Experimental and theoretical investigations of near infrared tomographic imaging methods and clinical applications.

Elizabeth M. C. Hillman
Department of Medical Physics and Bioengineering
University College London

Supervisors:
Dr Jeremy C. Hebden
Prof. David T. Delpy

Thesis submitted for the degree of Doctor of Philosophy (Ph.D.)
at the University of London

February 2002

ABSTRACT

This thesis describes investigations performed using a time-resolved optical tomographic medical imaging system. Optical tomography involves transilluminating a volume of tissue (e.g. an infant's head or an adult breast) using near infrared (NIR) light. The differing absorption characteristics of oxy- and deoxy- haemoglobin in this wavelength range can then be exploited to yield information about spatial variations in oxygenation and blood volume deep within the tissue. The distinctive scatter properties of tissues such as tumors, gray matter and white matter in the NIR range also provide diagnostic information.

A 32 channel time-correlated single photon counting instrument (MONSTIR), designed and built at University College London (UCL), is at first evaluated, and then used to perform a number of investigations into the capabilities and limitations of optical tomographic imaging. Simulations and image reconstructions are performed using the TOAST algorithm, developed at the UCL Department of Computer Science. The progression of the imaging system from laboratory prototype to clinical instrument is detailed.

Specifically, studies are presented demonstrating techniques for calibration of time-resolved data, along with validation of the TOAST forward model. Early phantom imaging trials are described with reference to problems associated with using 2D and 3D algorithms to reconstruct measured data. An investigation into the effects of static background structure on images of changes in properties is presented. Methods for derivation of functional parameters from time-resolved optical tomography are explored and demonstrated on clinical data acquired on the breast, forearm and neonatal head. An experimental study of the effects of non-scattering regions within diffusing media is presented, highlighting issues relating to refractive index mismatches. Finally, practical aspects of clinical data acquisition are addressed.

These investigations predominantly refer to general optical tomography issues rather than specifically relating to MONSTIR or TOAST. Findings and conclusions should be relevant to the wider biomedical optics community.

ACKNOWLEDGEMENTS

I would like to thank my supervisors Dr Jeremy Hebden and Prof. David Delpy for their support, encouragement and guidance throughout this PhD. I am indebted to Prof. David Delpy for originally inspiring me to pursue a career in Medical Physics and Bioengineering.

I would also particularly like to thank Prof. Simon Arridge, who engaged me in discussions that resulted in many of the ideas introduced in this thesis. I was both challenged and motivated by his encouragement and determination to strive for perfection.

This thesis would not have been possible without the contribution made by Dr Florian Schmidt, whose meticulous design and construction of the MONSTIR system provided such an excellent tool with which to perform the experiments described herein. Other members of the MONSTIR team who have contributed to aspects of this work include Dr Hamid Dehghani, Dr David Kirkby, Dr Martin Fry, Dr Adam Gibson, Dr Nick Everdell, Hylke Veenstra, Felipe Gonzalz, Rozarina Yusof, David Kaspar and Tara Bland.

I must express my gratitude to Dr Martin Schweiger and Prof. Simon Arridge for their work in conceptualizing, implementing and maintaining the TOAST image reconstruction algorithm, which was fundamental in much of this work. I would also like to credit Jason Riley for his contribution of 3D void-TOAST and his help in performing simulations and experiments relating to the void problem.

I would like to thank Dr David Boas, Dr Andreas Hielsher, Dr Joe Culver, Dr Harry Graber and Prof. Britton Chance for useful discussions relating to this work.

I thank my friends at the UCL Department of Medical Physics for their enduring support, for the biscuits, the discussions and the invaluable transfer of knowledge.

Finally I would like to thank Joy and Barry Hall, without whose boundless kindness, patience, generosity, food, lodgings and cups of mint tea, writing this thesis would have been impossible.

This work was generously supported by the Wellcome Trust. Additional funding was contributed by Action Research, EPSRC, The UCL Graduate School and Hamamatsu Photonics

PUBLICATIONS RESULTING FROM THIS WORK

Journal papers

***Hillman E. M. C**, Hebden J. C, Schweiger M, Dehghani H, Schmidt F. E. W, Delpy D. T, Arridge S. R. *"Time Resolved Optical Tomography of the Human Forearm":* Phys. Med. Biol. 46, 4, 1117-1130 (2001).

Hillman E. M. C, Hebden J.C, Schmidt F.E.W, Arridge S.R, Schweiger M, Dehghani H, Delpy D.T*, "Calibration techniques and datatype extraction for time-resolved optical tomography":* Rev Sci Instrum, 71, 9, 3415-3427 (2000).

Hebden J. C, Veenstra H, Dehghani H, **Hillman E. M. C**, Schweiger M, Arridge S. R, Delpy D. T, "*Three dimensional time-resolved optical tomography of a conical breast phantom":* Appl. Opt. **40** 19, 3278-3287 (2001).

Schmidt F.E.W, Hebden J.C, **Hillman E. M. C**, Fry M.E, Schweiger M, Delpy D.T , *"Multiple slice imaging of a tissue-equivalent phantom using time-resolved optical tomography":* Applied Optics, 39, 3380-3387 (2000).

Arridge S.R, Hebden J.C, Schweiger M, Schmidt F.E.W, Fry M.E, **Hillman E. M. C**, Dehghani H *"A method for 3D time-resolved optical tomography":* J. Imag. Sys. Tech**,** 11, 2-11 (2000).

Schmidt F.E.W, Fry M.E, **Hillman E. M. C**, Hebden J.C, and Delpy D.T*, "A 32-channel time-resolved instrument for medical optical tomography":* Review of Scientific Instruments 71, 256 - 265 (1999).

Hebden J.C, Schmidt F.E.W, Fry M.E, Schweiger M, **Hillman E. M. C**, Delpy D.T, Arridge S. R *"Simultaneous reconstruction of absorption and scattering images using multi-channel measurement of purely temporal data":* Optics Letters, 24, 534-536 (1999).

Conference proceedings

Hillman E. M. C, Arridge S. R. Hebden J. C, Delpy D. T, *"Processing multi-wavelength time-resolved optical tomography data to derive functional parameters"*: in Biomedical Tops, OSA Tech Digest. (in press 2002).

Hillman E. M. C, Dehghani H, Hebden J. C, Arridge S. R, Schweiger M, Delpy D. T, *"Differential imaging in heterogeneous media: limitations of linearization assumptions in optical tomography":* Proc SPIE 4250 (2001).

***Hillman E.M.C**, Hebden J.C, Schmidt F.E.W, Arridge S.R , Fry M.E, Schweiger M, Delpy D.T, *"Initial Clinical testing of the UCL 32 channel time-resolved instrument for optical tomography":* in Biomedical Tops, OSA Tech Digest. 100-102 (2000).

Hebden J. C, **Hillman E. M. C**, Gibson A, Everdell N, Yusof R. Md, Delpy D. T, *"Time-resolved optical tomography of the newborn infant brain: initial clinical results"*: in Biomedical Tops, OSA Tech Digest. (in press 2002).

Hebden J. C, Bland T, **Hillman E. M. C**, Gibson A, Everdell N, Delpy D. T, *"Optical tomography of the breast using a 32-channel time-resolved imager"*: in Biomedical Tops, OSA Tech Digest. (in press 2002).

Riley J. D, Arridge S. R, Chrysanthou Y, Dehghani H, **Hillman E. M. C**, Schweiger M, "*The Radiosity Diffusion Model in 3D*": SPIE Proc 4431, (2001).

Dehghani H, **Hillman E. M. C,** Schweiger M, Arridge S. R, and Delpy D. T, *"Optical tomography of a 3D multilayered head model",* Proc SPIE ***4250****, 45-52. (2001).*

Hebden J. C, Schmidt F. E. W, **Hillman E. M. C**, Arridge S. R, Schweiger M, and Delpy D. T. *"Optical tomography of the human forearm":* in Proceedings of the Inter-Institute Workshop on In Vivo Optical Imaging at the NIH, ed. A. H. Gandjbakhche (OSA, Washington DC, 2000) pp. 167-171.

Posters

Hillman E. M. C, Arridge S. R, Schweiger M, Dehghani H, Kirkby D, Delpy D. T. *"Evaluation of time-resolved optical tomography techniques using in-vivo and complex phantom studies",* Gordon Research conference (Lasers in Medicine and Biology) in Connecticut.

*Reproduced in appendix B (2.10).

TABLE OF CONTENTS

2. Optical tomography investigations..... 1.3—58

LIST OF ACRONYMS

ADF	Acquisition definition file
CCD	Charge coupled device
CFD	Constant fraction discriminator
CT	Computed tomography
DA	Diffusion approximation
DPF	Differential pathlength factor
EEG	Electroencephalograph
EIT	Electrical impedance tomography
FEM	Finite element method
FDM	Finite difference method
FMRI	Functional MRI
FWHM	Full width half maximum
HIE	Hypoxic ischaemic encephalopathy
ICG	Indocyanine green
IRF	Instrument response function
IVH	Intraventricular haemorrhage
MCP-PMT	Micro-channel plate photomultiplier tube
MONSTIR	Multi-channel opto-electronic near infrared system for time resolved image reconstruction
MRI	Magnetic resonance imaging
NICU	Neonatal intensive care unit
PC	Personal computer
PDE	Partial differential equation
PET	Positron emission tomography
PMDF	Photon measurement density function
PMT	Photomultiplier tube
PTA	Picosecond time analyser
RDS	Respiratory distress syndrome
RTE	Radiative transfer equation
TCSPC	Time correlated single photon counting
TPSF	Temporal point spread function
TOAST	Temporal optical absorption and scattering tomography
UCL	University College London

1. Introduction to Optical Tomography

Medical imaging is a widely expanding field of development and research. Techniques such as x-ray computed tomography (CT) and more recently magnetic resonance imaging (MRI) have revolutionised diagnosis and treatment of a wealth of illnesses. The contrast presented by a medical image is a consequence of the type of interactions occurring between the probing radiation and the different tissues that compose the body. X-rays yield images where bones are prominent since x-rays are far less absorbed by most other tissues in the body. MRI images examine the response of molecules to changes in magnetic filed, and ultrasound images depend on the acoustic mismatches between adjacent tissues.

Optical tomography is a new medical imaging technique that uses near infrared (NIR) light as the probing radiation. A major absorber of NIR in tissues is the haemoglobin in blood. The diagnostic potential of NIR was first identified by (Horecker, 1943), when the absorption characteristics of oxygenated and deoxygenated blood were found to be quite different at NIR wavelengths. Subsequently, it was found that the oxygenation status of brain tissue could be determined from measurements of its absorption at a number of NIR wavelengths (Jobsis, 1977). Changes over time of the concentrations of oxy- and deoxy- haemoglobin (and cytochrome C oxidase (Chance, 1966)) in the brains of neonates were first presented in (Ferrari *et al*, 1985) and (Ferrari *et al*, 1986).

A NIR *imaging system* that could provide maps of the oxygenation status of tissue within a living brain (or other part of the body) would reveal information about the functional and metabolic activity of tissues previously inaccessible. Applications considered in this thesis are imaging the neonatal brain, the adult breast and the forearm. The neonatal brain is a suitable application for optical tomography since new-born infants (particularly those born prematurely) can suffer permanent brain damage as a result of peri-natal illnesses and conditions which affect the supply of oxygen to brain tissue. Since their heads are small, NIR light can be used to probe functional parameters of tissue deep within the brain of a new-born infant. An advantage of using NIR light to create an imaging system is that it is safe, non

ionising and does not require a shielded room in which to acquire data (unlike x-ray and MRI methods), and so can be performed repeatedly at the bedside. Hence NIR can be used to monitor the progression of conditions and potentially allow new treatments to be developed, and the response to treatment to be monitored. Detecting abnormalities in the function of tissues should allow earlier intervention than most other imaging modalities which are unable to detect events that don't manifest as a gross change in tissue structure or composition. The only documented previous attempt to perform optical tomographic imaging of the neonatal brain using a time-resolved instrument is detailed in (Hintz *et al*, 1999). Breast imaging using NIR light would reveal information about the functional, absorbing and scattering properties of suspicious lesions. Current breast screening and diagnostic techniques are either invasive, uncomfortable or provide little specificity. Optical tomography offers the opportunity to identify the nature of a lesion without the need for compression of the breast or the use of ionising radiation or biopsy. Arm imaging allows the efficacy of optical tomography methods to be evaluated since functional changes in the arm can be readily induced and its size is similar to that of an infant's head.

Creating an imaging system that uses NIR light is not straightforward due to the scattering properties of tissue. Because the light quickly becomes diffuse as it traverses tissue, it is not possible to deduce the internal absorbing structure by simply looking at a number of projections (as is the case in x-ray imaging). Further, scattering structure will attenuate the light in a similar way to absorbing features. By looking at the intensity of light emerging from the tissue it is near impossible to separate the effects of absorption and scatter.

To overcome the problems of scatter in optical imaging (and also dynamic range issues relating to simple intensity measurements), a number of approaches have been investigated. The method described in this thesis involves measuring the temporal behaviour of light that has passed through tissue. This is achieved by using a pulsed laser as a light source rather than a simple continuous wave (CW) source. As the pulses propagate through tissue they become broader in time because, as the light scatters, some light takes a fairly direct path and some takes much longer paths, depending on the number of scatterers present. The absorbing features in the tissue also affect the shape of the pulse since the amount of light absorbed depends on the length of the path that it has travelled. By measuring the shape of the broadened pulses that have propagated through the tissue, we should be able to derive as much information as possible about the scattering and absorbing features within the volume.

An instrument to measure the shape of these pulses of light simultaneously at 32 positions was designed and built at UCL as described in (Schmidt, 1999). The instrument is known as the "multi-channel opto-electronic near infra-red system for time-resolved image reconstruction" (MONSTIR). This thesis describes the basic features of the MONSTIR

system along with methods developed for its calibration. Experiments are presented where MONSTIR is utilised to investigate the potential and constraints of optical tomography.

Producing images from *time-resolved* measurements is a complex process. An image reconstruction algorithm was developed using the finite element method (FEM) and the diffusion approximation. The algorithm can perform non-linear image reconstructions to find 3D distributions of the absorption and scatter properties within an object (Schweiger, 1994). The algorithm is known as the "time-resolved optical absorption and scatter tomography" (TOAST) algorithm. This thesis documents the use of this algorithm to reconstruct images from experimentally acquired data. In some cases, simpler algorithms and reconstructions of simulated data are also utilised.

The MONSTIR project began in 1996, with a team of around 5 people working on the system at any one time. Dr Martn Schweiger and Prof. Simon Arridge have been developing TOAST since around 1990. Over the past 3 years, the MONSTIR instrument has progressed from having only 8 detector channels, and it's first ever data acquisition, to being used in the University College Hospital neonatal intensive care unit to acquire data on two premature babies. The team project required each member to contribute to tasks such as designing and building phantoms, optical, electronic and hardware components, and fine-tuning, maintaining and upgrading the MONSTIR instrumentation.

This thesis is presented in two halves. The first half introduces the basic principals of optical tomography. This includes description of the MONSTIR measurement system and the TOAST image reconstruction algorithm. The second half of this thesis presents a series of investigations, which utilise MONSTIR and TOAST as tools to investigate various aspects of optical tomographic imaging in general. Each investigation identifies a particular aspect of data manipulation, data acquisition or image reconstruction that potentially presents an obstacle to successful clinical optical tomographic imaging. Contributions made by other members of the research team to aspects of these investigations are acknowledged where necessary, although in the most part the investigations were conceptualised and composed by the author. The implications of each study generate a framework within which clinical optical tomography may be successful. Developed techniques are demonstrated via studies of epoxy resin phantoms, the breasts of healthy adult volunteers, the exercising forearm, and ultimately experiments are described where data were acquired on the brains of two healthy premature infants.

This introductory section begins with a description of the optical properties of tissues (chapter 1.1). Firstly the basic optical properties of media are introduced, followed by description of how these relate to the optical properties of bulk tissues. The origin of optical

contrast in the chosen applications of neonatal brain imaging, breast imagng and arm imaging are then detailed. The relevant anatomy and pathology of common conditions are illustrated for each case.

Chapter 1.2 describes existing NIR techniques and instrumentation. These are presented in terms of the geometries used to acquire data, and then in terms of the types of instrumentation utilised. The way in which each type of measurement can be used to extract clinically relevant information is detailed. The MONSTIR system is then described, with reference to features that are relevant to subsequent investigations and studies.

Chapter 1.3 introduces the basis of optical tomographic image reconstruction. Methods commonly utilised to model the propagation of light through diffuse media are described. The way in which these models can then be used to reconstruct images is discussed with reference to the techniques utilised in subsequent investigations. Some specific features of TOAST are described, particularly the origin and characteristics of ‘datatypes’.

An introduction to the specific individual studies performed to investigate optical tomography methods can be found at the beginning of section 2.

1.1 Optical properties of tissues

This section introduces the basic theory of the interaction of light with tissue. Each optical property is described in turn. The bulk optical properties of composite tissues are then described along with the specific characteristics of water and haemoglobin. Finally the origin of contrast in clinical optical imaging is described with reference to brain, breast and arm imaging and their associated anatomy and pathologies.

1.1.1 Optical properties

The optical properties of a tissue are primarily described in terms of:

- Absorption
- Scatter
- Anisotropy
- Refractive index

All of these parameters will be dependent on the wavelength of the light probing the tissue. Additional optical characteristics of tissues include their fluorescence and inelastic scatter properties although the measurements and analysis in this thesis are limited to the effects of absorption and elastic scatter (and diffraction and reflection).

Below we describe the basic principles of absorption, scatter, anisotropy and refractive index. We will then relate these simple concepts to the optical properties of bulk tissues. The origins of optical contrast in the chosen optical tomographic imaging applications of neonatal brain imaging and breast imaging are then detailed.

1.1.1.1 Absorption

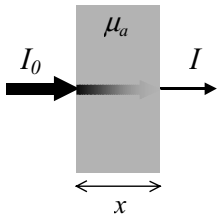

Figure 1.1.1 Attenuation of light through a non-scattering medium.

If a medium is non-scattering and illuminated with a collimated beam of light of intensity I_0 and wavelength λ the intensity of the emerging light will be:

$$I = I_0 e^{-\mu_a(\lambda)x}$$

[1.1.1]

where $\mu_a(\lambda)$ is the absorption coefficient of the medium and x is the width of the sample as illustrated in Figure 1.1.1. The absorption coefficient represents the probability per unit length of a photon being absorbed. The μ_a of a medium may be due to a number of absorbing substances (chromophores) mixed together. The individual extinction coefficients of each chromophore represent their absorption at a particular concentration. The absorption coefficient of a mixture of chromophores can be expressed as the sum of the products of the concentration of each chromophore c_n with its extinction coefficient ε_n.

$$\mu_a(\lambda) = \sum_n \varepsilon_n(\lambda) c_n$$

[1.1.2]

1.1.1.2 Scatter

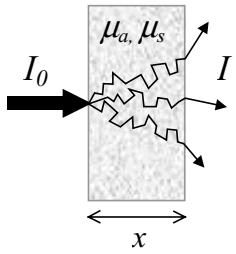

Figure 1.1.2 Attenuation of light through a scattering medium.

The scattering properties of a medium are described by its scattering coefficient μ_s. The scattering coefficient is the product of the number density of scattering particles and the scattering cross section of the particles. Therefore μ_s represents the probability per unit length of a photon being scattered. Compared to the non-scattering case above, if a medium is scattering the paths taken by the photons traversing it are no longer direct. This has two effects:

- We can no longer detect *all* the emerging photons unless the detector can collect over all angles and at all points on the surface of the medium.
- The photons will have travelled varying distances through the scattering media.

If the medium is absorbing (which is invariably true) the increased distance travelled by photons (known as the pathlength) will attenuate them in accordance with [1.1.1] (where x => DPFx, where DPF = the differential pathlength factor). The additional distance travelled by each photon (given by the DPF) will depend on how many scattering events it has encountered. DPF is a function of the scattering coefficient μ_s, the anisotropy of the scatter g, the absorption of the medium and also of the geometry of the medium. Attenuation through simple scattering materials can hence be described by the modified Beer Lambert law:

$$I = I_0 e^{-\mu_a(\lambda) DPF x + G}$$

[1.1.3]

where G represents losses due to the geometry. Since G is very difficult to quantify, most simple spectroscopy looks at changes in absorption given by:

$$\Delta A_{(2-1)} = Log\left(\frac{I_1}{I_2}\right) = DPFx(\Delta\mu_{a(2-1)})$$

[1.1.4]

where $\Delta A_{(2-1)}$ is the change in attenuation measured between state 2 and state 1 corresponding to an absorption change of $\Delta\mu_{a,(2-1)}$, assuming that all other terms are constant for the two measurements. DPF can be approximately calculated based on μ_a, μ'_s and the geometry of the object (Arridge *et al*, 1992) or determined, for example, via measurement of the mean-time taken for light to traverse a scattering medium <t> using:

$$DPF = \frac{c < t >}{xn}$$

[1.1.5]

where n is the refractive index of the material (see below) and c is the speed of light. Note however that the DPF is a function of μ'_s and μ_a (as is <t>), and is therefore wavelength dependent, which can make its use complicated in spectral analysis.

Light propagation through more complex scattering media can be described both analytically and numerically using the diffusion approximation to the radiative transfer equation as described in more depth in chapter 1.3.

1.1.1.3 Anisotropy

The scattering properties of bulk media are often described in terms of the reduced scattering coefficient μ'_s. When light scatters off a particular scatterer, the light generally emerges in a preferential direction, relative to its incident angle as shown in Figure 1.1.3.

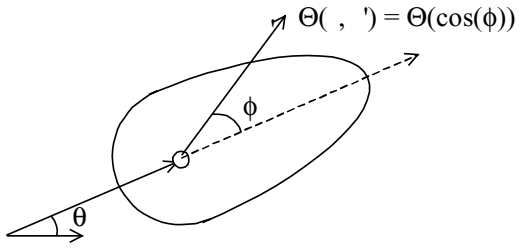

Figure 1.1.3 The scattering phase function Θ is often anisotropic.

The phase function Θ is a function of the angle between the incident and scattered light ϕ. It is assumed that Θ is independent of θ, the orientation of the scatterer, such that Θ is dependent only on the scalar product of ŝ and ŝ'. The mean cosine is given by:

$$g = \int_{4\pi} \Theta(\cos(\phi))\cos(\phi)d\phi$$

[1.1.6]

The scatter coefficient of material μ_s is a function of the number of scatterers. The anisotropy of the scatter given by the mean cosine or Landé g-factor *g* allows us to simplify directional effects by using the reduced scatter coefficient μ'_s given by:

$$\mu'_s = \mu_s(1-g)$$

[1.1.7]

The scattering cosine depends on the size, shape and refractive index mismatches of the scatterers. If g = 0 then the scatter is isotropic, and if g = 1, the scatter is entirely forward.

1.1.1.4 Refractive index

Refractive index describes the effect that a medium has on the speed of light traversing through it. The speed change when light passes between media of different refractive indices causes its direction to change as shown in [1.1.7].

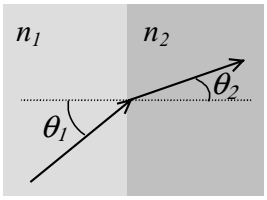

Figure 1.1.4 Refraction at the interface between two materials with different refractive indices.

Snell's law states that $n_1 \sin\theta_1 = n_2 \sin\theta_2$,

[1.1.8]

where *n* is the refractive index of the medium and θ is the angle between the propagation direction and the normal to the interface. The speed of light in a medium is given by:

$$c = \frac{c_{vacuum}}{n}$$

[1.1.9]

1.1.2 Optical properties of bulk tissues

The microscopic optical properties of tissue constituents do not directly translate to the macroscopic behaviour of light in the tissue. For example a red blood cell contains absorbing haemoglobin and scatters light due to its shape and the refractive index mismatch between it and the surrounding liquid. The bulk optical properties of whole blood will depend on how many blood cells are present, their orientation and the properties of the blood plasma.

Measuring the optical properties of tissues is difficult, particularly in-vivo. Methods of measuring and deriving optical properties from experimental measurements are detailed in section 1.2. Table 1.1.1 shows a selection of published optical properties measured ex-vivo using a number of different methods. These illustrate the wide variation in measured parameters due both to differing experimental preparation and measurement methods, and natural inter-subject variability.

The optical properties measured represent the *average* properties of the tissue i.e. the values of μ_a and μ'_s that would yield the measured data if the tissue were composed only of a substance with those properties. In reality the tissues are composed of many different components each of which has an effect on light propagation. For example, the amount, and oxygenation state of blood in the tissue will have a large effect on bulk absorption and scatter properties. Commonly ex-vivo samples are exsanguinated (blood is removed). Therefore, the optical properties of the tissues in vivo may be quite different to those listed below.

Tissue	Absorption coefficient μ_a (mm^{-1})	Reduced scatter coefficient μ'_s (mm^{-1})	g
Blood (arterial) (98%, 2m mol/L) (from measurements below)	0.398 (at 800 nm)	1 (at 800 nm)	-
Blood (venous) (75%, 2m mol/L) (from measurements below)	0.396 (at 800 nm)	1 (at 800 nm)	-
Blood partially oxygenated (Cheong *et al*, 1990)	1.55 (at 760 nm)	0.79 (at 760 nm)	0.9972
Brain (24 wk gest neonate) (van-der-Zee, 1992)	0.0215 ± 0.02 (at 800 nm)	0.748 (at 800 nm)	0.718
Brain (28 wk gest neonate) (van-der-Zee, 1992)	0.0373 ± 0.01 (at 800 nm)	0.673 (at 800 nm)	0.919
Brain (White matter 40 wk gest neonate) (van-der-Zee, 1992)	0.0373 (at 800 nm)	0.659 (at 800 nm)	0.981
Brain (Grey matter 40 wk gest neonate) (van-der-Zee, 1992)	0.0460 (at 800 nm)	0.529 (at 800 nm)	0.983
Brain (grey adult) (Cheong *et al*, 1990)	0.27± 0.02 (at 632.8 nm)	2.06± 0.2(at 632.8 nm)	0.94±0.004
Brain (white adult) (Cheong *et al*, 1990)	0.22± 0.02 (at 632.8 nm)	9.1± 0.5(at 632.8 nm)	0.82±0.01
Bone (pig skull) (Firbank *et al*, 1993)	0.025 (at 800 nm)	1.8 (at 800 nm)	-
Cerebrospinal fluid (CSF) (from water (Hale *et al*, 1973))	0.0022 (at 800 nm)	~0	-
Skin: Dermis (inc Epithelium) (Simpson *et al*, 1998)	0.013 (at 800 nm)	2 (at 800 nm)	-
Skin: Subdermis (Simpson *et al*, 1998)	0.008 (at 800 nm)	1.2 (at 800 nm)	-
Muscle (exsanguinated abdominal) (Simpson *et al*, 1998)	0.03 (at 800 nm)	0.7 (at 800 nm)	-
Breast (glandular) (Cheong *et al*, 1990)	0.05±0.01 (at 700 nm)	1.4±0.3 (at 700 nm)	--
Breast (adipose) (Cheong *et al*, 1990)	0.07±0.01(at 700 nm)	0.9±0.1 (at 700 nm)	--

Table 1.1.1 Published optical properties of common tissues (all ex-vivo)

The refractive index of each tissue is not shown since it is not easy to measure. It is generally assumed to be between that of lipid and water ($n \approx 1.4$).

Figure 1.1.5 shows that in the NIR range, melanin and haemoglobin are major absorbers compared to water. Scatter is also relatively low. (Note that absorption values of tissues at ~ 700nm in Table 1.1.1 will generally be significantly higher than they would be at 800nm). Water and haemoglobin spectra are considered in more detail below.

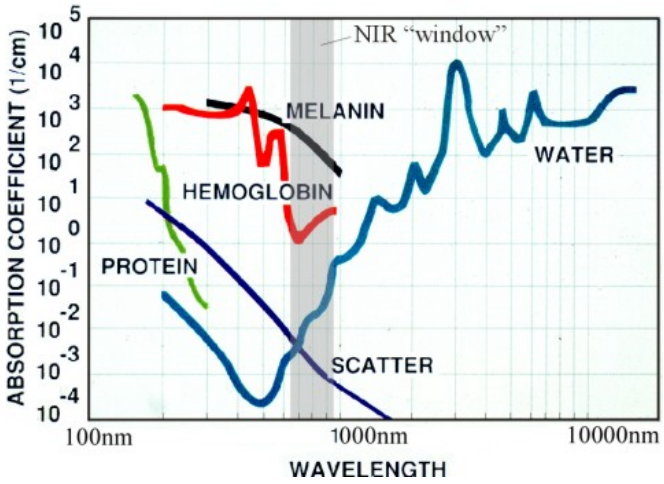

Figure 1.1.5 General optical properties of common tissues.

1.1.2.1 Water absorption

Water is a major constituent of tissues (Woodard *et al*, 1986). While it is not very absorbing in the NIR range (~0.0022mm^{-1} at 800nm) its wavelength dependence can affect spectroscopic measurements due to its abundance. Figure 1.1.6 demonstrates that over the 780nm-815nm wavelength range relevant to this thesis the water spectrum varies between 0.0021 and 0.0026mm^{-1}.

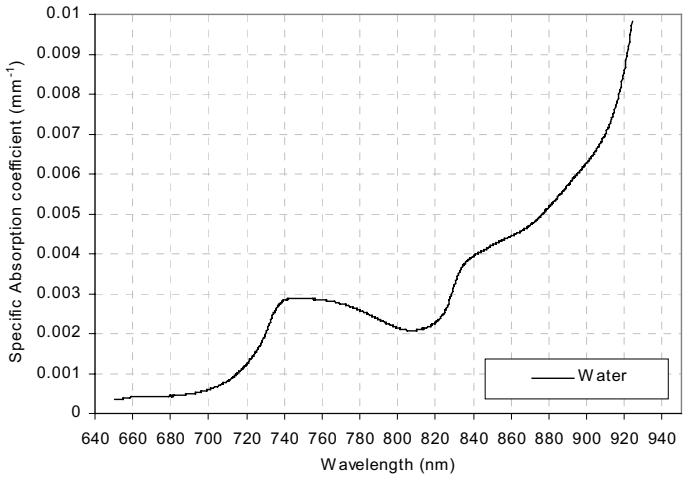

Figure 1.1.6 The absorption spectrum of waterin the NIR range measured by Veronica Hollis (Hollis et al, 2001)

1.1.2.2 Haemoglobin absorption

Healthy functioning of all tissues depends on an adequate blood supply, delivering sufficient oxygen. The blood content of a larger structure (e.g. muscle) will be significant, and the oxygenation of this blood will indicate functional activity. In the normal adult, blood is composed of ((Johnston, 1998), (Fox, 1999)):

- White blood cells and platelets <1% by volume (leucocytes - 1 per 100 rbc)
- Red blood cells. (~44% by volume (adult)) (1 cm^3 blood = 5 billion rbc)
- plasma (~55% by volume)

Red blood cells can be considered simply as small sacks containing haemoglobin. They deliver oxygen to tissues by attaching to oxygen in the lungs and becoming oxy-haemoglobin (HbO_2). At the tissue, the oxygen dissociates to leave deoxy- haemoglobin (Hb). The relative concentrations of oxy- and deoxy- haemoglobin in the blood tells us how well oxygenated the blood is. The oxygenation of blood in tissues is related to tissue blood supply and flow and the oxygen demand and usage of the tissue. A measure of the oxygenation of blood is the oxygen saturation (SO_2) given by:

$$SO_2 = \frac{[HbO_2]}{[HbO_2 + Hb]}$$

[1.1.10]

where [*y*] denotes the concentration of *y*. Over the NIR wavelength range, the absorption of blood is not only fairly significant (see Figure 1.1.5), but the absorptions of oxy- and deoxy-haemoglobin are different as shown in Figure 1.1.7.

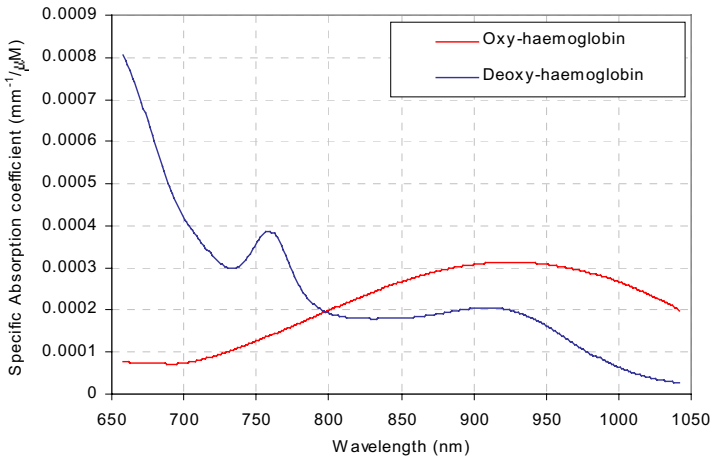

Figure 1.1.7 The absorption spectra of oxy- and deoxy-haemoglobin in the near infrared wavelength range.

These spectra can be converted into the absorption coefficient of whole blood by assuming a value for the concentration of haemoglobin (or haematocrit), and the saturation (which determines the relative concentrations of Hb and HbO_2). These values will depend on age, sex, state of health and the tissue that the blood is in. The haemoglobin concentration in whole blood is around 1.62 to 2.65 mM (assuming molecular weight of Hb = 68 000) and being higher for neonates than adults (and higher for men than women). Saturation of arterial blood is generally 94 to 99% and 55 to 70% for venous blood (Cope, 1991).

This means that if we can measure the absorption coefficient μ_a of a mixture of oxy- and deoxy-haemoglobin at two wavelengths, then we should be able to resolve their relative concentrations, and hence the SO_2 of the sample, given that we know the extinction coefficients of oxy- and deoxy- haemoglobin ($\varepsilon_{HbO_2}(\lambda)$ and $\varepsilon_{Hb}(\lambda)$). Using [1.1.2], for whole blood, we can see that:

$$\begin{aligned}\mu_{780nm} &= \varepsilon_{HbO_2}(780nm)\times c_{HbO_2} + \varepsilon_{Hb}(780nm)\times c_{Hb}\\ \mu_{820nm} &= \varepsilon_{HbO_2}(820nm)\times c_{HbO_2} + \varepsilon_{Hb}(820nm)\times c_{Hb}\end{aligned} \qquad [\ 1.1.11\ ]$$

So by simply solving simultaneous equations and using [1.1.10] we get:

$$SO_2 = \frac{\varepsilon_{Hb}(780nm)\times\mu_{820nm} - \varepsilon_{Hb}(820nm)\times\mu_{780nm}}{(\varepsilon_{Hb}(780nm)\times\mu_{820nm} - \varepsilon_{Hb}(820nm)\times\mu_{780nm}) + (\varepsilon_{HbO_2}(820nm)\times\mu_{780nm} - \varepsilon_{HbO_2}(780nm)\times\mu_{820nm})} \qquad [\ 1.1.12\ ]$$

This is a simplistic model that will work for whole blood if attenuation measurements along with the modified Beer-Lambert law are used (see section 1.1.1.2). Methods of spectroscopic analysis for optical tomographic imaging are explored in section 2.5, where the calculation of saturation and fractional blood volume within more complex tissues are considered.

1.1.3 The origin of contrast in clinical optical imaging

1.1.3.1 Neonatal brain

In terms of contrast, it is the *relative* values of absorption and scatter of different tissues at optical wavelengths that are important for imaging. For brain tissues, we can anticipate the contrast between different static structures from the relative optical properties of each tissue. Figure 1.1.8 illustrates the documented absorption and scatter properties of the major tissues in the neonatal brain from which we can infer contrast (at ~800nm, see Table 1.1.1). The values imply that the major absorber in brain images is likely to be grey matter, whereas the skull is likely to be the strongest component of a scatter image.

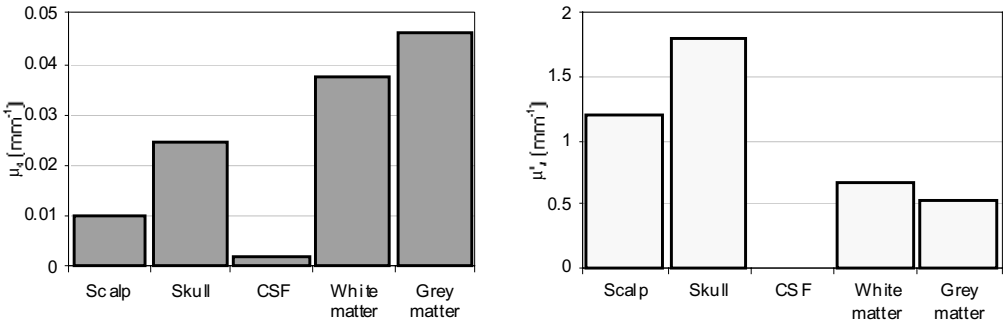

Figure 1.1.8 Approximate absorbing and scattering optical properties of the basic tissues of the head.

Figure 1.1.9 shows the skull of a term neonate. The plates of the skull are not fused and fontanels exist, composed of fibrous tissue rather than bone. The skull is soft and thinner than in an adult, and therefore likely to transmit NIR more readily than the adult skull. The face-to-cranium ratio is 1 to 8 rather than 1 to 2 in the adult. A plot of head circumference with gestational age is shown in appendix A (2.9.1) Typically, head diameter is ~11.2cm for a 40 week (term) male and 8.4cm for a 28 week gestation infant. Attenuation is exponential with distance ([1.1.3]) so infants with smaller heads will be more suitable for transillumination measurements.

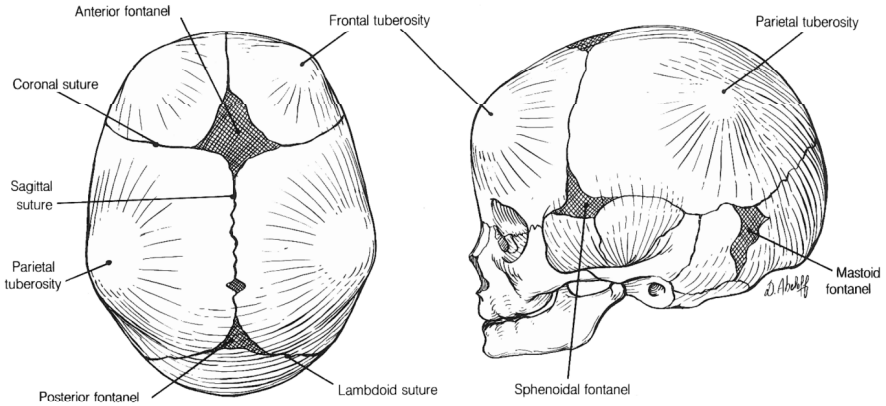

Figure 1.1.9 The skull of a term neonate (reproduced from (Hall-Craggs, 1995))

Figure 1.1.10 shows a cross section of the adult cerebral hemispheres. The brain of a neonate will have fewer convolutions in the peripheral grey matter (cortex) as shown in Figure 2.6.3 (section 2.6). White matter is composed of nerve fibres (axons) transmitting signals between the nerve cell nuclei (soma) that compose the grey matter. The axons are scattering which is why the white matter appears white. The brain is surrounded by cerebrospinal fluid (CSF), which also fills the ventricles (Figure 1.1.10). CSF and the layered structure of the head are described in more detail in section 2.6.

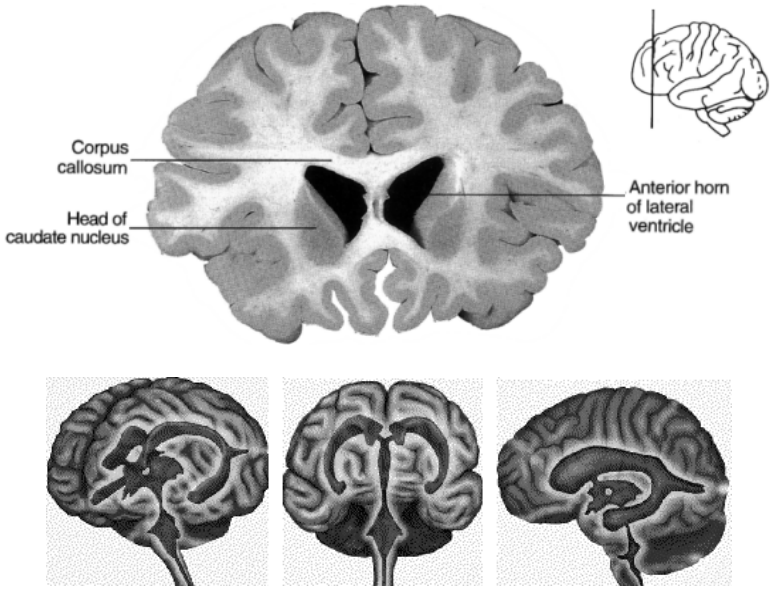

Figure 1.1.10 (top) A coronal slice of adult brain showing normal distribution of grey and white matter in the cerebral hemispheres and the lateral ventricles (reproduced from (Crossman, 1995) (p102))(bottom) The positions of the CSF filled ventricles within the adult brain

A detailed description of conditions likely to be encountered in the neonatal unit is given in (Hillman, 2001a). The different types of changes that occur in brain tissue with common illnesses can be summarised as:

	Caused by	Likely to manifest as
Haematomas/haemorrhages	Ruptured blood vessels	A collection of blood, usually clotted
Intra-ventricular haemorrhage (IVH)	Ruptured blood vessels around ventricles	Collection of blood around or in the ventricles
Ischaemia (low blood flow)	IVH (in periventricular regions); Infarct (stroke); birth asphyxia; poor autoregulation and perfusion; hydrocephalus	Region with decreased blood volume, decreasing oxygenation.
Hypoxia (low oxygenation)	Compromised lung function; asphyxia; and ischaemia.	Low oxygenation, blood flow may be normal, high or low if ischaemic.
Scatter changes	Degeneration of tissue e.g. following ischaemia in periventricular white matter	Changes in scatter and possible development of cysts.
Functional activation	Stimulus (e.g. visual, motor)	Slight increase in blood flow, volume and oxygenation in corresponding area of cortex.

