and perturbed measurement. The loca-

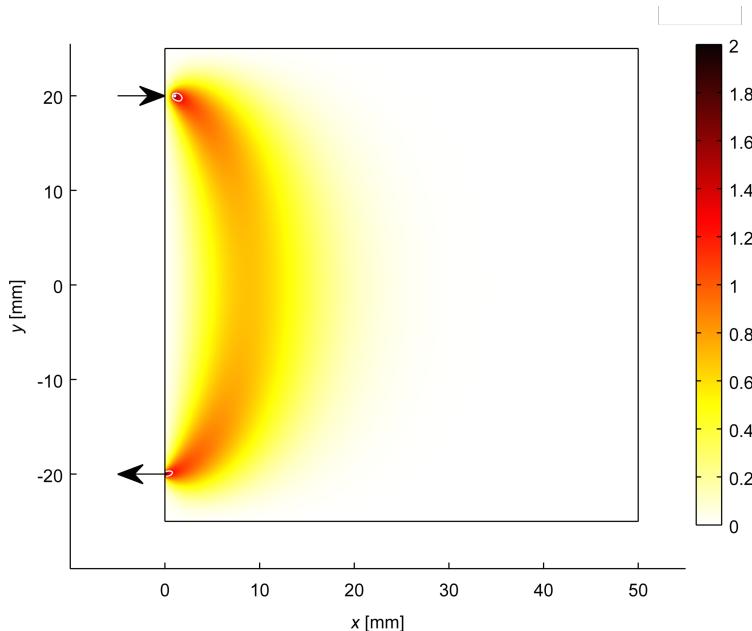


Figure 5.15: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using only the zero-lag Jacobian, equivalent to attempting a reconstruction using a single continuous wave source detector pair in the DOT problem.

tion of the acoustic focal regions is depicted by the plus symbols in the target reconstruction image of figure 5.14.

In each case the regularised system is solved by 5.2.3. The regularisation parameter was chosen by the l-curve method, though on occasion the point of maximum curvature was manually evaluated when the algorithm failed to choose the correct point. In the cases where no noise was added to the measurements, a regularised solution was still sought owing to the influence of slight numerical error in calculating the solution.

Reflection mode

We begin by evaluating the image reconstruction process in reflection mode. As per section 5.3.1 an isotropic point source is placed at position $(0.1, 20)$ mm to approximate a collimated source incident perpendicular to the boundary at $(0, 20)$ mm. A diffuse detector is located at $(0, -20)$ mm and integrates the outgoing correlation flux according to a Gaussian profile with full width half maximum of 0.1mm.

For completeness we begin by attempting to reconstruct the image using the zero-lag Jacobian. The result of the noiseless reconstruction is depicted in figure 5.15. Since the measurement is unperturbed by the acoustic field at this point in the correlation function, the Jacobian has a rank of one: all rows are identical to one another. The only information present

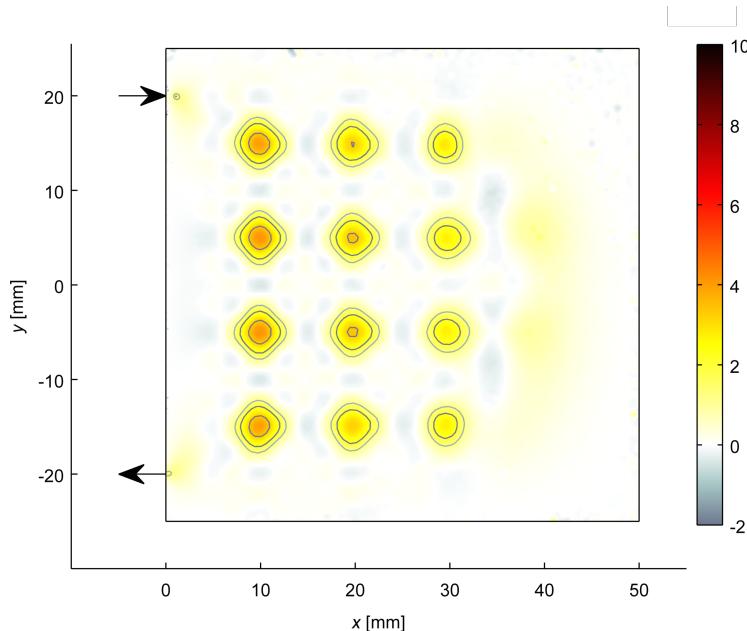


Figure 5.16: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using AC Jacobian absent of measurement noise.

in this inversion is the absorption sensitivity profile of figure 5.2; it is thus unsurprising that the reconstructed image is merely a scaled copy of this single set of linearly independent data.

We now consider a reconstruction based upon the AC measurement type and its associated Jacobian. In figure 5.16 we see that in the absence of noise the reconstruction produces an image which captures the features of the perturbation up to a depth of 30mm. Beyond this depth the sensitivity is so low that the solution has been forced to zero by regularisation. We also note that the reconstructed image has peaks close to the source and detector; we note from our analysis of the correlation measurement density functions of this measurement type that whilst a distinct region of sensitivity was found at the acoustic focus, the measurement was also sensitive near the source and detector location. The image demonstrates a spatial resolution comparable to the full-width half maximum of the acoustic field, this result is expected given the spatial sensitivity demonstrated in the associated correlation measurement density functions.

We now introduce 0.1% Gaussian noise onto the simulated measurement data, recalculate the regularisation parameter λ and reconstruct the image accordingly: the result is shown in figure 5.17. With added noise the quality of the image decays as regularisation smooths the regions of the image with low sensitivity which would otherwise, following inversion, dominate the reconstructed image in the presence of noise. Useful results are demonstrated in the image to a depth of over 20mm. Erroneous values near the source and detector locations

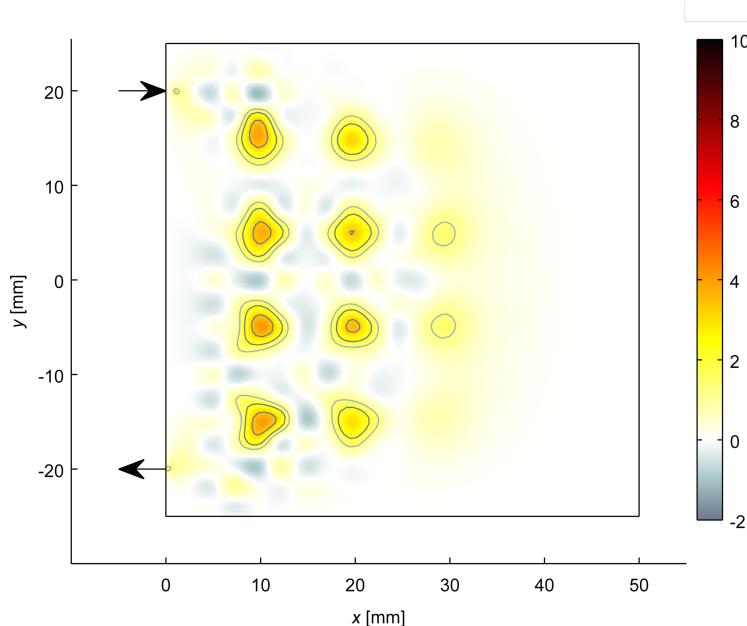


Figure 5.17: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using the AC Jacobian with 0.1% Gaussian measurement noise.

remain present in the image. Increasing the noise to 0.5% causes the regularisation process to remove any useful features of the perturbation from the image.

Finally we consider the reconstruction based upon the modulation depth measurement and its associated Jacobian. The noiseless reconstruction is shown in figure 5.18. The noiseless modulation depth reconstruction demonstrates good image quality throughout the reconstructed image. Owing to the nature of the measurement type, the source and detector locations are completely suppressed in this reconstruction. The image close to the source is in good agreement with the ideal reconstruction presented earlier, demonstrating an absorption change approaching the actual value $\Delta\mu_a = 10 \times 10^{-3} \text{ mm}^{-1}$. Even without noise, numerical inaccuracy is such that image decays into the region of lowest optical sensitivity.

In figures 5.19 and 5.20 we add 0.1% and 1% Gaussian noise to the simulated measurements and repeat the reconstruction process. These figures demonstrate that as the amount of noise in the measurement is increased, regularisation gradually suppresses the reconstructed images in the regions of least sensitivity. The modulation depth reconstruction demonstrates significantly improved tolerance to measurement noise than was demonstrated by the reconstruction based upon the AC measurement type. In the latter case, only 0.5% Gaussian noise could be added to the image before the resultant image became unusable, in the present case, noise of up to 1% can be added to the measurement whilst still resolving features of the image to a depth beyond 10mm.

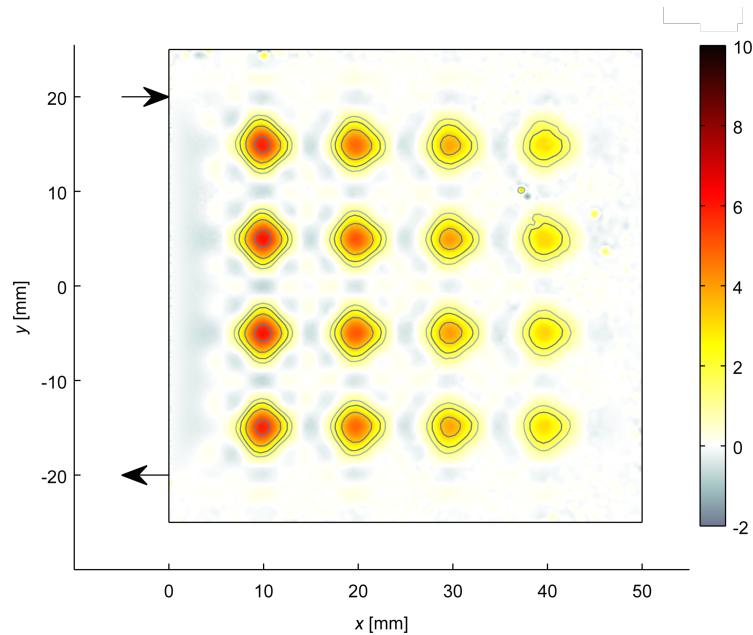


Figure 5.18: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using the UOT modulation depth Jacobian absent of measurement noise.

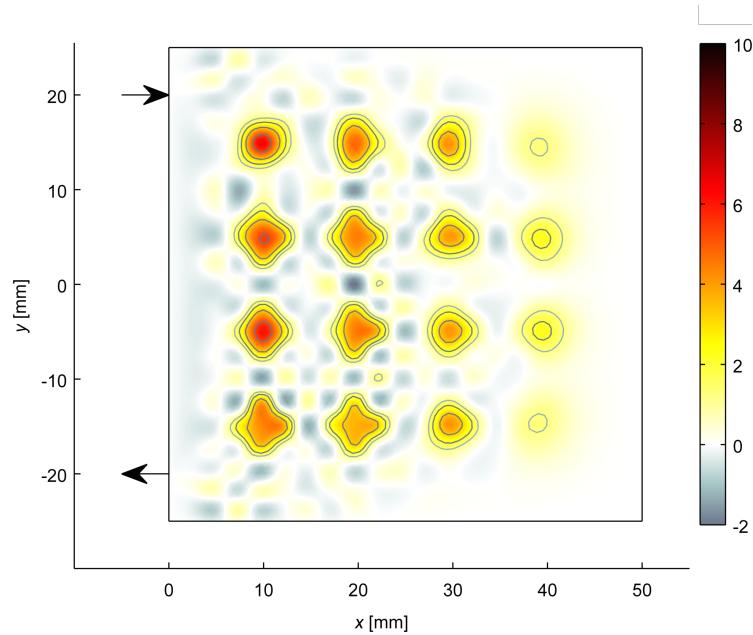


Figure 5.19: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using the UOT modulation depth Jacobian with 0.1% Gaussian noise.

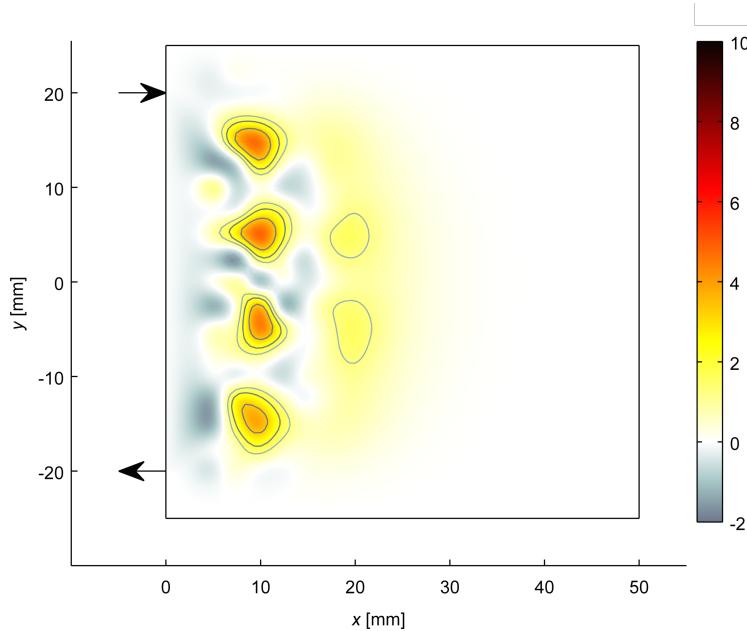


Figure 5.20: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using the UOT modulation depth Jacobian 1% Gaussian noise.

In the introduction to this chapter we noted that on many occasions a simple ‘direct-mapping’ approach was used for the generation of images in an acousto-optic experiment. We dismissed this approach in the case of a low-étendue system, or in a scenario where the optical sensitivity varies rapidly over the plane of the image. To demonstrate the improvement of the reconstructed images over that of a direct mapping approach, we generate an image by the direct mapping approach in the absence of noise by directly assigning the modulation depth difference measurement, $\Delta m(\mathbf{r})$, to pixels located at the centres of the acoustic foci.

Figure 5.21 demonstrates that the absorption perturbations along the axis of optical sensitivity (figure 5.10) can be distinctly located. Without compensating for the spatially varying optical sensitivity, those absorbers off axis are imperceptible. Moreover, the resulting unitless image provides no quantitative information as would be required if we wished to attempt a recovery of a clinically relevant parameter. It is a common feature of many hybrid imaging modalities that the relationship of the data to the desired parameter is often unclear. We also note that in the presence of noise a variety of rigorous regularisation strategies can be applied in a reconstruction to stabilise the recovery of the underlying parameters. In a direct mapping approach we may attempt to reduce the deleterious effects of the noise by some form of spatial filtering (for example, convolution with a smoothing kernel), but any such approach is somewhat ad-hoc.

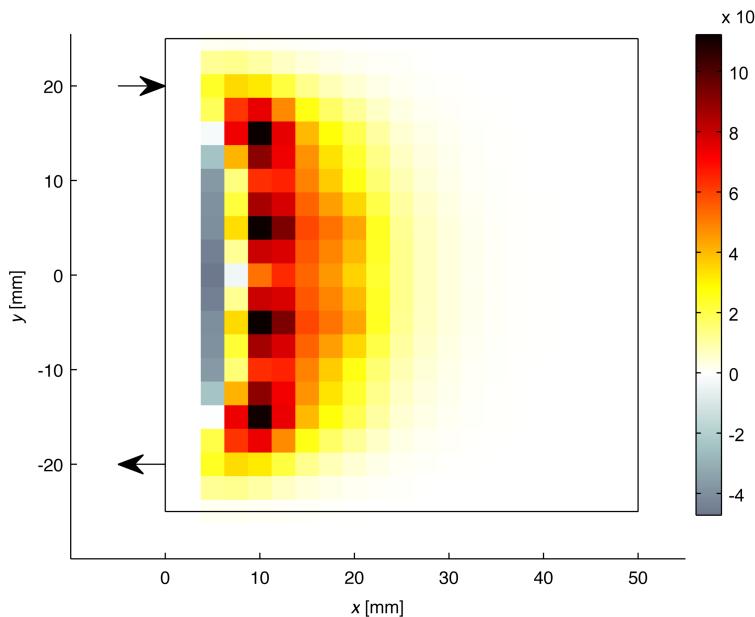


Figure 5.21: Direct mapping approach to reconstruction of an absorption perturbation using a modulation depth measurement in reflection mode without noise.