Table 1.1.2 Cause and physical manifestation of changes in brain tissue in neonates

Premature infants can suffer from a number of neurological conditions due to poor lung development (respiratory distress syndrome - RDS (Johnston, 1998) (p128)) and birth asphyxia (hypoxic ischaemic encephalopathy - HIE (Merenstein *et al*, 1998) (p596)) and as a result they are susceptible to subarachnoid and intraventricular haemorrhage (IVH). Premature infants also have generally poor regulatory control of blood pressure, blood flow and breathing. Term infants can also suffer from birth asphyxia, other traumatic injuries at birth or congenital abnormalities (Johnston, 1998) (p143).

The conditions detailed in Table 1.1.2 were identified since they represent changes in tissue that can potentially be detected using optical methods. Optical measurements are sensitive to both scatter and absorption. Multi-wavelength measurements can be used to determine haemoglobin concentration of tissues and its saturation. An example where an optical imaging system could play and important role in diagnosis and monitoring is IVH as shown in Figure 1.1.11. Not only could optical imaging potentially determine the grade of the haemorrhage, evaluate damage due to parenchymal haemorrhage, and monitor response to treatment, but continuous monitoring could potentially be used to characterise onset of the condition.

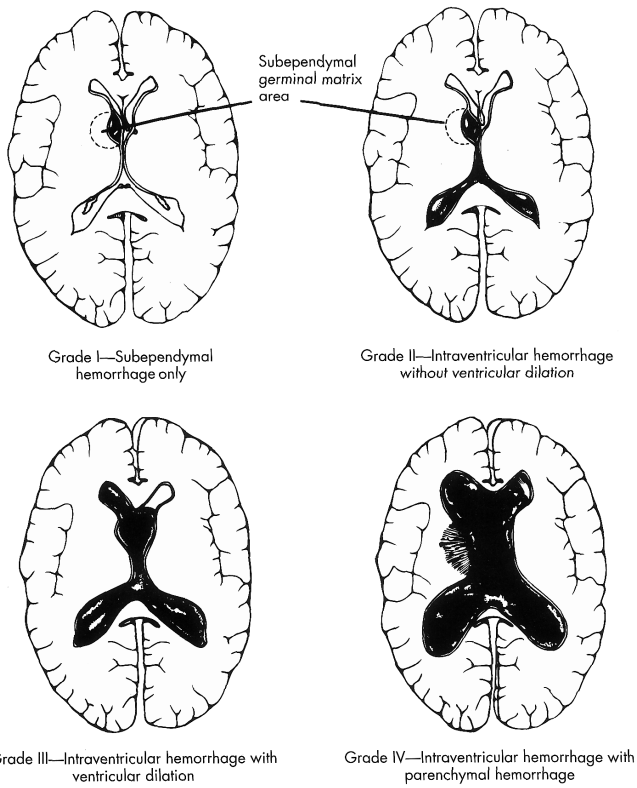

Figure 1.1.11 Periventricular - intraventricular haemorrhage grades 1-4(Merenstein *et al*, 1998)*(p559) (Hydrocephalus grades 3 and above)*

A common application of optical methods is to monitor functional responses to stimulus (see Table 1.1.2). This has a spatial aspect, since different areas of the cortex relate to different tasks as shown in Figure 1.1.12.

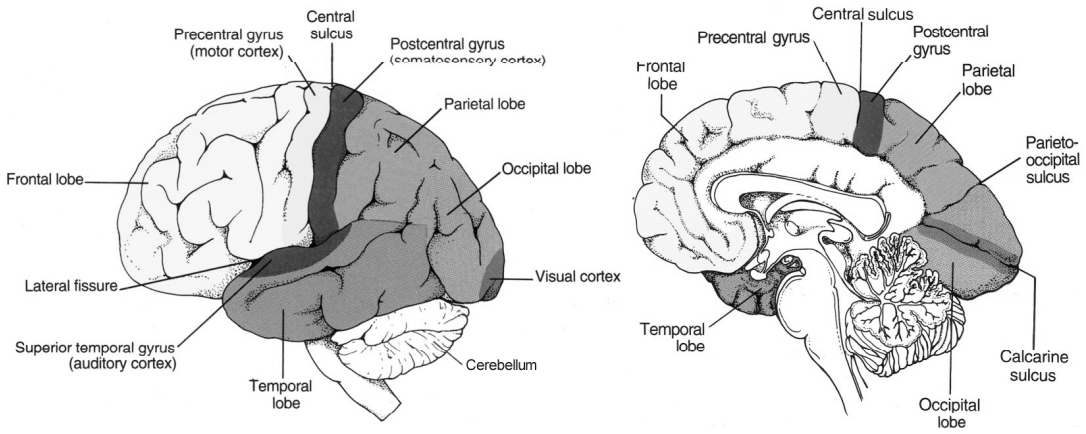

Figure 1.1.12 Regions of the adult cortex: ***Frontal Lobe****: Reasoning, planning, parts of speech, emotions, and problem-solving****, Motor cortex****: Co-ordination of complex movement, initiation of voluntary movement,* ***Somatosensory cortex:*** *Receives tactile information from the body,* ***Parietal Lobe:*** *perception of stimuli related to touch, pressure, temperature and pain ,****Occipital Lobe:*** *Visual processing,* ***Temporal Lobe****: Perception and recognition of auditory stimuli and memory,* ***Cerebellum:*** *(not strictly cortex) Learned processes, co-ordination.*

1.1.3.2 Breast

Optical tomography has the potential to identify the nature of suspicious lesions in the breast that have already been located by the patient or during screening. It is unlikely that optical methods will be used for routine screening in the immediate future since x-ray mammography is currently widely used and well established, despite its disadvantages of using ionising radiation and compression of the breast. Only following extensive testing and demonstration of accuracy and reliablitity could optical methods replace x-ray screening.

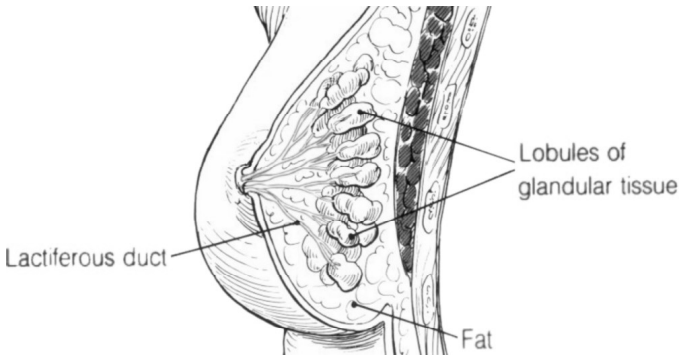

Figure 1.1.13 The breast is composed of ducts, lobes and adipose (fatty) tissue, along with connective structures. Most cancers are either associated with the lobes(Lobular) or the milk ducts (Ductal). Invasive or infiltrating tumours have broken through into the fatty tissue of the breast.

In the normal breast, glandular tissue has been found to be more scattering, but less absorbing than adipose tissue (see Table 1.1.1 and Figure 1.1.13). However, the optical properties of breast have been shown to depend on age and menstrual cycle (Cerussi *et al*, 2001). For imaging, optical contrast of malignant lesions in the breast may not have to rely on a recognisable change in μ_a or μ'_s in the tissue itself (Mourant *et al*, 1998). This is because malignant tumours are often associated with anomalous vasculature and different oxygenation to the surrounding tissue. In addition, common benign lesions have distinctive characteristics. These differences may be sufficient for diagnostic purposes as summarised in Table 1.1.3.

Condition	Type	Shape	Likely to manifest as
Cyst	Benign	Round and smooth	Low scatter
Blood filled cyst	Possibly malignant	Round and smooth	High / characteristic absorption, possibly low scatter.
Fibroadenoma	Benign	Round, mobile	High scatter, possibly high absorption, normal vasculature (common in younger women)
Fibrocystic/ fibroglandular changes	Benign	Boundaries not discrete	High scatter (common in older women)
Dormant tumour	Malignant	Small, within ducts or lobes	Possibly necrotic (de-oxygenated) core
Growing tumour	Malignant	Boundaries not discrete	Increased vasculature (hence increased absorption, scatter and anomalous oxygenation), may also have necrotic core.

Table 1.1.3 Different malignant and benign breast lesions and their potential optically detectable features.

Optical studies to characterise tumours in-situ are described in (Tromberg *et al*, 2000), (Fantini *et al*, 1998a). Tumours were found to predominantly manifest as regions of higher absorption than surrounding healthy tissue (due to angiogenesis (Weinstein, May 6 1994)). In general tumours were found to have lower oxygenation than surrounding tissue, although (Brizel *et al*, 1996) reports a correlation between tumour oxygenation and the presence of distant metastasises (secondary tumours). (Okunieff *et al*, 1993) suggest that lower oxygenation of tumours indicates effectiveness of radiation therapy. While the optical properties (including μ'_s) of the many types of malignancy that occur have not been categorically quantified, combined information about the scattering, absorbing, and functional characteristics of a suspicious lesion will clearly provide valuable diagnostic information.

1.1.3.3 Arm

Optical methods can be used to detect changes in blood volume, flow and oxygenation associated with muscle activation (Hillman *et al*, 2001b). The forearm is composed of densely packed muscles, two bones and skin, with veins around the perimeter and arteries running between the muscles as shown in Figure 1.1.14. Evaluation of functional changes in the arm with exercise is described in detail in chapter 2.5.5 .

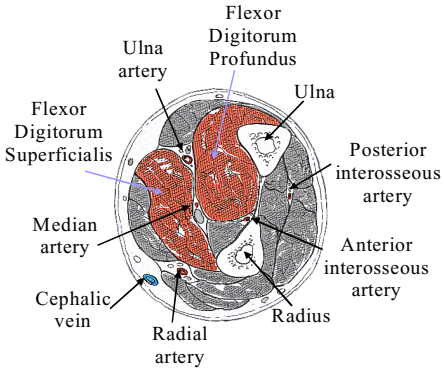

Figure 1.1.14 The basic anatomy of the forearm (modified and reproduced from (Hall-Craggs, 1995). Red regions represent the muscles activated during a simple finger flexing exercise.

1.2 Instrumentation and applications for optical techniques

Section 1.1.3 described that there are differences in the optical properties of tissues in the NIR range and that optical probing of tissues can be used to reveal anatomical, physiological and pathological features. The applications of NIR methods are therefore widely varied. This section presents a brief review of the different ways in which NIR light is applied to tissues first in terms of the geometry of the measurement, and then with reference to the type of light source and detection system utilised. The latter governs the parameters that can be usefully extracted from measurements. The former includes microscopy techniques, although only macroscopic geometries are detailed here. A more specific description is then given of the instrument that was used to make the measurements for investigations presented throughout this thesis, namely the UCL multi-channel time-resolved optical tomography system (MONSTIR).

1.2.1 Geometries

A NIR method will utilise a combination of a particular geometry, as described in section 1.2.1 and one of the source-detector types described in section 1.2.2 . The geometries commonly used to acquire NIR data on tissues are:

1) Single-point, for solely spectroscopic measurements e.g. of a tissue sample,
2) Global, for measurement of a large volume of tissue, e.g. an area of the brain,
3) Topographic, an imaging method where a number of measurements over the surface of the tissue map the optical properties of regions just below the surface.
4) Tomographic, imaging where measurements through a large volume of tissue are used to simultaneously quantifying the optical properties of different regions throughout the measured volume.

1.2.1.1 Single-point and global measurements

Simply placing a light source (or optical fibre carrying light) onto the surface of tissue, and detecting the transmitted light at some distance away via a fibre (or detector) can allow useful measurement of the optical properties of the tissue. The scattering nature of tissue means that a volume can be sampled in reflection or transmission geometries.

Figure 1.2.1 shows the sensitivity of intensity measurements made at the tissue surface to changes in the absorption coefficient within the tissue volume (see section 1.3.2.1). The depth

and size of the region sampled by a measurement depends on the shape of the object, the separation of the source and detector, and the optical properties of the tissue itself.

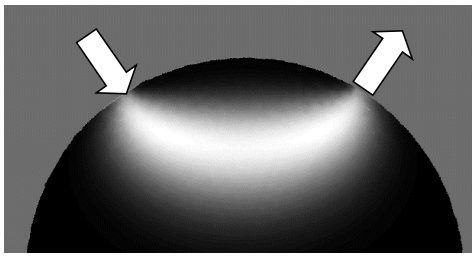

Figure 1.2.1 A source and detector placed onto the surface of tissue yield a measurement of the tissues volume's optical properties. The sensitivity of intensity measurements to absorption changes within the volume is shown. This sensitivity depends on the source-detector orientation and the background optical properties of the tissue.

Figure 1.2.2 (left) shows a single source and detector pair on an infant's head, and (right) shows the derived variation of oxy- and deoxy-haemoglobin in the brain during birth (Cope *et al*, 1988). This method is used for functional activation studies in infants and adults to evaluate responses to stimuli (e.g. visual, motor, see section 1.1.3.1) (Elwell *et al*, 1998), (Meek *et al*, 1995). Global measurements made on the breast have been used to characterise changes in optical properties with age and the menstrual cycle (Cerussi *et al*, 2001).

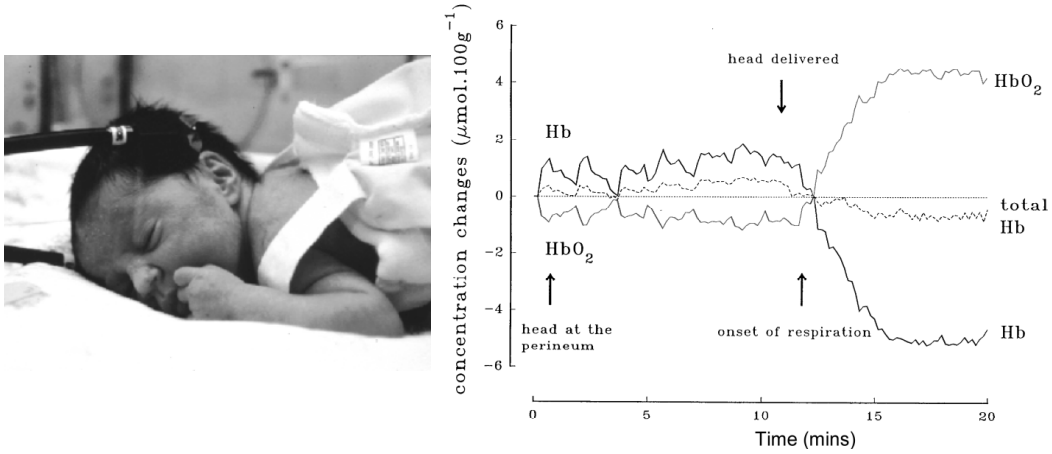

Figure 1.2.2 (Left), placement of optical fibres on a baby's head for global measurement. (Right) Plots of the temporal variation of cerebral Oxy- and Deoxy-haemoglobin during birth, measured using a single source / single detector set-up.

Single-point measurements are made in either reflection or transmission mode. In vivo, a probe composed of optical fibres can be used to measure the optical properties of specific tissues (e.g. brain tissue during surgery, through a needle into a lesion within the breast, or internal tissues via an endoscope (Bigio *et al*, 1996), (Mourant *et al*, 1995), (Mourant *et al*, 1998)). Ex-vivo measurements can utilise both modes to characterise excised tissue samples.

1.2.1.2 Topography

By utilising an array of sources and detectors placed on the surface of an object, the optical properties of the volumes underlying each source-detector pair can be mapped.

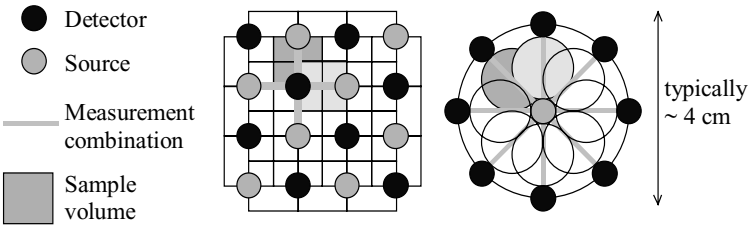

Figure 1.2.3 Different arrangements for arrays of sources and detectors for optical topography.

Figure 1.2.3 shows different source-detector arrangements for topography of the brain's cortex, as used by (Yamashita *et al*, 1996) (grid) and (Franceschini *et al*, 2000) (Circle). As indicated by Figure 1.2.1, the penetration depth of detected light depends on the distance between source and detector. The grid array in Figure 1.2.3 allows measurement of changes fairly close to the surface of the tissue (e.g. the cortex – through the skull) if close source-detector pairs are chosen. Measurement of deeper tissues can be achieved by examining signal between more distant points although, in general, signal-to-noise will worsen.

Compared to single-point and global measurements, creating a *map* of properties better demonstrates whether changes are occurring in areas where activation is expected, by simultaneously monitoring areas where no variation should be seen (Figure 1.2.4).

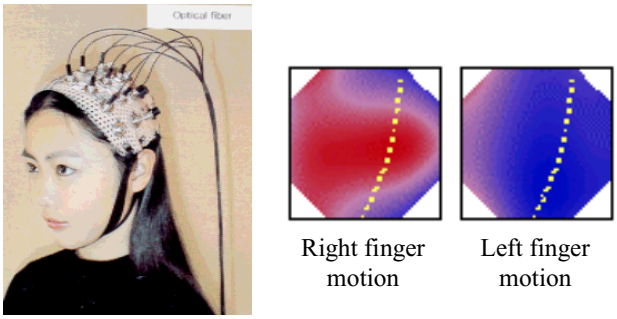

Figure 1.2.4 The grid geometry on a subject's head, and the topographical map for total haemoglobin seen over the left central sulcus (dotted line) for right and left finger tapping (Yamashita et al, 1996).

Problems with this method include calibration of individual detectors (if multiple detectors are used) and variations in the coupling coefficients for each source and detector due to bad contact and hair on the head. Where data must be acquired from multiple sources and / or detectors, a variety of protocols can be used:

- Serial acquisition (one source at a time, all detectors) can be slow, and since data are not acquired instantaneously time varying physiological processes may result in a spatially varying image. The advantage of the circular geometry shown in Figure 1.2.3 is that only one source is used.
- DC simultaneous acquisition (all sources on at once) requires assumptions about the amount of sources contributing to each detected signal (Vaithianathan *et al*, 2002).
- Frequency modulated simultaneous acquisition (all sources on at once). If each source is amplitude modulated at a different frequency (~KHz) parallel image acquisition is

possible since signal from each source to be separated via a Fourier transform of each measurement. (Yamashita *et al*, 1996).

Topographic techniques have been applied widely to brain monitoring applications, but also to exploring muscle activation for example on the forearm (Maris *et al*, 1994).

1.2.1.3 Tomography

Tomography involves employing a source - detector geometry that measures light that has travelled deep into the volume of tissue in order to produce an image of the internal structure. Figure 1.2.5 shows two possible arrangements for tomographic data acquisition.

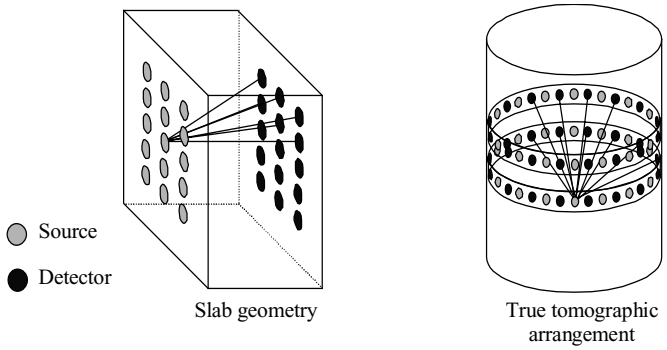

Figure 1.2.5 Source-detector arrangements for optical tomography

The slab geometry has been investigated by several researchers to image the (compressed) breast (Ntziachristos *et al*, 1998), (Fantini *et al*, 1998b), (Hebden *et al*, 1995) and (Rinneberg *et al*, 1998). In this geometry all detectors can be used for each source (if multiple detector channels are present), or a single source-detector pair can be raster scanned over both surfaces. In the tomographic arrangement, data are acquired in multiple directions, and generally all detectors acquire in parallel for each serially illuminated source. Raster scanning in a slab geometry is inefficient since it provides little depth resolution compared to acquiring data with multiple detectors for each source or better still using a tomographic arrangement. However, for the slab arrangement the dynamic range of the detectors need not be as large as for tomographic measurements where there are massive signal differences between near and far source-detector pairs.

Tomographic arrangements for breast imaging are being explored by (Hoogenraad *et al*, 1997), (McBride *et al*, in press 2001) and (Hebden *et al*, 2001). A comparison of the effects of geometry on image quality is given in (Pogue *et al*, 1999). Tomographic geometries for brain imaging have been investigated by (Eda *et al*, 1999) and (Hintz *et al*, 1999) and are also developed in this thesis. NIR tomographic methods have also been applied to imaging the finger to investigate the differing scattering properties of synovial fluid in arthritic joints (Xu *et al*, 2001), (Klose *et al*, 1998). Tomographic images of the arm have also been produced as described in (Hillman *et al*, 2001b) and (Graber *et al*, 2000).

1.2.2 Source and detector types

Description of NIR methods below is limited to applications that involve only absorption and elastic scattering of light (see section 1.1.1). Each source-detector type can in general be used in any of the geometries described above.

1.2.2.1 Continuous wave (CW)

If continuous wave (CW or DC) light is used as a light source for NIR methods then only the wavelength dependent attenuation of the emerging light can potentially be measured. Absorption changes can be calculated from attenuation changes using the modified Beer-Lambert law ([1.1.3]) for simple geometries. For more complex geometries and imaging applications, μ_a can be derived by comparing measured attenuation with an analytical or numerical model of light propagation (see chapter 1.3). μ'_s can also be calculated from attenuation, although a unique distribution of both absorption and scatter cannot be determined from a single attenuation measurement (Arridge *et al*, 1998c).

Sources composed of multiple wavelengths can yield absorption coefficients at multiple wavelengths. These data allow spectroscopic analysis of the tissue properties (although problems exist associated with extraction of μ_a when μ'_s and pathlength are wavelength dependent ((Kohl *et al*, 1998))).

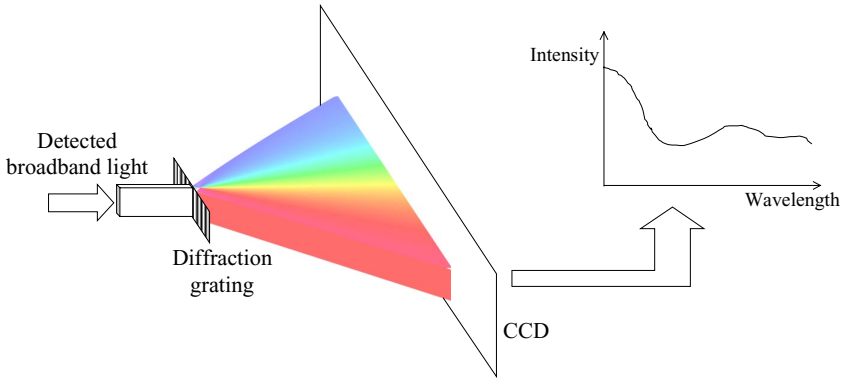

Figure 1.2.6 A CCD spectrometer uses a diffraction grating (more usually reflection) to split incoming white light into a spectrum, which is projected onto a CCD camera allowing the spectrum of the detected light to be recorded.

If a white light source is used, an almost continuous spectrum can be measured using a charge-coupled device (CCD) spectrometer as shown in Figure 1.2.6 (the spectra in Figure 1.1.6 and Figure 1.1.7 were measured in this way). By fitting a portion of the spectra of chromophores such as oxy- and deoxy- haemoglobin to a measured tissue spectrum, better accuracy in evaluation of chromophore concentrations can be achieved than if measurements are made at only a few wavelengths. A CCD spectrometer can record data at up to 100 Hz. Illumination intensities can be high since CW sources do not have to same safety limitations as lasers. Note however that detection of light through large thicknesses of tissue will be shot-noise limited. This technique is often used to acquire volume and single-point measurements

of tissue samples and time-series changes in functional parameters in the cerebral cortex (Meek *et al*, 1998), (Springett *et al*, 2000).

If a single detector is used (such as a photodiode or a photo-multiplier tube (PMT)), multi-wavelength measurements can made by either serially illuminating with monochromatic CW sources, or by frequency modulating (<20 kHz) a number of monochromatic CW sources and illuminating simultaneously. The attenuation of each wavelength can be derived from the amplitude of the Fourier transform of the data at each modulation frequency. A 32 channel tomographic optical imaging system which uses low-frequency modulated two-wavelength laser diode units for each source is detailed in (Schmitz *et al*, 2000). The gain of each individual photodiode is adjusted for each source position to account for the large dynamic range required for tomographic arrangements. The system can acquire image data sets at up to 20 Hz. A 256 channel CW tomographic breast imaging system is described in (Hoogenraad *et al*, 1997).

1.2.2.2 Frequency domain

If the light used to illuminate tissue is intensity modulated at ~MHz frequencies, in addition to being able to measure attenuation, a measure of the time taken for light to pass through tissue can be extracted. Comparison of the transmitted signal and the original signal reveals the phase difference between the two ($\Delta\phi$) (see Figure 1.2.7).

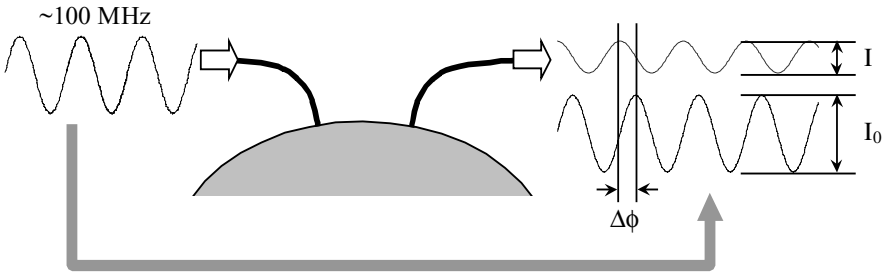

Figure 1.2.7 Frequency domain instrumentation measures amplitude and phase change during transmission of an intensity modulated signal through tissue.

Since the effects of absorption and scatter on phase are different, and different again from their effects on signal amplitude, using a frequency domain instrument offers the ability to separate absorption and scatter features. If the modified Beer-Lambert law is used to extract μ_a from attenuation, the phase part of the measurement can be used to evaluate pathlength (rather than requiring an estimate). As before, μ_a and μ'_s (global values and distributions) can also be extracted fitting data to analytic or numerical models. Since amplitude and phase are available, these models stand a better chance of determining μ_a and μ'_s than in the CW case. Signal-to-noise ratio (SNR) of attenuation measurements is likely to be better for frequency domain than CW measurements since a lock-in amplifier can be used

to reject DC signal e.g. due to room light. In reality phase information can be quite noisy and may be discarded.

The frequency chosen will govern the image (and temporal) resolution attainable, and the sensitivity to absorption and scatter although SNR is likely to worsen with higher frequencies (Boas, 1996). PMTs and photodiodes with a suitable frequency response can be used to detect emerging light. The frequencies used are usually around 100 MHz, although detection of frequencies up to a few GHz is now possible (Jacques *et al*, 1995), (Madsen *et al*, 1995), (Cerussi *et al*, 2001). The sources used are generally laser diodes (Pogue *et al*, 1997). Tissue is illuminated serially with different wavelengths if required. There are limitations on the amount of laser-light illumination allowable, however it is possible to acquire signal rapidly if sufficient photons are available. With large thicknesses of tissue, signal detection will be limited by shot noise.

Frequency domain instruments are also used in all detection geometries. A frequency domain global tissue measurement system is described in (Duncan *et al*, 1993), frequency domain topography systems are employed by (Franceschini *et al*, 2000) and (Fantini *et al*, 1998b). (Pogue *et al*, 1997) describes development of a frequency domain optical tomography system which utilises a number of laser diodes (of different wavelengths) modulated at 100 MHz, and a single PMT detector (which is serially illuminated by each detector fibre on a sliding stage).

1.2.2.3 Time-domain

A time-domain measurement represents measuring the time taken for individual photons to propagate through tissues. This can be achieved by illuminating tissue with a pulse of light, and using fast detectors to record how the pulse has broadened when it emerges. The broadened pulse (referred to as a temporal point spread function (TPSF)) represents a histogram of the time each photon took to travel from the source to the detector.

The integrated intensity (the total number of photons in the measured TPSF) represents a CW measurement. Frequency domain components can be extracted from the magnitude of the Fourier transform of a TPSF at the required frequency. A TPSF represents the tissue's impulse response function (IRF) which is the optimal measurement to characterise a system. MONSTIR is a time-resolved imaging system and is discussed in detail in section 1.2.3 .

μ_a and μ'_s can again be calculated for volume measurements using the modified Beer-Lambert law, measurement of the pathlength can readily be extracted from the mean-time of the detected photons ([1.1.5]). Fitting data to analytical and numerical models can also be used to derive global values and images of μ_a and μ'_s distributions (see chapter 1.3). Figure

1.2.8 illustrates how (in general) the optical properties of features inside an object will affect the shape of a TPSF.

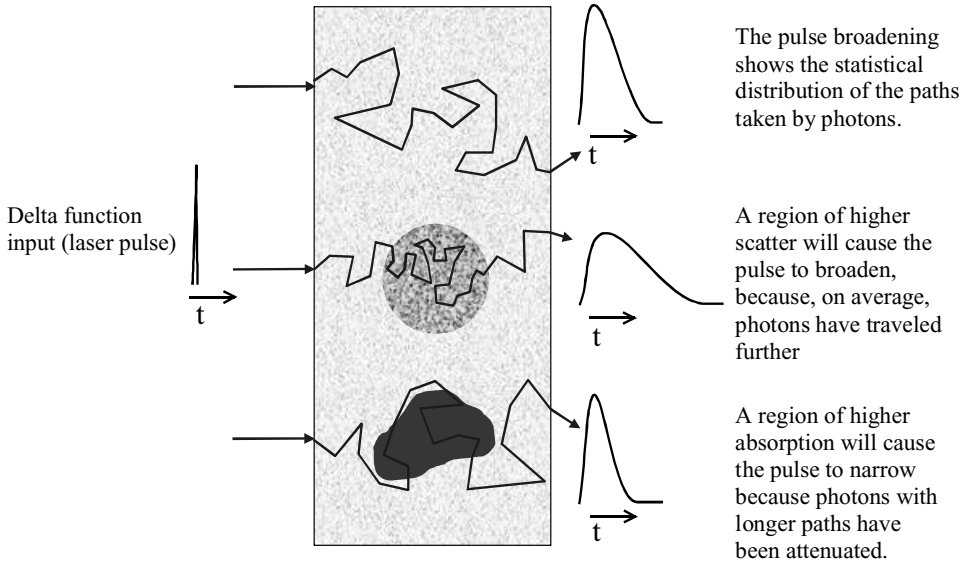

Figure 1.2.8 The broadening of a laser pulse passing through tissue is dependent on the internal structure of the object.

Pulsed lasers are used as the source light for time-resolved measurements. Such lasers would typically have a pulse width of a few ps, with a 40 – 80MHz repetition rate and around 40mW average power (although safe eye or skin exposure limits may require reduced power levels). Detection can be either analogue or digital. In an (analogue) streak camera, the broadened pulse of light is incident on a photocathode. The resulting electrons are swept across a phosphorescent screen that is imaged by a CCD camera. If the beam is swept at a frequency equivalent to the laser repetition rate, the CCD will record a line whose intensity profile is equivalent to the TPSF. A serial illumination single-channel optical tomography system which uses a streak camera is detailed in (Zint *et al*, 2001), in (Hebden *et al*, 1995) a streak camera is used as the detector in raster scanned transmission geometry tomography (see section 1.2.1.3). Pathlength ([1.1.3]) measurements made using a streak camera are evaluated in (Okada *et al*, 1995). However streak camera units are large, expensive and slow, and dark counts, limited illumination area and non-linearities make them unsuitable for use in a multi-channel system.

MONSTIR utilises digital time-resolved detection known as time-correlated single photon counting (TCPSC). Described in more detail below, TCSPC measures the time taken for individual photons to arrive at the detector and gradually builds up a histogram of their arrival times to form a TPSF. Such measurements can be slow and are inherently wasteful if many photons are available, but they are not affected by the dark counts and non linearities of the streak camera, and construction of a multi-channel TCSPC system is feasible.

Time-resolved measurements can distinguish between temporally correlated and temporally uncorrelated events. Temporally uncorreled noise such as thermal noise and background light should therefore be separable from measured signal. Increasing integration time for time-resolved measurements can thus improve SNR despite low photon flux. This is a significant advantage for tomographic optical imaging where very few photons will travel all the way through large tissue thicknesses (e.g. 10cm through breast). While time resolved techniques are rarely used for small volume measurements and topography due to the high cost of the instrumentation and the slow acquisition speed (~10 seconds per source position), they are well suited to tomographic imaging of large volumes of tissue (< 15 cm) that are not likely to vary significantly during the image acquisition period (~10 mins). (Hintz *et al*, 1998), (Eda *et al*, 1999) and (Schmidt *et al*, 2000a) describe time-resolved, true tomographic imaging systems. (Ntziachristos *et al*, 1998), (Grosenick *et al*, 1999) and (Rinneberg *et al*, 1998) describe transmission-geometry tomographic time-resolved optical mammography systems. A time-resolved topographic system is described in (Ntziachristos *et al*, 1999a).

1.2.3 The MONSTIR system

The MONSTIR system is a multi-channel *time-resolved* instrument for *tomographic* imaging of tissues as defined above. The design specifications and construction of the MONSTIR instrument, along with an initial system performance evaluation are extensively documented elsewhere (Schmidt, 1999), (Schmidt *et al*, 2000a). This section provides a brief overview of the MONSTIR system, detailing specific features of relevance to this thesis.

1.2.3.1 System overview

As described above (Figure 1.2.8), time-resolved instruments measure TPSFs. In order to create an image we need to measure a number of TPSFs between different sources and detectors in order to spatially sample the whole volume of the tissue. To achieve this, MONSTIR has 32 fibre bundles (for increased collection area, flexibility and low temporal dispersion) each leading to a separate detection system and 32 source fibres (multi-mode, graded index) multiplexed from the laser via a fibre switch. Thus up to 32 TPSFs can be acquired in parallel and by positioning the 64 fibres over the surface of an object (e.g. in a ring as shown below) and serially illuminating via each source fibre, 1024 TPSFs can be measured to constitute an image data set. The way that these TPSFs can then be used to create images is described in chapter 1.3.

Figure 1.2.9 shows the final design of the MONSTIR system. The basic photon-detection system is described below, followed by descriptions of a number of aspects of the system including the variable optical attenuators, the NIR laser and the computer control of hardware.

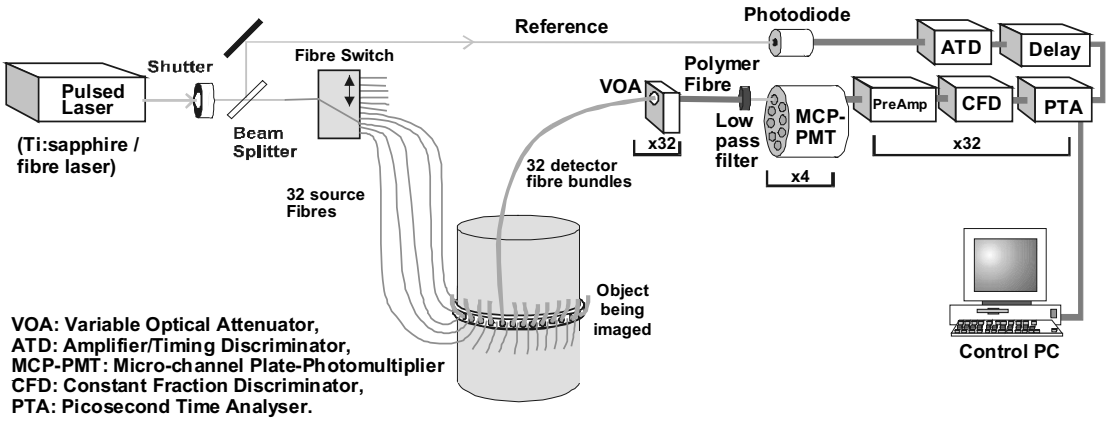

Figure 1.2.9 Schematic diagram of the MONSTIR system

1.2.3.2 Time correlated single photon counting (TCSPC) instrumentation.

MONSTIR utilises TCSPC to measure TPSFs. The laser light source in MONSTIR pulses at 80MHz. Every time a pulse enters the tissue being measured, it will emerge as a broadened pulse. Each fibre bundle will deliver the pulse corresponding to its position on the object to a fast detector. Each detector will only detect (less than) a single photon from each of these broadened pulses. TCSPC instrumentation converts signal corresponding to the detected photon into a digital pulse and then determines its temporal delay relative to a reference pulse. This temporal delay is recorded as a count in a histogram, which eventually, after multiple laser repetitions builds up into the TPSF, as shown schematically in Figure 1.2.10.

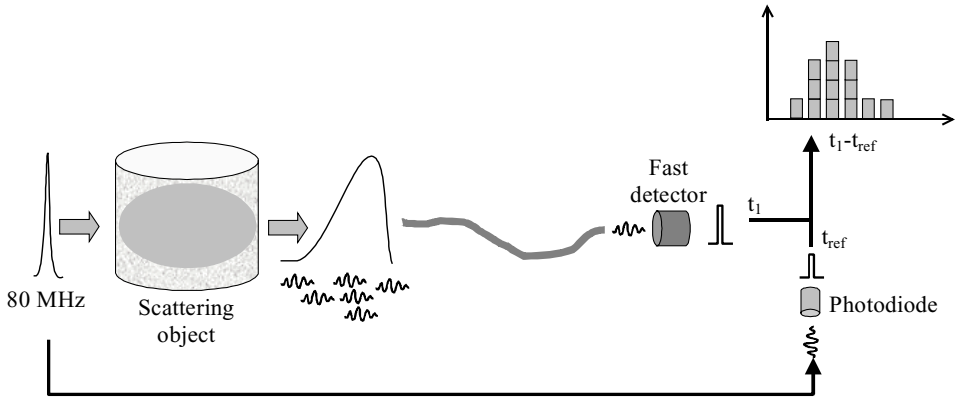

Figure 1.2.10 In TCSPC, a TPSF is built up as a histogram of the time-difference between a photon emerging from the scattering object and a reference photon.

Note that it is important that not more than one photon can be detected per incident broadened laser pulse. If two photons arrive within the same pulse, only the first will be detected. Pulse pile up error will result, where the measured TPSF becomes skewed to earlier times due to preferential detection of the earlier photon. In order for TCSPC to work well, less that 1 incident pulse in 100 should result in a detected photon. For an 80MHz laser this means detecting photons at 800KHz (see section 2.1.1.2). The TCSPC process in MONSTIR is illustrated in Figure 1.2.11.

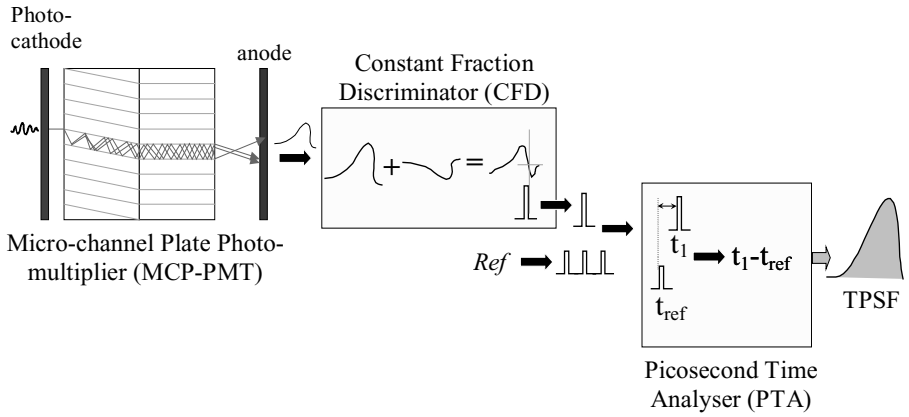

Figure 1.2.11 Schematic of the MONSTIR detection process

The fast detectors in MONSTIR are four 8-anode MCP-PMTs (Hamamatsu Photonics KK) in MONSTIR. These produce the electronic pulse in response to an incident photon and have lower temporal dispersion than a conventional PMT. The pulse is amplified and then converted into a digital NIM (nuclear instrument module) pulse, by one of 32 constant fraction discriminators (CFDs, from EG&G (Ortec)) which determine the temporal position of the arriving pulse independently of its magnitude. The NIM pulses then go to one of 32 picosecond time analysers (PTAs, from EG&G (Ortec)), which measure the time delay between a reference pulse and the NIM pulse from the MCP-PMT. The PTA builds up these times into a histogram, which is stored and then read out by a control PC. The laser repetition rate is usually 80MHz, so the temporal width of the histogram can be up to 12.5ns. Note that originally MONSTIR used a Ti:Sapphire laser. In this case a photodiode, illuminated via a beam splitter was used to generate the reference pulse. The new fibre laser has a digital sync output that can be used as a more robust reference signal. A typical TPSF measured using MONSTIR is shown in Figure 1.2.12.

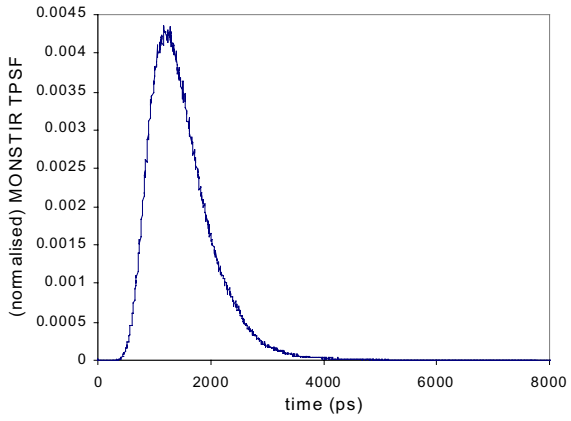

Figure 1.2.12 A typicl TPSF measured using the MONSTIR system

1.2.3.3 Variable optical attenuators

As described in section 1.2.1.3 above, a major problem for optical tomographic imaging methods is the required dynamic range of the detectors. In order to remain within the TCSPC limit of detecting photons at 80KHz, while still being sensitive to potentially very low incident photon rates MONSTIR requires active control of the light intensities incident on the detectors. MONSTIR has set of integrated variable optical attenuators (VOAs). It is critical that the attenuation setting for each channel is able to vary since any one detector will be exposed to signals over the whole dynamic range during one image acquisition as the source position varies. The optical properties of object being imaged and the chosen source-detector geometry will also require very different modifications to the detected signal intensity. The VOAs consist of thin metal blackened disks; each etched with a number of holes of varying sizes. Each disk is mounted on a stepper motor and positioned between an incoming detector bundle and a 4 mm-diameter polymer fibre leading to the MCP. (As shown in Figure 1.2.13).