Transmission mode

We now briefly consider the transmission mode scenario for the modulation depth experiment. As per section 5.3.1 an isotropic point source is placed at position $(0.1, 0)$ mm to approximate a collimated source incident perpendicular to the boundary at $(0, 0)$ mm. A diffuse detector of the same configuration as the reflection mode case is located at $(50, 0)$ mm.

The reconstructions of figures 5.22 through 5.24 demonstrate the modulation depth reconstruction in the case of measurements absent of noise, with 0.1%, and 1% Gaussian noise. Without noise, the reconstruction process provides an excellent approximation to the actual perturbed absorption profile. Under noise we see that regularisation smooths the regions of lower sensitivity which are now located off the transmission mode axis, echoing the results of the reflection mode measurements. Qualitatively, the transmission mode reconstructions appear to more accurately portray the actual absorption perturbation than their reflection mode counterparts: this is largely due to the more even optical sensitivity profile through the medium than in the reflection mode case.

5.4 Conculsions

The sensitivity functions and associated reconstructed images presented in this work provide considerable insight into the potential of autocorrelation based UOT. It has been shown that

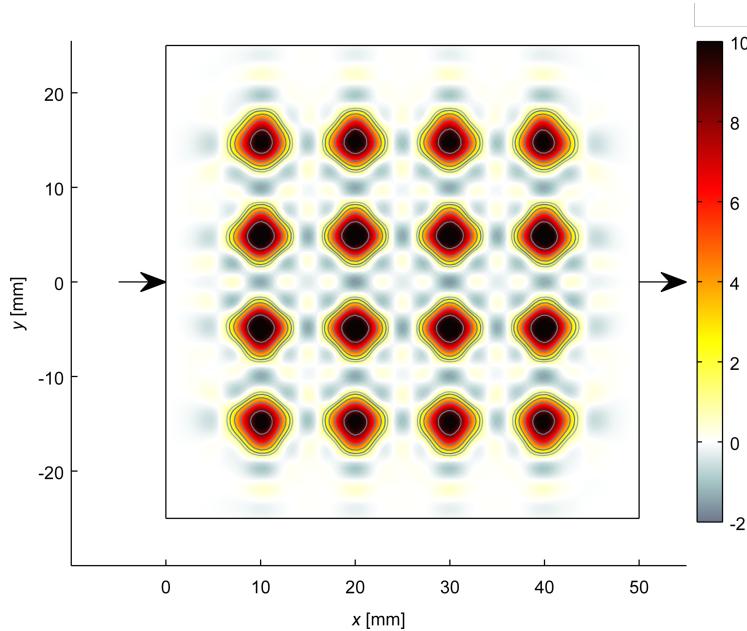


Figure 5.22: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using the UOT modulation depth Jacobian in a transmission mode configuration without measurement noise.

by moving the focus of an acoustic field in the medium, additional spatial resolution can be introduced into the UOT reconstruction: in the case of CW DOT, this would require that additional optical sources and detectors were employed, with associated cost and complexity. Whilst it is arguably unfair to directly compare a CW DOT image with the UMOT image, since the former was acquired with only two measurements compared to the latter's set of 372 measurements, it is nonetheless important to recognise the potential advantages of a system capable of reproducing such images with a small number of *optical* sources and detectors, owing to the implications in terms of cost and complexity. The low source-detector count demonstrated here suggests that a single (planar) optical source and point detector may be capable of reproducing images of optical parameters with sub-centimetre resolution, with the scanning of the acoustic focus achieved electronically by means of a phased array. Such a system would be mechanically straightforward and avoid many of the cost and complexity issues of a multi-channel DOT system.

Reconstructions based directly upon the AC measurement type are, in the absence of noise, capable of reproducing a fair representation of the underlying absorption perturbations. A disadvantage of this measurement type is the significant erroneous absorption indicated near the source and detector position: this is a direct consequence of the extreme sensitivity of the measurement in these regions. If an absorption perturbation were placed

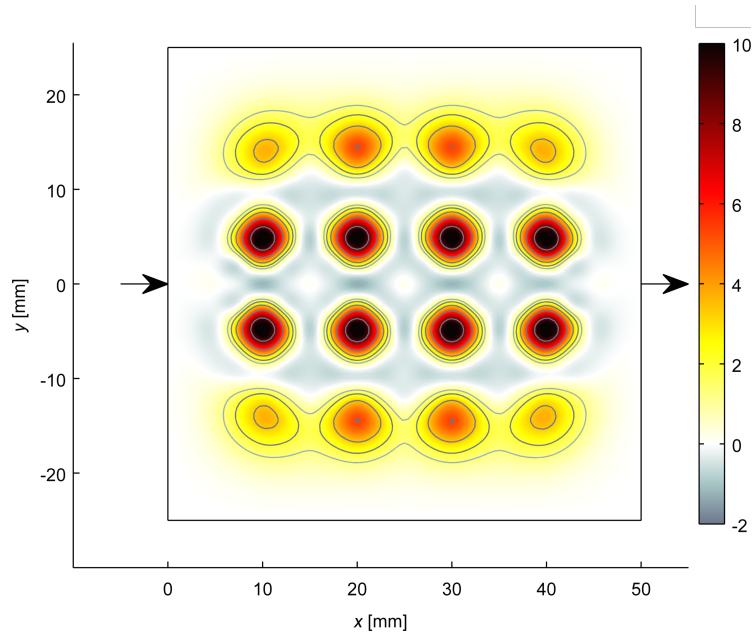


Figure 5.23: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using the UOT modulation depth Jacobian in transmission mode configuration with 0.01% Gaussian noise.

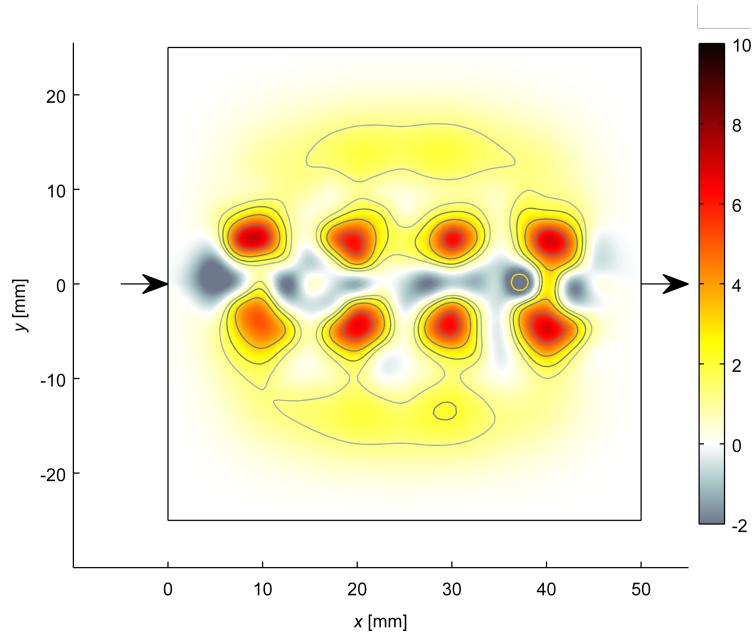


Figure 5.24: Reconstruction of the absorption perturbation $\Delta\mu_a$ [m^{-1}] using the UOT modulation depth Jacobian in transmission mode configuration with 1% Gaussian noise.

near the source or detector, the influence of this measurement change would dominate the resulting image.

In a reconstruction employing both the AC and DC measurements, regularisation would suppress the contributions from the AC measurements in which the perturbation from the DC value is below the noise floor. In this sense the enhancement of a DOT imaging system with ultrasound modulated correlation measurements will lead to a worse-case imaging resolution equivalent to that of CW DOT, with improved ultrasound modulated resolution where noise permits. The influence of the high regions of sensitivity near the source and detector remains a problem, however.

In contrast, the modulation depth measurement type completely suppresses sensitivity near the source and detector positions. Furthermore, the image reconstruction process is capable of producing acceptable results at levels of noise higher than in the AC measurement reconstruction. Suppressing sensitivity in regions close to the source and detector may be valuable in biomedical applications where physiological changes in a superficial region may otherwise dominate the measurement [130].

The amount of noise added to these reconstructions was chosen to explore the level at which the reconstructed images are degraded to the extent that they no longer provided any qualitatively useful information. In the case of the modulation depth measurement, accurate ($\pm 10\%$) reconstructions were achieved on the axis of optical sensitivity with additive noise of 1% applied to the simulated measurement data. The signal to noise ratio (SNR) of an actual experiment will vary significantly depending upon the chosen detection mechanism. In this work primarily consider autocorrelation based UOT which can employ photon counting detectors—such systems are inherently shot-noise limited. In this regime the available SNR is proportional to the square root of the detected signal level. Absolute values of 1% are not unreasonable, but this is highly dependent upon the properties and geometry of that which is being probed. In biological tissues, the optical input power may be maximised according to relevant safety limits, in this case decreasing detected signal levels (due to source-detector geometry, or the optical properties of the medium) will require that the integration time of the measurement be increased: this may potentially become (experimentally, or clinically) unviable. Interferometric methods such as digital holography utilise heterodyne mixing gain to achieve shot-noise limited performance [51], in this case the reference beam power must be increased to retain performance, with the requisite increases in cost and size of the coherent laser required for this function decreasing the utility of the technique.

In summary, we have demonstrated that standard techniques for image reconstruction in DOT can be successfully employed in UOT. The alternative measurement types available in UOT demonstrate sensitivity with a significantly improved spatial resolution over that of continuous-wave DOT. The modulation depth measurement type, and associated sensitivity functions, demonstrate particularly attractive qualities in suppressing the high sensitivity close to the source and detector, whilst also providing better noise immunity than direct use of the AC measurement data.

Chapter 6

Future work and summary

6.1 Future work

The methods and techniques presented in this work, and those which it references, represent a first step towards the clinical application of quantitative imaging in UOT. In this chapter we will consider both the challenges faced in moving towards clinical application, and some specific technical examples of how the development of the forward models and inversion techniques presented here may be extended.

6.1.1 Forward modelling

In this work we have considered two types of forward model which we may broadly categorise as statistical: Monte-Carlo sampling of the underlying transport equation, and deterministic: finite-element solutions to approximated forms of the transport equation. We consider the potential development of both.

Deterministic methods

In chapter 4 we derived a diffusion style model originally proposed by Sakadžić and Wang [118]. This model is closely related to earlier work by Wang in which the acousto-optic problem was treated in the style of diffusing wave spectroscopy [135], insofar as the phase increments are averaged and used within a diffusion style model of light transport.

We demonstrated that the accuracy of the model can be greatly improved when incorporated into a more advanced light transport model: we used an SP_3 model in this work. If the limitations of the model, in terms of the maximum permissible pressure amplitudes, and the requirement $(k_a l_{\text{tr}}) \gg 1$ can be accepted in a practical setting, then it would be of interest to explore higher order SP_N , or even P_N approximations. Such an investigation would be relatively straightforward given the extensive literature available in the field of diffuse optics, and the way that we have shown the AO phase increments to be integrated with the light transport model.

Even the simplest diffusion style models developed in this work represent a considerable computational burden for a reconstruction scheme, since a pair (forward and adjoint) of solutions must be formed *for each ultrasound location*, for each point of lag required for the solution. If the insonification generates an extra correlation absorption perturbation which is sufficiently small, it would be interesting to split the system matrix into the static component, absent of insonification, solve the system, and then approximate the effects of the small perturbation in correlation absorption due to the acoustic field by a Born series — that is to say that we would linearise the effects of the acoustic decorrelation on the optical field.

Alternatively, we could take an iterative approach (which with further investigation is likely to be related to the Born series). In our application we would replace the solution of equation 4.70,

$$\mathbf{A}(\tau)[x]\mathbf{G}(\tau) = \mathbf{q}, \quad (6.1)$$

with,

$$[\mathbf{A}_0[x] + \mathbf{H}(\tau)]\mathbf{G}(\tau) = \mathbf{q}, \quad (6.2)$$

where,

$$\begin{aligned} \mathbf{A}_0[x]_{ij} = \int_{\Omega} \sum_k \left[u_{D,k}(\mathbf{r}) D_k \nabla u_i(\mathbf{r}) \cdot \nabla u_j(\mathbf{r}) + u_{\mu,k}(\mathbf{r}) \mu_{a,k} u_i(\mathbf{r}) u_j(\mathbf{r}) \right] d^n \mathbf{r} \\ + \frac{1}{2A} \int_{\partial\Omega} u_i(\mathbf{r}) u_j(\mathbf{r}) d^{n-1} \mathbf{r}, \quad (6.3) \end{aligned}$$

and,

$$\mathbf{H}(\tau)_{ij} = \int_{\Omega} \sum_k u_{b,k}(\mathbf{r}) h_k(\tau) u_i(\mathbf{r}) u_j(\mathbf{r}) d^n \mathbf{r}. \quad (6.4)$$

The inverse of the matrix sum in the resultant equation can be found iteratively, the technique being described by Miller [96]. We recognise that in practice, solution of the matrix equation is not found by computing the inverse of the system matrix, and ostensibly this is required for the proposed method. Instead we anticipate that since the matrix need only be computed once, for all insonification profiles, a great deal of computational time would be saved compared to the present solution.

Either of these approaches would vastly decrease the amount of time required to form the Jacobian in a linear, or straightforward non-linear inversion procedure, although the effects on the accuracy of the resulting solution are unknown. Indeed, if the Jacobian is to be calculated, and regularisation to be applied, it may be appropriate to integrate these iterative approaches into the calculation of the SVD of the system matrix prior to application of the regularisation, since this may provide some indication of the level of iteration required before convergence is reached with respect to the noise floor of the measurements.

Statistical methods

We demonstrated in chapter 3, and in earlier work [73], that a significant amount of extra computation is required in a Monte-Carlo model of UOT in order to calculate the AO phase

increments. Given the encouraging results seen in coupling averaged phase increments with higher order light models (the SP_3 model of chapter 4), it may be interesting to incorporate these averaged increments in a Monte-Carlo light transport model. This would vastly increase the speed of simulation not only due to the reduced computation required, but also because the large acoustic dataset could be reduced in size (since propagation direction information is not required for averaged phase increments).

Even with such approximations, a Monte-Carlo approach to UOT image reconstruction will require another layer of parallelisation (i.e. a cluster of GPUs) to achieve the necessary simulation speed for use in an inversion process. Moreover, Monte-Carlo codes of the form described in this work, record the fluence in the domain which, if used in an inversion process, somewhat negates the extra accuracy afforded by the Monte-Carlo approach (the simplifications were detailed when we formulated an approximate adjoint method in chapter 3).

To properly exploit the transport qualities of a Monte-Carlo solution requires a model which records the R/CTE's field variable - (correlation) radiance. This could be achieved by extending each spatial point of the mesh used in the domain representation of the Monte-Carlo model in this work, with an associated angular mesh. If accelerated sufficiently, this would be of significant interest not only for UOT inversions, but also within the DOT and PAT communities.

6.1.2 The inverse problem

One of the most immediate investigations to be considered from the current work is to consider the recovery of the scattering coefficient, rather than just the absorption coefficient, in a UOT experiment. To do so would involve extending the derivation in chapter 5 to consider the gradient with respect to the scattering coefficient, and forming a compound Jacobian accordingly. A more subtle consideration is that of uniqueness — can certain independent changes in the scattering and absorption coefficients lead to the same measured boundary data? This problem is common to many inverse problems, and the author of a work which considers this problem in DOT [11] has suggested taking a similar approach in the UOT problem.

Given the substantial improvements demonstrated using the SP_3 model of chapter 4, it would be of value to reformulate the inversion procedure using this model. Such an approach has been reported for DOT [30], and good results were found in general, though instability in the solution was noted for higher order approximations. It is unclear without further investigation if such problems would be encountered in a UOT inversion.

Irrespective of a more developed understanding of uniqueness, the present work lends itself to be developed into a non-linear iterative inversion. This would permit absolute and quantitative imaging. A first step, which utilises the theoretical framework developed in this thesis, may be implementing a Gauss-Newton algorithm in which the Hessian is approxi-

mated using the Jacobian, which itself is repeatedly recalculated as the solution converges. We noted earlier that even the most computationally basic forward models presented in this work require significant computational effort, but that this would likely be viable if the proposed iterative approach to the addition of the acoustically induced correlation absorption were implemented. Alternatively the application of quasi-Newton methods such as the Broyden-Fletcher-Goldfarb-Shanno algorithm, or Krylov subspace methods such as conjugate gradients or GMRES could be explored.