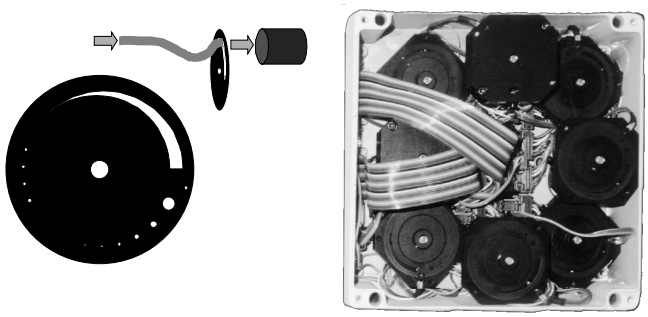

Figure 1.2.13 The variable optical attenuators. (left) the discs used to attenuate the light reaching the detectors, (Middle) a box with 8 discs mounted on 8 stepper-motors.

The VOAs can be programmed (using an acquisition definition file or ADF see section 1.2.3.5) according to the geometry of the sources and detectors to choose one of eight holes, the largest of which lets all light through and the smallest hole provides attenuation of ~ 10^5. The VOAs can also be set to zero transmission to protect the MCPs when the system is not acquiring data. The acquisition definition file for an imaging experiment can be estimated based on planned source-detector positions, or optimised empirically via an automated process that finds the VOA position corresponding to a predefined maximum count rate. This technique accounts for differing efficiencies of each source and detector channel and inter VOA variations, however it can take up to 20 minutes and must be performed with fibres in place on the subject.

Figure 1.2.14 shows the intensity profile of measurements made on a circular object as shown, (centre) calculated assuming no VOAs using the CW 3D infinite space Green's function (see section 1.3.1.2.1) and (right) measured using MONSTIR with the VOAs in their optimal positions.

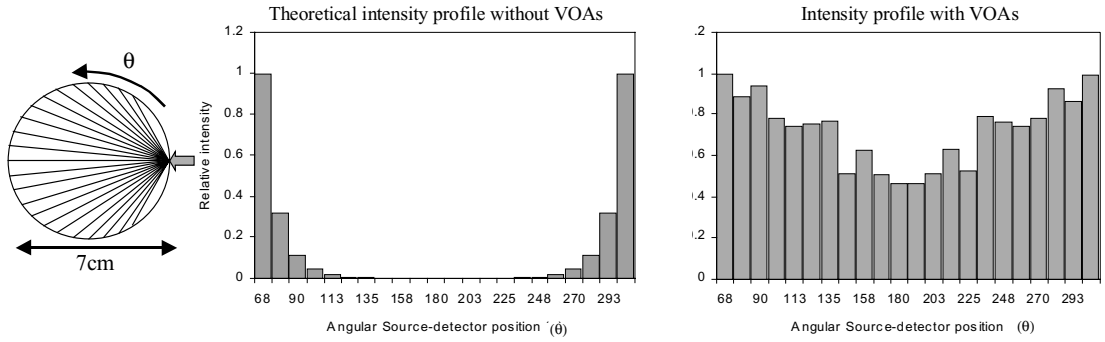

Figure 1.2.14 An attenuation range of 10^6 is seen in a tomographic arrangement around a 7 cm diameter circle with tissue-like optical properties (μ_a=0.01mm^{-1} and μ_s'=1mm^{-1}).(Left) the theoretical detected intensity profile around the circle(right) intensity measurements made using MONSTIR on a phantom with similar properties and optimised VOAs.

1.2.3.4 Laser source and dual wavelength acquisition

Since the detection system in MONSTIR does not distinguish between light of different wavelengths, a monochromatic pulsed light source is used. The original light source was a Ti:Sapphire laser which could be manually tuned to wavelengths from ~ 700nm to 850nm. Tuning was found to affect the output intensity and pulse characteristics. The Ti:Sapphire is sensitive to movement, environment, is often unstable and is certainly not portable.

Recently a new fibre laser (IMRA America, Inc.) was purchased which is more robust and very stable. While the laser itself cannot be manually tuned it was configured especially for the MONSTIR system to have two laser 'cavities', one producing pulses at 780nm and the other at 815nm. These outputs share the same 40MHz synchronising signal such that the laser can provide output in two different modes of operation (see Figure 1.2.15).

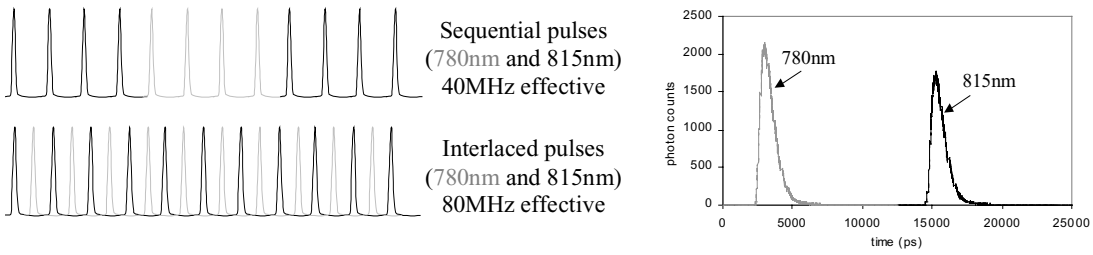

Figure 1.2.15 Two possible modes of operation of fibre laser to provide dual wavelength illumination. (top) sequential pulses, the laser switches the output to be one wavelength or the other at a predefined interval (~seconds). (bottom) exploiting the fact that the lasers are synchronised, an 80MHz output pulse train can be a created from both laser outputs by interleaving the pulses.

By configuring the MONSTIR system to measure the TPSF over a 25ns temporal window we can acquire two TPSFs for the interlaced mode and one per window in sequential mode. By acquiring two TPSFs (one at each wavelength) simultaneously in interlaced mode, we are effectively acquiring simultaneous dual wavelength data. The laser power illuminating tissue in interleaved mode is twice the power of the sequential mode output. Note that sequential mode illumination swaps between wavelengths on the order of seconds.

1.2.3.5 Computer control of hardware

A software package (MIDAS) was written by Florian Schmidt to control the hardware and data read-out of the MONSTIR system (Schmidt, 1999). This versatile program allows data acquisition protocols to be fully automated by means of an acquisition definition file (ADF). The ADF controls the VOA settings, the exposure time and order in which data are acquired. The TPSF data read out from the PTAs are written to a binary file along with header information about the measurement. All subsequent data processing is performed elsewhere.

MIDAS also controls safety aspects of the MONSTIR system. For example Peltier coolers around the MCP-PMTs (to reduce thermal noise) also require an external chiller. MIDAS ensures that the Peltiers cannot be turned on if the chiller is off. A hardware-based safety system (known as the inhibit system, see (Hillman, 2001a)) continuously checks the state of the MIDAS software by monitoring clock pulses issued by the software. If the software crashes the inhibit system takes over hardware control and returns the system to a safe state (MCP-PMT power off, shutter closed etc).

1.2.3.6 Final system

Figure 1.2.16 shows photographs of the completed system. It is around 2.2 m tall and requires an auxiliary water chiller, laser and appropriate interface device to hold the optical fibres on the surface of the subject (see section 2.7). Image reconstruction is performed elsewhere on a computer with superior processing power.

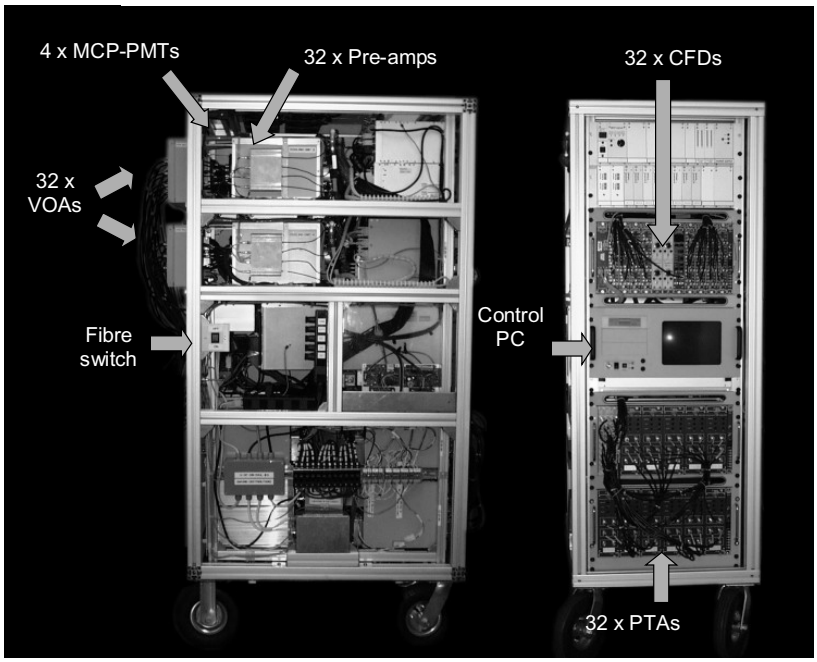

Figure 1.2.16 Photograph of the MONSTIR system

1.3 Image reconstruction

Conventional x-ray imaging does not require image reconstruction methods, since the image is a single projection through the tissue. X-ray computed topography (CT) is a more advanced imaging method that acquires data over a series of projection angles. Cross sectional images of the human body can then be generated by deducing the structures inside an object that would result in the measured projections. Although x-rays do scatter in tissue, in CT measurements x-rays are assumed to have travelled in straight lines. This makes image reconstruction fairly simple, via algorithms such as the Radon transform (Barret *et al*, 1981).

As was described in section 1.2, NIR data can be acquired in a variety of geometries to provide global measurements, surface mapping (topography) and full volume measurements (tomography). Unlike x-rays however, NIR light scatters heavily in soft tissue, which is advantageous for topographic mapping but detrimental to accurate cross-sectional image reconstruction. The path taken by NIR through a volume of tissue cannot be approximated as a straight line. Further, the paths which individual photons take through tissue will depend on the absorbing and scattering properties of the tissue itself. Since we must know the form of the photon propagation through the tissue in order to reconstruct images, we need to be able to model the behaviour of light.

This section will describe the way that models can be used to reveal the paths taken by light through an object. Once models have been formed, images can be reconstructed in a number of ways, from simple linear inversions to non-linear iterative methods. The implementation of additional aspects of a full image reconstruction algorithm will then be described with reference to the TOAST (temporal optical absorption and scatter tomography) algorithm which is used to reconstruct images from data measured using MONSTIR (section 1.2.3). Application of these reconstruction methods to real and simulated data will be explored later in section 2.

1.3.1 Modelling the behaviour of light in tissue

Light propagates through media in a deterministic and thus predictable way. The radiative transport equation (RTE) describes the transport of particles through a host medium. Models based on the RTE can be used for a diverse range applications from neutron diffusion through uranium fuel elements in nuclear reactors, to traffic flow (Duderstadt *et al*, 1979).

The radiative transfer (or Boltzmann transport) equation in the time domain is:

$$\left(\frac{1}{c}\frac{\partial}{\partial t}+\hat{s}\cdot\nabla+\mu_a(r)+\mu_s(r)\right)\phi(r,\hat{s},t)=\mu_s(r)\int_{s^{n-1}}\Theta(\hat{s}\cdot\hat{s}')\phi(r,\hat{s}',t)d\hat{s}'+q(r,\hat{s},t)$$

[1.3.1]

where c = speed of light in the medium, ϕ = number of photons per unit volume at position r at time t with velocity in direction $\hat{s}$ and $\Theta(\hat{s}\cdot\hat{s}')$ = normalised phase function (see Figure 1.1.3) and $cq(r,\hat{s},t)$ = the number of photons per unit volume per unit time sourced at position r at time t with velocity c in direction $\hat{s}$.μ_a and μ_s are defined in section 1.1.1 .

$$\text{Photon density } \Phi(r,t)=\int_{s^{n-1}}\phi(r,\hat{s},t)d\hat{s}$$

[1.3.2]

$$\text{Photon current } J(r,t)=\int_{s^{n-1}}\hat{s}\phi(r,\hat{s},t)d\hat{s}\,,$$

[1.3.3]

where s^n is the nth outer product of unit vector $\hat{s}$ with itself.

1.3.1.1 The diffusion approximation

In order to solve the radiative transport equation we need to make approximations and simplifications if we are to perform calculations within an acceptable amount of time. The P_N approximation considers the parameters ϕ and q and Θ in [1.3.1] as spherical harmonic expansions. The P_1 approximation yields:

$$\left(\frac{1}{c}\frac{\partial}{\partial t}+\frac{1}{3\kappa(r)}\right)J(r,t)+\frac{1}{3}\nabla\Phi(r,t)=q_1$$

[1.3.4]

Where the source function $$q_1(r,t)=\int_{s^{n-1}}\hat{s}q(r,\hat{s},t)d\hat{s}$$

[1.3.5]

and the diffusion coefficient $\kappa(r)$ for time-domain is:

$$\kappa(r)=\frac{1}{3(\mu_a(r)+\mu_s'(r))},$$

[1.3.6]

and for frequency domain $$\kappa(r)=\frac{1}{3(\mu_a(r)+\frac{i\omega}{c}+\mu_s'(r))}.$$

[1.3.7]

The diffusion approximation assumes that:

$$\frac{\partial J}{\partial t}=0 \text{ and } q_1=0$$

[1.3.8]

So within these assumptions, [1.3.4] becomes

$$J(r,t)=-\kappa(r)\nabla\Phi(r,t)\,.$$

[1.3.9]

Expressing [1.3.1] in terms of [1.3.3] and substituting [1.3.9] yields the time domain form of the diffusion approximation to the radiative transfer equation (or diffusion equation):

$$\mu_a(r)\Phi(r,t)+\frac{1}{c}\frac{\partial\Phi(r,t)}{\partial t}-\nabla\cdot\kappa(r)\nabla\Phi(r,t)=q_0(r,t)$$

[1.3.10]

the corresponding frequency domain form is:

$$\left(\mu_a(r)+\frac{i\omega}{c}\right)\hat{\Phi}(r,\omega)-\nabla\cdot\kappa(r)\nabla\hat{\Phi}(r,\omega)=\hat{Q}_0(r,\omega)$$

[1.3.11]

where $\hat{Q}_0(r,\omega)$ and $q_0(r,t)$ are source functions.

The diffusion approximation assumes that the source term is isotropic. Usually this means setting $q_1 = 0$, and using the isotropic function q_0 given by the zero'th order spherical harmonic. The diffusion approximation also assumes that the rate of change of photon current is negligible. This implies that the distribution of light within an object is instantaneous, which is erroneous for the time dependent case, but is usually justified by assuming $\mu_a << \mu'_s$.

1.3.1.2 Analytical Solutions to the diffusion approximation

1.3.1.2.1 Greens function solutions of the diffusion equation

Solutions to the diffusion equation can be found using the Green's operator $\mathcal{G}$ acting on the source function $Q_0(r,\omega)$ such that:

$$\Phi(r,\omega)=\hat{\mathcal{G}}[Q_0(r,\omega)]$$

[1.3.12]

and in the time domain $$\Phi(r,t)=\mathcal{G}[q_0(r,t)]$$

[1.3.13]

From the definition of a Green's operator the frequency domain Green's function solution $G^{(\Phi)}(r,r',\omega)$ is given by:

$$\hat{\mathcal{G}}[Q_0(r,\omega)]=\int_\Omega G^{(\Phi)}(r,r',\omega)Q_0(r',\omega)d^n r$$

[1.3.14]

and in the time domain:

$$\mathcal{G}[q_0(r,t)]=\int_{-\infty}^{\infty}\int_\Omega g^{(\Phi)}(r,r',t-t')q_0(r',t')d^n r dt'$$

[1.3.15]

where Ω is the region of interest. A useful property of Green's functions is that an operator acting on a Green's function $g(r,r',t-t')$ gives a delta function δ, at time t' (Boas, 1983). If $Q_0(r,\omega)$ or $q_0(r,t)$ are taken to be a delta function in space, the measured photon density $\Phi(r,\omega)$ or $\Phi(r,t)$ (equivalent to the system's impulse response) is then equal to the Green's function $G(r,\omega)$ or $g(r,t)$. Combining [1.3.13] and [1.3.15] we get:

$$\Phi(r,t)=\int_{-\infty}^{\infty}\int_\Omega g^{(\Phi)}(r,r',t-t')\,\delta\, d^n r dt'$$

[1.3.16]

Substituting [1.3.16] into [1.3.10] yields:

$$\mu_a(r)g(r,r',t,t') + \frac{1}{c}\frac{\partial g(r,r',t,t')}{\partial t} - \nabla\cdot\kappa(r)\nabla g(r,r',t,t') = \delta(r,r',t-t')$$

[1.3.17]

The Green's function solutions of the diffusion equation for 3 dimensional (3D) infinite space in the frequency domain [1.3.18], and time domain [1.3.19] are of the form:

$$G_{3D\_\inf}(\omega) = \frac{3(\mu_a + \mu_s')}{2d(2\pi)^{3/2}} e^{-d\sqrt{3(\mu_a+\mu_s')(\mu_a+\frac{i\omega n}{c})}}$$

[1.3.18]

$$g_{3D\_\inf}(t) = \frac{(3n(\mu_a + \mu_s'))^{3/2}}{(4\pi ct)^{3/2}} e^{-(\frac{\mu_a ct}{n} + \frac{3d^2 n(\mu_a+\mu_s')}{4tc})}$$

[1.3.19]

where d is the distance between the source and detector. In most cases measurements will be made on the boundary of a discrete domain. The measurable of the system is the boundary flux given by:

$$\Gamma(m,t) = \mathcal{B}\,\Phi(m)$$

[1.3.20]

where

$$\mathcal{B} = -\kappa\frac{\partial}{\partial v}$$

[1.3.21]

where m is the position on the boundary, v is the normal to the boundary and t is time. Since the source of photons in the system is a close approximation to δ function at time t_q and position r_q, then the measured flux can be approximated to the appropriate Green's function solution of the system (in the time domain) at m,

$$\Gamma(m,t) = -\kappa(m)\frac{\partial g_{\mu_s',\mu}(m, r_q, t-t_q)}{\partial v}$$

[1.3.22]

A plot of $\Gamma(m,t)$ v/s time with [1.3.19] substituted will represent the detected signal resulting from input of a spatially discrete delta function impulse (Figure 1.3.1). This is equivalent to the TPSF that would be measured by MONSTIR with source-detector separation d in an infinite medium with properties μ_a, μ_s' and n. MONSTIR measures the temporal response of tissue to a very short laser pulse.

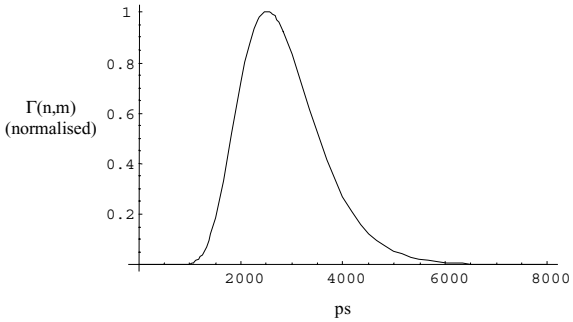

Figure 1.3.1 A plot of Γ(n,m) [1.3.22] calculated from the 3D infinite space Green's function [1.3.19] using μ_a=0.01 mm^{-1}, μ_s'=1 mm^{-1}, n=1.4, d=70 mm.

1.3.1.2.2 Green's function solutions for more complex geometries

To calculate Green's function solutions for more complex geometries than infinite space, methods similar to calculations of electric or magnetic fields can be used. For the case of a semi-infinite space (the space extending from one side of an infinite plane) with a pencil beam illuminating its surface, we assume that the photon fluence is zero at an extrapoltaed boundary at $z = 1/\mu'_s$ from the plane. To achieve this, an isotropic light source is assumed to be located at a distance z below the plane and a 'negative' source is added on the other side of the boundary, shown in Figure 1.3.2 as black and white circles respectvley. The distance z represents where we would expect photons to be initially scattered. (Patterson *et al*, 1989)

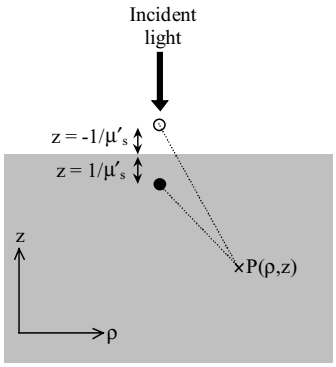

Figure 1.3.2 The geometry of real and negative sources either side of the boundary of a semi-infinite scattering medium, used to derive the Green's function solution to the diffusion equation for semi-infinite space.

So to obtain the semi-infinite Green's function solution, the contributions of the two sources are summed at the point P(ρ,z) as shown in [1.3.23]:

$$g_{3D\_semi-\inf}(t) = \frac{(3n(\mu_a + \mu'_s))^{3/2}}{(4\pi ct)^{3/2}} e^{-\frac{\mu_a ct}{n}} \left(e^{-\frac{((z-z_0)^2+\rho^2)(\mu_a+\mu_s)n}{4tc}} - e^{-\frac{((z+z_0)^2+\rho^2)(\mu_a+\mu_s)n}{4tc}} \right) \quad [\ 1.3.23\ ]$$

More complex derivations for geometries such as circles, cylinders, slabs and spheres have been performed and are detailed in (Arridge *et al*, 1992). Other analytical methods, including models of simple inclusions with different absorbing and scattering properties to the rest of the medium are described by (Boas, 1996), (Feng *et al*, 1995), (Boas *et al*, 1993a) and (Boas *et al*, 1993b). However these models are far too simplistic to model propagation of light through complex structures such as the brain.

1.3.1.3 Numerical solutions to the diffusion approximation

The diffusion equation [1.3.10] and [1.3.11] can be 'solved' using numerical methods instead of analytical ones as described above. A numerical method involves simply repeatedly applying the basic rules of propagation and interaction of photons as they travel through media. This approach is far more versatile than trying to derive an analytical solution. Numerical methods include Monte Carlo modelling, finite element method (FEM) and finite difference method (FDM).

1.3.1.3.1 Finite Element Method

The Finite element method is used to solve partial differential equations (PDEs), like the diffusion equation, in complex geometries. It is widely applied in engineering where it is used to model heat transfer, strains and stresses within structures, and fluid mechanics. FEM has been applied to a number of medical imaging reconstruction problems including Electrical impedance tomography (EIT) (Webster, 1990).

Where the region we want to examine is a domain Ω, we know that $\Phi(r,t)$ (the photon density as in [1.3.10] and [1.3.11]) will be continuous and approximately obey the diffusion equation according to the local optical properties within the domain. An example of the finite element method is to approximate $\Phi(r,t)$ by the piecewise polynomial function:

$$\Phi^h(r,t) = \sum_i^D \Phi_i(t) u_i(r)$$

[1.3.24]

This represents splitting the domain into a finite number of elements. Each element has nodes at its corners, so that the domain has D nodes. Effectively our aim is to estimate $\Phi_i(t)$ at each node i. The h superscript denotes that it is an approximate solution. $u_i(r)$ is a basis function that describes the way that the function is allowed to vary over an element. A simple form is for the basis to be piecewise linear, which means that the continuous function $\Phi(r,t)$ is estimated to vary linearly between nodes. In this case the PDE is replaced by a linear system of equations. Solving this linear system returns the approximation to $\Phi(r,t)$.

The solution can also be mapped to a local basis using [1.3.25]. A number of possible profiles for $v_p(r)$ are illustrated in Figure 1.3.3

$$u_i(r) = \sum_{p=1}^{N} u_p \mathrm{v}_p(r)$$

[1.3.25]

These bases generally have the effect of smoothing the solution u_i which is important for iterative methods described later (Schweiger *et al*, in preparation 2001).

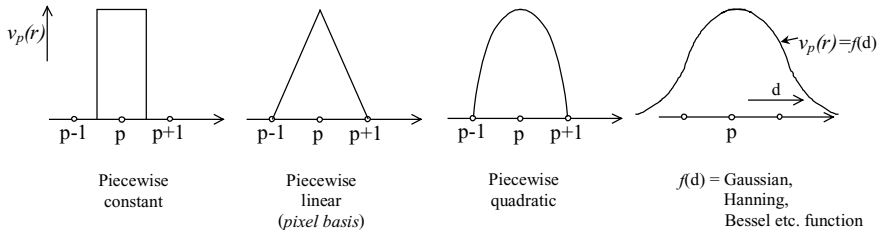

Figure 1.3.3 Different types of local basis functions implemented in TOAST

A mesh defines the locations of each node. The finite element mesh shown in Figure 1.3.4 is two-dimensional. It is also possible to create meshes and solve the FEM problem in 3D. Methods of generating meshes have been developed for a multitude of FEM applications.

Details about the mesh generation software commonly used in our application can be found in (Schöberl, 1997).

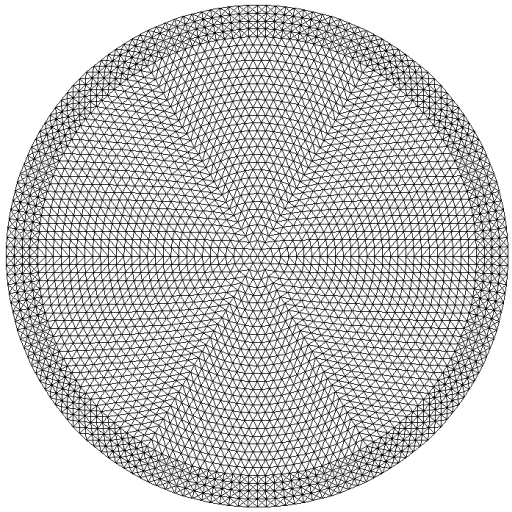

Figure 1.3.4 A typical mesh as used for finite element method modelling of light propagation through a circular domain. Each triangle is an element and each vertex is a node. The higher node density at the perimeter enables improved modelling of sources, since the field will vary most in this region.

To apply FEM to the diffusion equation [1.3.10], we multiply both sides by a test (or basis) function $u_j(r)$ and integrate over $d\Omega$.

$$\int_{\Omega} u_j(r)\left(\mu_a(r)\Phi^h(r,t)+\frac{1}{c}\frac{\partial\Phi^h(r,t)}{\partial t}-\nabla\cdot\kappa(r)\nabla\Phi^h(r,t)\right)d\Omega=\int_{\Omega}u_j(r)q_0(r,t)d\Omega$$

[1.3.26]

Integrating by parts, substituting [1.3.24] into [1.3.26] and rearranging leads to the matrix form shown in [1.3.27] (time domain) and [1.3.28] (frequency domain).

$$\left[\mathbf{K}(\kappa)+\mathbf{C}(\mu_a c)\right]\Phi+\mathbf{B}\frac{\partial\Phi}{\partial t}=Q$$

[1.3.27]

$$\left[\mathbf{K}(\kappa)+\mathbf{C}(\mu_a c)+i\omega\mathbf{B}\right]\hat{\Phi}(\omega)=\hat{Q}(\omega)$$

[1.3.28]

Where the matrix elements are given by:

$$K_{i,j}=\int_{\Omega}\kappa(r)\nabla u_j(r)\cdot\nabla u_i(r)d\Omega$$

$$C_{i,j}=\int_{\Omega}\mu_a(r)cu_j(r)u_i(r)d\Omega$$

$$B_{i,j}=\int_{\Omega}u_j(r)\nabla u_i(r)d\Omega$$

$$Q_j(t)=\int_{\Omega}u_j(r)q_0(r,t)d\Omega$$

$$\Phi=\begin{bmatrix}\Phi_1(t)\\ \Phi_2(t)\\ \vdots\\ \Phi_D(t)\end{bmatrix}$$

[1.3.29]

The *i, j* subscripts refer to each node in the domain, *i,j* =1→*D* (where *D* is the number of nodes in the mesh). An extra term is also required to accurately model the boundary elements on ∂Ω, where $\hat{v}_\perp$ is the normal to the surface:

$$A_{i,j} = \int_{\partial\Omega} u_i(r) u_j(r) d^{n-1} \hat{v}_\perp$$

[1.3.30]

The matrices **K**, **C**, **B** and **A** are $D \times D$ in size and sparse (have few non-zero entries). They are usually sparse since they are only non-zero when *i* and *j* refer to the nodes on the same element. The form of [1.3.28] can be compared to [1.3.12], such that (including the boundary element term):

$$\left[\mathbf{K}(\kappa) + \mathbf{C}(\mu_a c) + \mathcal{B}\mathbf{A} + i\omega\mathbf{B}\right]^{-1} \hat{Q}_0 = \hat{\mathbf{G}}[\hat{Q}_0(r,\omega)] = \Phi ,$$

[1.3.31]

where $\hat{\mathbf{G}}$ is the discrete representation of the Green's operator $\mathcal{G}$ described in section 1.3.1.2.1. This means that rather than deriving the analytical Green's function for the geometry and structure that we are trying to model (as demonstrated in section 1.3.1.2.2), we can create the operator $\hat{\mathbf{G}}$ from the integrals **K**, **C, Q** and **B**. Since these integrals depend on the mesh, geometry, optical properties throughout the mesh, and the basis, they provide a tailor-made operator to model a given object.

1.3.1.3.2 Finite Difference method

Finite difference methods (FDM) can be considered as a special case of FEM, where a regular grid is used instead of a more versatile mesh. Solutions methods of FDM are considerably easier than those of FEM. (Arridge, 1999).

1.3.1.4 Numerical solutions to radiative transfer and hybrid models

In some cases, the diffusion approximation is not valid. These include when the measurements are over distances comparable to the scattering length ($1/\mu'_s$) (e.g. imaging finger joints (Klose *et al*, 1998), small tissue samples (Dunn *et al*, 2000), and for tissues where the condition $\mu_a << \mu'_s$ does not hold (e.g. CSF see section 2.6). In these cases, an alternative model must be used if the propagation of light is to be properly modelled.

1.3.1.4.1 Monte Carlo Modelling

Monte Carlo involves establishing a model of a particular measurement geometry and optical property distribution. Photons are then 'launched' one by one into the system. The probability of absorption or scattering events occurring is given by the μ_a and μ_s of the local region. Photons that reach the 'detector' (given by whether the photon's path has lead it to the position of the detector, and the probability of the photon having been absorbed on the way) constitute the signal that would be measured. Monte Carlo modelling is considered the 'gold

standard' of modelling techniques and, if sufficient photons are considered, should yield the same results as an analytic solution (if one exists) or an experimental measurement. However the process of Monte Carlo modelling is very slow and computationally intensive.

1.3.1.4.2 Radiative transfer and hybrid models

Comparative studies of the diffusion approximation (DA) and the radiative transfer equation (RTE) given in (Hielsher *et al*, 1998) deduce that within regions where $\mu'_s >> \mu_a$ the two models produce very similar results with the DA being considerably easier to solve. However the DA, unlike the RTE, is unable to correctly model propagation through regions where $\mu'_s \approx \mu_a$ or $\mu'_s < \mu_a$. Image reconstruction methods that incorporate the RTE are being developed (Dorn, 2000),(Dorn, 2001). A hybrid FEM model that uses radiosity and the DA to account for non-scattering regions in diffusing media is described in (Arridge *et al*, 2000a) and evaluated in detail in chapter 2.6.

1.3.2 Image reconstruction using models

If a measurement is made on an object with a particular μ_a and μ'_s, we can determine μ_a and μ'_s by comparing predictions of a model with the measurements. If values of μ_a and μ'_s can be found that result in modelled parameters that match the measurements, it is then assumed that these values of μ_a and μ'_s correspond to those of the object. Clearly the accuracy of this technique will depend on how accurately the model represents the object, how well experimental errors have been removed from the measured data, and whether there is sufficient information in the model to describe unique values of μ_a and μ'_s.

Image reconstruction involves deducing the internal distribution of μ_a and μ'_s that correspond to a set of measurements made using MONSTIR. The first step is to use the model to create a matrix relating changes (or perturbations) in μ_a and μ'_s at points within the object to perturbations in measurements.

1.3.2.1 Perturbation theory and PMDFs

The perturbation approach to optical imaging is the approximation that a change in state in an object results in a change in the measurement that can be expressed as a series based on known functions describing the original state of the object. A measurement $Y_{n,m}(\mathbf{x},\omega)$, made at frequency ω, using source *n* and detector *m* on an object with optical properties $\mathbf{x}$ can be expressed (via a Taylor series) as:

$$Y_{n,m}(\mathbf{x}_1(r),\omega) = Y_{n,m}(\mathbf{x}_0(r),\omega) + \frac{\partial Y_{n,m}(\mathbf{x}_0(r),\omega)}{\partial \mathbf{x}}[\mathbf{x}_1(r)-\mathbf{x}_0(r)] + \frac{1}{2!}\frac{\partial^2 Y_{n,m}(\mathbf{x}_0(r),\omega)}{\partial \mathbf{x}^2}[\mathbf{x}_1(r)-\mathbf{x}_0(r)]^2 .....$$

[1.3.32]

where $$\mathbf{x}(r) = \begin{pmatrix} \mu_a(r) \\ \kappa(r) \end{pmatrix},$$

[1.3.33]

representing the distribution of optical properties over the object. $\mathbf{x}_0(r)$ denotes the initial state and $\mathbf{x}_1(r)$ a perturbed state. If we ignore terms in [1.3.32] above the first derivative, we get:

$$J_{n,m}(\mathbf{x}_0(r),\omega) = \frac{Y_{n,m}(\mathbf{x}_1(r),\omega) - Y_{n,m}(\mathbf{x}_0(r),\omega)}{\mathbf{x}_1(r) - \mathbf{x}_0(r)} = \frac{\partial Y_{n,m}(\mathbf{x}_0(r),\omega)}{\partial \mathbf{x}}$$

[1.3.34]

Values of $J_{n,m}(\mathbf{x}_0(r),\omega)$ plotted over the volume / area of the object of interest are called photon measurement density functions (PMDFs). These show the magnitude of the change in the measurement $Y_{n,m}(\mathbf{x},\omega)$ that will result from unit change in optical properties $\mathbf{x}$ at position *r* for a given source detector pair *n,m*. PMDFs for a circular mesh for a single source-detector geometry, and measurement $Y_{n,m}(\mu_a,0)$ (integrated intensity) are plotted in Figure 1.3.5.

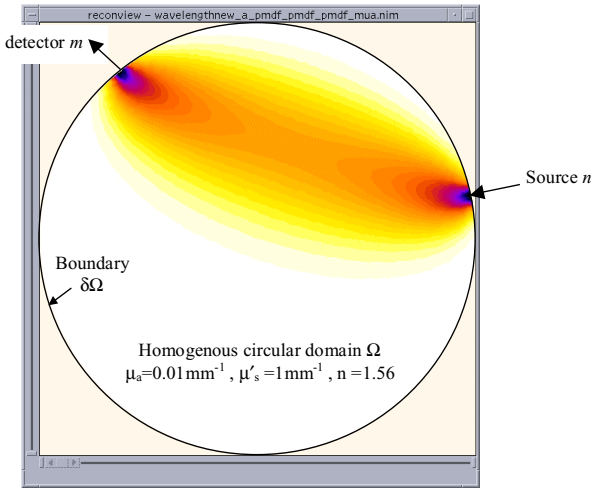

Figure 1.3.5 A plot of the photon measurement density function for source n, detector m, and intensity as the measurement type, and μ_a as the optical property varying on the circular mesh shown in Figure 1.3.4. All values are negative so hotter colours = numbers closer to zero

A matrix composed of a full set of derivatives $J_{n,m}(\mathbf{x}_0(r),\omega)$ for all source-detector combinations at all mesh nodes is **J**, the Jacobian of the forward operator (sometimes called a weight or sensitivity matrix). **J** can be evaluated analytically (using the Born approximation (Arridge, 1999)) using:

$$J_{n,m}(\mathbf{x}_0(r),\omega) = \begin{pmatrix} G_m^+(r,\omega)G_n(r,\omega) \\ \nabla G_m^+(r,\omega)\cdot\nabla G_n(r,\omega) \end{pmatrix}.$$

[1.3.35]

where $G_m^+(r,\omega)$ and $G_m(r,\omega)$ are the forward and adjoint Green's functions for source *n* and detector *m* and the appropriate geometry and optical properties. Reciprocity implies that the steady state field of photons between a given source and detector will be the same if the source and detector positions are swapped. The adjoint form of the Green's function represents propagation in the opposite direction to the normal Green's function (Arridge,

1999). This assumption allows the computational effort of calculation of the Jacobian to be significantly reduced

In order to evaluate these derivatives $J_{n,m}(\mathbf{x}_0(r),\omega)$ for complex geometries and heterogeneous distributions of optical properties, the Green's functions in [1.3.35] can be substituted for the FEM-derived discrete Green's operator $\hat{\mathbf{G}}$ described in section 1.3.1.3.2 [1.3.31] (Arridge *et al*, 1995a). Other methods of calculating **J** include Monte Carlo modelling (Graber *et al*, 1993), (Eda *et al*, 1999) (see section 1.3.1.4.1) and empirical measurement (using a tank of scattering liquid with a grid of positions r where a small absorber (and / or scatterer) can be placed and the change in measurement recorded) (Sevick *et al*, 1994). The effects of inaccurate calculation of $J_{n,m}(\mathbf{x}_0(r),\omega)$ are investigated in chapter 2.4.

1.3.2.2 Linear image reconstruction

The simplest form of image reconstruction possible with optical tomographic measurements is known as back-projection. This method combines the approach of x-ray CT image reconstruction with the weight functions or PMDFs described above in section 1.3.2.1 which account for the non-linear propagation of light compared to x-rays.

Re-writing [1.3.34], we find that the change in measurement can be expressed as the product of the Jacobian and the change in optical properties:

$$\Delta Y_{n,m}(\mathbf{x}_0(r),\omega) = J_{n,m}(\mathbf{x}_0(r),\omega)\Delta\mathbf{x}(r)$$

[1.3.36]

in the limit where Δx is small. In order to determine the optical properties $\Delta\mathbf{x}$ corresponding to the measurement $\Delta Y_{n,m}(\mathbf{x}(r),\omega)$ we must invert [1.3.36]. More generally where $\mathbf{M}_{n,m}$ is derived from $Y_{n,m}(\mathbf{x}(r),\omega)$, (see section 1.3.3) we must solve:

$$\mathbf{M}^{(1)}{}_{n,m} - \mathbf{M}^{(0)}{}_{n,m} = J_{n,m}{}^{M}(\mathbf{x}_0(r))[\mathbf{x}_1(r) - \mathbf{x}_0(r)]$$

[1.3.37]

A number of methods exist to achieve this inversion. Linear reconstructions performed in this thesis use Tikonov regularisation given by:

$$\Delta\mathsf{x} = [\mathsf{J}^T\mathsf{J} + \tilde{\lambda}\mathsf{I}]^{-1}\mathsf{J}^T\Delta\mathsf{M},$$

[1.3.38]

where $\Delta\mathbf{x}$ is the matrix of nodal optical properties sought, $\Delta\mathbf{M}$ is the matrix of changes in measurement, **J** is the Jacobian and $\tilde{\lambda} = \lambda F_{max}$ where F_{max} is the maximum main diagonal element value of the matrix $\mathbf{J}^T\mathbf{J}$ and λ is a regularisation parameter which determines the accuracy of the match between the model and the data (Dehghani *et al*, 1999a).

1.3.2.3 Non-linear image reconstruction (TOAST)

Linear reconstructions are limited in use since they can only account for small changes in optical properties. If an absolute measurement only ($Y_{n,m}(\mathbf{x},\omega)$, or $\mathbf{M}^{(1)}{}_{n,m}$) is to be reconstructed rather than a *change* in measurement (to yield an image of the *absolute* optical properties of the object being imaged), a non-linear approach is required.

Differences between linear and non-linear reconstructions are described in detail in chapter 2.4 (section 2.4.1). In general, a non-linear method uses an iterative scheme that updates the Jacobian for each estimate of the absolute optical properties $\mathbf{x}(r)$. Starting from an initial estimate of $\mathbf{x}_g(r)$, the model is used to simulate $M^{(g)}{}_{n,m}$ corresponding to an estimated measurement. The difference between this estimate and the measurements being reconstructed ($\mathbf{M}_{n,m}$) is used to determine a better estimate of the optical properties $\mathbf{x}_{g+1}(r)$ and so on.

$$\mathbf{M}^{(1)}{}_{n,m} - M^{(g)}{}_{n,m} = J_{n,m}{}^{M}(x_g(r))[x_{g+1}(r) - x_g(r)]$$
$$\mathbf{M}^{(1)}{}_{n,m} - M^{(g+1)}{}_{n,m} = J_{n,m}{}^{M}(x_{g+1}(r))[x_{g+2}(r) - x_{g+1}(r)].....$$

[1.3.39]

Iteration continues until $x_{g+N}(r)$ minimises

$$\sum_{n,m}\left(\frac{\mathbf{M}^{(1)}{}_{n,m} - M^{(g+N)}{}_{n,m}}{\sigma_{n,m}}\right)^2$$

[1.3.40]

where $\sigma_{n,m}$ is the standard deviation on measurement $\mathbf{M}_{n,m}$. The method for computing the updated estimate and the method of inversion will affect the way in which the reconstruction proceeds and converges (Hielscher *et al*, 1999), (Hielscher *et al*, 2000),(Arridge *et al*, 1998a), (Arridge *et al*, 1998b), (Schweiger *et al*, 1996). Reconstructions can be constrained by using a-priori information, for example about tissue boundaries (Schweiger *et al*, 1999a), (Dorn, 2001), (Kolehmainen *et al*, 2000). TOAST is a non-linear diffusion approximation, FEM-based algorithm.