6.1.3 Towards clinical application

Detection mechanisms, and data types

In chapter 2 we explored the large number of detection mechanisms that have been applied to, or developed for, UOT. Each of these techniques have advantages and disadvantages which lend themselves to different applications in biomedical imaging, and it is thus desirable that the development of an image reconstruction technique be applicable to as many existing and future techniques as possible.

The first difficulty which is encountered is that the different techniques each generate data of different forms. In UOT, the fundamental measurement is that of the power-spectral density (PSD) of the optical flux emanating from the medium being probed. In contrast, consider the data generated by three common detection mechanisms.

- Autocorrelation based techniques such as those considered in this work record the time (lag) domain intensity autocorrelation. This can be transformed to power-spectral density, but absolute units are only available if the intensity autocorrelation function is un-normalised, and if the optical system is properly characterised to permit application of the Seigert relationship.
- Interferometric techniques are typically capable of determining the flux shifted by the acoustic frequency (or its harmonics). An absolute measurement will require knowledge of the heterodyne gain in the system.
- Photo-refractive measurement systems, depending upon their configuration, may measure the total ‘unmodulated flux’, or alternatively, the ‘modulated flux’. This measurement is scaled by the mixing gain of the system.

Indeed, the only techniques in which the PSD can be directly measured are those which employ some form of extremely narrow-band filtering, such as spectral-hole burning.

Raw power spectral-density may not be the best candidate for a standard measurement-type, not least because some measurement techniques are unable to measure it, but also because in a clinical setting numerous practical matters will affect the absolute quantity of flux collected. This is akin to the preference for normalised measurement types in fields such as diffuse optical tomography.

A more sensible standard would be that of the normalised power spectral density, where in the absence of the acoustic modulation the power in the DC component of the flux is normalised to unity. In the CW case, this is equivalent to knowledge of the n^{th} harmonic modulation depths employed in this work, and others. Most of the techniques discussed in this work can be modified to provide this data type.

The presence of a (*de-facto*, or otherwise) standard data type would permit a focused approach to the development and evaluation of forward models in UOT, in addition to the benefits inherent in developing a common language between researchers employing the various techniques.

Imaging, sensing, and the UOT acoustic source

Future clinical UOT devices may include imaging systems which seek to recover quantitative images of optical coefficients (or chromophore concentrations) within a three-dimensional volume, or simpler sensing devices which can be tuned to recover a point-wise measurement. Reconstruction will be essential in both, but practical matters suggest different approaches may be taken.

It is likely that a clinical UOT imaging system would be based around an existing diagnostic ultrasound system, since such devices contain the hardware to safely generate the requisite acoustic field distributions whilst simultaneously providing extra clinically relevant information in the form of a standard echography image. In this scenario the UOT acoustic source would take the form of a short pulse of moderate amplitude with a centre frequency in the 1–10 MHz range. Lateral scanning would be achieved by beam-forming by the ultrasound system, and depth resolution would be achieved by time-gating: this approach has been demonstrated using the photo-refractive detection mechanism. In these physical circumstances the diffusion style approach to the forward modelling and reconstruction of UOT images may be applicable. The computational efficiency of these techniques, which are only applicable at higher acoustic frequencies, are key to allowing volume reconstructions with spatial resolution equivalent to that which can be offered by the spatio-temporal acoustic field distribution provided by such hardware.

A point-wise sensing system may be developed from the ground up with an acoustic transducer suitable for the specific application. For example, in a device designed to achieve measurement at significant depth it may be prudent to choose a lower operating frequency (to minimise attenuation), and higher peak power levels than would be employed in traditional diagnostic ultrasound. One may also choose to exploit acoustic non-linearity to directly generate higher harmonics in the UOT measurement which are confined to smaller volumes in the tissue. Each of these design choices suggest that diffusion style models may be ill-suited, and that a transport model should be employed in the reconstruction process to properly capture the underlying physics. The vastly increased computation effort required to execute such models may be offset in this case by the requirement to reconstruct parameters in a

single, fixed, volume.

Acoustic modelling

The acoustic component of UOT receives little attention in the literature. This is perhaps because the focused or pulsed transducers employed experimentally can be characterised by hydrophone measurements in media of similar acoustic properties to the (optically inhomogeneous) UOT phantoms which they scan. Knowledge of the acoustic field distribution is critical for image reconstruction in UOT—moving towards clinical application will require greater consideration of the acoustic properties within the biological tissues of interest.

In the simplest cases of imaging soft tissues, the principle concern will be to account properly for attenuation as the ultrasound propagates. For a single transducer pulsed measurement made at moderate pressures, it may be sufficient simply to rescale the measured UOT data according to the distance travelled by the pulse at the instant the data is recorded. Measurements with a focused (pulsed, or CW) field may require some level of simulation to ensure that the region of interest is appropriately insonified.

The very presence of contrast in ultrasound images of soft tissues informs us that such media are acoustically inhomogeneous. Typical acoustic reflection coefficients between soft tissues soft tissues should not cause problems in maintaining suitable amplitudes for driving acousto-optic modulation, but as per the case of absorption, simulation may be required to properly understand the effects on a given focal region. In the case of an imaging system integrated with an ultrasound scanner, one possibility might be to use the information from a standard ultrasound scan to attempt to approximate the acoustic impedance mismatches, and thus the distortion of the acoustic field driving the UOT measurement.

6.2 Summary

Ultrasound modulated optical tomography is a hybrid imaging modality with numerous potential clinical applications. A great deal of effort has been expended by the research community in optimising the spatial selectivity of the technique, and in the development of more sensitive detection mechanisms. Significantly less effort has been applied to the problem of modelling and image reconstruction, especially when compared with more mature areas of research, such as diffuse optical tomography.

One reason for this disparity may be that the promise of UOT lies in the use of an acoustic field to control the spatial resolution of the technique, and thus unlike areas which rely upon image reconstruction to achieve their results, such as DOT, this aspect of the research has been of a lower priority. In fact, image reconstruction is essential if we wish to recover quantitative information regarding the optical properties of the medium.

By employing some of the latest techniques in parallel computation, we developed a fast, flexible and efficient reference forward model by which we could evaluate the models of chapter 4 which were more suited for use as part of an image reconstruction process. Following

the work of others in the field, we improved the accuracy of a diffusion-style approximation for UOT by integrating the acousto-optic modulation with a more accurate model of light transport. Finally, we employed the diffusion-style approximation to the forward model in an image reconstruction method where we followed many of the approaches taken in DOT — not only benefiting from the rigorous methods of this field, but hopefully permitting the appreciation and development of this work by a wider audience.

Bibliography

- [1] Milton Abramowitz and Irene A Stegun. *Handbook of Mathematical Functions*. With Formulas, Graphs, and Mathematical Tables. Courier Dover Publications, 1964.
- [2] B J Ackerson, R L Dougherty, N M Reguigui, and U Nobmann. Correlation transfer: Application of radiative transfer solution methods to photon correlation problems. *Journal of thermophysics and heat transfer*, 6(4):577–588, 1992.
- [3] Erik Alerstam, Stefan Andersson-Engels, and Tomas Svensson. White Monte Carlo for time-resolved photon migration. *Journal of Biomedical Optics*, 13(4):041304–041304–10, 2008.
- [4] 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):060504–060504–3, 2008.
- [5] T J Allen and P C Beard. Light emitting diodes as an excitation source for biomedical photoacoustics. *SPIE BiOS*, 8581:85811F–85811F–9, February 2013.
- [6] Moritz Allmaras and Wolfgang Bangerth. Reconstructions in ultrasound modulated optical tomography. *Journal of Inverse and Ill-posed Problems*, 19(6).
- [7] Raphael Aronson and Noel Corngold. Photon diffusion coefficient in an absorbing medium. *Journal of the Optical Society of America A*, 16(5):1066, 1999.
- [8] Simon R Arridge. A finite element approach for modeling photon transport in tissue. *Medical Physics*, 20(2):299–309, 1993.
- [9] Simon R Arridge. Photon-measurement density functions. Part I: Analytical forms. *Applied Optics*, 34(31):7395–7409, 1995.
- [10] Simon R Arridge. Optical tomography in medical imaging. *Inverse Problems*, 15(2):R41–R93, January 1999.
- [11] Simon R Arridge and W R B Lionheart. Nonuniqueness in diffusion-based optical tomography. *Optics Letters*, 23(11):882–884, 1998.

- [12] Simon R Arridge and John C Schotland. Optical tomography: forward and inverse problems. *Inverse Problems*, 25(12):123010, December 2009.
- [13] Simon R Arridge and M Schweiger. Photon-measurement density functions. Part 2: Finite-element-method calculations. *Applied Optics*, 34(34):8026–8037, 1995.
- [14] Simon R Arridge and M Schweiger. Image reconstruction in optical tomography. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 352(1354):717–726, June 1997.
- [15] Simon R Arridge, Hamid Dehghani, Martin Schweiger, and Eiji Okada. The finite element model for the propagation of light in scattering media: A direct method for domains with nonscattering regions. *Medical Physics*, 27(1):252, 2000.
- [16] Guillaume Bal and John C Schotland. Inverse Scattering and Acousto-Optic Imaging. *Physical Review Letters*, 104(4):043902, January 2010.
- [17] Michele Bassani, Fabrizio Martelli, Giovanni Zaccanti, and Daniele Contini. Independence of the diffusion coefficient from absorption: experimental and numerical evidence. *Optics Letters*, 22(12):853, 1997.
- [18] P Beard. Biomedical photoacoustic imaging. *Interface Focus*, 1(4):602–631, June 2011.
- [19] Bruce J Berne and Robert Pecora. *Dynamic Light Scattering*. With Applications to Chemistry, Biology, and Physics. Courier Dover Publications, 2000.
- [20] M Bertero and P Boccacci. *Introduction to Inverse Problems in Imaging*. Taylor & Francis, December 2010.
- [21] 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.
- [22] Florian J Blonigen, Alex Nieva, Charles A DiMarzio, Sébastien Manneville, Lei Sui, Gopi Maguluri, Todd W Murray, and Ronald A Roy. Computations of the acoustically induced phase shifts of optical paths in acoustophotonic imaging with photorefractive-based detection. *Applied Optics*, 44(18):3735, 2005.
- [23] D A Boas and A G Yodh. Spatially varying dynamical properties of turbid media probed with diffusing temporal light correlation. *Journal of the Optical Society of America A*, 14(1):192, 1997.
- [24] D A Boas, D H Brooks, E L Miller, C A DiMarzio, M Kilmer, R J Gaudette, and Quan Zhang. Imaging the body with diffuse optical tomography. *IEEE Signal Processing Magazine*, 18(6):57–75, 2001.

- [25] David Boas. A fundamental limitation of linearized algorithms for diffuse optical tomography. *Optics Express*, 1(13):404, 1997.
- [26] Emmanuel Bossy, Lei Sui, Todd W Murray, and Ronald A Roy. Fusion of conventional ultrasound imaging and acousto-optic sensing by use of a standard pulsed-ultrasound scanner. *Optics Letters*, 30(7):744–746, April 2005.
- [27] A Bratchenia, R Molenaar, and R P H Kooyman. Acousto-optic spectroscopy as a tool for quantitative determination of chemical compounds in tissue: a model study. In *Biomedical Optics (BiOS) 2007*, pages 64371P–64371P–10. SPIE, February 2007.
- [28] Aliaksandr Bratchenia, Robert Molenaar, Ton G van Leeuwen, and Rob P H Kooyman. Acousto-optic-assisted diffuse optical tomography. *Optics Letters*, 36(9):1539, 2011.
- [29] Jenghwa Chang, Raphael Aronson, Harry L Graber, and Randall L Barbour. Imaging diffusive media using time-independent and time-harmonic sources: dependence of image quality on imaging algorithms, target volume, weight matrix, and view angles. In *Photonics West '95*, pages 448–464. SPIE, May 1995.
- [30] Michael Chu and Hamid Dehghani. Image reconstruction in diffuse optical tomography based on simplified spherical harmonics approximation. *Optics Express*, 17(26):24208–24223, 2009.
- [31] D L Colton and R Kress. *Integral equation methods in scattering theory*. John Wiley & Sons Inc, New York, May 1983.
- [32] 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, 1997.
- [33] B T Cox, J G Laufer, and P C Beard. The challenges for quantitative photoacoustic imaging. *SPIE BiOS: Biomedical Optics*, 7177:717713–717713–9, February 2009.
- [34] Ben Cox, Jan G Laufer, Simon R Arridge, and Paul C Beard. Quantitative spectroscopic photoacoustic imaging: a review. *Journal of Biomedical Optics*, 17(6):0612021–06120222, June 2012.
- [35] R L Dougherty, B J Ackerson, N M Reguigui, F Dorri-Nowkoorani, and U Nobbmann. Correlation transfer: Development and application. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 52(6):713–727, 1994.
- [36] T Durduran, A G Yodh, B Chance, and D A Boas. Does the photon-diffusion coefficient depend on absorption? *Journal of the Optical Society of America A*, 14(12):3358–3365, 1997.

- [37] D J Durian. The diffusion coefficient depends on absorption. *Optics Letters*, 23(19):1502, 1998.
- [38] Daniel S Elson, Rui Li, Christopher Dunsby, Robert Eckersley, and Meng-Xing Tang. Ultrasound-mediated optical tomography: a review of current methods. rsfs.royalsocietypublishing.org.
- [39] Q Fang and D A Boas. Tetrahedral mesh generation from volumetric binary and grayscale images. *Biomedical Imaging: From Nano to Macro*, 2009.
- [40] Qianqian Fang. Mesh-based Monte Carlo method using fast ray-tracing in Plücker coordinates. *Biomedical Optics Express*, 1(1):165–175, 2010.
- [41] 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, 2009.
- [42] Shechao Feng, Fan-An Zeng, and Britton Chance. Photon migration in the presence of a single defect: a perturbation analysis. *Applied Optics*, 34(19):3826, 1995.
- [43] Martin Frank, Axel Klar, Edward W Larsen, and Shugo Yasuda. Time-dependent simplified PN approximation to the equations of radiative transfer. *Journal of Computational Physics*, 226(2):2289–2305, October 2007.
- [44] K Furutsu and Y Yamada. Diffusion-Approximation for a Dissipative Random Medium and the Applications. *Physical Review E*, 50(5):3634–3640, November 1994.
- [45] Sanjiv Sam Gambhir. Molecular imaging of cancer with positron emission tomography. *Nature Reviews Cancer*, 2(9):683–693, September 2002.
- [46] E M Gelbard, J Davis, and J Pearson. Iterative solutions to the P_1 and double P_1 equations. *Nuclear science and engineering*, pages 36–34, 1959.
- [47] A P Gibson, J C Hebden, and Simon R Arridge. Recent advances in diffuse optical imaging. *Physics in Medicine and Biology*, 50(4):R1–R43, February 2005.
- [48] R Graaff and JJ Ten Bosch. Diffusion coefficient in photon diffusion theory. *Optics Letters*, 25(1):43–45, 2000.
- [49] R Graaff, M H Koelink, F F M de Mul, W G Zijistra, A C M Dassel, and J G Aarnoudse. Condensed Monte Carlo simulations for the description of light transport. *Applied Optics*, 32(4):426–434, 1993.
- [50] Gabriele Gratton, Monica Fabiani, Paul M Corballis, Donald C Hood, Marsha R Goodman-Wood, Joy Hirsch, Karl Kim, David Friedman, and Enrico Gratton. Fast

and Localized Event-Related Optical Signals (EROS) in the Human Occipital Cortex: Comparisons with the Visual Evoked Potential and fMRI. *NeuroImage*, 6(3):168–180, October 1997.

[51] M Gross, P Goy, and M Al-Koussa. Shot-noise detection of ultrasound-tagged photons in ultrasound-modulated optical imaging. *Optics Letters*, 28(24):2482–2484, December 2003.

[52] Sonny Gunadi and Terence S Leung. Spatial sensitivity of acousto-optic and optical near-infrared spectroscopy sensing measurements. *Journal of Biomedical Optics*, 16(12):127005, 2011.

[53] Jenni Heino, Simon Arridge, Jan Sikora, and Erkki Somersalo. Anisotropic effects in highly scattering media. *Physical Review E*, 68(3):031908, September 2003.

[54] L G Henyey and J L Greenstein. Diffuse radiation in the galaxy. *Astrophysical Journal*, 93(1):70–83, January 1941.