1.3.3 Datatypes

The total information measurable in optical tomography at a single wavelength is the full temporal profile of a broadened delta-function-like laser pulse that has passed through the tissue. The MONSTIR system (described in section 1.2.3) measures this broadened pulse or TPSF at 5 ps (picosecond = 10^{-12} s) intervals over a temporal window up to 12500 ps wide. This represents a frequency range of 0 to 100 GHz in 40MHz intervals. The measured TPSF $y_{n,m}(t)$ is equivalent to the boundary flux $\Gamma_{n,m}(\mathbf{x}(r),t)$ [1.3.20].

1.3.3.1 Full time-dependent problem

Using the FEM model described in section 1.3.1.3.1 we could, in theory, calculate the whole TPSF, evaluating $\Gamma_{n,m}(t)$ for each time step, and then minimise the error between it and the measured TPSF $y_{n,m}(t)$ over some time-gate or over the whole TPSF window. In this case, we would be minimising:

$$\frac{1}{2}\sum_{n=1}^{S}\sum_{m=1}^{M_n}\sum_{t=0}^{T_{n,m}}\left(\frac{\Gamma_{n,m}(t)-y_{n,m}(t)}{\sigma_{n,m}(t)}\right)^2$$

[1.3.41]

However, solving for the entire TPSF (or even a number of time-steps) is computationally intensive. It is also likely that there is some redundancy in calculating the full TPSF (Arridge *et al*, 1998b). This method is not implemented TOAST, although other researchers have produced good results utilising the full TPSF (Gao *et al*, 2002) who also explore fitting to the normalised TPSF.

Instead, TOAST *directly* simulates integral transforms of the full temporal profile, known as datatypes. These datatypes need to be correspondingly extracted from measured TPSFs such that the difference between measured and simulated datatypes can be minimised using [1.3.40]. Different datatypes characterise particular features of the TPSF and thus hold differing information e.g. Intensity = its magnitude, mean-time = its position, variance = its width, skew = its asymmetry etc. Describing a TPSF in terms of a number of datatypes instead of using the whole profile reduces redundancy. Direct calculation of datatypes reduces the number of forward model calculations required. Choice of datatypes requires the following considerations:

- Can they be calculated simply within the TOAST forward model?
- Can they be extracted reliably from the measured data (and calibrated accordingly)?
- What are the dependencies of each datatype on μ_a and μ'_s and what are their sensitivities to noise?
- Which datatype or combination of datatypes provides a unique description of the object?

The datatypes found to be simply calculable using the TOAST forward model are described below. Considerations for choosing datatypes are then detailed in terms of uniqueness, information content and noise tolerance. Techniques for extracting datatypes from measured TPSFs, and calibration considerations are investigated in chapters 2.1 and 2.2.

1.3.3.2 Mellin-Laplace transforms

The general form of the Mellin–Laplace Transform is:

$$\mathcal{M}_{p,u}[\Gamma_{n,m}(t)] = \frac{1}{N}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)t^u e^{-pt}dt$$

[1.3.42]

where $\Gamma_{n,m}(t) = \Gamma_{n,m}(\mathbf{x}(r),t)$ for the given optical property distribution $\mathbf{x}(r)$ and where u is the order of the transform. Most of the datatypes currently implemented in TOAST are simply special cases of this form, as summarised in Table 1.3.1.

Datatype $\mathcal{M}_{p,u}$	u	p	N
Integrated Intensity	0	0	1
Complex form	0	iω	1
Meantime	1	0	E
2nd moment	2	0	E
3rd moment	3	0	E
4th moment	4	0	E
Laplace	0	s	E
Mellin Laplace$_1$	1	s	E
Mellin Laplace$_2$	2	s	E
Mellin Laplace$_3$	3	s	E
Mellin Laplace$_4$	4	s	E

Temporal moments (rows: Meantime, 2nd moment, 3rd moment, 4th moment)

s = Laplace Coefficient (rows: Laplace, Mellin Laplace$_1$, Mellin Laplace$_2$, Mellin Laplace$_3$, Mellin Laplace$_4$)

Table 1.3.1 Values of p and u in [1.3.42] corresponding to different datatypes

where:

$$E = \int_{t=0}^{\infty}\Gamma_{n,m}(t)dt$$

[1.3.43]

The same datatypes can also be derived from:

$$\hat{\mathcal{M}}_{p,u}[Y_{n,m}(\omega)] = \frac{i\dfrac{\partial^u Y_{n,m}(\omega)}{\partial\omega^u}\Big|_{\omega=-ip}}{Y(0)}$$

[1.3.44]

where $Y_{n,m}(\omega)$ is the frequency domain form of $\Gamma_{n,m}(t)$. TOAST calculates the moments directly by applying this differential operator to [1.3.31]. TOAST can calculate datatypes of any order, by starting from intensity ($\hat{\mathcal{M}}_{0,0}[Y_{n,m}(\omega)]$ (N=0)) and incrementing *u* (Schweiger *et al*, 1995), (Schweiger *et al*, 1997a), (Arridge *et al*, 1995b).

Note that it is possible to extract information equivalent to that measured by a frequency domain system by taking the Fourier transform of a TPSF and calculating the amplitude and phase at a given frequency. A Fourier transform can be described by the generalised integral transform in [1.3.42] so 'complex' datatype in Table 1.3.1 refers to this approach.

1.3.3.3 Central Moments

TOAST can also derive *central* moments (variance, skew and kurtosis) using:

$$Central\mathcal{M}_{p,u}[\Gamma_{n,m}(t)] = \frac{1}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)(t-<t>)^{u}e^{-pt}dt$$

[1.3.45]

where $<t>$ is mean-time ($\mathcal{M}_{0,1}[\Gamma_{n,m}(t)]$) . As shown in appendix A (2.9.2), central moments can be expressed in terms of 1st, 2nd, 3rd etc temporal moments, and thus expressed in terms of $Y_{n,m}(\omega)$ using [1.3.44]. Parameters u and p for central moments are summarised in Table 1.3.2:

Central Datatype	u	p	E
2nd Central moment (Variance)	2	0	N
3rd Central moment (Skew)	3	0	N
4th Central moment (Kurtosis)	4	0	N

Table 1.3.2 Central datatypes extracted from measured TPSFs

1.3.3.4 Uniqueness

Uniqueness of a set of measurements refers to whether there is a single unique distribution of optical properties that corresponds to the data. While erroneous data with low dependence on changes in optical properties may result in (incorrect) solutions which are equally likely, uniqueness is an intrinsic problem that occurs even with perfect data.

TOAST attempts to simultaneously solve for μ_a and μ'_s, so two images are produced. If integrated intensity at a single wavelength only is used as the input data, TOAST will be unable to determine whether structure is absorbing or scattering since a distribution of each could result in the same boundary data (Arridge *et al*, 1998c). While this result has not been proven for other datatypes, it is likely that a combination of a number of datatypes (with differing dependencies on μ_a and μ'_s) will be required to achieve separation of μ_a and μ'_s features.

1.3.3.5 Implementation of multi-datatype reconstructions

In linear notation, if we want to reconstruct μ_a using a single datatype we will use a Jacobian $J_{\mu a}^{datatype}$ that relates changes in μ_a to changes in that datatype. Where $\mathbf{M}_{datatype}$ is the measurement:

$$[\Delta\mathbf{M}_{Mean}] = [J_{\mu_a}^{Mean}][\Delta\mu_a]$$

[1.3.46]

If we wish to reconstruct for both μ_a and μ'_s simultaneously from single datatype (e.g. mean-time) we use:

$$\left[\Delta \mathbf{M}_{Mean}\right] = \left[J_{\mu_a}^{Mean} \quad J_{\mu_s'}^{Mean}\right] \begin{bmatrix} \Delta\mu_a \\ \Delta\mu_s' \end{bmatrix}$$

[1.3.47]

If we with to use more than one datatypes in a dual parameter reconstruction, we solve:

$$\left[\Delta \mathbf{M}_{Intensity} \quad \Delta \mathbf{M}_{Mean}\right] = \begin{bmatrix} J_{\mu_a}^{Intensity} & J_{\mu_s'}^{Intensity} \\ J_{\mu_a}^{Mean} & J_{\mu_s'}^{Mean} \end{bmatrix} \begin{bmatrix} \Delta\mu_a \\ \Delta\mu_s' \end{bmatrix}$$

[1.3.48]

This can be extended further to reconstructing (for example) data acquired at different times or wavelengths simultaneously as detailed in chapter 2.1. Non-linear algorithms will also use these types of matrices to simultaneously solve for multiple parameters. If too many input data sets are used (e.g. the time-points on a TPSF) image reconstruction can become very computationally intensive (Gao *et al*, 2002).

Note that in fact, when more than one data set or multiple parameters are being derived, the terms must be normalised so that they have the form:

$$\begin{bmatrix} \dfrac{\Delta \mathbf{M}}{mean\left|\Delta \mathbf{M}\right|} \end{bmatrix} = \begin{bmatrix} \dfrac{\mathbf{x}_o J_{x,\lambda}^{M}(\mathbf{x}_o)}{mean\left|\Delta \mathbf{M}\right|} \end{bmatrix} \begin{bmatrix} \dfrac{\Delta \mathbf{x}}{\mathbf{x}_o} \end{bmatrix}$$

[1.3.49]

If this is not done, in the case of mean-time and intensity, mean-time will dominate the inversion and the solution will be virtually independent of the intensity information.

1.3.3.6 Datatype information content v/s noise tolerance

Choice of datatype is incredibly important in optical tomography. When reconstructing using models, datatypes from experimental measurements must be well calibrated since the model will try to account for all features of the data provided, including erroneous measurements. Some datatypes are more tolerant of noise on the measured TPSF than others (Hebden *et al*, 1998), while the sensitivity of datatypes to changes in μ_a and μ'_s also varies widely. Just as if noisy data are reconstructed, the wrong solution will result, if a datatype is well calibrated, but also insensitive to μ_a and μ'_s, the image reconstruction may not converge.

The datatypes explored early in the development of MONSTIR were intensity, mean-time, variance and Laplace as they were shown to be effective when reconstructing simulated data (Schweiger *et al*, 1999b), (Schweiger *et al*, 1997b), (Arridge *et al*, 1993). An analysis of the effects of experimental noise and systematic errors on extraction of datatypes from MONSTIR TPSFs is presented in chapter 2.2 which reveals that variance and higher central moments can be badly affected by systematic noise.

Determining the information content of datatypes is not straight-forward. Plots of the boundary data profiles resulting from absorbing, scattering and both absorbing *and* scattering

perturbations on each of the datatypes detailed above are shown in appendix A (2.9.3). These imply that Mellin-Laplace transforms with p = s > 0 of differing order u all show a very similar percentage change in response to the perturbations. This suggests that when used together they will contribute little more information to the reconstruction than just one datatype. Complex amplitude and intensity are found to be very similar, as are complex phase and mean-time. Amplitude and phase at 50MHz were found to be very similar to those at 100MHz. All datatypes showed opposing signal change with scatter and absorbing perturbations except for amplitude and intensity, suggesting that a combination of normalised (all those with N=E in Table 1.3.1 and Table 1.3.2, and phase which self-normalises) and unnormalised (amplitude and intensity) datatypes may provide a good description of internal structure.

Normalised datatypes were explored since the absolute magnitude of a TPSF (intensity) is very sensitive to coupling changes between sources and detectors and the object being imaged (Delpy *et al*, 1997). Dynamic range reduction VOAs (see section 1.2.3.3) also introduce errors into derivation of unnormalised datatypes. A measure of the equivalence of datatypes to each-other can be demonstrated by considering how the temporal moments (p=0), the central moments, or the Mellin-Laplace (s = p >0) moments represent the coefficients in the Taylor series expansion of $Y_{n,m}(\omega)$ ($Ln(Y_{n,m}(\omega))$ for central moments) around $\omega = -ip$. This is demonstrated in appendix A (2.9.4). The resulting series will approximate the Fourier transform of a TPSF. Intensity represents the scaling factor of each term. Mean-time was found to always contribute to the imaginary part of $Y(\omega)$, suggesting it should be used alongside a real datatype to ensure both parts of the Fourier transform are represented in the image reconstruction. The number of terms in each series determines the range over which it converges. The central moment series (using only intensity, mean-time, variance and skew) was found to converge well up to higher than 14 GHz. While series with as many as 6 terms in the ordinary moments and Mellin-Laplace (s>0) cases diverged after ~ 2GHz and 6GHz respectively.

The advantage of *time-resolved* optical tomography over frequency domain (FD) or continuous-wave (CW) imaging is the additional information available from a TPSF (along with improved tolerance to shot noise). If time-resolved optical tomography is to prove the superior method of imaging diffuse media, the additional information available in the TPSF and the accuracy of the measurement must outweigh its cost and the long acquisition time. Most images from experimental data presented in this thesis are generally processed using mean-time, and sometimes variance or intensity, within the bounds of calibration limitations. Possible alternatives to datatypes are discussed in section 2.2.7.1.

2. Optical tomography investigations

Throughout the development of MONSTIR, experimental studies have been performed to evaluate the system and explore methods of data acquisition, calibration and image reconstruction. Chronologically this began with imaging very simple epoxy resin cylindrical phantoms containing small regions with higher absorption or scatter properties (Hebden *et al*, 1999), (Hillman *et al*, 2000a), (Schmidt *et al*, 2000a). These phantoms were gradually made more complex, with wavelength dependent dyes and smaller features distributed throughout the phantom (Hillman *et al*, 2000b), (Arridge *et al*, 2000b). The first clinical evaluation of MONSTIR was performed using a single ring of sources and detectors around the arm of a volunteer. This experiment was then extended to determine whether functional changes could be seen when a simple exercise was performed during data acquisition (Hillman *et al*, 2001b). Following the design and manufacture of a specialist fibre holder, phantoms representing the geometry of the breast were imaged (Hebden *et al*, 2001). Following shortly after, the breasts of a number of adult volunteers were imaged (Hebden *et al*, 2002a). Finally, the MONSTIR system was evaluated by acquiring data on a 31 week gestation neonate.

During the gradual progression of the project, a number of specific aspects of optical tomographic imaging that required investigation became apparent. The second half of this thesis addresses a number of these important issues. The imaging trials described above play a role within many of the studies, and are used to illustrate the concepts being explored.

Calibration was found to be a major obstacle in successful image reconstruction of data acquired using all of MONSTIR's 32 independent detector channels. Chapter 2.1 describes the sources of noise and error on data measured using MONSTIR. The types of calibration measurements developed to evaluate and correct for systematic errors are described. Data processing techniques for application of calibration measurements are then detailed, along with methods developed for reduction of the effects of stochastic noise. This work is described in (Hillman *et al*, 2000a). As a result of this investigation it was possible to define a

number of distinct imaging paradigms for acquisition and processing of data. The characteristics and required protocol for each is summarised.

Besides accounting for the relative systematic differences between MONSTIR's channels, it had never been demonstrated that the data acquired using MONSTIR corresponded well to data simulated using TOAST. Mismatches could be a result of instrument non-linearities, or errors in the implementation of the TOAST FEM model. Chapter 2.2 presents an experiment designed to quantifiably evaluate the accuracy of the absolute calibration methods developed in chapter 2.1, while validating the TOAST forward model. The match between data acquired using MONSTIR and values simulated using TOAST is explored for both raw TPSFs and for all implemented datatypes. While the match between some measured and simulated parameters was found to be good, a number of datatypes (such as skew) were found to be very sensitive to systematic errors and noise such that they were unlikely to prove useful in image reconstruction.

Preliminary phantom imaging trials are described in chapter 2.3. Specifically some of the assumptions commonly made by optical tomography researchers are addressed such as the use of 2D reconstruction algorithms to reconstruct measured data. The limitations of correcting data such that a 2D algorithm can be used are demonstrated and discussed. The accuracy of absolute and simple difference imaging paradigms for different combinations of datatypes are explored, and the limitations on quantitation, resolution and absorption / scatter separation are illustrated. Much of this work is detailed in (Hillman, 2001a). The use of a suitable homogenous phantom to acquire calibration data in the form of a reference data set was found to yield images far superior to those achievable via absolute calibration of data.

The apparent usefulness of a reference data set prompted the desire to use the object being imaged as its own reference. This is possible if data are acquired before and after some change is induced in an object, or if data are acquired at two or more wavelengths (two possible paradigms identified in chapter 2.1). However if the reference object is not homogenous, it was found that resulting images could be erroneous. Chapter 2.4 addresses the, often overlooked, inherent problems associated with image reconstruction of *changes* in data where the reference state is heterogeneous. Simulations and wavelength-difference experimental data are used to demonstrate that lack of knowledge about the static optical properties of an object will detrimentally affect images generated from such difference data. The origin of the error is derived and a technique for reducing its effect is proposed and demonstrated. This work is described in (Hillman *et al*, 2001c). The significance of the identified error to optical tomography and NIR spectroscopy in general is discussed.

Ultimately the goal of optical tomography is to generate maps of the functional characteristics of tissue. Since preceding chapters identify the limitations on absorption and scatter image accuracy, it is necessary to determine whether this accuracy is sufficient to generate clinically useful images of functional parameters. Preliminary arm imaging results are presented in appendix B ((Hillman *et al*, 2000b), (Hillman *et al*, 2001b)). Chapter 2.5 derives a number of methods by which multi-wavelength optical tomography data can be combined to derive functional parameters such as absolute oxygen saturation and fractional blood volume. The accuracy of deriving these parameters from reconstructed maps of *absolute* absorption (at two wavelengths) are investigated using a simulation, first for noise-free data and then for data with noise added. The efficacy of both linear and non-linear image reconstruction methods are explored and the effect of quantitative accuracy in absorption images correlated with the resulting functional parameters derived. More complex methods of deriving functional information from 3-wavelength data are then explored in a similar way. Results are found to be promising. While quantitative accuracy in the derived parameters is not perfect, areas of low and high saturation and fractional blood volume are well localised. Data acquired on the breast, and the exercising forearm are then used to further evaluate the techniques developed. While the imaging results are preliminary (since many approximations were made in their derivation) the parameters retrieved are physiologically reasonable, and a repeatable functional change is identified in the exercising forearm. Simulation results suggested that differing convergence of non-linear reconstructions of single-wavelength absorption images could cause errors in derived functional parameters. Methods of constraining absorption solutions by forcing the scatter solution to be the same for both wavelengths are explored and found to improve the stability of reconstructions. The technique is also shown to improve absorption / scatter separation for reconstructions that use only a single datatype. A further technique explored is the reconstruction of images of saturation, volume and scatter directly, rather than reconstructing absorption separately first. Possible linear implementation of this method is described and demonstrated. Preliminary parts of this work are described in (Hillman *et al*, 2002).

A major potential obstacle to successful brain imaging using optical tomography is the presence of non-scattering cerebrospinal fluid (CSF) around and within the brain, since the diffusion approximation is not valid for non-scattering regions (voids). Models have been developed to account for CSF, although none have been accurately validated experimentally. Chapter 2.6 describes the void problem in detail and the theoretical basis of void-TOAST. An experiment is then detailed which utilises a resin, intralipid and glass phantom to evaluate the accuracy of the 3D void-TOAST algorithm. Early analysis of these data is presented in (Riley *et al*, 2001). Detailed analysis of the experimental results revealed an important relation

between the effect of void regions and the refractive index mismatch between the void area and the surrounding tissue. Measurements of the change in data when an absorbing object is placed beyond a non-scattering layer are compared with the signal measured without the clear layer. Images are reconstructed from these data, which demonstrate that the refractive index mismatch between the non-scattering and scattering regions will contribute to the accuracy of reconstructed images. Implications for imaging the neonatal brain are discussed, although it is emphasised that only limited interpretation of the experimental results is appropriate due to the simplistic geometry utilised. The accuracy and applicability of the 3D void-TOAST model are also discussed.

Finally, clinical implementation of optical tomography presents major practical challenges relating to data acquisition protocols and patient positioning. Chapter 2.7 describes the ongoing development of the practical aspects of data acquisition, in particular methods of holding fibres in place and calibrating data in the clinical environment. Methods developed for data acquisition on a neonate are described. Manufacture of a fibre holding helmet for neonatal brain imaging is illustrated within a description of the imaging experiments performed in the UCL neonatal intensive care unit (NICU) on two healthy premature infants. The results of these studies are presented with reference to achievable signal-to-noise. Very preliminary imaging results are also shown. The fibre holder design and the performed experimental protocol are then evaluated, and future modifications and extensions proposed. Fibre holder designs for breast imaging are then described, and their development detailed. Potential improvements to breast fibre-holders and the experimental protocol used for subjects imaged so far are discussed. Aspects of fibre-holder design and clinical imaging experiments are described in (Hebden *et al*, 2002a), (Hebden *et al*, 2002b) and (Hebden *et al*, 2002c).

A summary and discussion of the results of each investigation is given at the end of each chapter. An overall summary of ideas for further work is presented in Chapter 2.8.

2.1 Calibration of time-resolved data and datatype extraction

The time-resolved data measured using MONSTIR suffers from a number of different systematic and stochastic errors. In order to reconstruct images from measured data TOAST must be able to simulate data that will match the measured data (see section 1.3.2). Therefore it is important to ensure that errors in the data read-in to TOAST are small enough that the image reconstruction can converge to the correct solution (Hillman *et al*, 2000a).

This chapter extends the system evaluation detailed in (Schmidt, 1999) to describe the specific sources of noise and error that affect data measured using MONSTIR (and in many cases time-resolved systems in general). Calibration measurements that can be made to attempt to account for measurable systematic errors are then detailed. Data processing methods developed for application of calibration measurements, for reducing the effects of remaining noise, and for deriving datatypes suitable for TOAST are then presented. Finally a summary of the different imaging paradigms developed, along with their implementation and limitations is given.

2.1.1 MONSTIR system performance and noise characteristics

The raw TPSFs measured by MONSTIR are affected by:

- Errors due to convolution of the TPSF with the system's (source-detector dependent) impulse response function (IRF)
- Shot noise (α Intensity$^{1/2}$)
- Source-detector and intensity dependent cross-talk
- Stochastic noise (e.g. detector thermal noise)
- Temporal drift and jitter in detection electronics

These errors in the raw TPSF will then undergo a transform from raw data to an extracted datatype (such as mean-time or total intensity) before being read-in to TOAST (see section 1.3.3) where the datatypes then undergo a transform into one or more images. Therefore the relation between simple noise on TPSFs to image noise is very complex.

The approach adopted here was to try to develop a calibration system that produced extracted datatypes that agreed as closely as possible with simulated data.

Careful evaluation of the performance of the MONSTIR is well documented in (Schmidt, 1999). During the development of MONSTIR substantial noise reduction and improvements

in temporal stability have been achieved. However certain aspects of the system will always limit performance and their origin must be considered if corrections are to be applied to data.

Figure 2.1.1 shows a typical (normalised) TPSF measured using MONSTIR shown on a linear and a log scale. A simulated TPSF corresponding to the same geometry and optical properties is shown for comparison. The origins of the five features identified in Figure 2.1.1 are described below.

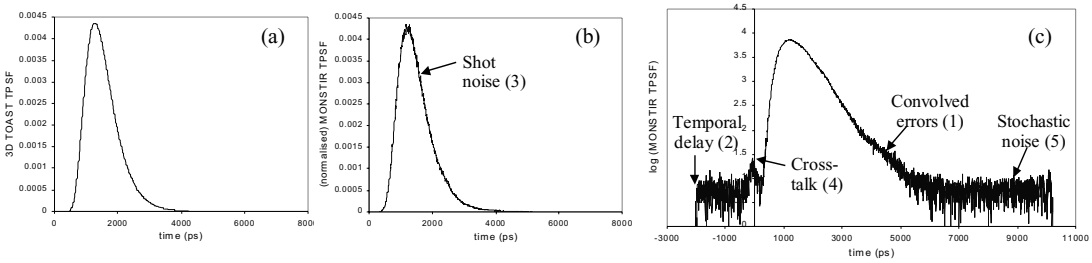

Figure 2.1.1 The types of error affecting measured TPSFs. (a) shows the TOAST simulated TPSF for the geometry, (b) shows the measured TPSF. (c) is (b) on a log scale. Arrows indicate features discussed in sections (1) and (2):2.1.1.1, (3):2.1.1.2, (4): 2.1.1.3 and (5): 2.1.1.5

2.1.1.1 Convolved errors and the system's impulse response function

The impulse response function (IRF) of a system is defined as its output when a delta function is the input. For MONSTIR the IRF is the slightly temporally broadened and attenuated pulse that would be recorded if a perfect delta function of laser light were the input. The impulse response of a measurement system will convolve with any signal measured by the system. Thus, a TPSF measured using MONSTIR will be broader and have decreased magnitude compared to the true TPSF of the object being imaged. Since the image reconstruction assumes a delta function input to the tissue and that the emerging broadened pulse is detected perfectly without any subsequent broadening, the IRF must be eliminated from measured TPSFs.

Measuring the IRF is difficult. It represents the measurement of a TPSF when a particular source fibre and a detector bundle are in direct contact with each other. A tool designed for measurement of the system IRF is shown in Figure 2.1.2. It holds a source and detector a fixed distance apart, so that the source light diverges onto a thin sheet of white paper. This will add a slight broadening to the measurement, but allows the detector bundle to be more uniformly illuminated and accounts for the fact that, in practice, a fibre bundle will be illuminated by diffuse light. The 7 mm separation adds a 23 ps delay. Additional broadening in the measured IRF will result from the fact that the laser pulse itself is not a true delta function, but has finite width. An additional signal which may appear in a measured IRF is a contribution from (non-convolved) cross-talk and additive noise within the system (see sections 2.1.1.3, 2.1.1.4 and 2.1.1.5). This demonstrates that it is never going to be possible to measure the exact IRF of the system.

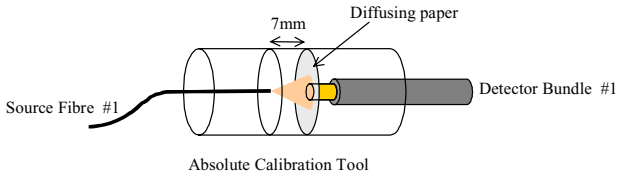

Figure 2.1.2 The 'Absolute calibration tool' - designed to allow measurement of the MONSTIR impulse response function.

A typical IRF measurement is shown in Figure 2.1.3 on both linear on log scales. Note that the delay on the time-scale is arbitrary.

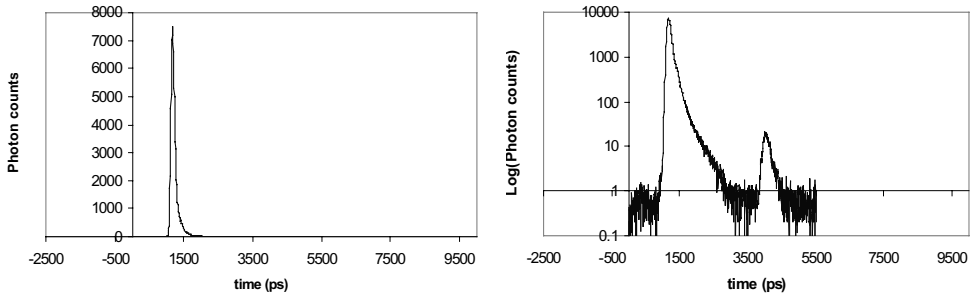

Figure 2.1.3 (Left) a typical IRF measurement on a linear scale and (right) on a log scale.

The main features of the IRF are its finite width, and the small secondary peak ~ 3000 ps after the main peak. This corresponds to a round-trip distance of ~ 60 cm in a medium with a refractive index of 1.5, and represents an optical reflection in the polymer fibres linking the MCP-PMTs and the VOAs, which are ~30 cm long (see section 1.2.3.3).

Note however that the IRF shown is for source 1 and detector 1. It would be different if one of the other detectors had been used since the relative delay, broadening, attenuation and secondary reflection will be unique to a particular detection channel. This is because each channel will consist of a series of fibres and cables each with differing lengths and dispersion properties. Coupling between the object and the fibre bundle and then the bundle and the detector will affect attenuation. Each MCP-PMT output pulse will be depend on the channel's photocathode and gain and will require different threshold and walk settings on the CFDs to determine which photons are counted. This can affect broadening, delay and the photon detection efficiency. The secondary reflection magnitude and position is unique to each detector channel since it depends on the coupling, polishing and the exact length of each polymer fibre.

The IRF will also depend on which source fibre is used. The length and attenuation and dispersion properties of the delivery fibres will all differ slightly. Coupling into and out of the fibre switch will also introduce an extra unique attenuation of the delivered pulse.

Therefore, when measurements are made on an object where many or all source-detector combinations are used, calibration for the impulse response must consider the appropriate correction for the specific source-detector pair used for each measurement.

The TPSF in Figure 2.1.1c represents the true TPSF of the object convolved with an impulse response similar to that shown in Figure 2.1.3. The IRF has caused the TPSF to have a shifted time-delay (see arrow (2)). The secondary reflection, clearly visible in the IRF manifests, as a 'hump' in the tail of the measured TPSF (depending on the TPSF width) as shown by arrow (1). While broadening of the TPSF due to the IRF is difficult to distinguish from the simulated TPSF both the secondary reflection and the broadening will radically affect certain extracted datatypes as will be shown in chapter 2.2.

2.1.1.2 Shot noise limitations

Since shot noise is proportional to the square root of the number of photons in a TPSF, the more photons acquired, the better the signal to noise ratio (SNR) becomes. Arrow (3) in Figure 2.1.1b indicates the shot noise spikes visible on the top of the TPSF relative to the simulated TPSF (which represents the result where shot noise = 0). Note that shot noise, while being normally distributed will not be temporally un-correlated in a TPSF measurement, since the amount of shot noise is proportional to the number of photons in each time-bin, and will therefore be greatest around the peak of the TPSF. While SNR will improve with increased numbers of photons, there are limits on rate at which photons can be detected, and also on the amount of time that each acquisition can reasonably take.

As described in section 1.2.3.2, TCSPC involves measuring a TPSF by detecting a maximum of one photon per incident laser pulse. In fact it is important to ensure that the probability of two photons arriving at the detector from a single laser pulse is negligible, since this would skew the TPSF in favour of earlier times as only the first photon would be detected. To ensure that this coincidence occurs in less than 1 in 10^4 pulses, it is necessary to ensure that we detect no more than one photon per 100 laser pulses. The lasers (both the Tsunami, and the interlaced pulse train from the fibre lasers) currently pulse at 80 MHz, which suggests that we could detect photons at 800kHz for TCSPC. However the detection electronics (in particular dead-time in the PTAs) limits the maximum detected count rate possible to 250 kHz.

The efficiency of the MONSTIR system determines how many of the photons arriving at the detector contribute counts to the TPSF. The total ‘exposure time' required for a given object depends on:

- The efficiency of the MONSTIR system

- The size and optical properties of the object
- The allowable incident light intensity (for skin and eye safety)
- The number of photons required in the collected TPSFs

Additional acquisition time is required between each successive source illumination. During this time the PTAs must be read out (~1.5 seconds), the VOAs must be moved to their new positions and the fibre switch must activate the next source (~1.5 seconds).

Shot noise on a TPSF will affect extracted datatypes differently. Results published in (Hebden *et al*, 1998) indicate that the coefficient of variation of most datatypes is less than 0.2% if a typical TPSF contains 10^6 photons.

Assuming a detection efficiency of 0.04 (as calculated in (Schmidt, 1999)), the exposure time required to acquire 10^6 photons through a 9 cm object with optical properties μ_a = 0.01 mm^{-1} and μ'_s= 1 mm^{-1}, at MONSTIR's hardware limited count rate of 250 kHz, will be 4 seconds. For the fibre-laser (where 2 TPSFs are acquired simultaneously) the exposure time will be twice this if both TPSFs are to contain 10^6 photons. This assumes that 10 mW or more is incident on the tissue (calculated using the 3D infinite space Green's function - see section 1.3.1.2.1). However in practice we find significantly higher incident powers are required to achieve a photon detection counts of ~ 100kHz, requiring 10 to 15 seconds exposure time per source. So total acquisition time (for one wavelength, 32 sources, on a typical object) is hardware limited to 32 × (4 + 3) = 224 seconds = 3.7 minutes, although around 10 minutes is more usual. Additional time is required to perform calibration measurements prior to image data acquisition. Setting up the optode placement on the subject, loading software, and system and laser warm-up time will also add to the whole time taken to acquire an image data set. The reasonable time for acquisition of a full image data set will be limited by the amount of time the subject can remain still. It must also be considered that more than one image data set may be required.

If more photons are required, or a time of ~ 10 minutes per scan is unsatisfactory for a clinical study, improvement to required exposure time may be achieved by regulating the detection efficiencies of each detector (e.g. using appropriate VOA settings - see section 1.2.3.3), and regulating the source fibre delivered intensities. The exact loss in each source fibre is difficult to quantify since they are easily damaged. Variations in delivered light intensities may be up to 300% (see Figure 2.1.9). The highest value must be less than the limit for subject exposure and the lowest will yield the worst SNR. If all illumination intensities are the same, the maximum SNR can be achieved for minimum illumination. There is no way to decrease dead time in the PTAs, although more modern photon counting instrumentation now offers improved count rates.

2.1.1.3 Cross-talk in the fibre-switch

The measured TPSF shown in Figure 2.1.1 has a small pre-peak indicated by arrow (4). The origin of this 'pre-peak' has been identified as cross-talk in the fibre-switch used to switch laser light from one delivery fibre to the next (see Figure 1.2.9). This cross-talk is not a convolved error and therefore will not represent part of the system IRF. The peak was found to have a fairly constant magnitude, such that if the main TPSF were heavily attenuated (e.g. on the opposite side of an object to the source), the pre-peak would be very large in comparison. Where TPSFs were from detectors quite close to the source, the pre-peak would be negligible. Its broadening suggested it was a diffuse signal, but was not wide enough to be a TPSF from another channel. The mechanism causing the cross-talk is shown in Figure 2.1.4.

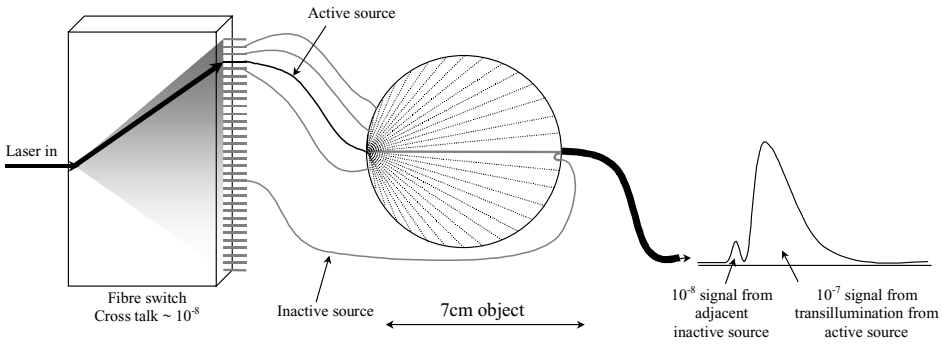

Figure 2.1.4 The origin of the pre-peak, via fibre switch light leakage into inactive source fibres.

Light leakage into de-activated channels of the DiCon (Berkeley, CA) VX500 fibre switch is quoted to be $< 10^{-8}$. However if the light from an active source is attenuated by $> 10^{-6}$ (e.g. at large source-detector separations, or in heavily absorbing or scattering objects), then light from inactive sources close to the detector will have a significant magnitude compared to the main TPSF. This signal will be slightly broadened since it has passed a short distance through the object, and is subject to convolution by the system IRF. It arrives before the main peak since the distance it has travelled is shorter than that of photons from the active source.

The pre-peak is tolerable if it is small, since post-processing can remove signal arriving prior to the 'ballistic photon time' (the time taken for unscattered photons to cover the source-detector spacing given the object's refractive index). However the *identifiable* pre-peak (or sometimes double pre-peaks) is only the signal from the closest inactive sources. Those inactive sources farther from the detector will also be adding to the TPSF which, while small, may skew measurements such as the mean photon flight time. Further, TCSPC was chosen since improved SNR can be achieved with increased exposure time, allowing measurements to be made in cases where attenuation is high. In the clinical setting, inherent limitations of the instrument to cope with attenuations above a certain level could cause erroneous data to be recorded (since attenuation would not be known for a given subject). This is demonstrated by the TPSF acquired on a human forearm, shown in Figure 2.1.5. Without subjective

examination of data prior to datatype extraction and image reconstruction, pre-peaks can overwhelm data and prevent convergence to the correct image solution.

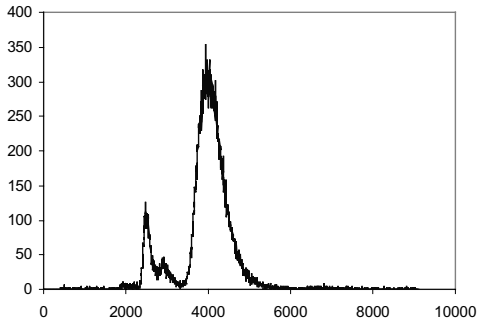

Figure 2.1.5 The TPSF collected by source 1 and detector 14 on the human forearm, without the secondary shutters. The source and detectors are on opposite sides of the arm. Double pre-peaks can be seen from the two closest inactive sources.

In order to overcome this problem, an auxiliary bank of shutters was constructed. This extra unit was intended to reduce signal from the in-active sources to allow better tolerance of highly attenuating objects. Figure 2.1.6 shows how these shutters function.

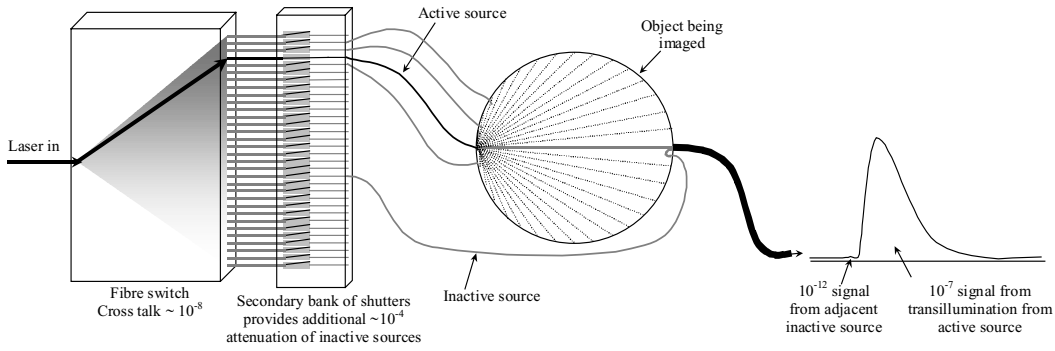

Figure 2.1.6 A secondary array of shutters should reduce the amount of light reaching the object via inactive sources.

It is noteworthy that there were many problems with light-leakage between fibres *within* the secondary shutter unit. In addition there was a variable delay in the time taken for each shutter to open, reducing exposure time for given sources during data acquisition. These problems required multiple modifications and drastically delayed successful integration of the secondary shutter unit into MONSTIR.

2.1.1.4 MCP-PMT cross-talk

Additional features are occasionally seen in the TPSF acquisition window. These appear to correspond to cross-talk originating from the MCP-PMTs. The intensity of light reaching an 8-anode MCP-PMT from one detector bundle may be significantly larger than the signal arriving via an adjacent detector bundle, mounted on the same MCP-PMT. The photocathode is common to all 8 channels, and the light shielding between adjacent 'channels' is not perfect. There are also likely to be some electrostatic effects inside the MCP-PMTs contributing to the cross-talk peak. The cross-talk peaks seem to be predictable. They are much broader than the

fibre-switch peaks (see Figure 2.1.5) since they are entire TPSFs from adjacent detector channels. The magnitude of the MCP-PMT cross-talk depends on a factor that is channel specific and on the magnitude of the signal in the interfering channel. It may be possible to 'background subtract' this cross-talk, if a calibration measurement is made of the signal in all detector channels when each detector is illuminated individually (the VOAs being used to block light entering from all other channels). In general MCP-PMT cross-talk is difficult to see since it is small in relation to the main TPSF, or its delay is similar to that of the measured TPSF so it superimposes the main signal. However Figure 2.1.7 shows a TPSF where MCP-PMT cross-talk is clearly visible as a large peak at around 10000ps.

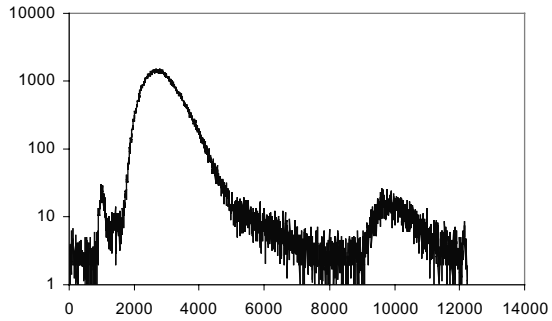

Figure 2.1.7 A TPSF collected on a cylindrical liquid phantom by source 12 and detector 22. The pre-peak is due to cross-talk from the fibre switch. The large peak after the TPSF is caused by cross-talk at the MCP-PMT.

2.1.1.5 Stochastic (dark counts and thermal noise)

The MCP-PMT units are surrounded by an insulating case that contains Peltier cooling units. If the MCP-PMTs are operated at room temperature with the VOAs closed and the laser off, they will register detected counts at around 2000Hz. In practice the Pelitier units (and surrounding water cooling) maintain the MCP-PMTs at ~ -20^{o}C during data acquisition. This reduces the dark count rate to less than 50Hz. These noise-related counts represent the (temperature dependent) probability of electrons spontaneously leaving the photocathode without an incident photon. This noise is therefore temporally un-correlated or stochastic. Threshold and walk settings on the CFD will affect the amount of thermal noise pulses that are counted as genuine photons. Ambient room light may also cause temporally un-correlated noise on TPSFs, although a low pass (optical) filter will reduce the amount of room light that can reach the photocathode. Stochastic noise manifests as a constant background offset on measured TPSFs as indicated by arrow (5) in Figure 2.1.1c.

2.1.1.6 Drift and jitter

While not strictly a form of noise, TPSFs measured with a particular channel have been shown to change their magnitude and temporal characteristics with time. Over a long time scale this is referred to as drift, jitter defines the case when jumps occur rapidly. This

phenomenon is of concern if calibration measurements are to be made that need to be valid throughout a series of data acquisitions.