[55] Nam T Huynh, Diwei He, Barrie R Hayes-Gill, John A Crowe, John G Walker, Melissa L Mather, Felicity R A J Rose, Nicholas G Parker, Malcolm J W Povey, and Stephen P Morgan. Application of a maximum likelihood algorithm to ultrasound modulated optical tomography. *Journal of Biomedical Optics*, 17(2):026014, 2012.

[56] Akira Ishimaru. *Wave Propagation and Scattering in Random Media*. Wiley-IEEE Press, February 1999.

[57] Akira Ishimaru and S T Hong. Multiple scattering effects on coherent bandwidth and pulse distortion of a wave propagating in a random distribution of particles. *Radio Science*, 10(6):637–644, 1975.

[58] F F Jobsis. Noninvasive, infrared monitoring of cerebral and myocardial oxygen sufficiency and circulatory parameters. *Science*, 198(4323):1264–1267, December 1977.

[59] Qianqian Fang David R Kaeli. Accelerating mesh-based Monte Carlo method on modern CPU architectures. *Biomedical Optics Express*, 3(12):3223–3230, December 2012.

[60] Willi A Kalender. X-ray computed tomography. *Physics in Medicine and Biology*, 51(13):R29–R43, July 2006.

[61] Marleen Keijzer, Willem M Star, and Pascal R M Storchi. Optical diffusion in layered media. *Applied Optics*, 27(9):1820–1824, 1988.

[62] M Kempe, M Larionov, D Zaslavsky, and A Z Genack. Acousto-optic tomography with multiply scattered light. *Journal of the Optical Society of America A*, 14(5):1151–1158, 1997.

- [63] Andrew Kensler and Peter Shirley. Optimizing Ray-Triangle Intersection via Automated Search. In *2006 IEEE Symposium on Interactive Ray Tracing*, pages 33–38. IEEE.
- [64] Alwin Kienle and Michael S Patterson. Determination of the optical properties of turbid media from a single Monte Carlo simulation. *Physics in Medicine and Biology*, 41(10):2221–2227, January 1999.
- [65] A D Klose and E W Larsen. Light transport in biological tissue based on the simplified spherical harmonics equations. *Journal of Computational Physics*, 220(1):441–470, 2006.
- [66] 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):6096–6100, November 2010.
- [67] Puxiang Lai, Xiao Xu, and Lihong V Wang. Ultrasound-modulated optical tomography at new depth. *Journal of Biomedical Optics*, 17(6):0660061–0660066, June 2012.
- [68] E W Larsen, J M McGhee, and J E Morel. The simplified P_{sub N} equations as an asymptotic limit of the transport equation. *Transactions of the American Nuclear Society; (United States)*, 66, January 1992.
- [69] Edward W Larsen. Asymptotic solution of neutron transport problems for small mean free paths. *Journal of Mathematical Physics*, 15(1):75, 1974.
- [70] Edward W Larsen and J E Morel. Asymptotic solutions of numerical transport problems in optically thick, diffusive regimes II. *Journal of Computational Physics*, 83(1):212–236, July 1989.
- [71] Edward W Larsen, J E Morel, and Warren F Miller, Jr. Asymptotic solutions of numerical transport problems in optically thick, diffusive regimes. *Journal of Computational Physics*, 69(2):283–324, April 1987.
- [72] Terence S Leung. Light propagation in a turbid medium with insonified microbubbles. *Journal of Biomedical Optics*, 18(1):015002–015002, January 2013.
- [73] Terence S Leung and Samuel Powell. Fast Monte Carlo simulations of ultrasound-modulated light using a graphics processing unit. *Journal of Biomedical Optics*, 15(5):055007–055007–7, September 2010.
- [74] W Leutz and G Maret. Ultrasonic modulation of multiply scattered light. *Physica B: Condensed Matter*, 204(1-4):14–19, January 1995.
- [75] A Lev and B G Sfez. Direct, noninvasive detection of photon density in turbid media. *Optics Letters*, 27(7):473, 2002.

- [76] A Lev and B G Sfez. Pulsed ultrasound-modulated light tomography. *Optics Letters*, 28(17):1549–1551, 2003.
- [77] S Lévéque, A C Boccara, M Lebec, and H Saint-Jalmes. Ultrasonic tagging of photon paths in scattering media: parallel speckle modulation processing. *Optics Letters*, 24(3):181–183, February 1999.
- [78] Sandrine Lévéque-Fort. Three-Dimensional Acousto-Optic Imaging in Biological Tissues with Parallel Signal Processing. *Applied Optics*, 40(7):1029, 2000.
- [79] Hui Li and Lihong V Wang. Autocorrelation of Scattered Laser Light for Ultrasound-Modulated Optical Tomography in Dense Turbid Media. *Applied Optics*, 41(22):4739–4742, 2002.
- [80] Jun Li, Geng Ku, and Lihong V Wang. Ultrasound-modulated optical tomography of biological tissue by use of contrast of laser speckles. *Applied Optics*, 41(28):6030–6035, October 2002.
- [81] Youzhi Li, Philip Hemmer, Chulhong Kim, Huiliang Zhang, and Lihong V Wang. Detection of ultrasound-modulated diffuse photons using spectral-hole burning. *Optics Express*, 16(19):14862–14874, 2008.
- [82] Youzhi Li, Huiliang Zhang, Chulhong Kim, Kelvin H Wagner, Philip Hemmer, and Lihong V Wang. Pulsed ultrasound-modulated optical tomography using spectral-hole burning as a narrowband spectral filter. *Applied Physics Letters*, 93(1):11111, July 2008.
- [83] André Liemert and Alwin Kienle. Light diffusion in a turbid cylinder I Homogeneous case. *Optics Express*, 18(9):9456, 2010.
- [84] André Liemert and Alwin Kienle. Light diffusion in a turbid cylinder II Layered case. *Optics Express*, 18(9):9266, 2010.
- [85] Jorge Ripoll Lorenzo. *Principles of diffuse light propagation: light propagation in tissues with applications in biology and medicine*. World scientific, New Jersey, NJ, 2012.
- [86] Iván Lux and László Koblinger. *Monte Carlo particle transport methods: neutron and photon calculations*. CRC Press, Boca Raton, FL, 1991.
- [87] F MacKintosh and Sajeev John. Diffusing-wave spectroscopy and multiple scattering of light in correlated random media. *Physical Review B*, 40(4):2383–2406, August 1989.
- [88] G D Mahan, W E Engler, JJ Tiemann, and E Uzgiris. Ultrasonic tagging of light: Theory. *Proceedings of the National Academy of Sciences*, 95(24):14015–14019, November 1998.

- [89] G Maret and P E Wolf. Multiple light scattering from disordered media. The effect of brownian motion of scatterers. *Zeitschrift für Physik B Condensed Matter*, 65(4):409–413, 1987.
- [90] Eduardo Margallo-Balbás and Patrick J French. Shape based Monte Carlo code for light transport in complex heterogeneous Tissues. *Optics Express*, 15(21):14086–14098, 2007.
- [91] Fay A Marks, Harold W Tomlinson, and Glen W Brooksby. Comprehensive approach to breast cancer detection using light: photon localization by ultrasound modulation and tissue characterization by spectral discrimination. In *OE/LASE'93: Optics, Electro-Optics, & Laser Applications in Science & Engineering*, pages 500–510. SPIE, September 1993.
- [92] George Marsaglia and Arif Zaman. A New Class of Random Number Generators. *The Annals of Applied Probability*, 1(3):462–480, August 1991.
- [93] Fabrizio Martelli, Daniele Contini, Adriana Taddeucci, and Giovanni Zaccanti. Photon migration through a turbid slab described by a model based on diffusion approximation. II. Comparison with Monte Carlo results. *Applied Optics*, 36(19):4600, 1997.
- [94] Fabrizio Martelli, Samuele Del Bianco, Andrea Ismaelli, and Giovanni Zaccanti. *Light Propagation Through Biological Tissue and Other Diffusive Media*. Theory, Solutions, and Software. SPIE-International Society for Optical Engineering, January 2010.
- [95] Michele Martinelli, Adam Gardner, David Cuccia, Carole Hayakawa, Jerome Spanier, and Vasan Venugopalan. Analysis of single Monte Carlo methods for prediction of reflectance from turbid media. *Optics Express*, 19(20):19627, 2011.
- [96] Kenneth S Miller. On the Inverse of the Sum of Matrices. *Mathematics Magazine*, 54(2):67–72, January 1981.
- [97] Michael I Mishchenko. Vector Radiative Transfer Equation for Arbitrarily Shaped and Arbitrarily Oriented Particles: a Microphysical Derivation from Statistical Electromagnetics. *Applied Optics*, 41(33):7114, 2002.
- [98] Michael I Mishchenko. Maxwell's equations, radiative transfer, and coherent backscattering: A general perspective. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 101(3):540–555, October 2006.
- [99] Michael I Mishchenko. Poynting–Stokes tensor and radiative transfer in discrete random media: the microphysical paradigm. *Optics Express*, 18(19):19770–19791, 2010.