MONSTIR's 32 separate detection channels all suffer from both drift and jitter. We can measure drift and jitter using the phantom shown in Figure 2.1.8. MIDAS (see 1.2.3.5) can be configured to automatically acquire a set of TPSFs every 80 seconds. Historically, the peak position of the TPSF was assessed using a 3rd order polynomial fit (to avoid spurious shot-noise peaks). Later, the mean value of the TPSF was used to evaluate the system stability. Additional parameters such as the integrated intensity and the full width half maximum (FWHM) are also calculated.

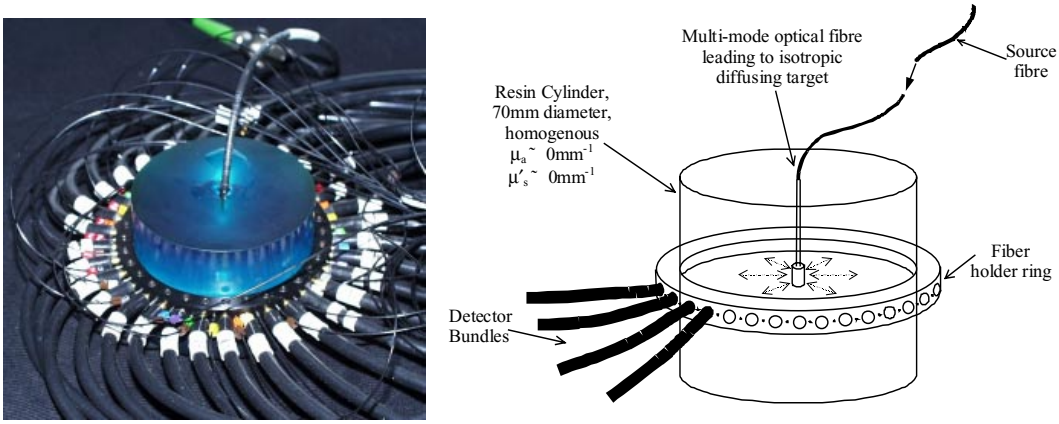

Figure 2.1.8 Tool used for drift tests. A small diffuser in the centre of the clear resin phantom is connected to a source fibre. The diffuser simultaneously illuminates all detectors with approximately the same intensity and temporal delay.

Figure 2.1.9 shows a 'drift test' acquired using the Ti:Sapphire laser. It shows instability in mean-time, FWHM and total counts. Large jitter can be seen in the mean-time plot for channels 24 and 25. Slight drift can be seen in total counts and the mean-time. The mean-times are shown as changes from their initial value, but the FWHM and total counts give an indication of the variability between channels. Although the distribution of light within the phantom is not entirely isotropic (so there may be different illumination characteristics for each detector) FWHMs are between 110 ps and 230 ps, and total counts between 200 000 and 900 000 (for 10 second exposure time). Jitter is occasionally seen in the FWHM and total counts.

Figure 2.1.9 highlights the differing dependencies of datatypes on system performance, since the width of the TPSF (indicated by the FWHM) is evaluated for image reconstruction via calculation of the variance (see section 1.3.3.3). Intensity measurements would be calculated from the total number of counts in a TPSF.

The origin of drift is likely to be temperature change. When the system is first turned on, a clear trend can be seen in the total counts and mean-time plots (note that mean-time calculation can depend on SNR, especially if simple calculation does not subtract background

counts). As the Peltiers units gradually cool the MCP-PMTs, SNR will improve. Also, variations in the length and refractive index of fibres, and the conduction properties of electrical cables, along with slight variations in alignment are likely to be responsible for drift. Making alignments more stable (e.g. addition of a diffuser to the reference photodiode) along with shielding instrumentation from the chiller fan and ensuring a more stable temperature environment following a sufficient warm-up period results in fairly drift-free data.

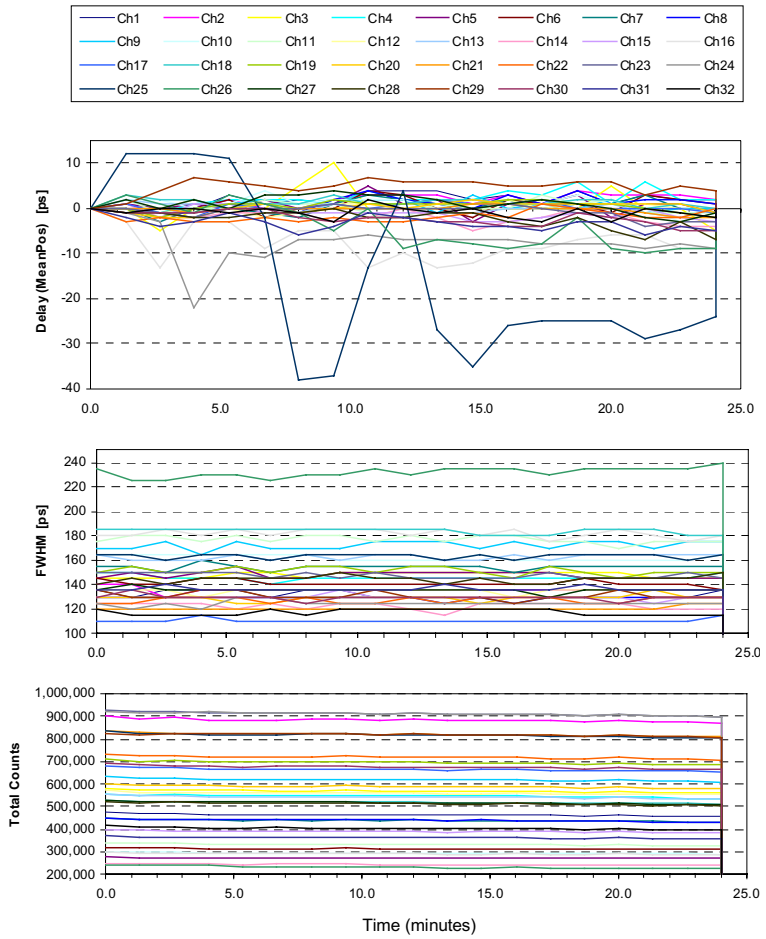

Figure 2.1.9 Drift Plots of data from MONSTIR, acquired over 25 minutes.

The major cause of jitter is thought to be insufficient warm-up time. It has been found that the CFDs and PTAs are optimal when warmed up for 7-10 hours prior to data acquisition. The Peltier units, MCP-PMTs, the laser and the chiller need to be on for 2-3 hours before data is acquired. Other aspects of the system that may affect jitter include:

- Faulty connections, e.g. Lemo plugs into CFDs being loose.
- Faulty PTA units - certain units, (despite warm-up) consistently produce jitter in mean-time drift tests, occasionally double peaks are seen in channel 24 and recently channel 13.
- Faulty CFD units - channels grouped together by a common quad CFD unit can suffer from jitter. FWHM jitter may originate from unstable settings of walk and threshold on the CFDs.

- INHIBIT (see section 1.2.3.5), causing MCP shut down can result in discrete jumps in mean-time and other data following re-start.
- Historically some apparent jitter on mean-time in drift tests may have been caused by poor fitting to the TPSF, the use of 'peak position' or the dependence of extracted mean-time on the calculation window size.

It may be possible to remove jitter from data by looking at extracted data and identifying discrete jumps, although subjective data analysis is not desirable in a clinical setting. A second-generation system would benefit greatly from having channels that are not entirely independent, where any drift and jitter would occur in all channels simultaneously and hence be much easier to identify and eliminate with automated post-processing. However sufficient warm-up, along with integration of the new fibre laser have been found to radically reduce drift and jitter problems.

2.1.2 Calibration measurements

From the descriptions of system errors and noise detailed above we can devise a simplified model of how a measured TPSF will differ from the true TPSF of the object being measured. An experimentally measured TPSF, represented by y(t) can be described as the convolution of the ideal TPSF D(t), and the IRF specific to each source detector pair I(t) + cross-talk and stochastic noise *f*(t). The composition of a measured TPSF is shown schematically in Figure 2.1.10.

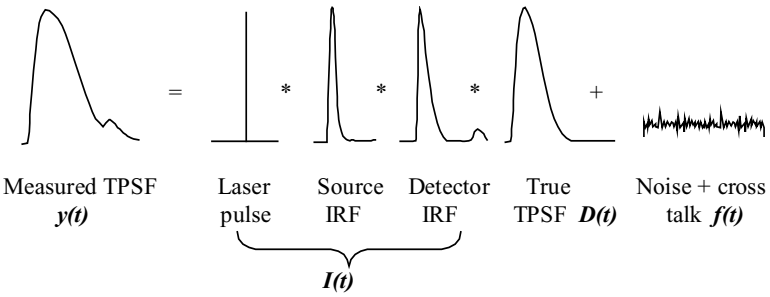

Figure 2.1.10 Measured TPSFs are a convolution of the true object TPSF, the source IRF, detector IRF and the laser pulse, plus background noise and cross-talk.

In Figure 2.1.10, * denotes a convolution. We can write this model as:

$$y_{n,m}(t) = D_{n,m}(t) * I_{n,m}(t) + f_{n,m}(t) \qquad [\,2.1.1\,]$$

where the system IRF is given by:

$$I_{n,m}(t) = laserpulse(t) * Source_n IRF(t) * Detector_m IRF(t) \qquad [\,2.1.2\,]$$

In order to calibrate data measured by MONSTIR we need to devise ways of measuring or accounting for $I_{n,m}(t)$ and $f_{n,m}(t)$ such that suitable corrections can be applied to the raw data

$y_{n,m}(t)$ to extract $D_{n,m}(t)$. The techniques devised to obtain calibration measurements are each described below. The earliest attempts at calibration were performed using three dedicated tools to measure relative source and detector channel properties. Subsequently, a device which fits onto the end of a source and detector pair was designed to obtain in-situ calibration measurements (termed a 'monstode'). Other techniques explored included the use of a reference measurement and using averaging to derive calibration factors. These techniques were more specifically aimed at correcting for $I_{n,m}(t)$ since relatively simple data processing techniques can be used be used to reduce the contribution of $f_{n,m}(t)$ (as described in section 2.1.3.2).

Considering the whole IRF as the optimal measurement for calibration is useful because measurements of simple parameters such as the relative delays in each channel, only provides calibration for datatypes that depend on delay. For example, the greatest potential source of uncertainty in mean flight time is the variation in the lengths of the many optical fibres and electrical cables that make up the system. Reflections and variations in TPSF *shape* due to the finite instrument response of the system also contribute error, but to a lesser degree. By contrast, variance (the central moment about the mean) is not affected by differences in delay, but is heavily dependent on the TPSF shape and hence the width of the IRF, and the secondary reflection peak.

A factor that heavily affects the effectiveness of calibration is the drift and jitter of the system as described in section 2.1.1.6. The optimal calibration system will require minimal time and effort, and preferably not require the optical fibres to be removed from the subject being studied. This would mean that multiple calibrations could be acquired to compensate for the continually changing properties of the system (as described in section 2.1.1.6).

2.1.2.1 Source, detector and absolute calibration tools

The first calibration system developed focused on measuring the system's IRF in a more indirect way than using the absolute calibration tool shown in Figure 2.1.2, since there are 1024 unique values of the system's IRF $I_{n,m}(t)$. This involved exploiting the fact that every TPSF measured using MONSTIR will be the convolution of the true TPSF of the object being measured and the (object independent) system IRF (as shown in Figure 2.1.10).

To facilitate measurement of the system IRF, $I_{n,m}(t)$, two calibration tools were constructed as shown in Figure 2.1.11 and Figure 2.1.2 (in section 2.1.1.1). The source and detector *relative calibration tool* (Figure 2.1.11) is an epoxy resin cylinder of diameter 70mm. The same phantom is used for drift tests and is shown in Figure 2.1.8. An optical fibre is coupled to a small scattering cylinder embedded in the centre.

In configuration A, a source fibre from the laser is connected to the central scatterer, providing isotropic illumination of the 32 detector bundles positioned in the central plane. This is termed the *relative detector calibration*.

In configuration B, a detector bundle is coupled to the fibre leading to the central scatterer, which then effectively collects any light incident on the central scatterer from each of the 32 source fibres positioned in the central plane. This is termed the *relative source calibration*.

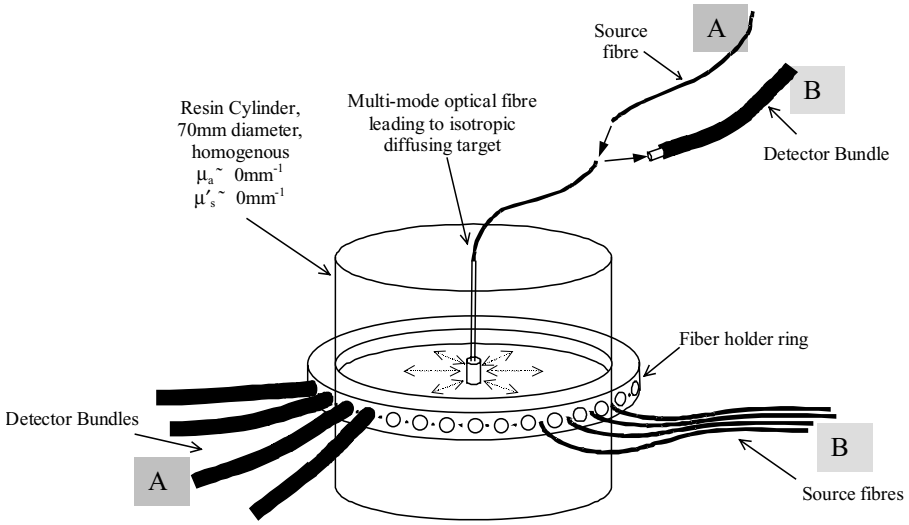

Figure 2.1.11 The relative calibration tool for IRF measurement for sources (configuration B) and detectors (configuration A) Figure 2.1.8 shows a photograph of this tool.

The *absolute calibration* tool shown in Figure 2.1.2 in section 2.1.1.1 holds a source fibre and a detector bundle at a known distance apart, shielded from extraneous light and utilising the beam divergence to illuminate the whole bundle. A small piece of diffusing paper is positioned in front of the bundle to ensure realistic excitation of all the modes in its constituent multi-mode fibres.

Calibration measurements made using these tools would be required each time an image data set is acquired in order to compensate for the system drift.

Each measurement can be decomposed into its constituents in the same way as the TPSF shown in Figure 2.1.10.

1) Relative detector calibration:

MONSTIR measures all 32 TPSFs from the detectors in a planar ring, equidistant from the scatterer in the centre of the phantom, which is delivering light via a single source fibre. All other source fibres are inactive. This provides the temporal response of each detector channel to the same incident signal.

$$Detcal_{spare,m} = laserpulse * Source_{spare}IRF * Detector_{m}IRF * Tool_{A}TPSF \quad [\,2.1.3\,]$$

2) Relative source calibration

MONSTIR measures the TPSF in detector channel 1 for each serial illumination of the scattering target by a source fibre. This provides the temporal characteristics of each incident source, when collected by the same detector channel.

$$Srccal_{n,1} = laserpulse * Source_n IRF * Detector_1 IRF * Tool_B TPSF$$

[2.1.4]

3) Absolute calibration

MONSTIR measures the TPSF resulting from detector 1 being illuminated by source 1.

$$Abscal_{1,1} = laserpulse * Source_1 IRF * Detector_1 IRF * Tool_{Abs} TPSF$$

[2.1.5]

where $Tool*TPSF$ refers to the IRF of the calibration tool itself. We want to be able to derive an expression for $I_{n,m}(t)$ [2.1.2] in terms of these three measurement sets. By rearranging [2.1.3], [2.1.4] and [2.1.5] and substituting into [2.1.2] we obtain:

$$I_{n,m}(t) = Srccal_{n,1} * Detcal_{spare,m} * Abscal_{1,1} \otimes Srccal_{1,1} \otimes Detcal_{spare,1} \otimes Tool_{Abs} TPSF$$

[2.1.6]

Where ⊗ denotes a de-convolution. Thus the IRF of all 1024 (32×32) source-detector combinations can be calculated by convolving and de-convolving TPSFs from the 65 measurements described above. Methods for performing this deconvolution are described in detail in section 2.1.3.1.

An advantage of this method is that the properties of the relative calibration tool, the $Tool_{A\ or\ B} TPSF$ term, should cancel out, eliminating factors such as the tool's wavelength dependence. However, as well as being a complex 3-step process, the method suffers from the disadvantage that it is not measuring truly diffuse light, which may affect the IRF. A further disadvantage is that the anisotropy of the central scatterer in the relative calibration tool will mean that calibrations for intensity of the TPSF (i.e. measurement of the absolute magnitude of each IRF) will not be accurate (see section 2.1.2.4). Finally, the $Tool_{Abs} TPSF$ term in [2.1.6] is not actually available from the measurements made. It represents the delay and broadening of a laser pulse passing through the absolute calibration tool. From the tool's design we can assume that broadening will be negligible, and the delay will be ~ 23ps (due to the geometry). The way in which $I_{n,m}(t)$ and this extra term are accounted for in extracting datatypes is described in section 2.1.4.1

2.1.2.2 Monstodes

The source, detector and absolute calibration tools described in section 2.1.2.1 above suffer from the major disadvantage that the sources and detectors have to be removed from the object being imaged in order to acquire calibration data. Further, if the object being imaged is not cylindrical, the fibres have to be moved into the 70mm diameter fibre-holder

ring, which is an intricate and time consuming process. The monstode (described in detail in section 2.7.2.2.1) offers an alternative to calibration tools and allows much more practical measurements of the system's IRF to be made. Figure 2.1.12 shows a monstode, and illustrates how a detector fibre held within the same monstode as an active source fibre can be used to measure back-reflected light.

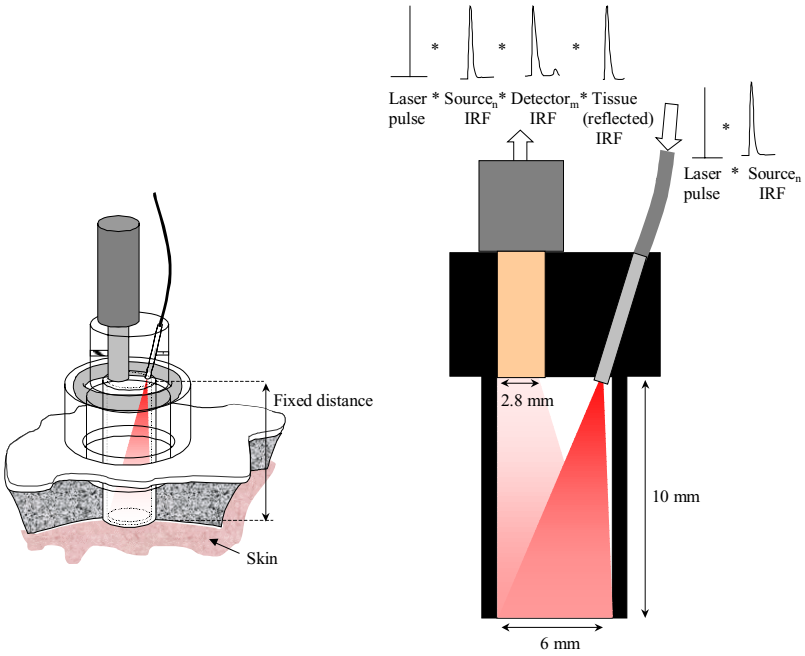

Figure 2.1.12 The monstode (See section 2.7.2.2.1) can be used for calibration measurements, by measuring light reflected from the skin

The detected signal represents the IRF of the source and detector convolved with the reflection TPSF of the tissue. A set of calibration data is acquired by serially illuminating each source (with heavy attenuation), and acquiring a TPSF with the detector corresponding to the illuminated source. Using the notation introduced above:

$$MonstodeCal_{n,n} = Laserpulse * Detector_n IRF * Source_n IRF * Tissue_{n,n} TPSF \quad [\,2.1.7\,]$$

However, there will only be 32 measurements: for source-detector pairs (1,1), (2,2), (3,3) etc. This is not sufficient information to derive all 1024 IRF combinations *n,m*. Hence, in addition, we must still acquire a set of $Srccal_{n,1}$ measurements independently.

Examination of data acquired using the relative calibration phantom for source calibration (section 2.1.2.1), has shown that the source IRFs are far more stable than detector IRFs. Drift in the source IRF may be caused by temperature variations changing the properties of the source fibres. Other changes may be due to fibre flexure and fibre damage (which may significantly affect the accuracy of the intensity of the IRF). However during the course of an experiment, it is likely that source IRFs will remain stable. In fact, since the new fibre-laser has been added to MONSTIR, a source-calibration measurement appears to remain similar over a time-scale of weeks. If we can measure $Srccal_{n,1}$ independently (using a

modified version of the relative calibration tool shown in Figure 2.1.11, since all 32 monstodes will not fit around a 70mm diameter cylinder), we can retrieve the IRF measurements using:

$$I_{n,m}(t) = MonstodeCal_{n,n} * Source_m IRF \otimes Source_n IRF \otimes Tissue_{n,n} TPSF$$

[2.1.8]

which from [2.1.4], can be evaluated using:

$$I_{n,m}(t) = MonstodeCal_{n,n} * Srccal_{n,1} \otimes Srccal_{m,1} \otimes Tissue_{n,n} TPSF$$

[2.1.9]

Clearly we are still relying on source calibration measurements made indirectly using a non-scattering phantom. However there is significantly less manipulation required here, and less terms which need to cancel compared to [2.1.6].

Note that again however there is a term in [2.1.9] which is unknown: $Tissue_{n,n}TPSF$. This represents the broadening and attenuation of the back-reflected pulse due to scattering of light in the tissue. The varying paths of photons within the monstode, causing broadening, along with the distance from the source to the tissue and back to the detector are not included in this term. This is because, since we are measuring the calibration in-situ, these effects will also affect all measurements made with the monstode and so will cancel when the calibration is applied. The properties of $Tissue_{n,n}TPSF$ will depend on the following factors:

- The absorption and scatter of the local tissue (which may vary with hair, freckles etc.)
- The specular reflectivity of the tissue (which depends on roughness of the skin and the refractive index mismatch between air and tissue).

The broadening of a laser pulse that is reflected from a scattering surface can be described analytically using (Patterson *et al*, 1989):

$$R(\rho,t) = \left(\frac{\mu_a + 3\mu'_s}{4\pi c}\right)^{3/2} \frac{1}{\mu'_s t^{5/2}} e^{-\left(\mu_a ct + \frac{(\rho^2 + (1/\mu'_s)^2)(\mu_a + 3\mu'_s)}{4ct}\right)}$$

[2.1.10]

where ρ is the source-detector separation. Although this is derived from Green's functions, uses the diffusion approximation, and is thus unsuitable for use in true modelling of the IRF of tissue for the monstode geometry, it does demonstrate that there is a dependence on the absorption and scatter properties of the tissue. Monte Carlo modelling of the monstode geometry has indicated that the dependence of a monstode calibration on surface optical properties is small (Hebden *et al*, 2002c). However it may be necessary to include some measure of $Tissue_{n,n}TPSF$ in calibration. For a given monstode this could potentially be evaluated using:

$$MonstodeCal_{1,1} \otimes Abscal_{1,1} = Tissue_{1,1} TPSF \otimes Tool_{Abs} TPSF$$

[2.1.11]

where $Tool_{Abs}TPSF$ is the fairly well characterised TPSF of the absolute calibration tool. While this correction wouldn't be performed on all 32 $MonstodeCal_{n,n}$ measurements, it may provide an estimate of the general properties of the tissue.

The dependence of the $Tissue_{n,n}TPSF$ on the tissue optical properties will be affected by how much of the detected light has actually entered the tissue and is scattered, and how much is reflected due to the refractive index mismatch between tissue and air. The largest angle possible between the source and detector in the monstode geometry is 17^o = 0.3 rad. In the limit when angle between the source and detector is small, and assuming that the refractive index of skin is 1.4 (and air has n=1), the specular reflectance r is given by:

$$r = \left(\frac{n-1}{n+1}\right)^2 = \left(\frac{1.4-1}{1.4+1}\right)^2 = 0.028 = 2.8\%$$

[2.1.12]

From [2.1.11], considering the geometry of the monstode, approximately 0.6% of light entering the tissue will be scattered back to the detector. Thus only 100x0.6/2.8=18% of the light being detected has been *into* the tissue. This factor further reduces the potential influence of the type of tissue under the detector on the accuracy of a calibration performed using a monstode.

2.1.2.3 Reference measurements and difference imaging

A reference measurement implies a measurement made on an object fairly similar to the object being imaged. For example:

- A homogenous phantom, with similar properties to the object being imaged.
- The object itself at some different state (e.g. a liquid phantom with and without an embedded object, or the brain before and after a haemorrhage)
- The same object, but imaged at a different wavelength.

Such reference measurements can be combined with image data to yield a form of calibrated data. This can be explained by the introduction of an important property of convolved functions that can be exploited during manipulation of time-resolved data. As described in section 1.3.3 measured TPSFs are inevitably used to derive datatypes. These are extracted via integral transforms such as:

$$\mathcal{M}_{p,u}[y_{n,m}(t)] = \frac{1}{N}\sum_{t=T_1}^{t=T_2} y_{n,m}(t) t^u e^{-pt}$$

[2.1.13]

where

$$N = \sum_{t=T_1}^{t=T_2} y_{n,m}(t)$$

[2.1.14]

and where p and u specify the datatype being calculated. It is well known that the Fourier transform of a convolved function is equal to the product of the Fourier transform of each function:

$$Y_{n,m}(\omega) = \int y_{n,m}(t)e^{-i\omega t}dt = F[y_{n,m}(t)] = F[D_{n,m}(x) * I_{n,m}(v)] = F[D_{n,m}(x)] \times F[I_{n,m}(v)]$$

[2.1.15]

The most usual way to perform a de-convolution is to take to Fourier transform of the data and divide by the Fourier transform of the IRF, and then transform the result back into the time domain and then derive datatypes. However, it can be shown that, like Fourier transforms, certain of the integral transforms used to extract datatypes also result in simple relations between transformed convolved functions. In particular we find that when p=0 and u=1 in [2.1.13]:

$$\mathcal{M}_{0,1}[y_{n,m}(t)] = Mean[y(t)] = \frac{\sum_t y(t) \times t}{\sum_t y(t)} = \frac{\sum_{t=(x+v)} D(x) * I(v) \times t}{\sum_{t=(x+v)} D(x) * I(v)}$$

[2.1.16]

that

$$Mean[y(t)] = Mean[D(t)] + Mean[I(t)]$$

[2.1.17]

Intensity also proves to have simple relations with convolved functions:

$$N = Intensity[y(t)] = \sum_t y(t) = Intensity[D(x)] \times Intensity[I(v)]$$

[2.1.18]

Note that not all datatypes can be manipulated in this way. A more thorough analysis of these relations is given in section 2.1.3.1.

The results above reveal additional possible calibration methods. The use of a reference measurement is commonplace amongst optical tomography researchers. This is because most simplified image reconstruction techniques (generally linear – see section 1.3.2.2) utilise the difference between two measurements made on similar objects, or the same object in different states. However, due to the convolution relations shown above, depending on the datatype being used, a significant advantage of using a reference measurement is that some measurement (and model) errors will cancel automatically, calibrating the difference data.

In general, the IRF of a particular source-detector combination does not change (apart from when affected by temporal drift, fibre damage and, possibly, VOA settings and the wavelength of light being used). Hence *any data* acquired with that source detector pair will affected by their unique IRF. It is feasible then that we can use two sets of image data (that use the same source-detector pairs *n,m*) to directly eliminate all $I_{n,m}(t)$ by acquiring:

$$y^{Object1}{}_{n,m}(t) = D^{Object1}{}_{n,m}(x) * I_{n,m}(v)$$

$$y^{Object2}{}_{n,m}(t) = D^{Object2}{}_{n,m}(x) * I_{n,m}(v)$$

[2.1.19]

If we extract datatypes e.g. intensity, and calculate the ratio (or the difference of the log) of the two data sets, we can use the relations derived above to yield:

$$\frac{Intensity[y^{Object1}(t)]}{Intensity[y^{Object2}(t)]} = \frac{Intensity[D^{Object1}(t)] \times Intensity[I(t)]}{Intensity[D^{Object2}(t)] \times Intensity[I(t)]} = \frac{Intensity[D^{Object1}(t)]}{Intensity[D^{Object2}(t)]} \quad [\ 2.1.20\ ]$$

and similarly for mean flight time, taking the difference, we get:

$$Mean[y^{Object1}(t)] - Mean[y^{Object2}(t)] = Mean[D^{Object1}(t)] - Mean[D^{Object2}(t)] \quad [\ 2.1.21\ ]$$

In both cases the $I_{n,m}(t)$ term cancels out. The cancellation of $I_{n,m}(t)$ between two datatypes using this method is generally better than results achieved using the calibration tools and monstodes described in sections 2.1.2.1 and 2.1.2.2. This may be due to the fact that by using a reference measurement we are directly measuring *every* IRF for all n,m combinations of data, rather than trying to compose all of the n,m pairs from individual indirect measurements. However note from [2.1.20] and [2.1.21] that we are left with data which does not truly represent $D^{Object1}(t)$. This is an important aspect of using two measurements to eliminate the IRF. The various different ways to account for the extra $D^{Object2}(t)$ term, and the different types of image reconstruction possible using difference data are described in section 2.1.4 .

2.1.2.4 Intensity calibration

Measurement of the IRF should allow calibration of all datatypes. However the accuracy with which we can measure the features of the IRF determine how readily each datatype can be calibrated. As a consequence, calibrating for intensity is very difficult. The indirect measurements described in section 2.1.2.1 measure the shape of the IRF quite well, but fail to provide information about the true attenuation due to each source and detector channel. This is due in part to the fact that the positions of the VOAs during data acquisition will not be accounted for by the calibration measurements. Calibrating for VOA attenuation may be difficult since each VOA provides slightly different attenuation for each setting (see section 1.2.3.3). Repeatability of VOA attenuation has however been shown to be quite reliable, so measurement of VOA attenuation factors for all 8 VOA settings for all 32 channels is feasible.

Intensity measurements have been successfully utilised for image reconstruction thanks to direct measurement of the IRFs achieved by using a second reference object as described above in section 2.1.2.3. The applicability of using a reference measurement is limited however since a suitable reference object must be available, the surface coupling between both objects must be similar, and the VOA settings for each source-detector pair must be the same for both acquisitions.

A method for calibrating intensity data suggested by (Boas *et al*, 2001a) considers coupling variations, and potentially other experimental errors, as unknowns that can be found as part of the image reconstruction process. In the future, TOAST could be modified to attempt to calculate the 64 coupling coefficients *and* the optical property distribution that corresponds to the measured data. To date however, no reliable method has been found to reliably calibrate *absolute* intensity.

2.1.3 Data processing and datatype extraction

TOAST, and most other reconstruction methods, require data in the form of datatypes. As introduced in section 1.3.3 these are parameters extracted from measured TPSFs. The advantage of using datatypes is that they represent the fundamental characteristics of a TPSF as a number of single parameters. Creating a model to fit these few parameters, rather than finding a fit to the entire TPSF, greatly reduces the computational expense of image reconstruction.

Once TPSF data has been measured, we must extract datatypes while accounting for systematic errors and noise. Techniques developed to improve the calibration and datatype extraction process are described below. Initially we described in detail the deconvolution-equivalent processing developed from examining the effect of extracting datatypes from convolved functions (as introduced in section 2.1.2.3). This method can be used to apply IRF calibrations using measurements described in section 2.1.2 . We then describe the other data processing methods and considerations required to remove other sources of noise from TPSFs prior to extraction of datatypes.

2.1.3.1 Techniques for deconvolution and datatype extraction

Since a TPSF measured by MONSTIR is the convolution of the true TPSF and the system's IRF (plus noise), we can calibrate by performing a deconvolution using the IRF measurements described in sections 2.1.2.1 and 2.1.2.2. This generally involves performing a Fourier transform on the measured TPSF, dividing by the Fourier transform of the IRF and then transforming the result back into the time-domain (with suitable filtering). This takes advantage of the relation:

$$F[y_{n,m}(t)] = F[D_{n,m}(t) * I_{n,m}(t)] = F[D_{n,m}(t)] \times F[I_{n,m}(t)] \quad [2.1.22]$$

where

$$y_{n,m}(t) = D_{n,m}(t) * I_{n,m}(t) \quad [2.1.23]$$

However this kind of deconvolution is time-consuming, computationally intensive, and may impose certain errors on the resulting TPSF due to the necessary filtering. In fact, in order to extract $I_{n,m}(t)$ from calibration measurements, multiple convolutions and

deconvolutions are required ([2.1.6] and [2.1.9]), potentially amplifying errors. Note that if we are assuming that $y_{n,m}(t)$ is a TPSF measured by MONSTIR, we are neglecting the additional noise term $f(t)$ (see [2.1.1]). The following analysis assumes that the effects of this additive term have already been removed from data and calibration TPSFs via pre-processing methods described later in section 2.1.3.2.

An alternative to using a Fourier transform to perform deconvolution exists. It was found that certain of the transforms used to extract datatypes from measured TPSFs had similarly simple relations with convolved functions. In practice, the datatypes most frequently employed for image reconstruction to date have been the mean flight time, variance (about the mean), integrated intensity, and the normalised Laplace transform. These are simple to extract, were implemented early on in the development of TOAST and shown to be effective in simulations (Schweiger *et al*, 1999b).

Where $y_{n,m}(t)$ is the discrete form of $\Gamma_{n,m}(t)$ defined in [1.3.20] in 1.3.1.2, we can see that extraction of datatypes from measured TPSFs can be achieved by evaluating:

$$\mathcal{M}_{p,u}[y_{n,m}(t)] = \frac{1}{N}\sum_{t=T_1}^{t=T_2} y_{n,m}(t) t^u e^{-pt} \qquad [2.1.24]$$

or

$$Central\mathcal{M}_u[y_{n,m}(t)] = \frac{1}{N}\sum_{t=T_1}^{t=T_2} y_{n,m}(t)(t - < t >)^u \qquad (u > 1), \qquad [2.1.25]$$

where

$$N = \sum_{t=T_1}^{t=T_2} y_{n,m}(t)$$

apart from complex amplitude and phase, and integrated intensity where N=1. $y_{n,m}(t)$ is the measured TPSF, which is currently sampled every 5 ps. T_1 and T_2 are the minimum and maximum times over which the datatypes are calculated. Values of p and u govern which datatype is calculated are summarised (see Table 1.3.1 in 1.3.3).

The discrete form of the convolution in [2.1.23] is:

$$y(t) = \sum_v D(t-v)I(v) . \qquad [2.1.26]$$

From [2.1.24], the mean flight time of y(t) is:

$$\mathcal{M}_{0,1}[y_{n,m}(t)] = Mean[y(t)] = \frac{\sum_t y(t) \times t}{\sum_t y(t)} . \qquad [2.1.27]$$

If we substitute [2.1.26] into [2.1.27] we obtain:

$$Mean[y(t)] = \frac{\sum_t y(t) \times t}{\sum_t y(t)} = \frac{\sum_t \sum_v D(t-v)I(v)t}{\sum_t \sum_v D(t-v)I(v)} = \frac{\sum_x \sum_v D(x)I(v)(x+v)}{\sum_x \sum_v D(x)I(v)} = \frac{\sum_x D(x) \times x}{\sum_x D(x)} + \frac{\sum_v I(v) \times v}{\sum_v I(v)} \qquad [2.1.28]$$

where *x=(t-v)*, which leads to:

$$Mean[y(t)] = Mean[D(x)] + Mean[I(v)]\,.$$

[2.1.29]

So rather than de-convolving the IRF from a TPSF and then calculating mean flight time, the above analysis implies that we can get the same result from simply calculating the mean-time of the IRF and subtracting it from the mean-time of the measured TPSF. Further, as shown in appendix A (2.9.5), the same is true for the 2nd and 3rd central temporal moments (variance and skew) where:

$$Central\mathcal{M}_2[y_{n,m}(t)] = Variance[y(t)] = \frac{\sum_t y(t)\times(t-\langle t\rangle)^2}{\sum_t y(t)} = Variance[D(x)] + Variance[I(v)]$$

[2.1.30]

and

$$Central\mathcal{M}_3[y_{n,m}(t)] = Skew[y(t)] = \frac{\sum_t y(t)\times(t-\langle t\rangle)^3}{\sum_t y(t)} = Skew[D(x)] + Skew[I(v)]$$

[2.1.31]

For the normalised Laplace transform and integrated intensity datatypes, it can be shown (see appendix A (2.9.5)) that the ideal datatype D(t) is *multiplied* by a factor equal to the datatype of the IRF (I(t)):

$$Intensity[y(t)] = \sum_t y(t) = Intensity[D(x)]\times Intensity[I(v)]$$

[2.1.32]

$$\mathcal{M}_{s,0}[y_{n,m}(t)] = Laplace[y(t)] = \frac{\sum_t e^{-st} y(t)}{\sum_t y(t)} = Laplace[D(x)]\times Laplace[I(v)]$$

[2.1.33]

where s = the Laplace coefficient. Simple relations also exist for the complex amplitude and phase of convolved functions (see appendix A (2.9.5)).

$$Amp[\mathrm{Y}] = Amp[\mathrm{D}]Amp[\mathrm{I}]$$

$$Phase[\mathrm{Y}] = Phase[\mathrm{D}] + Phase[\mathrm{I}]$$

[2.1.34]

Note that real and imaginary parts of $y_{n,m}(t)$ can be calculated using:

$$\begin{aligned}\mathcal{M}_{(i\omega),0}(y_{n,m}(t)) &= \sum_{t=T_1}^{t=T_2} y_{n,m}(t)e^{-i\omega t}\\ &= \sum_{t=T_1}^{t=T_2} y_{n,m}(t)(\cos(\omega t) + i\sin(\omega t))\\ &= \sum_{t=T_1}^{t=T_2} y_{n,m}(t)\cos(\omega t) + i\sum_{t=T_1}^{t=T_2} y_{n,m}(t)\sin(\omega t)\end{aligned}$$

[2.1.35]

While all of the datatypes detailed above have simple relations with convolved functions, making them easy to calibrate and ideal for difference imaging (see section 2.1.2.3), there are

certain datatypes that do not have simple relations e.g. Mellin Laplace (where p>0 and *u*>0 in [2.1.24]) or kurtosis (where *u*=4 in [2.1.25]). A general rule (derived in the appendix A (2.9.5), demonstrates that the relation between datatypes of the form shown in [2.1.24] can be deduced using the following series:

$$\begin{aligned}
&\mathcal{M}_{p,x}[y] \\
&= \mathcal{M}_{p,0}[D]\mathcal{M}_{p,x}[I] \\
&+ x \times \mathcal{M}_{p,1}[D]\mathcal{M}_{p,(x-1)}[I] \\
&+ \frac{x(x-1)}{2!} \times \mathcal{M}_{p,2}[D]\mathcal{M}_{p,(x-2)}[I] \\
&+ \frac{x(x-1)(x-2)}{3!} \times \mathcal{M}_{p,3}[D]\mathcal{M}_{p,(x-3)}[I] \\
&+ \frac{x(x-1)(x-2)(x-3)}{4!} \times \mathcal{M}_{p,4}[D]\mathcal{M}_{p,(x-4)}[I].....
\end{aligned}$$

[2.1.36]

Note that $\mathcal{M}_{0,0}[\ ]$ is 'normalised intensity' which = 1. This series predicts for example that mean-time $\mathcal{M}_{0,1}[y]=\mathcal{M}_{0,0}[D]\mathbf{x}\mathcal{M}_{0,1}[I]+\mathcal{M}_{0,1}[D]\mathbf{x}\mathcal{M}_{0,0}[I]$, in agreement with [2.1.29] and also that $\mathcal{M}_{0,3}[y]=\mathcal{M}_{0,3}[I]+3\mathcal{M}_{0,1}[D]\mathbf{x}\mathcal{M}_{0,2}[I]+3\mathcal{M}_{0,2}[D]\mathbf{x}\mathcal{M}_{0,1}[I]+\mathcal{M}_{0,3}[D]$, the 3rd temporal moment, is not simply composed of the sum of the 3rd moment of the two convolved functions, but contains two extra terms containing the 2nd and 1st moments also. This series also holds for Mellin-Laplace datatypes where p = s >0 where s is the Laplace coefficient. Note however that $\mathcal{M}_{s,0}[\ ]$ is equal to the Laplace datatype such that all Mellin-Laplace datatypes with u>0 will have extra terms in their convolution relations.

The corresponding series for central moments (u >1) is shown in [2.1.37]. Again $Central\mathcal{M}_0[\ ]=1$, and note also that $Central\mathcal{M}_1[\ ] = 0$.

$$\begin{aligned}
&Central\mathcal{M}_x[y] \\
&= Central\mathcal{M}_0[D]Central\mathcal{M}_x[I] \\
&+ \frac{x(x-1)}{2!} \times Central\mathcal{M}_2[D]Central\mathcal{M}_{(x-2)}[I] \\
&+ \frac{x(x-1)(x-2)(x-3)}{4!} \times Central\mathcal{M}_4[D]Central\mathcal{M}_{(x-4)}[I] \\
&+ \frac{x(x-1)(x-2)(x-3)(x-4)(x-5)}{6!} \times Central\mathcal{M}_6[D]Central\mathcal{M}_{(x-6)}[I].....
\end{aligned}$$

[2.1.37]

This demonstrates that central moments are simple sums up until u = 4 (kurtosis), where an additional term appears (see appendix A (2.9.5.8)).

So for certain datatypes, we can incorporate deconvolution of the system's IRF into the datatype extraction process by exploiting the relations derived. This can be done either by using the same relations to extract $I_{n,m}(t)$ from independent calibration methods (as described

in sections 2.1.2.1 and 2.1.2.2) or by using a reference measurement to cancel out common $I_{n,m}(t)$ between two sets of datatypes (see section 2.1.2.3).

2.1.3.2 Techniques for removal of non-convolved noise

The effectiveness of the convolution relations for datatypes described in section 2.1.3.1 depends on the influence of the additive noise term *f*(t) in [2.1.1]. Stochastic noise, shot noise and cross-talk are not convolved functions. Higher order moments such as skew are heavily dependent on such noise and its effects may overwhelm the corrections calculated for individual source - detector combinations. Background noise on calibration TPSFs may have far more effect on the datatype derived than the shape or temporal position of the actual TPSF, hence limiting the usefulness of the subsequent calibration. Although shot noise, stochastic noise and cross-talk can be reduced via hardware alterations (see section 2.1.1), processing for remaining non-convolved features is necessary.

2.1.3.2.1 Low pass filtering

High frequency noise removal using a low pass filter is inappropriate as a method of removing stochastic noise from TPSFs. This is because a low pass filter is simply a multiplication of the Fourier transform by a chosen window, which in turn represents the convolution of a constant function (the inverse Fourier transform of the window) in the time domain. As demonstrated in section 2.1.3.1, in most cases, convolving TPSFs with any constant function will simply change the subsequently derived datatype by a constant amount. Hence the absolute values of some datatypes (variance, intensity and Laplace) will be adversely affected by low pass filtering whereas mean and skew will be unchanged, the effects of the stochastic noise will remain. However, potentially, high frequency spikes due to shot noise could be removed using a median filter, which may help to reduce spurious errors in higher order datatypes.

2.1.3.2.2 Background subtraction

A background offset intensity will influence the calculation of datatypes. For example, the existence of apparent photon counts at flight times shorter than the ballistic photon time (where t = source detector separation / (c / n)), where c = speed of light and n is the refractive index) will produce an effect on datatypes not accounted for by the forward model. A background subtraction scheme will reduce the effects of a constant offset, such as that due to the level of temporally un-correlated noise on a particular TPSF (dependant on thermal noise from the MCP-PMT and ambient light). To avoid features such as the pre-peak (section 2.1.1.3) representative background noise can be sampled at long photon flight times (i.e. following the region containing the main TPSF, where the tail of the TPSf is < noise). For datatype extraction from MONSTIR TPSFs, the last 160 data points are averaged into 16 bins

(10 data points = 50 ps). The lowest value bin is then assumed to represent the background level. This method avoids the possibility of including the reflection peak (section 2.1.1.1). This background level value is subtracted from every data point in the TPSF automatically and before extraction of datatypes.

The low frequency background variation due to cross-talk in the 8-anode MCPs (section 2.1.1.4), may continue to affect datatypes following 'constant background' subtraction. However it may be possible to adequately predict the magnitude and temporal position of such cross-talk, if the fraction of cross-talk occurring can be measured, and assumed to remain constant. However, note that if a cross-talk correction based on measured values is to be applied it must be subtracted *before* constant-background subtraction is performed.

2.1.3.2.3 Dynamic windowing

The position of a TPSF in the data acquisition temporal window depends solely on the length of the fibres and cables for a particular channel compared to the length of the cables delivering the reference pulse to the PTAs. So a length of recorded data representing a 12500ps temporal window may have the TPSF in the centre, or perhaps shifted to the left or right. The position of the TPSF in the window, and its corresponding unique time-scale is accounted for in calibration, although central datatypes such as variance are unaffected by relative time-scale shifts. However, when we calibrate, we combine calibration TPSFs acquired on a number of different channels (e.g. all source calibrations are acquired using detector 1). We need to devise a way to ensure that the position of the TPSF in its temporal window does not adversely affect calibration of extracted datatypes. Figure 2.1.13 shows a typical TPSF along with a system IRF measured using the tool shown in Figure 2.1.2. Appropriate selection of the range over which to calculate a datatype from a TPSF can reduce the effects of the stochastic noise, shot noise and cross-talk either side of the TPSF which isn't accounted for in modelled datatypes.

Figure 2.1.13 demonstrates that measurements from MONSTIR imply the arrival of photons before the ballistic photon time. Also, while the TPSF tail decays exponentially with increasing time, there is a clear point in this log-plot where background noise overwhelms any signal from the TPSF tail. Furthermore, not all channels have the same amount of data before or after the TPSF as each other. If we calculate TPSFs including these regions, we will extract datatypes that will never match those modelled by TOAST. This is particularly important for the central moments (variance and skew), which are more dependent on photons far from the mean of the TPSF (as shown in Figure 2.2.16 in section 2.2.6.2.1).

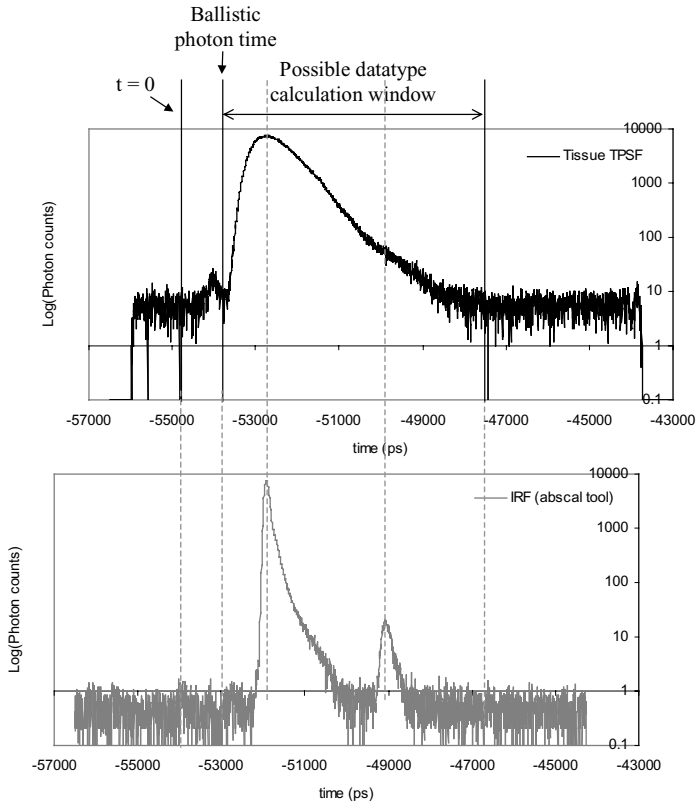

Figure 2.1.13 Plots of a typical TPSF (from a scattering phantom) and an IRF (measured using the absolute calibration tool). A temporal window is indicated that would allow datatype extraction to avoid the effects of background noise. The IRF has been aligned to demonstrate how the temporal window would also be used to generate calibration datatypes so that convolved errors will cancel.

For effective application of the calibration procedures described in sections 2.1.2.1 and 2.1.2.2, and for difference imaging as described in section 2.1.2.3, the derivation of datatypes from the main data set and the calibration or reference data sets should be equivalent. A temporal window for data measured on an object must contain as much as possible of the broadened TPSF. A window for the IRF that is to be used effectively to provide the de-convolved true TPSF must include the main features of the IRF. In theory these features will fit into a window of width v, where t=x+v, where t is the window used for the measured TPSF and x is a window that would contain the whole de-convolved TPSF. In practice however it is easier to use a constant temporal window for datatype calibration, especially since absolute calibration requires so many convolutions and deconvolutions. Using a constant window means calculating all TPSFs involved in derivation of a datatype over a temporal window that is equivalent, for each given source-detector pair. This ensures that any reflections or other temporal features specific to the source - detector pair are included in the calculation of calibration values if they have influenced the datatype values derived from the actual data TPSF, as shown in Figure 2.1.13. Note however that this technique assumes that the noise either side of the IRF, or a constant background, has been adequately accounted for. Figure 2.1.13 illustrates the need for background subtraction *before* datatype extraction for calibration or before deconvolution proper so that we calculate:

$$y(t)^{meas} - f_y(t) = D(x) * (I(v)^{meas} - f_I(v))$$

[2.1.38]

In order to use a constant temporal window, we must define the window as the time either side of a fixed point. Although it requires the assumption that the MONSTIR TPSFs are approximately symmetric, the mean flight time has been chosen as this anchor point, since it is easier and more reliable to calculate than peak position. The TPSFs in Figure 2.1.13 have been aligned in this way, indicating the window that will be used for datatype extraction from both TPSFs. We need an iterative procedure for derivation of the mean flight time datatype since the mean of the TPSF is required to deduce the window over which to calculate a better estimate of the mean flight time. Once the optimal window for a particular TPSF has been deduced, the same distance either side of the mean flight time of each calibration TPSF is used to calculate the 'calibration datatypes'.

Once the raw mean flight times have been calculated, it is possible to deduce absolute time = 0 using [2.1.6] or [2.1.9]. This allows further refinement of the window since we can eliminate the inclusion of any apparent photon counts prior to either the entrance time of photons (t = 0) or, if necessary, the ballistic photon time. Since most of the noise on TPSFs will be normally distributed (except cross-talk), its effects should cancel during calculation of most datatypes. However, particularly for higher order moments, spurious peaks a distance away from t = 0 (or t = mean for central moments) could overwhelm the resulting datatype. Windowing considerations for specific datatypes include:

- For Laplace, the chosen coefficient of s = 0.005 ps^{-1} results in the datatype being particularly sensitive to early photons, and hence the pre-peak, so the ballistic photon time is used to establish the window for calculation.
- Since skew is an odd function, the use of a symmetric window either side of the TPSFs mean should aid cancellation of remaining additive stochastic noise. If a pre-peak is present, a section of background noise from the end of the TPSF could be used to mimic the datatype calculation window extending an equal distance before and after the mean.
- Since variance is an even function, it is desirable to calculate the datatype over as small a window as possible, while not truncating the tail of the TPSF (see section 2.2.6.2.1). However there is also the possibility of extrapolating the tail of the TPSF or using a model fit to better account for the truncation of the TPSF tail by noise.
- Intensity and mean flight time should be least susceptible to normally distributed noise although inadequate background subtraction and contributions from pre-peaks and low frequency noise components such as MCP cross-talk suggest that using a reduced temporal window for their calculation should be considered.

2.1.4 Overview of different imaging paradigms

The variety of processing and calibration methods described above need to be used in a number of combinations to yield data suitable for the multiple types of image reconstruction possible from time-resolved data. Requirements are listed below.

2.1.4.1 Absolute data for absolute imaging, using calibration tools

Data are acquired on the object and then immediately on the calibration tools detailed in section 2.1.2.1. A source calibration, detector calibration and absolute calibration are required. Calibration can be achieved by performing background subtraction and cross-talk correction followed by using [2.1.6] in a full deconvolution and then extracting datatypes. Alternatively raw datatypes can be extracted from background-subtracted and windowed calibration and measurement TPSFs and then combined using:

$$\mathcal{M}[D_{n,m}] = \mathcal{M}[y_{n,m}] - (\mathcal{M}[Srccal_{n,1}] + \mathcal{M}[Detcal_{spare,m}] - \mathcal{M}[Srccal_{1,1}] - \mathcal{M}[Detcal_{spare,1}] + \mathcal{M}[Abscal_{1,1}] - \mathcal{M}[Tool_{Abs}TPSF]) \quad [\,2.1.39\,]$$

The $\mathcal{M}[Tool_{Abs}TPSF]$ term represents the datatype of the TPSF of the absolute calibration tool which should have negligible broadening and a delay of ~ 23ps. A table of appropriate values for each datatype is shown in appendix A (Table 2.9.1, section 2.9.6).

This process will yield absolute datatypes, suitable for non-linear reconstruction using TOAST, yielding images of absolute scatter and absorption. It is suitable for extraction of mean-time, variance, skew and phase. For Laplace the addition and subtraction in [2.1.39] are replaced with multiplication and division (or the log of datatypes are used). Note that Mellin-Laplace and higher order moments must be extracted from properly de-convolved TPSFs (see section 2.1.3.1).

Intensity and complex amplitude extracted in this way will not be suitable for image reconstructions since the VOAs will have been used during data acquisition but not during calibration measurements (see section 2.1.2.4). It is not currently possible to accurately quantify the effect of each individual VOA setting on the number of photons detected. Further, the calibration tools described in 2.1.2.1 were not designed for intensity calibration. The central scatterer is unlikely to emit light isotropically when used in the detector calibration, and it will not collect light in the same way for each source during a source calibration.

2.1.4.2 Absolute data for absolute imaging, using monstode calibration

Data are acquired on the object to be imaged, and then a monstode calibration is performed in-situ. This involves serially illuminating each source (with high attenuation of the incident beam), and collecting back-reflected light using the detector within the same

monstode as the illuminated source. A source calibration measurement must also be made, although it is less likely to be susceptible to system drift and jitter and so doesn't need to be acquired immediately. An absolute calibration measurement can also be made if an estimate of the effect of the optical properties of the object's surface is required (see [2.1.11]). Calibration can be achieved by performing background subtraction and cross-talk correction followed by using [2.1.9] in a full deconvolution and then extracting datatypes. Alternatively raw datatypes can be extracted from background-subtracted and windowed calibration and measurement TPSFs and then combined using:

$$\mathcal{M}[D_{n,m}] = \mathcal{M}[y_{n,m}] - (\mathcal{M}[MonstodeCd_{n,n}] + \mathcal{M}[Srccal_{n,1}] - \mathcal{M}[Srccal_{m,1}] - \mathcal{M}[Tissue_{n,n}TPSF] \quad [\,2.1.40\,]$$

The resulting datatypes are exactly equivalent to those extracted using the calibration tools as described in 2.1.4.1. The only additional considerations are that the $Tissue_{n,n}TPSF$ term may impose some variable not accounted for by calibration (e.g. due to hair, freckles etc.). This may affect some datatypes very differently to others. The other difference is that monstode measurements are *potentially* suitable for extraction of calibrated intensity data since the IRF is measured much more directly, and without moving the fibres to a different location (which also results in very good mean-time calibrations). Intrinsically monstodes also illuminate a much larger area than just placing a fibre in contact with the skin, so movement and variations in coupling should have less effect on data. However the monstode calibration currently still requires use of the source calibration phantom, and hence will have certain problems with accuracy of source intensity calibration. Further, monstode measurements cannot currently account for VOA settings used during data acquisition. The using of monstode calibration measurements for wavelength-difference imaging is discussed in section 2.1.4.5.

2.1.4.3 Difference data for absolute imaging using an arbitrary reference phantom

Data are acquired on the object of interest and then on a reference object as soon as possible afterwards. The same VOA settings and source-detector combinations should be used (which limits the allowable optical properties and shape of the reference phantom in itself). No external calibration is required although the availability of a t=0 estimate will aid calculation of datatypes over a suitable window (section 2.1.3.2.3). Data must be processed to remove cross-talk and background noise, then raw datatypes can be extracted from both the object and reference sets.

If the optical properties of the reference are very similar to those of the object of interest, the un-calibrated data can be used directly in a linear or non-linear reconstruction. For a linear reconstruction the difference in the two data sets is used as input data (see section 1.3.2.2). An

initial estimate of the optical properties of the reference object is required to generate suitable PMDFs for the linear inversion. An estimate may be available from manufacture of the reference phantom, or could be calculated by using averaging as described in appendix A (2.9.7). The resulting images should represent absolute absorption and scatter distributions within the object (assuming that the reference object is well characterised and homogenous). Similarly for a non-linear reconstruction, the two sets of data can be read into TOAST (under FILE and REFFILE). The estimated optical properties of the reference phantom should be used as the initial starting parameters for the reconstruction. Again, absolute images should result. A further advantage of this technique, besides the automatic calibration of raw data, is that errors in the FEM model (such as optode and mesh boundary positioning errors) also cancel to a certain degree. The difference between two models is minimised to match the difference between two measurements. This is shown schematically in Figure 2.1.14.

The construction of a reference phantom for a neonatal head could potentially be incorporated into the fibre holder manufacturing process as described in section 2.7. Using a homogenous reference phantom in this way has been shown to be effective in imaging the adult breast. Since the breast will conform to the conical fibre holder, a conical reference phantom provides a suitable reference. Note that since the reference phantom had wavelength dependent absorption, estimates of its background optical properties were required at both wavelengths in order to obtain multi-wavelength absolute images (see section 2.5.4).

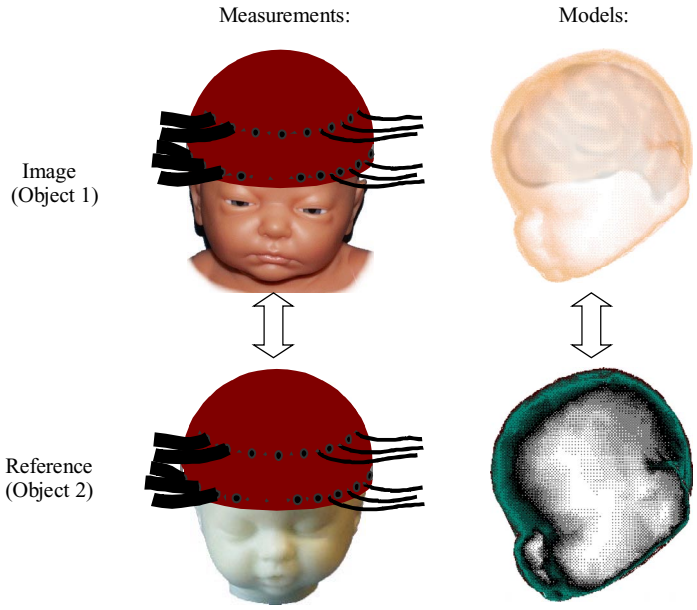

Figure 2.1.14 Using a homogenous phantom as a reference, and accounting for its properties by simulating $D^{Object2}(t)$, helps to eliminate mesh and model errors as well as removing the IRF from measured data. TOAST compares the difference between the measurements to the difference between the models.

If the optical properties of the reference are quite different to those of the object of interest, TOAST may still be able to derive absolute images. The limiting factor is that the

initial conditions must equal the reference object's properties, and if these are too far from the object's the solution may not converge. If this is the case, a simulated set of data representing the reference phantom's properties and geometries could be added on to the reference data sets. The starting parameters of the reconstruction can then be much closer to the estimated properties of the object of interest, although for this to be effective, a very good estimate of the optical properties of the reference phantom is required.

Alternatively, if the reference is a homogenous phantom with circular symmetry, and data is of the form shown in [2.1.19], $D^{Object2}_{n,m}(t)$ and hence $D^{Object1}(t)$ can be derived by using the average (appropriate) datatype for each source-detector spacing *(d)*, along with averaged indirect calibration measurements. The average datatype for a given source-detector spacing on an axially symmetric object is given by:

$$Average_{\,d}[\mathcal{M}[y(t]] = \frac{1}{32}\sum_{a=1}^{32}\mathcal{M}[y^{d,a}(t)] = \frac{1}{32}\sum_{a=1}^{32}\mathcal{M}[D^{d,a}(t)] + \frac{1}{32}\sum_{a=1}^{32}\mathcal{M}[I^{d,a}(t)] \qquad [\,2.1.41\,]$$

Where $y^{d,a}(t)$ is one of the 32 TPSFs acquired with a source-detector pair that is a distance *d* apart (assuming a circular planar arrangement). Since the phantom is cylindrical and homogeneous:

$$\frac{1}{32}\sum_{a=1}^{32}\mathcal{M}[D^{d,a}(t)] = \mathcal{M}[D^{d}(t)]\,, \qquad [\,2.1.42\,]$$

and the average of all $I_{n,m}(t)$ terms is just the average correction for all sources and detectors. So using [2.1.6] and [2.1.17] (for mean-time), we can derive absolute values of $D^{Object1}(t)$ from:

$$\begin{aligned}\mathcal{M}[D^{(1)\,d}(t)] = {} & \mathcal{M}[y^{(1)\,d}(t)] - \mathcal{M}[y^{(2)\,d}(t)] + Average_{\,d}[\mathcal{M}[y^{(2)}(t]] \\ & - \frac{1}{32}\sum_{n=1}^{32}\left[\mathcal{M}[Srccal_{n,1}] + \mathcal{M}[Detcal_{spare,n}]\right] \\ & + \mathcal{M}[Srccal_{1,1}] + \mathcal{M}[Detcal_{spare,1}] - \mathcal{M}[Abscal_{1,1}]\end{aligned} \qquad [\,2.1.43\,]$$

Since the *average* of the indirect calibration measurements are used, all the values of $D^{Object2}_{n,m}(t)$ for a given *d* will be the same so this method would not be prone to the spurious errors likely to affect the calibration methods described in section 2.1.2.1.

Datatypes suitable for difference imaging relative to an arbitrary reference phantom are all those where a simple relation between convolved functions exists (see section 2.1.3.1, note if the relation is multiplicative, the log of the input data should be used). Intensity difference data can be utilised in reconstructions, although problems can occur relating to laser power fluctuations, variations in the time taken for (faulty) secondary shutters on the fibre switch to

open, and difference in coupling between the reference phantom and the object of interest (especially if an arm or breast are being imaged).

2.1.4.4 Difference data for state-difference imaging

We define state-difference imaging as the case where data are acquired on the same object before and after some change has occurred, or is induced, in the object's optical properties. This may be the arm in a relaxed and then tensed state, the breast before and after the injection of a contrast agent, an infant's brain before and during a hyper-capnic swing, or a liquid phantom with moveable inclusions. As with a reference phantom, calibration is not explicitly required, although it is useful to determine a suitable calculation window for datatypes. As with data from a reference phantom, background subtraction and cross-talk compensation should be applied before extraction of raw datatypes from both sets of data.

It is likely that the properties of the two states will be similar, so a linear or non-linear reconstruction can be used. An estimate of the optical properties of one of the states is required to generate PMDFs (for linear reconstruction) or for use as initial conditions in a non-linear reconstruction. Note that the initial condition parameters should subsequently be subtracted from non-linear images if the changes in properties with state-change are required. Estimates of the average optical properties of one of the states could be obtained via data averaging as described in appendix A (2.9.7).

The resulting images should represent the *changes* in absorption and scatter between one state and the other. Note however that this is only an approximation since if one state isn't homogenous or very well characterised the images may be distorted or quantitatively incorrect due to intrinsic limitations of using difference data (unrelated to calibration). A study of this phenomenon is presented in chapter 2.4.

As with using measurements on a reference phantom, most datatypes are suitable for state-difference imaging, including intensity, but with the same limitations as before. Coupling between the object and the fibres should be more consistent here, although movement during a state change may affect coupling and even the shape of the object (e.g. in tensing the arm).

2.1.4.5 Difference data for wavelength-difference imaging

Wavelength-difference imaging involves acquiring data at two different wavelengths. This is useful to evaluate spectral properties of tissue, potentially leading to derivation of functional parameters (see chapter 2.5). The integration of the new dual-wavelength pulse-train laser into the MONSTIR system now makes it very simple to acquire dual-wavelength data-sets simultaneously during a single image acquisition (see section 1.2.3.4). The resulting

data is in the form of two halves of a 25ns acquisition window, the first 12.5ns representing 780nm data and the second half 815nm. Background subtraction and cross-talk compensation are still required. Datatypes can then be extracted, again calibration measurements are useful to determine a good position for the datatype calculation window.

All of the previous types of processing described are entirely applicable to data acquired using the dual wavelength system. However the data needs to be treated carefully if the *difference between the two wavelengths* is required. In the same form as the reference-object data as shown in [2.1.19] we acquire data of the form:

$$y^{\lambda_1}{}_{n,m}(t) = D^{\lambda_1}{}_{n,m}(t) * I^{\lambda_1}{}_{n,m}(t)$$

$$y^{\lambda_2}{}_{n,m}(t) = D^{\lambda_2}{}_{n,m}(t) * I^{\lambda_2}{}_{n,m}(t)$$

[2.1.44]

However state-difference and reference-measurement imaging described in sections 2.1.4.3 and 2.1.4.4 (rightly) made the assumption that the IRF term should cancel between the two measurements. For the case of wavelength-difference imaging, the IRFs of the system are fundamentally different at the two wavelengths:

1) Within the laser, the laser pulses are not exactly 12.5ns apart. An additional optical delay within the unit adds around 280ps onto one wavelength.
2) The intensity of each wavelength output appears to be unstable, particularly when the laser is first turned on, such that the ratio between the two wavelengths is not necessarily constant.
3) The pulse profile of the output from each of the lasers may be slightly different.
4) The fibres and MCP-PMTs have slightly different properties at the two wavelengths.

The delay between the two pulses appears to be fairly easy to characterise, and the properties of the system are thought to have only a small effect on the system IRF, however these aspects are quite difficult to evaluate since a truly wavelength independent object is required to make test-measurements. Measurements are currently being made to evaluate the severity of the intensity fluctuations of the two wavelengths. It may be possible to monitor the intensity of each wavelength real-time within the fibre laser via beam splitters and photodiodes. Another option is to use a dual wavelength calibration measurement (such as with the monstodes) prior to using data to derive wavelength-difference images. The calibration would only need to account for the differing powers of the pulses at each wavelength and their relative delays since all other typical IRF features are likely to cancel since difference data is still being used. However, monstode measurements will potentially introduce an additional wavelength dependent factor from the spectral properties of the surface used for back-reflecting the calibration light.

If correction factors were developed, wavelength-difference data could potentially be used in the following way. It could be reconstructed using linear or non-linear image reconstruction (since changes should be relatively small). An estimate of the properties of the object at one wavelength would be required to generate PMDFs or as initial parameters for TOAST.

Resulting images should represent the difference in absorption and scatter between the two wavelengths (if TOAST images have initial values subtracted). Again however, the limitations of difference imaging are likely to cause images to not be entirely accurate (see chapter 2.4). Clearly the datatypes available for wavelength-difference imaging will depend on the eventual stability achieved with the laser and the stability of the wavelength dependent properties of the system. Since the data are acquired simultaneously, wavelength-difference data would virtually unaffected by problems such as drift and coupling variations.

2.1.5 Summary and discussion

This chapter has introduced a number of aspects of the MONSTIR system performance including noise and systematic error characteristics. Techniques for calibrating measured data and extracting datatypes were introduced. The resulting distinct paradigms for reconstructing images from measured data were then described. While the effectiveness of calibration has been inferred from a study of the noise characteristics of raw data, this chapter has not demonstrated the effectiveness of each method in practice. Chapter 2.2 details an experiment designed to validate absolute calibration methods to determine the inherent limitations of datatype extraction from time-resolved data. The effectiveness of the different imaging paradigms are demonstrated and evaluated on a number of simple phantoms in chapter 2.3 and then more specific aspects of each method are explored in subsequent chapters.

2.2 Evaluation of MONSTIR accuracy using simple fluid phantoms

Chapter 2.1 predicted the likely experimental inaccuracies of the MONSTIR system by exploring the noise characteristics of raw data. This chapter describes an experiment that was designed to test how well calibrated data from MONSTIR matches data simulated using TOAST. A homogenous phantom was constructed and independently characterised such that measurements could be compared with simulated values to determine which datatypes and calibration methods are most likely to allow accurate image reconstruction. Fundamentally these data also allow verification that measurements made using MONSTIR agree with data generated using TOAST.

For very early imaging experiments it was difficult to determine the quality of extracted data. Its 'smoothness' (lack of discrete jumps and spikes) and the patterns of plotted data could be used to test whether calibration was effective in correcting for the differences between channels. However there was no way to test whether the measured values were actually correct or mutually consistent since the properties of the phantoms were not well characterised. It was also not possible to validate that the FEM modelling in TOAST was actually producing datatypes that were realistic when a poor data match could be due to inaccurate properties of the object being modelled (as well as inaccuracies in extracting datatypes from measured data). Early imaging trials used mean-time, variance and Laplace (with $s = 0.005ps^{-1}$) since simulations implied that such a combination of three temporal moments was effective in separating absorption and scatter (Schweiger *et al*, 1997b). However in practice results using these datatypes were often noisy and quantitatively poor. (see section 2.3)

2.2.1 Test solutions

A medium was required that could be well-characterised and constrained to a controlled geometry. A common scattering liquid used for biomedical optics experiments is called Intralipid (Choukeife *et al*, 1999), (van-Staveren *et al*, 1991). It is a clinical nutrient suspension composed mainly of water and soybean oil. It has a similar shelf-life to UHT milk. It comes in 20% and 10% forms. The general assumption is that a 1% solution of intralipid (e.g. 1 part 20% intralipid, 19 parts distilled water) has a μ'_s of 1 mm^{-1} at ~ 800 nm. However, the scattering and absorbing properties and the particle size (and hence anisotropy) of intralipid have been shown to vary with: batch, whether 10% or 20% solution is used, and the

age of the sample (Choukeife *et al,* 1999). Therefore it is unreasonable to assume very accurate knowledge of the scatter properties of a sample only from knowledge of the intralipid concentration. Intralipid must only be diluted using distilled water, since impurities may affect the emulsification properties of the oil and water. Its density is approximately that of water (allowing its mass to be used for accurate volume measurement). Its properties vary slightly with temperature. Its absorption is assumed to be approximately the same as water (0.0022 mm^{-1} at 800 nm).

Intralipid was chosen as our test material since it would conform to the shape of its container, allowing its optical properties to be measured in different geometries. It can also be diluted and have dye added to change the optical properties of the sample. The dye chosen was a sample of water-soluble dye known as ICI S109564 that has good absorption in the near infrared and is not thought to react with intralipid. Its optical properties are not well defined (due to changing concentration of the stock dye with time).

Deriving μ_a and μ'_s of the intralipid mixtures from measurements made by MONSTIR would not independently test the accuracy of the measurements. Hence, a method to independently measure the optical properties of a number of mixtures of dye and diluted intralipid needed to be established. The mixtures then needed to be transferred to a container with a geometry that was suitable for measurement with MONSTIR, which also provided an adequate challenge to TOAST.

2.2.2 Cylindrical hollow phantom design

A cylindrical geometry was chosen for the measurements made with MONSTIR. This was because most of the early validation phantoms imaged using MONSTIR were cylindrical (see section 2.3), and ring-shaped fibre holders were available. The final design of a hollow cylindrical phantom that could hold intralipid solutions is shown in Figure 2.2.1.

The main part of the phantom is a length of black PVC drainpipe (3 inch inner diameter). Plugs were machined to fit the bottom and top of the tube to exclude light. In order for the fibres to be held, the drainpipe had to separate into two halves, sandwiching the fibre holder ring between them. Coupling rings were made to hold the fibre holder ring in line with the two pieces of drainpipe to make a rigid tube. However the fibres must pass through holes in the fibre holder ring. These holes, and the joints between the drainpipe pieces, would leak if the tube were filled with intralipid. Hence a sheath of very fine rubber was made, to line the phantom. Latex was chosen since a sheath can be made by pouring a little latex into the tube, coating the inside, and pouring away the excess. After drying with hot air, a second layer of latex is then created to ensure any small holes due to bubbles are filled. Silicon oil spray is

required to stop the latex from sticking to itself when removed. The sheath created is strong, a perfect mould of the inside of the tube, and very thin.

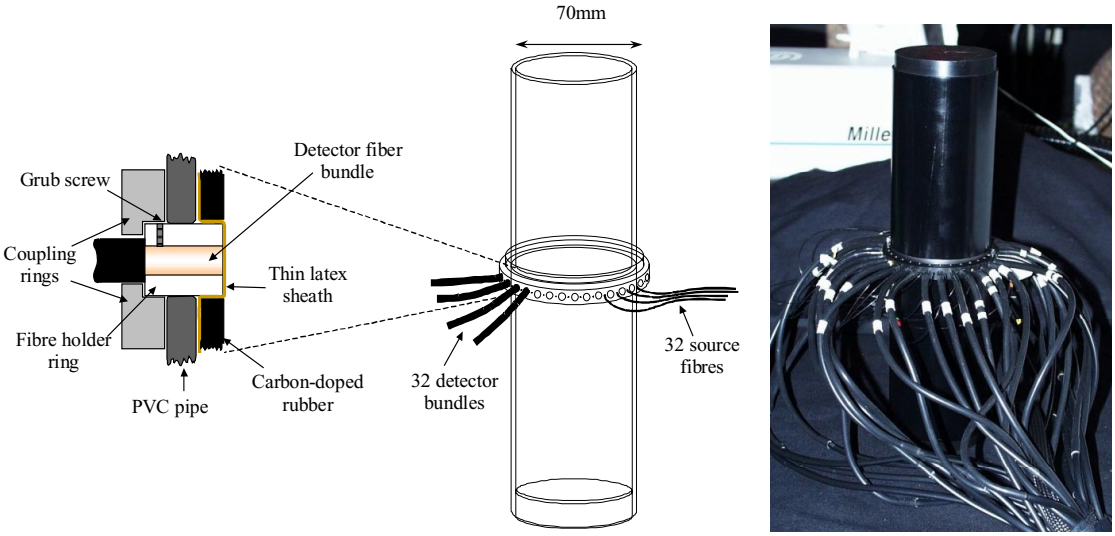

Figure 2.2.1 The 'drainpipe' phantom, used to hold scattering and absorbing solutions for testing MONSTIR.

Although the latex sheath is slightly scattering when wet, it shouldn't attenuate the light significantly. However there is a slight possibility that the sheath may cause reflections and light-piping (see section 2.6). Adding dye to the sheath was tried, but it was feared that the dye would leak into the intralipid solution. A layer of heavy-duty carbon-doped rubber was found to be an excellent near infrared absorber, allowing the simple boundary condition for TOAST that all photons hitting the boundary are absorbed (see section 1.3.1.3.1). The black rubber sheets were cut as rectangles that fitted snugly *inside* the sheath, one above the fibre holder ring and the other below. The sheath was then only in direct contact with the intralipid when over the sources and detectors. This is shown schematically in Figure 2.2.1 (left).

2.2.3 Measuring the Optical properties of solutions

2.2.3.1 Measuring the μ_a of dye

Measuring the μ'_s of solutions is difficult. It can be made easier if we have prior knowledge of the absorption properties of the solution. Also, an estimate of the absorption of the dye solution is necessary for calculating the amount of dye required to create solutions with suitable optical properties. The μ_a of a non-scattering solution can be calculated using the Beer Lambert equation (see section 1.1.1.1):

$$I = I_0 e^{-\mu_a x} \qquad [\ 2.2.1\ ]$$

The experimental set-up used to measure variation in intensity I with distance x is shown schematically in Figure 2.2.2. Collimated white light passes through a glass cylinder and into

the integrating sphere, where it is evenly distributed over the inner surface of the sphere such that a fibre mounted in the wall of the sphere collects a constant fraction of the light. The fibre carries the light to a CCD spectrometer system (see Figure 1.2.6, section 1.2.2.1), which records the spectra of the light at a rate of >1 Hz. A medical infusion pump is used to accurately add dye solution to the tube at a constant and well-defined rate.

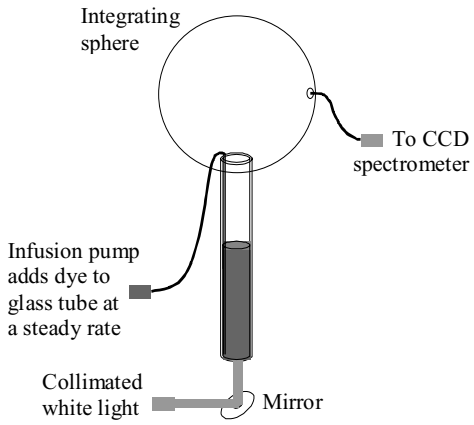

Figure 2.2.2 The method used to measure the absorption coefficient of a non-scattering dye solution. By pumping the dye into a cylindrical tube at a steady rate and acquiring continuously, the variation in intensity with dye thickness is easily derived using the surface area of the tube and the flow rate.

A 1:20 solution of the stock dye was mixed using distilled water. A 1:1000 dilution of this solution was then made and tested in the above set-up (since attenuation through the 1:20 solution was too great to allow accurate measurement).

The area of the cylinder, minus the area of the delivery tube (which leads to the base of the cylinder) was then used to convert time to distance *x:*

$$\frac{1}{x} = \frac{1}{time} \times \frac{surface\ area}{flow} \qquad [\ 2.2.2\ ]$$

Using [2.2.1], we can then derive μ_a of the (1:20000) dye solution using the gradient of the log plot of intensity versus distance. A plot of the data is shown in Figure 2.2.3. Error analysis of the resulting μ_a value is shown on the right (errors were added in quadrature). μ_a was found to be $0.0271 mm^{-1} \pm 0.0005 mm^{-1}$.

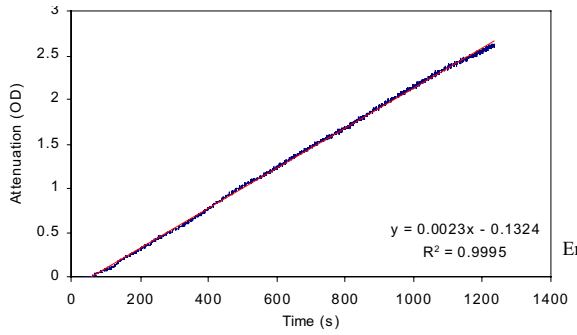

Gradient =	$0.00227 s^{-1}$	$\pm 1 \times 10^{-6} s^{-1}$
Inner diameter of tube =	2.065cm	$\pm 0.01 cm^{-1}$
Diameter of little tube =	1.8mm	$\pm 0.1 mm^{-1}$
Total time =	1232s	± 1s
Surface area =	$3.32 cm^2$	$\pm 0.03 cm^2$
Flow =	$0.0278 ml\ s^{-1}$	$\pm 0.0004 ml\ s^{-1}$
μ_a	$0.0271 mm^{-1}$	$\pm 0.0005 mm^{-1}$

Error on gradient = $\Delta m = \sqrt{\frac{\sigma}{N(x^2 - \bar{x}^2)}}$ σ = standard deviation, N = number of points

Figure 2.2.3 Plot of the attenuation of light with time as dye is added to the tube as shown in Figure 2.2.2. The least-squares fit line is shown, and error analysis performed on the final value of μ_a derived

2.2.3.2 Measuring μ′s of absorbing and scattering solutions

Simple analytic expressions for intensity exist for simple geometries. For a semi-infinite medium, uniformly illuminated with a uniform collimated beam over the whole surface, the intensity within the medium a distance *x* from the surface is given by [2.2.3].

$$I = I_0 e^{-\mu_{eff} x} \qquad [\,2.2.3\,]$$

where μ_{eff} is the effective attenuation coefficient given by:

$$\mu_{eff} = \frac{1}{x}\log\left(\frac{I}{I_0}\right) = \sqrt{3\mu_a(\mu_a + \mu_s')} \qquad [\,2.2.4\,]$$

This geometry has a simple relation between μ_a, μ_s' and distance *x* (Patterson *et al*, 1991). It assumes that an isotropic detector is available, and that it is possible to mimic a medium with infinite extents. This was attempted using the experimental set-up shown in Figure 2.2.4.

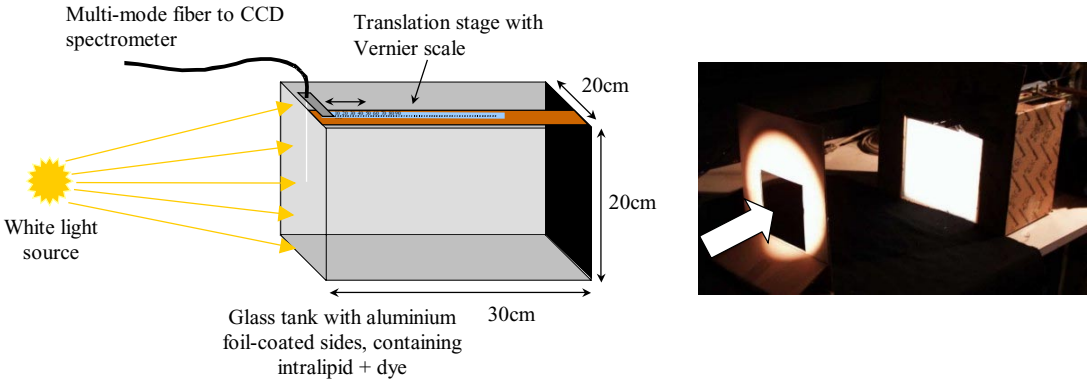

Figure 2.2.4 A glass tank with four walls coated with aluminium foil mimics infinite extents. The front surface is uniformly illuminated using a white light source. The photo on the right shows illumination of the source via a mask to reduce stray light.

Illuminating the front surface with approximately collimated light, and having the sides of the tank silvered, means that light reaching the sides should be reflected back, as if there were an identical medium next to it, as shown in Figure 2.2.5. Although the aluminium foil is not a perfect reflector, the glass itself will ensure total internal reflection over some angles, and the foil does ensure that light leakage (in and out) is minimised.

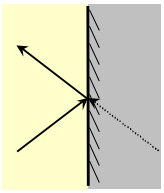

Figure 2.2.5 Silvering the sides of the tank mimics infinite extents.

On top of the tank there is a sliding stage with a Vernier scale. This stage holds a wide multi-mode rigid fibre with its tip suspended half way up the height of the tank. This means

that all boundaries are the same distance away from the fibre, minimising their effect on model accuracy. The model assumes that the fibre collects light isotropically. The fibre used has a high numerical aperture, and is within a scattering media, so since we are only using relative measurements, the limited sample angle of the fibre should cancel.

The fibre carries light to a CCD spectrometer system (see Figure 1.2.6 in section 1.2.2.1). This measures the entire spectrum of the light detected over the near infrared range. Slit width and integration time determines the sensitivity of the CCD spectrometer, so it was ensured that all measurements for a given solution were performed using the same system settings. Measurements were repeated as the translation stage was moved away from the illuminated surface in steps of 5 – 10 mm, depending on the attenuation. Once the photon detection limit had been reached the translation stage was moved back towards the illuminated surface in steps to provide repeat measurements for the solution. Each solution was measured in turn. This technique provided repeated measurements of I/I_0, which were used to derive μ_{eff} and ultimately μ_a and μ'_s of each solution, as described below.

2.2.3.3 Data analysis of μ_{eff} measurement

Values of μ_{eff} were calculated for each of the four solutions by plotting attenuation v/s distance from the illuminated surface and evaluating the gradient of a least squares fit to the data. The plots for sets 1 and 3 are shown in Figure 2.2.6. The fits were only performed over the linear range of the data. When the detector fibre is very close to the illuminated surface, the light may not be truly diffuse. At distances far from the illuminated surface the effects of finite boundaries and the contribution of extraneous room light will affect linearity.