- [100] Tomas Möller and Ben Trumbore. Fast, minimum storage ray-triangle intersection. *Journal of Graphics Tools*, 2(1):21–28.
- [101] Stephen P Morgan, Nam T Huynh, Haowen Ruan, and Felicity R Rose. Ultrasound-Modulated Optical Tomography. In *Optical Techniques in Regenerative Medicine*. CRC Press.
- [102] Todd W Murray, Lei Sui, Gopi Maguluri, Ronald A Roy, Alex Nieva, Florian Blonigen, and Charles A DiMarzio. Detection of ultrasound-modulated photons in diffuse media using the photorefractive effect. *Optics Letters*, 29(21):2509–2511, November 2004.
- [103] D J Pine, D A Weitz, P M Chaikin, and E Herbolzheimer. Diffusing wave spectroscopy. *Physical Review Letters*, 60(12):1134–1137, 1988.
- [104] G C Pomraning. Asymptotic and variational derivations of the simplified PN equations. *Annals of Nuclear Energy*, 20(9):623–637, September 1993.
- [105] Samuel Powell and Terence S Leung. Highly parallel Monte-Carlo simulations of the acousto-optic effect in heterogeneous turbid media. *Journal of Biomedical Optics*, 17(4):045002–04500211, 2012.
- [106] Samuel Powell and Terence S Leung. Image reconstruction techniques for ultrasound-modulated optical tomography. In *SPIE BiOS*, pages 857819–857819–10. SPIE, March 2013.
- [107] Scott Alan Prahl. *Light Transport in Tissue*. PhD thesis, The University of Texas at Austin.
- [108] V Raman and K S Venkataraman. Determination of the Adiabatic Piezo-Optic Coefficient of Liquids. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 171(945):137–147, May 1939.
- [109] F Ramaz, B Forget, M Atlan, A C Boccara, M Gross, P Delaye, and G Roosen. Photorefractive detection of tagged photons in ultrasound modulated optical tomography of thick biological tissues. *Optics Express*, 12(22):5469–5474, November 2004.
- [110] Steffen G Resink, Albert C Boccara, and Wiendelt Steenbergen. State-of-the art of acousto-optic sensing and imaging of turbid media. *Journal of Biomedical Optics*, 17(4):0409011–04090110, April 2012.
- [111] H Ruan, M L Mather, and S P Morgan. Pulsed ultrasound modulated optical tomography utilizing the harmonic response of lock-in detection. *Applied Optics*, 52(19):4755, 2013.

- [112] Haowen Ruan, Melissa L Mather, and Stephen P Morgan. Pulse inversion ultrasound modulated optical tomography. *Optics Letters*, 37(10):1658–1660, 2012.
- [113] Haowen Ruan, Melissa L Mather, and Stephen P Morgan. Pulsed ultrasound modulated optical tomography with harmonic lock-in holography detection. *Journal of the Optical Society of America A*, 30(7):1409–1416, 2013.
- [114] Sava Sakadžić and Lihong Wang. Ultrasonic modulation of multiply scattered coherent light: An analytical model for anisotropically scattering media. *Physical Review E*, 66(2):026603, August 2002.
- [115] Sava Sakadžić and Lihong V Wang. High-resolution ultrasound-modulated optical tomography in biological tissues. *Optics Letters*, 29(23):2770–2772, 2004.
- [116] Sava Sakadžić and Lihong V Wang. Modulation of multiply scattered coherent light by ultrasonic pulses: An analytical model. *Physical Review E*, 72(3):036620, September 2005.
- [117] Sava Sakadžić and Lihong V Wang. Correlation transfer equation for ultrasound-modulated multiply scattered light. *Physical Review E*, 74(3), 2006.
- [118] Sava Sakadžić and Lihong V Wang. Correlation transfer and diffusion of ultrasound-modulated multiply scattered light. *Physical Review Letters*, 96(16), 2006.
- [119] Sava Sakadžić and Lihong V Wang. Correlation transfer equation for multiply scattered light modulated by an ultrasonic pulse. *Journal of the Optical Society of America A*, 24(9):2797, 2007.
- [120] 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*, 29(10):2110, 2012.
- [121] J M Schmitt. Optical coherence tomography (OCT): a review. *Selected Topics in Quantum Electronics, IEEE Journal of*, 5(4):1205–1215, 1999.
- [122] John C Schotland, John C Haselgrave, and John S Leigh. Photon hitting density. *Applied Optics*, 32(4):448–453, 1993.
- [123] M Schweiger, Simon R Arridge, and D T Delpy. Application of the finite-element method for the forward and inverse models in optical tomography. *Journal of Mathematical Imaging and Vision*, 3(3):263–283, 1993.
- [124] M Schweiger, Simon R Arridge, and M Hiraoka. The finite element method for the propagation of light in scattering media: boundary and source conditions. *Medical Physics*, 1995.

[125] E M Sevick, J K Frisoli, C L Burch, and J R Lakowicz. Localization of absorbers in scattering media by use of frequency-domain measurements of time-dependent photon migration. *Applied Optics*, 33(16):3562, 1994.

[126] Jan Sikora, Athanasios Zacharopoulos, Abdel Douiri, Martin Schweiger, Lior Horesh, Simon R Arridge, and Jorge Ripoll. Diffuse photon propagation in multilayered geometries. *Physics in Medicine and Biology*, 51(3):497–516, January 2006.

[127] C Rebecca Simpson, Matthias Kohl, Matthias Essenpreis, and Mark Cope. Near-infrared optical properties of ex vivo human skin and subcutaneous tissues measured using the Monte Carlo inversion technique. *Physics in Medicine and Biology*, 43(9):2465–2478, September 1998.

[128] Michael Stephen. Temporal fluctuations in wave propagation in random media. *Physical Review B*, 37(1):1–5, January 1988.

[129] P Surya Mohan, T Tarvainen, M Schweiger, A 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, 2011.

[130] Ilias Tachtsidis, Terence S Leung, Anchal Chopra, Peck H Koh, Caroline B Reid, and Clare E Elwell. False positives in functional near-infrared topography. *Advances in experimental medicine and biology*, 645:307–314, 2009.

[131] Tanja Tarvainen. *Computational Methods for Light Transport in Optical Tomography*. PhD thesis, University of Kuopio.

[132] Tanja Tarvainen, Marko Vauhkonen, Ville Kolchmainen, and Jari P Kaipio. Hybrid radiative-transfer-diffusion model for optical tomography. *Applied Optics*, 44(6):876–886, 2005.

[133] Lihong Wang, Steven L Jacques, and Xuemei Zhao. Continuous-wave ultrasonic modulation of scattered laser light to image objects in turbid media. *Optics Letters*, 20(6):629, 1995.

[134] Lihong Wang, Steven L Jacques, and Liqiong Zheng. MCML—Monte Carlo modeling of light transport in multi-layered tissues. *Computer Methods and Programs in Biomedicine*, 47(2):131–146, July 1995.

[135] Lihong V Wang. Mechanisms of ultrasonic modulation of multiply scattered coherent light: a Monte Carlo model. *Optics Letters*, 26(15):1191, 2001.

[136] Lihong V Wang and Geng Ku. Frequency-swept ultrasound-modulated optical tomography of scattering media. *Optics Letters*, 23(12):975–977, 1998.

- [137] Ashley J Welch and Martin J C van Gemert. *Optical-Thermal Response of Laser-Irradiated Tissue*. Springer Science+Business Media, 2010.
- [138] Gang Yao, Shuliang Jiao, and Lihong V Wang. Frequency-swept ultrasound-modulated optical tomography in biological tissue by use of parallel detection. *Optics Letters*, 25(10):734–736, 2000.