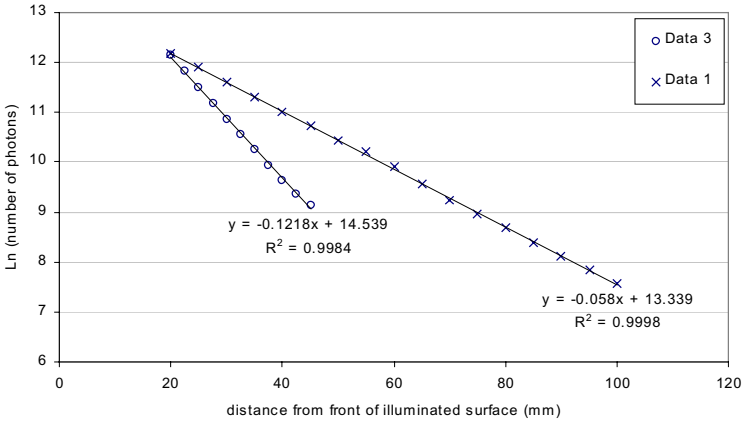

Figure 2.2.6 Plots of ln(number of photons) v/s distance from illumination for two of the solutions measured, and the corresponding least squares fir lines.

If we re-arrange [2.2.4] we get:

$$\mu_a = \frac{1}{2}\left(-\mu'_s + \sqrt{\frac{4\mu_{eff}^{\ 2} + 3\mu'^{\,2}_s}{3}}\right)$$

[2.2.5]

One method for determining μ_a and μ'_s of a liquid involves adding constant fractions of dye to a scattering solution, measuring its μ_{eff} after each addition and then solving for μ'_s and $\Delta\mu_a$ from the measured μ_{eff} values. The values of μ_a and μ'_s chosen for the four solutions of intralipid and dye in this experiment were intended to be suitable for this method.

However analysis of the data implied that the fractions of dye added could not have been entirely constant unless scatter properties varied widely between each solution. Since the concentration of intralipid was thought to be the most stable aspect of the experiment, an alternative method for calculating μ_a and μ'_s was devised.

The major errors in this experiment would be introduced by the fact that once a sample had been mixed and measured, 1.5 litres were removed from the (10 litre) tank and refrigerated for subsequent measurement with MONSTIR. This 1.5L then needed to be replaced in the tank, and then dye added to increase the μ_a of all the liquid in the tank. To improve accuracy of dye addition, a large dilute batch of dye was prepared (from the 1:20 solution, potentially introducing further errors if μ_a is to be calculated from values measured as described in 2.2.3.1). Since adding dilute dye then requires larger volumes of dye to be added to increase μ_a, the main solution will become less scattering. Therefore we need to add a very small amount of extra scatterer with the dye. Because we wish to store liquid that has actually been measured, dye was also added to 1.5L extra of each solution (mixed with intralipid as one batch to start), which was then used to replace the removed solution prior to the next dye addition. This complicated process ensured that each batch that was measured was the exact batch that was subsequently measured using MONSTIR. However the exact amount of dye added is clearly not accurate. Therefore it was decided to calculate the μ_a of each solution assuming no prior knowledge of the dye concentrations added.

Values of μ_a were generated using the four values of μ_{eff} measured from the gradient of log-scale plots of the detected intensity at 800 nm v/s distance for each solution (see Figure 2.2.6) and an estimate of μ'_s. The μ_a values calculated vary by ~ ± 20% when changes are made to the μ'_s value (± 10%). With $\mu'_s = 0.5$ mm^{-1} the μ_a value for the first sample, which was just diluted intralipid, agrees well with the absorption coefficient of water (0.0022 mm^{-1}). However the subsequently derived values of μ_a for the rest of the solutions indicate that the value of μ_a of the dye, calculated via the method described in section 2.2.3.1, is too high, or that large errors were made in measuring the volume of the dye added (as suspected). If we adjust μ'_s to $\mu'_s = 0.4$ mm^{-1}, although the values for μ_a appear more like what we had intended (via dilution of the dye), the μ_a of pure intralipid is overestimated (0.0028 mm^{-1}). Since the first measurement is most accurate (since it was the least attenuating), and since the

absorption of intralipid is likely to be stable, the values found using μ'_s=0.5 mm^{-1} are quoted although error analysis was performed to account for potential dilution errors of the intralipid.

A further possible source of error is that, although we can calculate the theoretical μ_a and μ'_s properties of the solution according to the amount of intralipid and dye mixed, we cannot account for interactions between the dye and the intralipid. On a microscopic scale, the oil droplets in the intralipid may interact with the dye, such that the dye distribution is not homogenous, thus affecting the absorption coefficient. Measurements were made using MONSTIR on the same solutions as were measured here. Hence it is reasonable to assume that the values measured are relevant, and not the values implied by the dye concentrations added etc. The values derived for each solution are shown in Table 2.2.1.

	μ_{eff} measured (mm^{-1})	μ'_s (mm^{-1})	μ_a derived (mm^{-1})
Mix 1	0.057880 ± 0.0001	0.500 ± 0.013	0.00222 ± 0.00001
Mix 2	0.099209 ± 0.0001	0.500 ± 0.013	0.00648 ± 0.00002
Mix 3	0.126906 ± 0.0001	0.500 ± 0.013	0.01052 ± 0.00003
Mix 4	0.146476 ± 0.0001	0.500 ± 0.013	0.01392 ± 0.00004

Table 2.2.1 μ_{eff}, μ_a and μ'_s values derived from independent measurements of the optical properties of three intralipid + dye mixtures.

Errors on μ'_s were derived by adding potential dilution errors in quadrature. Errors on the gradient for μ_{eff} used standard deviation for return measurements and the formula shown in Figure 2.2.3. Since the μ_{eff} errors are small compared to those in μ'_s we can calculate the range of μ_a values that would have been calculated if the μ'_s error were propagated thus:

	μ'_s = 0.5 – 0.0013	μ'_s = 0.5	μ'_s = 0.5 + 0.0013
μ_a (mix 1)	0.00217	0.00222	0.00228
μ_a (mix 2)	0.00632	0.00648	0.00665
μ_a (mix 3)	0.01026	0.001052	0.01078
μ_a (mix 4)	0.01358	0.01392	0.01427

Table 2.2.2 Range of μ_a and μ'_s values derived from independent intralipid measurements (± 1 standard deviation)

These values give us the range of values that should correspond to the true values of the four mixes.

2.2.4 Experimental method for MONSTIR data acquisition

The solutions were prepared and measured as detailed in section 2.2.3.2. The bottles of solution were refrigerated overnight before being measured using MONSTIR the next day.

Following adequate warm up, the MONSTIR system was used to acquire 'absolute' calibration measurements using the tools and methods described in section 2.1.4.1. Each solution was then placed into the 'drainpipe' phantom (Figure 2.2.1). The VOA settings used attempted to regulate the intensity being detected by each channel. For mixes 1 and 2, 26 detectors were active (3 either side of the source were shut off), and the exposure time was 15

seconds per source. For mixes 3 and 4, 24 detectors were active (4 either side shut off), with an acquisition exposure of 15 seconds per source. Full image data sets were acquired (all possible source-detector combinations).

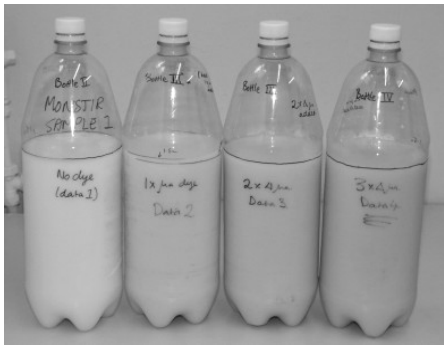

Figure 2.2.7 The four intralipid solutions used to assess the accuracy of MONSTIR (1 to 4 left to right)

2.2.5 Modelling data using TOAST

TOAST was used to model both datatypes and TPSFs for the sets of optical properties shown in Table 2.2.2 using the mesh shown in Figure 2.2.8 which accurately represents the geometry of the drainpipe phantom shown in Figure 2.2.1. The mesh has 124485 nodes and 80576 quadratic elements and it has a waist at the level of the sources and detectors to represent the true geometry of the phantom (generated by Hamid Dehghani). Data were simulated for the same source-detector geometry, but for only one source, for increased speed and because the liquid is homogenous. The following datatypes were simulated: Intensity, mean-time, variance, skew, complex phase (50 MHz and 100 MHz) and Mellin Laplace (u = 0 to 3 inclusive).

Mesh:

```
124485 nodes
80576 quadratic elements
```

forward data simulated using:

```
ALGORITHM = CG <Conjugate gradient>
SOURCE_TYPE = ISOTROPIC
SOURCE_PROFILE = GAUSSIAN
SOURCE_WIDTH = 2
BOUNDARY_CONDITION = ROBIN
CG_PRECON_MODE = INCOMPLETE_CH
CG_ERRORLIMIT = 1e-12
CG_PRECON_FILLIN = 6
```

TPSFs generated using:

```
1000 points, 5ps per point.
```

Figure 2.2.8 The mesh used for forward calculation of datatypes and TPSFs for the liquids detailed in Table 2.2.1 in the phantom shown in Figure 2.2.1.

2.2.6 MONSTIR data:TOAST model comparison

2.2.6.1 TPSF comparison

Figure 2.2.9 shows TPSFs generated using TOAST. TPSFs are shown for two different source-detector separations: source 1 and detector 6 (separation = 36 mm = 31°), and source 1 and detector 16 (separation = 70 mm = 87°). The variation corresponding to the uncertainty in μ'_s is shown by the grey lines either side of each black TPSF (from Table 2.2.2).

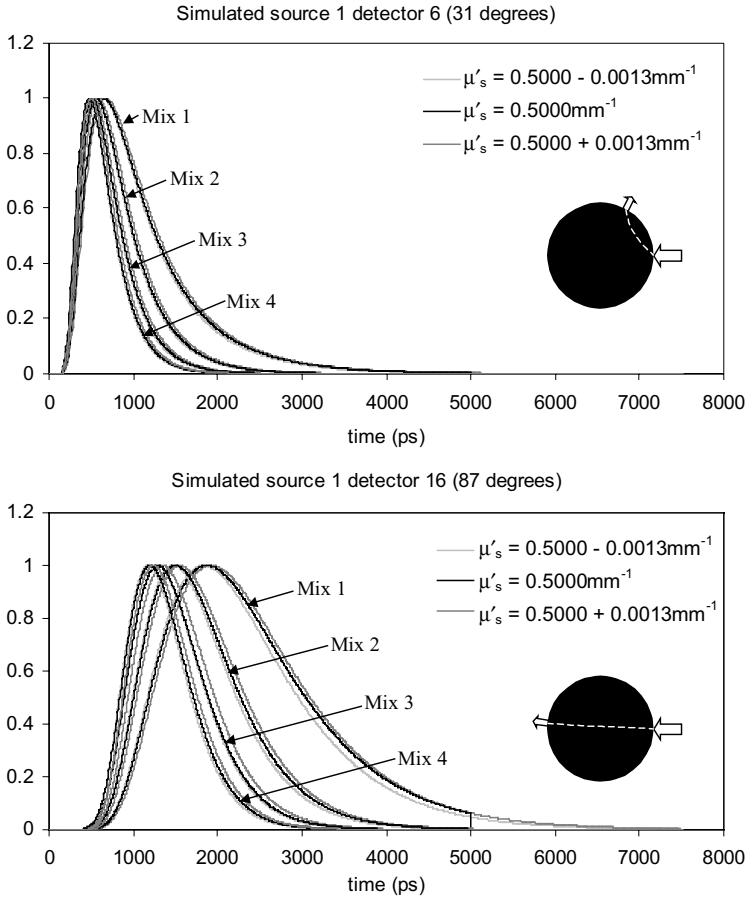

Figure 2.2.9 Simulated TPSFs corresponding to the properties of the intralipid mixes detailed in Table 2.2.2 for two different source-detectors geometries (as indicated by the black circles). Grey lines imply the variation with the uncertainty in μ'_s.

Figure 2.2.10 shows TPSFs measured by MONSTIR superimposed on the simulated TPSFs. The MONSTIR TPSFs are those for the corresponding source-detector pairs for each mix. The measured TPSFs shown are raw and not de-convolved. Therefore we would not expect the TPSFs to match perfectly with the simulated ones. The reflection bump described in section 2.1.1.1 is not really visible in the shape of the tail since the TPSFs are quite narrow (since $\mu'_s = 0.5$ mm^{-1}). However, the measured TPSFs will be slightly broader than the true TPSFs that they represent since the IRF is still affecting the data. The TPSFs have been time-scale calibrated using the absolute calibration measurements detailed in section 2.1.2.1 using [2.1.39]. This equation provides the mean-time of the true TPSF, but since the TPSFs shown

in Figure 2.2.10 have not had the IRF shape de-convolved, mean-time = <true TPSF> + <IRF>. In order for the raw TPSFs peaks to agree with where the peak of the de-convolved TPSF would be, we must shift the time axis by <IRF>-IRF peak. This is the difference between the mean value and the peak value of the absolute calibration tool measurement (section 2.1.1.1) and ≈ 80 ps. The TPSFs have been normalised to a peak value of 1. For the measured TPSFs, the maximum value of a mean-filtered version of the TPSF was used to determine the normalisation to avoid effects from stochastic noise.

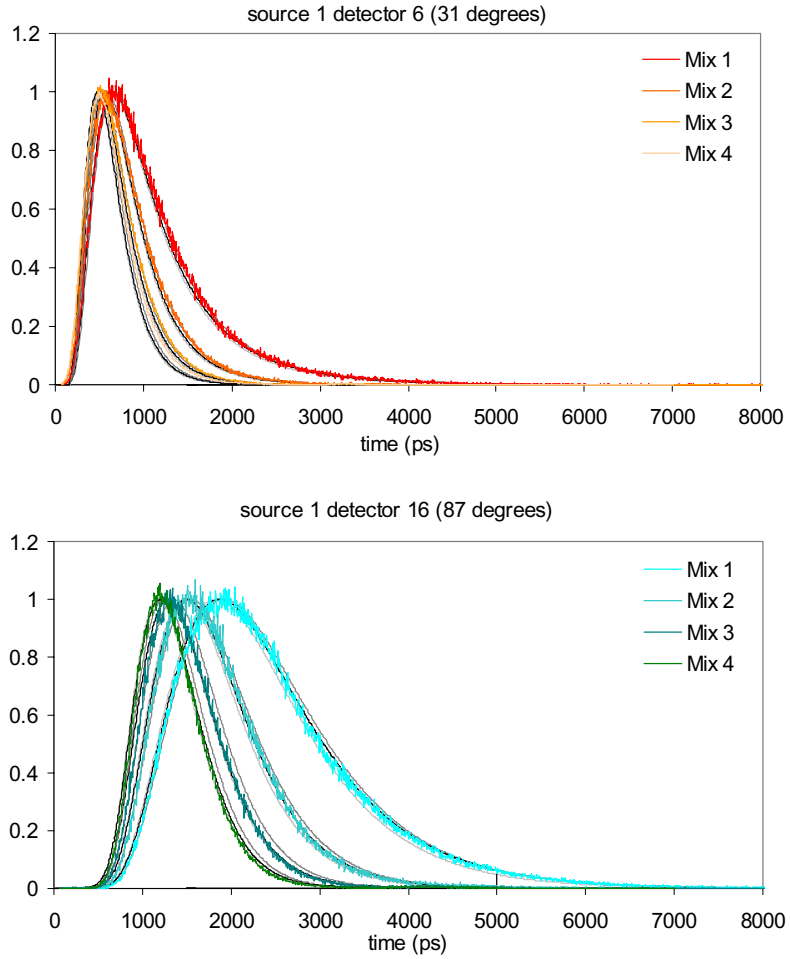

Figure 2.2.10 TPSFs measured using MONSTIR compared with those simulated using TOAST for liquids of four different absorptions, at two different source-detector separations.

These raw TPSFs are shown for simple qualitative comparison with the simulated ones. In order to quantitatively analyse the match between the model and measured data we need to compare the extracted calibrated datatypes (which is where a good match is of prime importance for successful image reconstruction).

2.2.6.2 Datatype comparison

Datatypes were extracted from MONSTIR TPSFs and calibrated using the absolute calibration measurements described in 2.1.2.1 and the methods detailed in 2.1.4.1.

deconvolution equivalent methods (see section 2.1.3.1) were used to extract mean-time, variance, skew and complex phase (at 50 and 100MHz). For Laplace and Mellin Laplace (s = 0.005 ps^{-1}), datatypes were calculated from TPSFs that had been de-convolved from the IRF using [2.1.6]. Dynamic windowing and background subtraction were also performed. The mean-average datatype for each source-detector separation was then calculated as well as the standard deviation. For 32 sources and detectors in a symmetric geometry this corresponded to taking the mean of 64 measurements for each source-detector spacing. Corresponding datatypes were simulated using TOAST as described in section 2.2.5 .

Figure 2.2.11 shows measured and simulated absolute mean-time, variance and skew for all source-detector combinations for mix 1 and mix 4 (see Table 2.2.2). The grey lines either side of each black line indicate the variation in the simulated datatype corresponding to the uncertainty in μ'_s and μ_a of the solutions measured. Error bars show the standard deviation n the mean of each datatype for each source-detector separation.

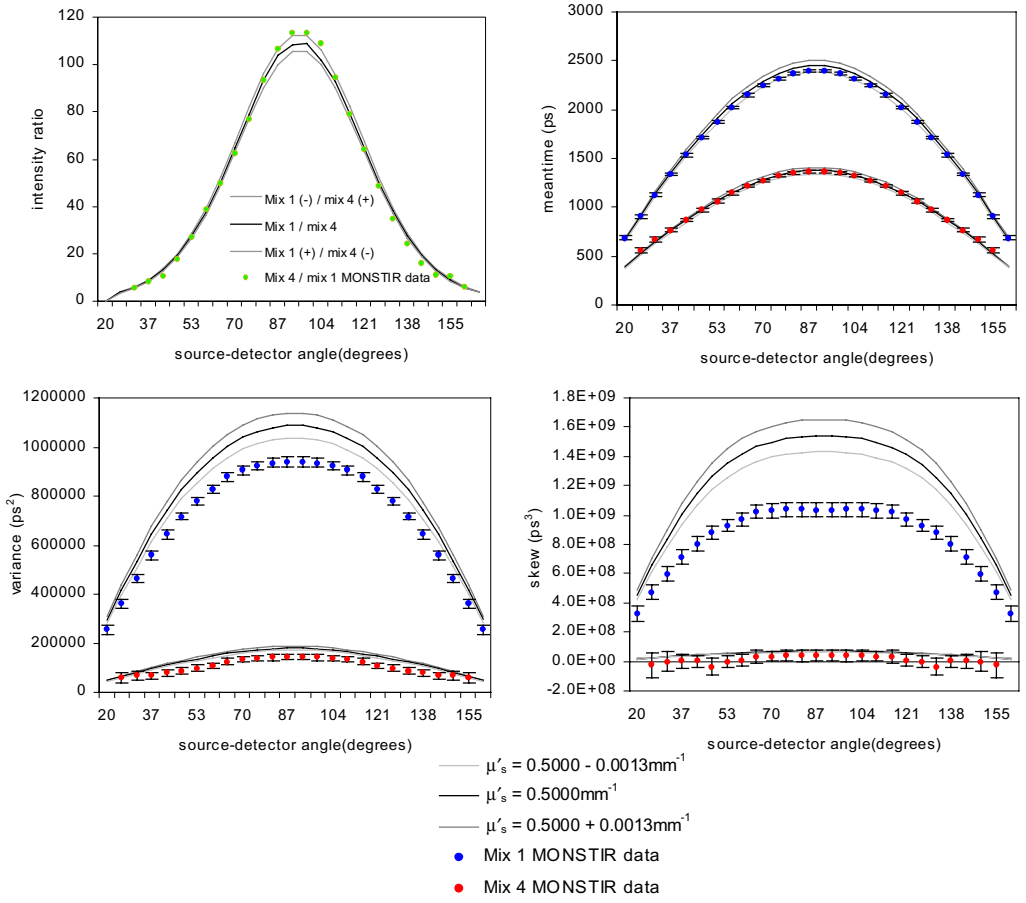

Figure 2.2.11 Comparison between simulated datatypes and those extracted from MONSTIR TPSFs using calibration procedures described above, for mixes 1 and 4.

Due to intensity calibration problems as described in section 2.1.2.4, we cannot compare *absolute* intensity values with those simulated. However the ratio between intensities for mix 1 and mix 4 will self-calibrate (2.1.2.3) since the same VOA settings were used for both acquisitions. Note that an ND filter correction was also applied to the intensity ratio.

Figure 2.2.12 shows plots of phase simulated using TOAST and evaluated from measured TPSFs (using [2.1.35]), for 100 MHz and 50 MHz for mixes 1 and 4. Complex amplitude (not shown) would yield very similar results to integrated intensity (shown in Figure 2.2.11).

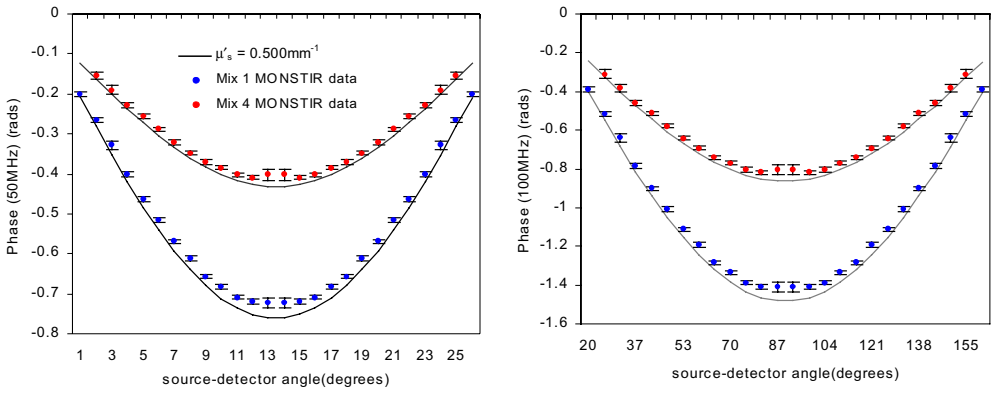

Figure 2.2.12 Phase extracted from MONSTIR TPSFs compared to phase from TOAST for mixes 1 and 4, for 100 MHz and 50 MHz

Following full de-convolution of MONSTIR TPSFs and the calibration IRF derived using [2.1.6], Mellin Laplace datatypes for u = 0 to 3, s = 0.005 ps^{-1} were extracted from data acquired on mixes 1 and 4. These are compared to the same datatypes derived using TOAST in Figure 2.2.13.

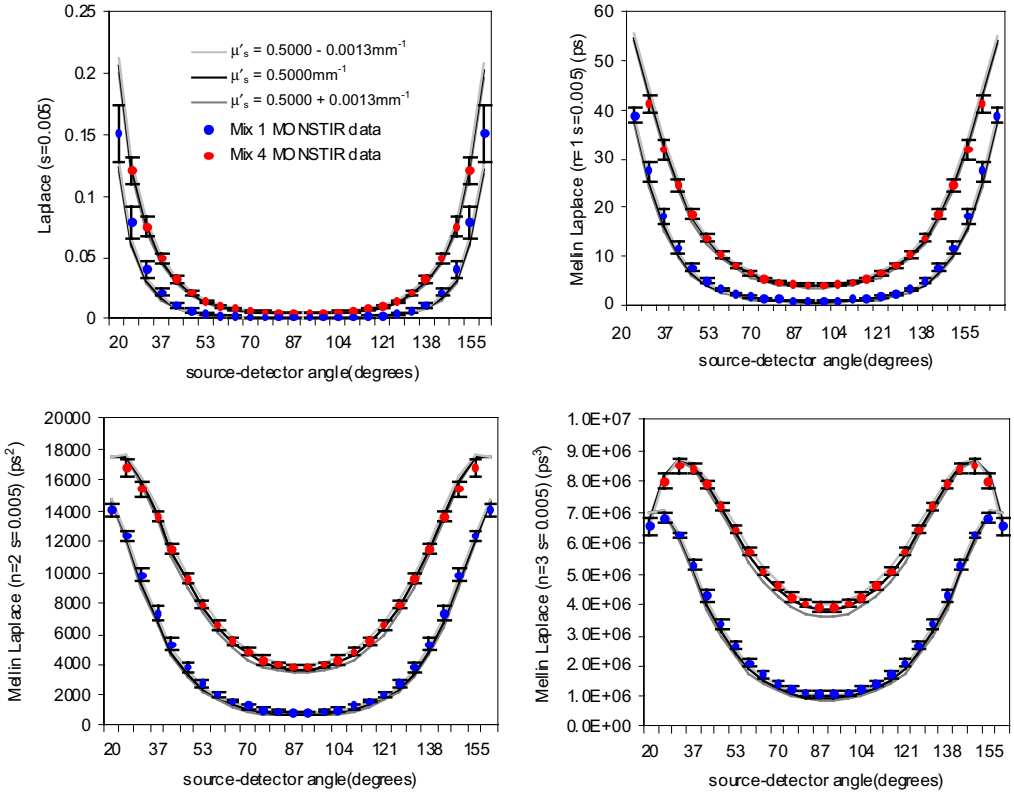

Figure 2.2.13 Mellin Laplace datatypes extracted from MONSTIR TPSFs that have had the system IRF fully de-convolved, compared to simulated values from TOAST (n=u)

We can quantitatively compare these results by evaluating an error norm based on the difference between the simulated values and the MONSTIR datatypes. Although the error norm used as the measure of convergence in TOAST is usually referred to in the form of [1.3.40], in practice the measurement itself is often used as the standard deviation (since multiple measurements are not available in a conventional image data-set). If we assume that the values of μ_a and μ'_s for the solutions measured are correct, then an image reconstruction of the data extracted from the measured TPSFs would be minimising an error norm of the form:

$$\sqrt{\frac{1}{N}\sum_{i}^{N}\left(\frac{M_i - \mathrm{M}_i}{\mathrm{M}_i}\right)^2}$$

[2.2.6]

where $\mathbf{M}_i$ is a MONSTIR measurement and M_i is the corresponding datatype calculated by the TOAST forward model. If the match between the data and the simulation were perfect this error norm would be zero. Figure 2.2.14 shows values of this error norm for each of the datatypes shown above.

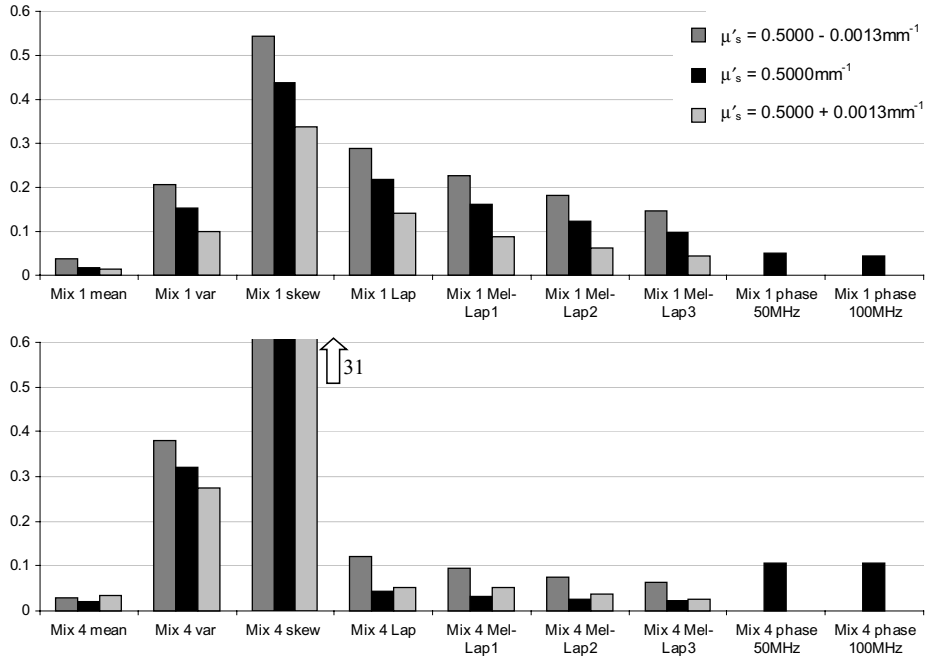

Figure 2.2.14 Error norms for each of the datatypes shown in Figure 2.2.11, Figure 2.2.12 and Figure 2.2.13

It is possible that the mismatch between measured datatypes and simulated ones is due to inaccurate evaluation of optical properties. However the plots above demonstrate that a good match has been made between calibrated mean-time data and simulated data. Similarly phase and the higher Mellin-Laplace orders indicate quite a good fit especially for mix 4. This supports the assumption that the optical properties of the solutions are indeed a good estimate. However error norms are very large for skew, and large for variance and lower orders of Mellin-Laplace.

These results are significant since if a datatype with a large inherent error were to be used in a reconstruction (particularly in combination with other datatypes), the solution would not be able to converge correctly. Further, if there were structure in the object being imaged, its contribution to the error norm would only be small and would be overwhelmed by the magnitude of the mismatch between the absolute values of the datatypes.

The grey lines either side of the black line on the plots in Figure 2.2.11, Figure 2.2.12 and Figure 2.2.13 demonstrate the sensitivity of each datatype to small changes in μ_a and μ'_s. The intrinsic data information content of each datatype is discussed briefly in section 1.3.3.6. We can see from these plots that in many cases datatypes that have a large error norm are also those with a high sensitivity to changes in optical properties (e.g. skew and variance). This means that the signal representing structure would in fact be larger within these measurements so convergence may not need to be as accurate as for less sensitive datatypes.

In general, for mix 1, the fact that all datatypes have a lower error norm for the highest value of μ'_s implies that mix 1 is likely to have had properties more similar to those corresponding to a slightly higher scatter. In mix 4 the error norms seem to be smallest for $\mu'_s = 0.5\text{mm}^{-1}$. The origin of the datatype specific mismatches and the implications for each datatype are discussed briefly below.

2.2.6.2.1 Moments

The fit between the simulated and measured intensity ratio shown in Figure 2.2.11 supports the findings earlier that systemic errors in intensity will cancel if the ratio of two similar data-sets is used. For mean-time the fit between simulated and measured data is very good implying that it is a very good datatype in terms of ease of calibration, although its sensitivity to μ_a and μ'_s is not as good as for variance and skew. However the large error norms for variance and skew cannot be explained purely by their sensitivity to optical properties. This can be seen if we look at the mix 4 data in more detail as shown in Figure 2.2.15.

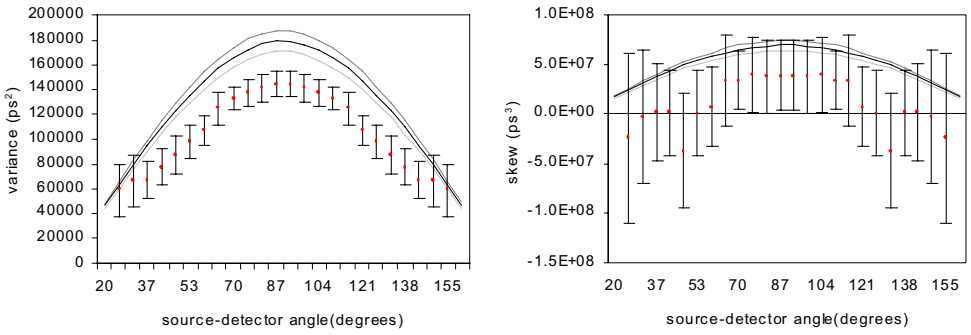

Figure 2.2.15 Simulated and measured datatype comparison for variance and skew for mix 4 only. (key as for Figure 2.2.11)

The standard deviations on the skew and variance data for mix 4 are very large and the profiles of the datatypes are distorted. Some of the values derived for skew are even negative.

We can explore the sensitivities of moments to types of noise by looking at plots of the windows over which effectively we are sampling the TPSF y(t) (from [2.1.24]). Figure 2.2.16 shows that extraction of intensity is achieved by simply summing the number of photons in a TPSF. Calculating the mean represents multiplying each y(t) value by *t* and summing all of these products. Similarly in extracting variance we are multiplying the TPSF with a parabola, centred on the mean. Skew represents weighting the TPSF by a cubic function.

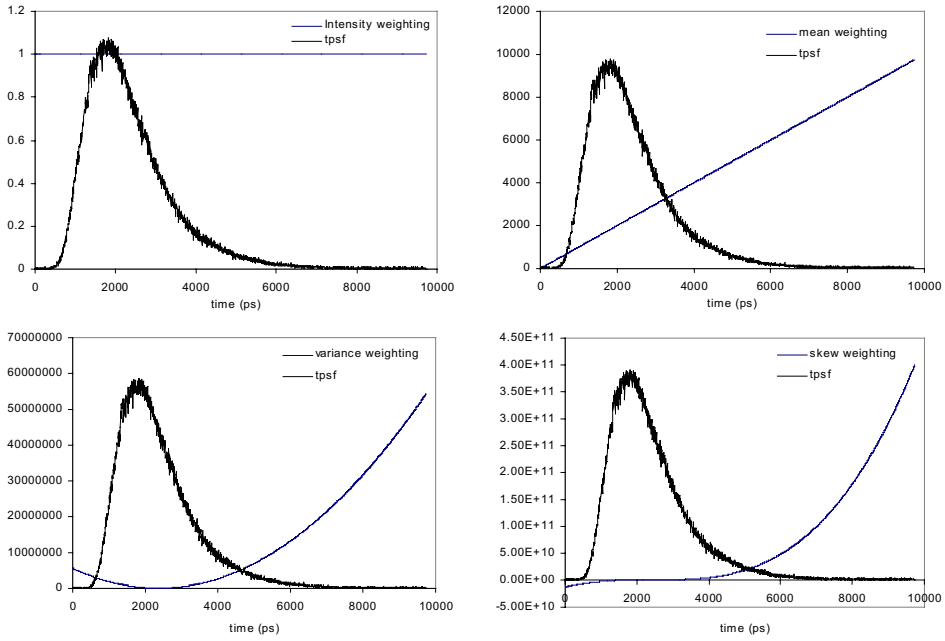

Figure 2.2.16 Plots showing the weighting applied to different parts of the TPSF during extraction of intensity, mean, variance & skew.

These weightings result in mean-time giving us general information about the TPSF position in time. Variance provides a measure of the width of the TPSF about its mean, and skew provides a measure of the asymmetry of the TPSF about its mean. However, Figure 2.2.16 also demonstrates that, as (central) moment order increases, increasing weighting is given to the beginning and tail end of the TPSF, where signal to noise is worst. As a consequence, naturally we would expect variance, and even more-so skew to be highly dependent on the level of stochastic noise on the TPSF.

It is also clear that the finite sample window of a MONSTIR TPSF will heavily influence the higher moments (as for intensity and mean also). The effect of a finite sample window can be demonstrated by calculating mean-time and variance from the TPSFs generated by TOAST (as shown in Figure 2.2.9). These have a finite sample window of 5000 ps for the whole TPSF, and discrete sampling of 5 ps, yet they have used exactly the same numerical model as

that which derived the datatypes shown in Figure 2.2.11. The datatypes extracted from the MONSTIR data were calculated from 5000 ps *after* the mean-time of each TPSF as described in the dynamic windowing section (2.1.3.2.3). Figure 2.2.17 shows the datatypes extracted from the MONSTIR TPSFs, the simulated TOAST datatypes, and datatypes calculated from TPSFs generated by TOAST for mixes 1, 2 and 4. (Only simulated values for $\mu'_s = 0.5\text{mm}^{-1}$ are shown).

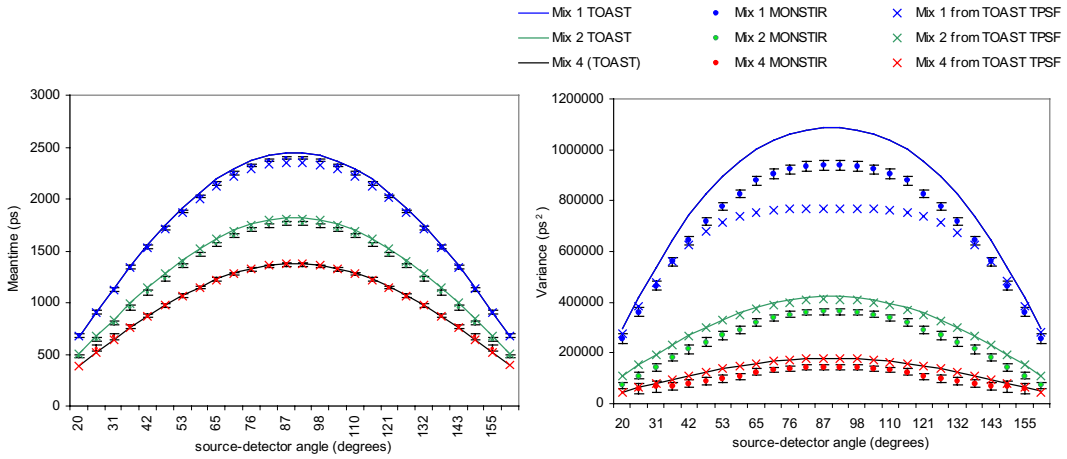

Figure 2.2.17 Comparison between datatypes derived by TOAST, datatypes extracted from MONSTIR TPSFs and datatypes extracted from TPSFs generated by TOAST to demonstrate the effect of a finite acquisition window on datatype accuracy.

Where the TPSFs are broadest (mix 1 and source-detector angle ~90°) mean-times calculated from the truncated simulated TPSFs (crosses) are lower than those derived by the full forward model (line). The mean-times from MONSTIR (dots) are between the two. For mixes 2 and 4 the mean-time from TOAST TPSFs are very close to the TOAST datatypes implying that, for these narrower TPSFs (see Figure 2.2.9) that tail truncation is having little effect. The MONSTIR TPSFs are not as truncated as the TOAST TPSFs which may explain why the mix 1 MONSTIR mean-times are closer to the simulated values, although errors in the mean-times from MONSTIR are also likely to be due to mismatches in the optical property estimates.

For variance, the truncation of the TOAST TPSFs has had a dramatic effect on the calculated values for the far source-detector separations for mix 1. As with mean-time the mix 1 MONSTIR variances are between the correct and truncated simulated values, since MONSTIR truncation is not as severe. However for the narrower TPSFs for mixes 2 and 4, the variances calculated from the truncated TPSFs are very close to those calculated using the full forward model while the MONSTIR variances quite a lot lower. The error norm for variance and skew was also a lot larger for mix 4 compared to mix 1. Figure 2.2.15 showed that this mismatch is unlikely to be due to optical property mismatches, particularly since the difference between simulated and measured values is more of an offset, while an optical

property mismatch would be produce an error proportional to source-detector spacing (see [2.9.69] (appendix A)). Note also that having used [2.1.40] to calibrate variance, the correction factors derived are constant for all mixes. This indicates that, where the TPSF is narrow, effects relating to stochastic noise and cross talk at large and small times are likely to be affecting the applicability of the calibration factors. Where TPSFs themselves are quite narrow, the effects of a non-zero background begin to overwhelm the values of variance and skew.

2.2.6.2.2 Phase

The match between simulated and measured phase was shown to be quite good with the error norm for mix 1 being lower than for mix 4 (Figure 2.2.12). The method used to extract phase (given by [2.1.35]) is not optimal; an FFT could have been used that included better filtering and windowing. We can explore the sensitivity of complex datatypes again by looking at the weightings applied to parts of the TPSF when the datatypes are calculated. Figure 2.2.18 shows the sensitivity of real and imaginary part of the Fourier transform at 50 MHz and 100 MHz.

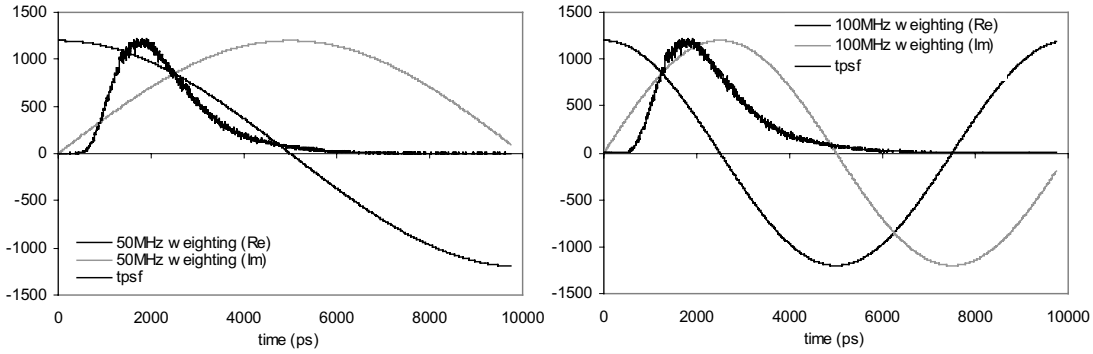

Figure 2.2.18 Plots showing the weighting applied to different parts of the TPSF during extraction of complex datatypes (Real and imaginary)

Unlike variance and skew, the extraction of complex datatypes does not weight the larger or smaller times any more than the main TPSF although clearly different frequencies will emphasise features at different positions on the TPSF. We have not seen a marked difference in the error on phase with frequency but the two frequencies chosen are quite similar. If higher or lower frequencies are to be considered in the future this analysis should be extended to explore whether errors in extraction worsen with increasing frequency (as is the case with frequency domain systems (Boas, 1996).

2.2.6.2.3 Mellin Laplace

The Mellin-Laplace datatypes extracted from MONSTIR TPSFs have fairly large error norms, decreasing with increasing order u. However it can be seen from the plots of the datatypes (Figure 2.2.13) that there is an almost exponential dependence of Laplace on

distance. Indeed we found in section 2.1.3.1 that Laplace, like intensity has a mutiplicative relation with convolved functions. This means that TOAST should really work with the Ln(Laplace) datatype rather than Laplace. If we calculate the error-norm of the log of the Laplace data we get a much smaller value. Figure 2.2.19 shows the Laplace plots on a log scale and the corresponding error norms for mix 1 and 4 if the log of Laplace is used. Ln(Laplace) looks much like phase, but has a lower error norm for mix 4 than mix 1. This is the opposite of the moments.

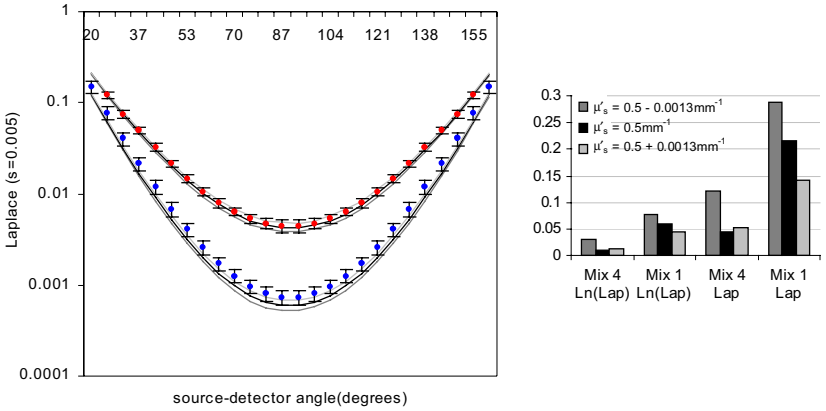

Figure 2.2.19 Laplace data on a log scale and the corresponding error norms for Laplace and Ln(Laplace) (compare with Figure 2.2.14)

Figure 2.2.20 shows the weighting applied to a TPSF when Mellin Laplace are calculated.

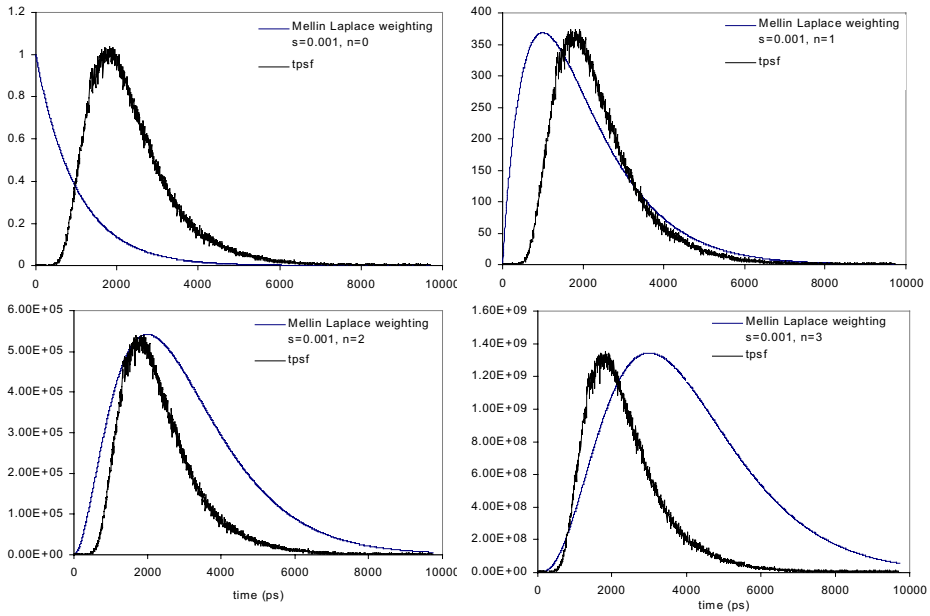

Figure 2.2.20 Plots showing the weighting applied to different parts of the TPSF during extraction of Mellin Laplace (n=0-2, s=0.001).

An interesting feature is the exponential decay (given by s) of the weighting at lower and higher times (except for Laplace which is highest at t=0). This may explain why the fit to simulated data is better for the narrower TPSFs of mix 4 than the wider TPSFs of mix1. The

signal contributed from the TPSF is large compared to the contribution from noise (the opposite of variance). However the grey lines in Figure 2.2.13 demonstrate that none of the Mellin_laplace orders are as sensitive to μ_a and μ'_s changes as variance or skew.

The shape of the Mellin Laplace weighting also depends on the chosen value of s. Figure 2.2.21 shows the change in the shape of the weighting with increasing s (for u=2), this shape change is very similar to that seen with changing u in Figure 2.2.20.

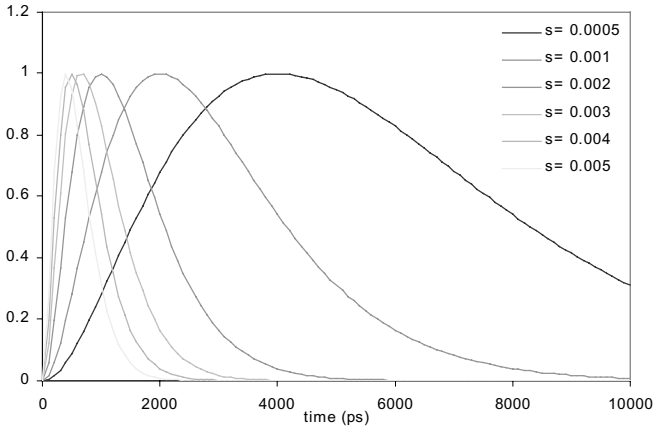

Figure 2.2.21 Plot showing the weighting applied to different parts of the TPSF during extraction of Mellin Laplace for u=2 and s=0.0005-0.005).

While in general the error norms for mix 4 were comparable with mean-time, they were high for mix 1. While the exponential decay of the weighting explains why Mellin Laplace is more tolerant of the effects that overwhelm variance and skew, it is not clear whether they represent a robust datatype. Further, the different Mellin-Laplace datatypes seem quite similar in their μ_a and μ'_s dependence (as shown in Figure 2.2.13 for n=0 to 3, this was also implied by results shown in appendix A (Figure 2.9.4 and Figure 2.9.8). The weighting changes with changing s are similar to those with changing u implying that radically different dependencies won't result from varying u and s. Further, the relations between datatypes of convolved functions derived earlier in section 2.1.3.1 ([2.1.36]) show that calibration of Mellin Laplace is complicated by lengthy, and potentially inaccurate de-convolution procedures. An approximate method of calibrating Laplace and Mellin Laplace by accounting for the temporal delay of the system's IRF without accounting for its breadth is explored in (Hillman, 2001a). This approximation yields Mellin-Laplace values that agree fairly well with simulated results. However this implies that Mellin-Laplace is not very sensitive to the shape of the system's IRF, and may hence not be sensitive to subtle changes in TPSFs relating to μ_a and μ'_s structure.

2.2.7 Summary and discussion

This chapter detailed an experiment that was designed to evaluate the quality of extracted and calibrated measured data, in addition to providing validation for the TOAST forward model (note that this experiment did not utilise or test the iterative image *reconstruction* part of TOAST).

Four liquid phantoms were measured using an independent method to characterise their optical properties. Each phantom was then measured using MONSTIR and the resulting calibrated and extracted absolute data were compared with simulated values based on the nominal optical properties derived from the independent method. It was concluded that the match between measured and simulated data depended on:

- The accuracy of the optical property estimate for a given solution.
- The accuracy of the model (e.g. the dependence of derived datatypes on meshing errors and approximations)
- Which datatype is chosen and its tolerance to the different types of noise on a measured TPSF and the effectiveness of calibration.

In general skew and variance demonstrated poor agreement with simulated data, particularly with narrow TPSFs. This could be attributed to:

- Their strong sensitivity to μ_a and μ'_s changes
- Their increased weighting at high and low times and hence their sensitivity to background noise and truncation of the tail of the TPSF.

Future work could investigate the true effectiveness of replacing noise at the end of a TPSF by extrapolating the tail and then calculating variance and skew. Phase was found to have reasonable agreement between measured and simulated values, although techniques for extracting phase from TPSFs could undoubtedly be improved. If different frequencies than those investigated are to be used (50MHz and 100MHz) similar evaluation of noise tolerance would be required since it is likely to be frequency dependent.

Preliminary studies of the information content of the various datatypes were described in section 1.3.3.6. From the results for Mellin-Laplace it is clear that more work is required to properly evaluate this datatype.

As a result of findings described in this chapter, most subsequent image reconstructions from experimental data were performed using mean-time absolute data and, where applicable, intensity and mean-time changes only. The effect of using a limited set of datatypes is explored in section 2.5.6.1. Where simulations are presented within the following chapters, in most cases mean-time and intensity are used as examples. Early images reconstructed using variance and Laplace are presented in chapter 2.3.

2.2.7.1 Possible future alternatives to datatypes

Much work is still required to determine whether the full potential of time-resolved imaging can be realised via extraction of datatypes that are both robust to noise and contain sufficient information to describe the internal properties of an object. It is possible that a change in approach is required to make TOAST more tolerant of the inevitable errors in input data. Possible alternatives to extracting datatypes in the current fashion would be:

- To use only the equivalent of Fourier domain data (exploring optimal combinations of frequencies)
- To include better estimates of the standard deviation of measurements in the image-reconstruction error norm (rather than using the data). This would allow a positive weighting to be applied to data that is thought to be more accurate. These standard deviations could be calculated from repeated measurements or determined from a model based on the characteristics of individual sources, detectors, and the number of photons in each measured TPSF.
- To perform reconstructions based on the entire TPSF (or a number of time-gates) rather than extracting datatypes. (Gao *et al*, 2002) reports improvements in images reconstructed using this method although far more input data values are used making the reconstruction less ill-posed but far more computationally intensive.
- It may be possible to derive datatypes via the TOAST forward model that are in fact parameters of a fit to the TPSF. Although moments etc. are shown to represent coefficients in a polynomial fit to the frequency domain form of a TPSF (see appendix A (2.9.4)), these datatypes could be extracted by fitting a function to a selected portion of the measured TPSF. This may be more robust than needing to calculate transforms based on the TPSF extending to infinite times, as required currently.
- There may be other integral transforms of a TPSF that could be derived fairly simply via an FEM forward model that would provide complementary information to existing datatypes.

The future development of datatypes will depend on how the primary focus of optical tomography methods evolves. If difference-imaging techniques are determined to be sufficient for clinical applications, calibration requirements will be less severe. Continuing development of instrumentation will also govern the ultimate limitations of calibration and noise compensation.

2.3 2D and 3D image reconstruction of phantom data

Once the construction of the MONSTIR system was complete, a series of imaging trials were performed. Cylindrical resin phantoms were made using a recipe described in (Firbank *et al*, 1995a). The phantoms consisted of epoxy resin (Araldyte), white paint (composed of titanium dioxide particles) and NIR absorbing dyes. Optical properties were determined from the concentrations of dye and scatterer added to the mixture. Most phantoms were machined into cylinders 70mm in diameter. Each phantom had a number of inclusions. To begin with these were rods, and later discrete cylindrical inclusions were embedded. The inclusions had scatter or absorption (or both) properties of various multiples of the background properties of the phantoms. The purpose of these phantoms was to demonstrate the ability of the MONSTIR system and TOAST to reconstruct the correct absorption and scatter distributions from measurements.

The first images from the MONSTIR system were produced from measurements made on a cylindrical phantom with three rods embedded. One had five times the background absorption, one had five times the background scatter and the third had twice both background scatter and absorption. These images were reconstructed from measurements made using only 8 detectors and two sources arranged within a single plane. A full image data set was acquired by twisting the phantom within a ring holding the sources and detectors. The images were reconstructed on a 2D mesh and used mean-time and variance datatypes. It was possible to distinguish between the absorbing and scattering structure (Hebden *et al*, 1999).

While the images produced in this way showed structure corresponding to the locations of the rods in the phantom, quantitation and model convergence was poor. Investigations were undertaken to explore the validity of the variety of assumptions made in reconstructing these early images. Investigations into the quality of calibration of the chosen datatypes were presented in chapter 2.2. In this chapter we investigate the use of a 2D model to reconstruct measured data since this was suspected to introduce errors. While a 3D model has now been implemented in TOAST, the findings detailed are relevant since few researchers within the field of optical tomography use true 3D models for image reconstructions.

Some examples of early validation imaging experiments are also presented where reconstructions were performed using both 2D and 3D models. Both absolute data and difference data (where a homogenous phantom provides a reference measurement) were used for reconstructions. Performance is compared in each case. Analysis of the effects of imaging an axially symmetric phantom (e.g. with rods rather than discrete inclusions) is also

presented. Note that a full chronological catalogue of the images produced by MONSTIR are presented elsewhere in (Hillman, 2001a).

2.3.1 Reconstructing real data using a 2D mesh

Development of the TOAST reconstruction software naturally began with the optimisation of 2D meshes and reconstruction methods, prior to development of the 3D methods now available. The majority of the early experiments conducted with MONSTIR utilised sources and detectors confined to a single plane (as shown in Figure 2.7.1 – section 2.7.1). As a consequence it was tempting to suppose that reconstruction could be achieved using an algorithm that is likewise limited to 2D.

However, during any real experiment, necessarily photons are free to migrate in all three dimensions within a scattering medium, so measurements made at the surface will inevitably be affected by photons that have travelled through structure located above and below the plane of interest. The data plots shown in Figure 2.3.1 illustrate this.

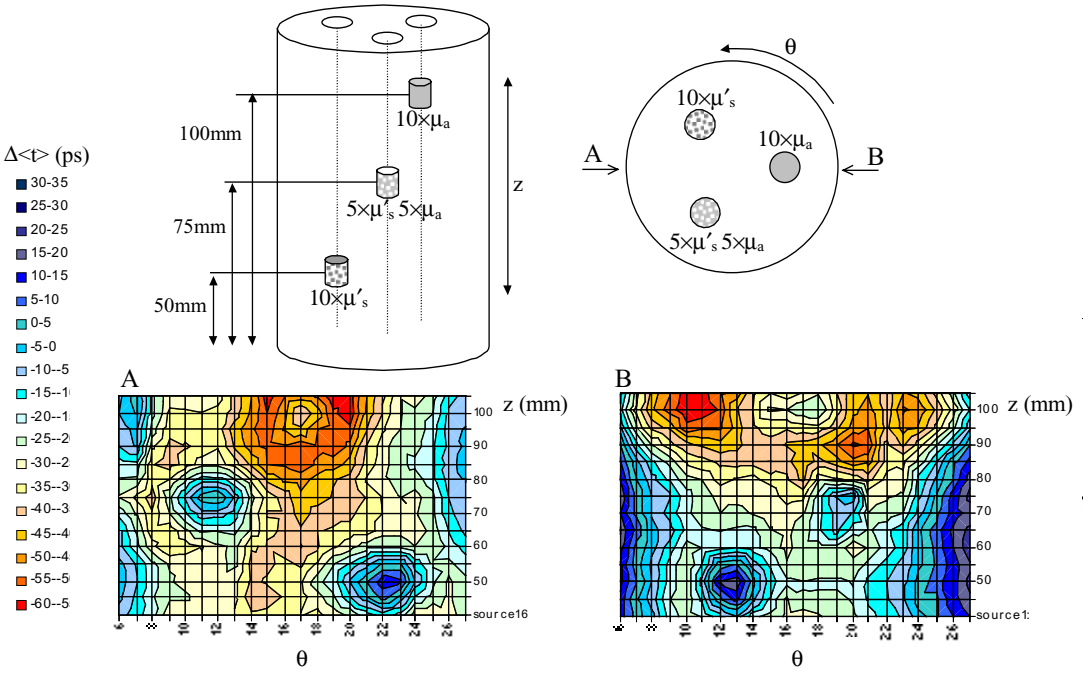

Figure 2.3.1 Plots of the difference in mean-time between the 14 planar acquisitions on the 3D Multi-level phantom, and a single slice acquisition on the homogenous phantom. The orientation of the phantom for each plot A, B is shown. The features of the phantom are easily identified from these data alone.

Mean-time data were acquired on a 14cm high cylindrical phantom (with background μ_a = 0.01mm^{-1} ± 0.001mm^{-1} and μ'_s = 1mm^{-1} ± 0.2mm^{-1}) using 32 sources and detectors interspersed in a planar ring. 14 separate image data sets were acquired by translating the fibre holder up the phantom in 5mm steps. The phantom had three discrete cylindrical inclusions (each 10mm high and 10mm in diameter) at different heights. The top inclusion had 10xμ_a, the middle inclusion, 5xμ_a and 5xμ'_s, and the lower inclusion 10xμ'_s. In addition a single data

set was acquired on a homogenous reference phantom. Plots show the detected mean-time difference between the structured and homogenous phantoms for each opposing detector illuminated by (A) source 1 and (B) source 16 (in each of 14 vertical positions). Out-of-plane structure will obviously affect planar images reconstructed in 2D since data acquired in planes without anomalies still show evidence of the perturbations above or below (e.g. z=65mm).

However out-of-plane structure is not the only factor that affects the suitability of using a 2D model to reconstruct measured data. In fact, regardless of the effects of out-of plane structure, the statistical distribution of photon flight times between two points for 2D and 3D models are sufficiently dissimilar to produce a significant difference in the forward model predictions of the datatypes. This is illustrated in Figure 2.3.2, which shows TPSFs generated using 2D and 3D FEM models for a homogeneous disc and a cylinder of equal diameter (70 mm). Both models employ the same reduced scattering coefficient ($\mu_s' = 1.0\ mm^{-1}$), absorption coefficient ($\mu_a = 0.01\ mm^{-1}$), and source-detector separation (d = 70 mm). The differences between the TPSFs suggest that datatypes derived from a 2D FEM model are likely to be significantly different from those acquired experimentally. Similar differences are also observed between datatypes and TPSFs derived using analytic solutions for more general geometries. The Green's function solutions for infinite space yield different values if the solution is sought for 3D space rather than 2D space (see section 1.3.1.2.1).

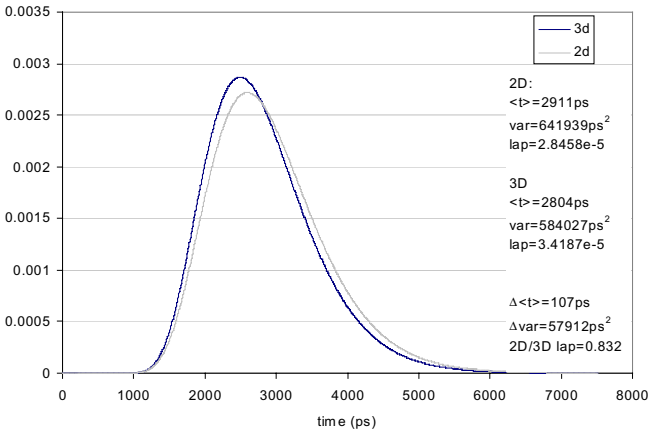

$$G_{3D\_inf}(\omega) = \frac{3(\mu_a + \mu_s')}{2d(2\pi)^{3/2}} e^{-d\sqrt{3(\mu_a+\mu_s')(\mu_a+\frac{i\omega n}{c})}}$$

$$G_{2D\_inf}(\omega) = \frac{3(\mu_a + \mu_s')}{2d(2\pi)^{3/2}} K_0\left[d\sqrt{3(\mu_a+\mu_s')(\mu_a+\frac{i\omega n}{c})}\right]$$

(where K0[] is the modified Bessel function of order zero)

Figure 2.3.2 (left) TPSFs generated using finite element models of a cylinder (3D) and a disc (2D). Three datatypes extracted from the TPSFs are shown on the right. $\mu_A = 0.01mm^{-1}$, $\mu_s' = 1mm^{-1}$, d=70mm and n=1.4.(right) the infinite space Green's function solutions of the DA in 2D and 3D which also yield differing results for the same properties (Arridge et al, 1992)

The requirement of successful image reconstruction is that TOAST can simulate data corresponding to the measurements. Clearly this will not be possible if the data read-in is 3D and the TOAST model is 2D. However the fact remains that 2D models are very attractive since computation and hence image reconstruction is substantially faster. Many other researchers continue to assume that 2D models are sufficient for reconstructing data with specific source-detector geometries (see section 2.3.4). Three possible approaches to using a 2D model were explored:

1) Reconstructing the measured data using a 2D model and assuming the data mismatch to be insignificant.
2) Deriving and applying an ad-hoc correction factor to measured data to allow it to be reconstructed in 2D.
3) Using difference data where only the changes in data between an object with structure and a homogenous reference object were reconstructed.

2.3.1.1 2D:3D corrections

An ad-hoc correction factor can be derived to allow experimental data to be more readily reconstructed using a 2D mesh although it will in no way correct for effects from out-of-plane objects. A preliminary exploration of the feasibility of deriving a correction factor was conducted by Schweiger and Arridge (Schweiger *et al*, 1998) who employed data simulated using a 3D FEM model. Image reconstructions based on the uncorrected 3D data and a 2D algorithm exhibited a significant ring artefact. The simulated data were then adjusted by re-scaling each datatype by a factor dependent on the source-detector separation. These factors were generated by dividing the datatype derived from a 3D FEM model of a homogenous cylinder by the same datatype generated from a 2D FEM model of a homogenous disc of the same diameter and optical properties. This can be expressed as the assumption that:

$$\frac{\mathbf{M}^{3D}{}_{n,m} \times M^{2D}{}_{n,m}}{M^{3D}{}_{n,m}} \approx \mathbf{M}^{2D}{}_{n,m}$$

[2.3.1]

Where $\mathbf{M}^{3D}{}_{n,m}$ is the measured data, $M^{*D}{}_{n,m}$ is a 2D or 3D simulated datatype and $\mathbf{M}^{2D}{}_{n,m}$ represents a 2D-corrected measurement suitable for reconstruction using a 2D model. This crude correction method produced images with significantly fewer artefacts. Without a correction, in non-linear reconstruction notation (see section 1.3.2.3) TOAST will be solving:

$$\mathbf{M}^{3D}{}_{n,m} - M^{2D(g)}{}_{n,m} = J_{n,m}{}^{M}(x^{2D}{}_{g}(r))[x^{2D}{}_{g+1}(r) - x^{2D}{}_{g}(r)]$$

[2.3.2]

Since the approach of using 'absolute' calibrated data from MONSTIR is fairly unique, other researchers do not encounter major image artefacts. This is because they utilise difference imaging, where the data reconstructed are the differences between a reference phantom $\mathbf{M}^{3D(0)}{}_{n,m}$ and the object of interest $\mathbf{M}^{3D(1)}{}_{n,m}$. As described in section 2.1.2.3 using difference data has advantages relating to calibration and modelling error tolerance (as well as a number of constraints and disadvantages). In terms of a 2D:3D correction, using difference data in 2D reconstructions means solving an equation of the form:

$$\mathbf{M}^{3D(1)}{}_{n,m} - \mathbf{M}^{3D(0)}{}_{n,m} = J_{n,m}{}^{M}(x^{2D}{}_{0}(r))[x^{2D}{}_{1}(r) - x^{2D}{}_{0}(r)]$$

[2.3.3]

(using linear reconstruction notation - section 1.3.2.2).

Effectively the assumption is being made that:

$$\mathbf{M}^{3D(1)}{}_{n,m} - \mathbf{M}^{3D(0)}{}_{n,m} \approx \mathbf{M}^{2D(1)}{}_{n,m} - \mathbf{M}^{2D(0)}{}_{n,m}$$

[2.3.4]

This corresponds well to the intuitive approach taken by Schweiger and Arridge (Schweiger *et al*, 1998) since their method used a correction factor that involved dividing by a 2D model and multiplying by a 3D model. Examination of the magnitude of the changes in datatypes calculated in 2D and 3D revealed relations similar to those derived for extracting datatypes of convolved functions described in section 2.1.3.1. It was found that *adding* the correction factors to moments, and multiplying by corrections for Laplace yielded the best results. This means that for Ln(Laplace), Ln(Intensity) and moments we would be solving:

$$\mathbf{M}^{3D(1)}{}_{n,m} - M^{3D(0)}{}_{n,m} = J_{n,m}{}^{M}(x^{2D}{}_{0}(r))[x^{2D}{}_{1}(r) - x^{2D}{}_{0}(r)]$$

[2.3.5]

In the non-linear case iteration g would be given by:

$$\mathbf{M}^{3D(1)}{}_{n,m} - M^{3D(0)}{}_{n,m} + M^{2D(0)}{}_{n,m} - M^{2D(g)}{}_{n,m} = J_{n,m}{}^{M}(x^{2D}{}_{g}(r))[x^{2D}{}_{g+1}(r) - x^{2D}{}_{g}(r)]$$

[2.3.6]

Therefore we are correcting 3D data by effectively simulating a 3D-reference measurement. Specific corrections can be applied to absolute data using [2.3.1] (or its moment equivalent). However the correction can also be quickly applied within TOAST by:

- Simulating the required datatypes on a homogenous medium with the same shape as the object of interest (3D), with optical properties fairly close to those of the object being imaged (to create $M^{3D(0)}{}_{n,m}$)
- Reading these data in as a REFFILE (as if they were a reference measurement for difference imaging).
- Setting the initial optical property conditions in TOAST equal to the optical properties of this reference set (so that TOAST derives a corresponding $M^{2D(0)}{}_{n,m}$ as shown in [2.3.6])
- Reconstructing on a 2D mesh with the same shape in the plane where the measurements are made.

Limitations of this technique are that the sources and detectors must be in the same plane. A 3D homogenous mesh of the object and a suitable solver are required for a single forward solution. Any structure out of the plane of interest will affect images, so the object being imaged must be very simple. Using a reference measurement instead of effectively simulating one has many advantages relating to calibration and model error cancellation, however it is not always possible to measure a suitable reference in practice (section 2.1.4).

Data acquired on the 3D Multi-level phantom shown in Figure 2.3.1 were used for a systematic examination of the 2D:3D correction procedure as described below.

2.3.1.2 Imaging tests of the ad-hoc 2D:3D correction for absolute data

Figure 2.3.2 showed that there is an intrinsic difference between forward data simulated on a 2D and a 3D mesh. Figure 2.3.3 shows a reconstruction of simulated mean, variance and Laplace (s=0.005 ps^{-1}) data. The datatypes were simulated on a homogenous 2D mesh (μ_a = 0.01mm^{-1}, μ'_s = 1mm^{-1}) and then the ad-hoc correction factors, approximately equal to the difference between 2D and 3D datatypes, were added to the simulated data. Images were then reconstructed on the original 2D mesh. This is equivalent to reconstructing simulated 3D data on a 2D mesh. The resulting artefacts represent the typical nature of the 2D:3D artefact.

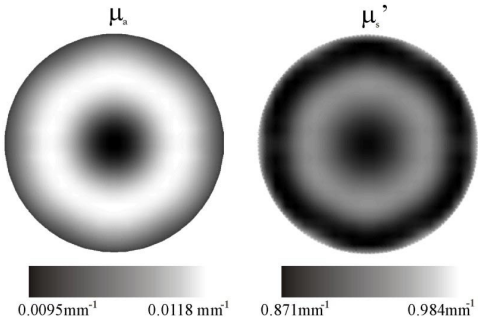

Figure 2.3.3 2D Image reconstructions from simulated 2D data - with correction factors added to mimic 3D data. Thus these artefacts represent the typical 2D:3D artefact.

Experimental data were acquired on the central plane of the 3D Multi-level phantom (which contains the both absorbing and scattering feature), using 32 sources and 22 detectors per source. Calibration measurements were also obtained using the absolute calibration procedure described in section 2.1.4.1. Mean-time, variance and Laplace (s=0.005ps^{-1}) were extracted from the measured data. Figure 2.3.4 (top) shows images reconstructed on a 2D circular mesh from these data without the ad-hoc 2D:3D correction being applied.

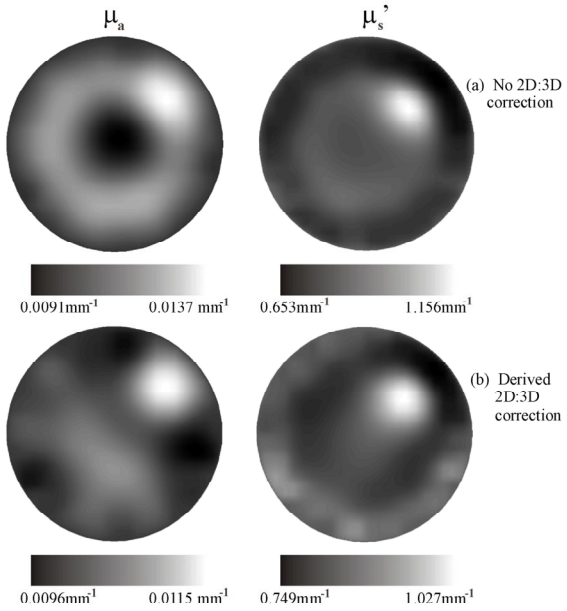

Figure 2.3.4 (top) images reconstructed from calibrated absolute data, using a 2D mesh. (bottom), the same data reconstructed on the same mesh, following ad-hoc 2D:3D correction.

The top images in Figure 2.3.4 represent the case where TOAST is trying to find a 2D distribution that yields boundary datatypes that match the absolute measured (3D) data. The observed artefacts are very similar to those demonstrated in Figure 2.3.3. Note also that the position of the anomaly appears to be different in the μ_a and μ'_s images.

Figure 2.3.4 (bottom) shows images that *have* had the ad-hoc correction factors applied directly to the data prior to image reconstruction (along with calibration). Clearly the 2D:3D ring artefacts have been reduced. Both features have also moved more towards their correct positions. Nevertheless artefacts remain due to the difficulties of calibrating absolute measurements.

2.3.1.2.1 Image summary (2D:3D corrected absolute images)

These results were published in (Hillman *et al*, 2000a) (submitted December 1999).

Phantom	3D Multi-level
Mesh	Circular 3781 nodes 7392 linear elements (2D)
Basis	Pixel 24×24
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iterations	8
Sources	32
Detectors per source	22
Datatypes	Mean, Var, Laplace ($s = 0.005$ ps^{-1})
Calibration	De-conv equivalent
2d 3d correction applied?	No (top) Yes (bottom)
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	10 secs
Wavelength	800 nm

Table 2.3.1 Properties of images shown in Figure 2.3.4

2.3.1.3 Imaging tests of 2D:3D mismatch compensation using difference data

Assuming there is no structure out of the plane, a 2D mesh can be used to reconstruct difference measurements between a homogenous reference phantom and the object of interest if we assume that [2.3.4] holds. The effectiveness of the 2D:3D compensation should be equivalent to the correction applied to the absolute images shown in Figure 2.3.4. However in addition, other model errors and systematic errors on data should be better accounted for in difference measurements.

In addition to the data acquired on the 3D Multi-level phantom used for the absolute images above data were also acquired, using exactly the same acquisition protocol, on a homogenous cylindrical phantom with similar but not identical background optical properties. ‘Raw’ mean-time data were extracted from the TPSFs collected on both phantoms, *without* calibration measurements being applied (in fact the calibration measurements were utilised, but purely to determine the optimal window over which to calculate datatypes see section 2.1.3.2.3). The raw data were read-in to TOAST as ‘FILE’ and ‘REFFILE’, effectively providing TOAST with data in the form of [2.3.3]. Figure 2.3.5 shows the images produced

by TOAST (the 11th iterations) when data are reconstructed (top) on a 2D mesh and (bottom) using a cylindrical 3D mesh. The 3D image shown represents the reconstructed values in the plane of the fibre holder. For the 3D image, 11 iterations required 23 hours on a 700 MHz Pentium III computer. The 2D images took 6 minutes to perform 11 iterations.

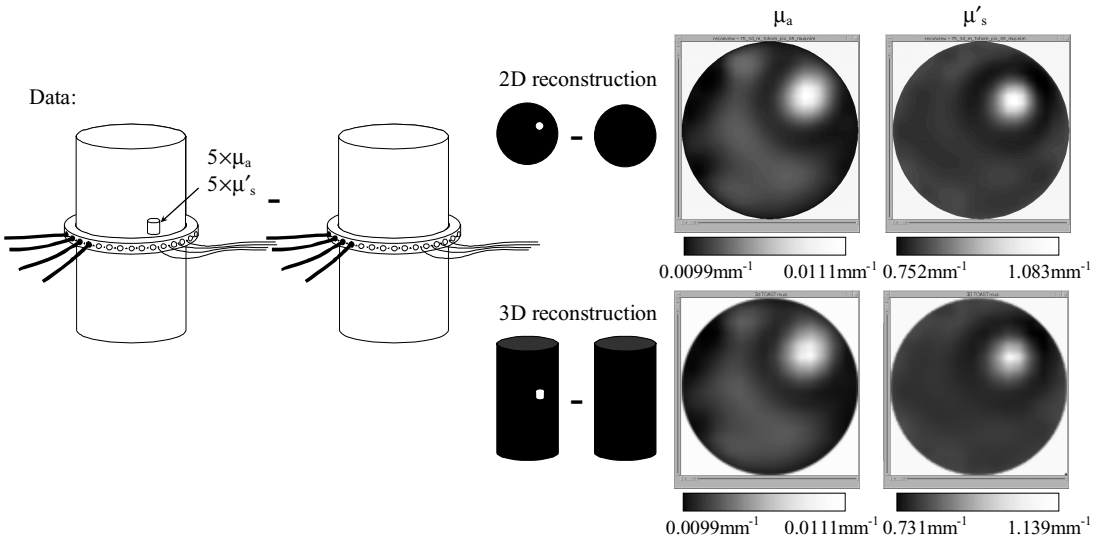

Figure 2.3.5 Images produced from the difference in mean-time between the 3D Multi-level phantom and the homogenous phantom. (top) images reconstructed using a simple 2D mesh, (bottom) the same data reconstructed on a correct 3D cylindrical mesh.

2.3.1.3.1 Image summary (difference)

The 2D:3D comparison in Figure 2.3.5 was published in (Hillman *et al*, 2001c) (submitted Jan 2001).

Phantom	3D Multi-level - Homogenous	3D Multi-level - Homogenous
Mesh	Circular 3781 nodes 7392 linear elements (2D)	(3D) Cylindrical 140 mm tall, 31525 nodes, 21600 quadratic elements. (24 layers, 950 nodes in each)
Basis	Pixel 24×24	Pixel 16×16×5
Starting parameters: μ_a	0.01 mm^{-1}	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}	1 mm^{-1}
Iteration	11	11
Sources	32	32
Detectors per source	22	22
Datatypes	(raw) Mean	(raw) Mean
Calibration	Difference (homog reference)	Difference (homog reference)
2d 3d correction applied?	No	No
Simultaneous	μ_a and μ'_s	μ_a and μ'_s
Acquisition time (per source)	10 secs	10 secs
Wavelength	800 nm	800 nm

Table 2.3.2 Properties of images shown in Figure 2.3.5

2.3.1.4 Discussion of 2D:3D corrections

Clearly the ad-hoc 2D:3D correction applied to absolute data reduces the artefacts relating to 2D:3D mismatch seen in uncorrected images (Figure 2.3.4). However, manipulating datatypes prior to image reconstruction may have repercussions on image accuracy and convergence particularly if datatypes are not necessarily quantitatively correct anyway (e.g. variance – see section 2.2.6.2.1).

Difference imaging was quite successful at producing 2D images from measured data (Figure 2.3.5). Note however that whether absolute corrections or difference data are used to compensate for 2D:3D mismatches, both are strictly limited in usefulness since they require the regions above and below the plane of interest to be homogenous. As a method of re-constructing clinical data, particularly brain data, using a 2D model is likely to yield unacceptable results regardless of attempts to perform a 2D:3D correction. Arm images reconstructed from data where a 2D-3D correction has been applied are shown in chapter 2.5.

The 2D and 3D difference images shown in Figure 2.3.5 agree remarkably well with each-other, both qualitatively and (especially for μ_a) quantitatively. The likely reasons why this method has been so effective are:

- that there is no structure out of the plane (within the areas probed by the PMDFs),
- that systematic and modelling errors have been very well cancelled out by using difference data (as suggested in section 2.1.2.3).
- that the effect of the 3D propagation of photons through the homogenous layers above and below the plane of interest has somehow cancelled, since it is approximately the same for both data sets despite the slight differences in the optical properties of the phantoms.

The quantitative solution is far from the target. However, low iterations are shown, and peak μ_a and μ'_s values do not necessarily indicate the quantitative accuracy of an image, since any blurring will always reduce the peak value in the image. These effects are explored in appendix A (2.9.8) and chapter 2.5.

The fact that the difference imaging results for 2D and 3D imaging (shown in Figure 2.3.5) were so similar suggested that subsequent phantom images could be reconstructed on both 2D and 3D meshes. In general 2D reconstructions are used for speed and evaluating image-data quality. Convergence of 2D reconstructions is also often better than for 3D. This might be due to the fact that the number of unknowns (given by the number of nodes in the mesh) is a lot smaller for 2D problems. Comparison between 2D and 3D results for a number of phantoms is presented below.

2.3.2 Multi-level phantom 2D and 3D image reconstructions

Figure 2.3.1 showed plots of data acquired at 14 levels on the multi-level phantom (which contains 3 inclusions at different heights). Data were also acquired on a homogenous reference phantom. These data can be reconstructed in a number of ways:

- 2D:3D corrected absolute data can be used to reconstruct 14 images, each representing a 2D slice through the phantom (or use difference data relative to the homogeneous phantom to compensate for 2D:3D mismatches).
- All of the absolute data (or a subset) can be reconstructed on a cylindrical 3D mesh.
- Difference data corresponding to each level subtracted from the homogenous phantom measurement can be reconstructed on a 3D cylindrical mesh.

2.3.2.1 Multi-slice 2D imaging

Figure 2.3.6 shows 14 absolute 2D:3D corrected images of the multi-level phantom shown in Figure 2.3.1. Each image was reconstructed separately using a 2D mesh from only data acquired in the single plane represented by the image. The primary motivation for this study was to establish whether the data were of high enough quality to attempt a full 3D image reconstruction. However it is also useful to establish whether images of an object with simple 3D structure can be produced using multiple 2D slices. While out-of plane structure will affect images, having multiple successive 2D images may allow the true positions of the inclusions to be determined.

Data were collected using all 32 sources in the fibre holder ring, and the 22 opposite detectors. Mean-time, variance and Laplace ($s = 0.005ps^{-1}$) were extracted from the measured TPSFs. Absolute calibration was performed using the calibration measurements and methods described in section 2.1.4.1. A 2D:3D correction was applied to each single-plane data set as detailed in section 2.3.1.2. Certain source-detector combinations had to be removed prior to image reconstruction due to drifts and jitter in system IRFs during acquisition of the 14 sets of data (only 2 calibration measurements were made: one before and one after the 14 data sets were acquired).

Iteration 5 for μ_a and 7 for μ'_s were chosen. The images are scaled between the global maximum and minimum of all 14 images. The target images are shown to the left. If the reconstruction were perfect, all images would be black besides the slices at 50mm, 75mm and 100mm which would have the μ_a and μ'_s distributions indicated (although perhaps the images ± 5mm either side would also include evidence of the 10mm high inclusions).

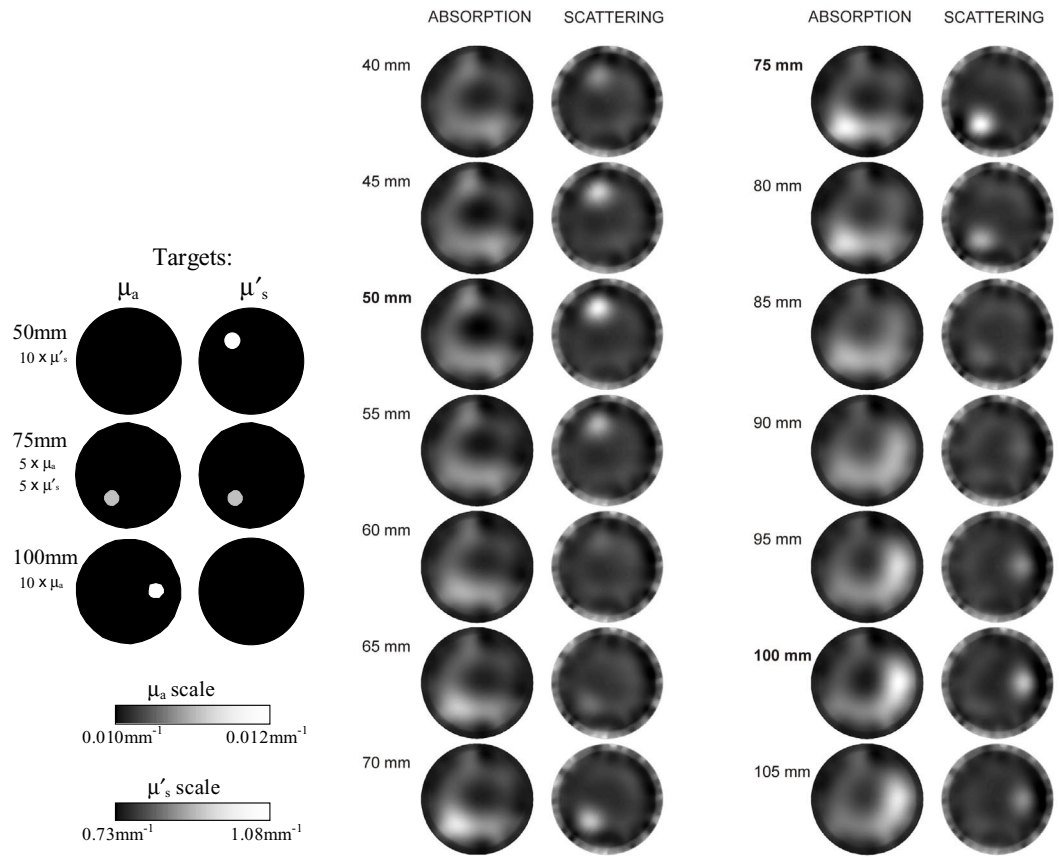

Figure 2.3.6 Multi-slice images of the 3D phantom shown in Figure 2.3.1. Data were acquired at each height and each image was reconstructed from only the data from that plane using a 2D mesh with a 2D:3D correction.

2.3.2.1.1 Image summary (2D multi-slice)

These results were published in (Schmidt *et al*, 2000b) (submitted October 1999).

Phantom	3D Multi-level
Mesh	Circular 3781 nodes, 7392 linear elements (2D)
Basis	Forward mesh
Starting parameters: μ_a	$0.01mm^{-1}$
Starting parameters: μ'_s	$1mm^{-1}$
Iteration	5 for μ_a, 7 for μ'_s
Sources	32 (per image)
Detectors per source	22
Datatypes	Mean, Cvar, lap ($s=0.005ps^{-1}$)
Calibration	De-conv equivalent
2d 3d correction applied?	Yes
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	15 secs
Wavelength	780nm

Table 2.3.3 Properties of images shown in Figure 2.3.6

2.3.2.1.2 Discussion of 2D multi-slice images

There are a number of artefacts in the multi-slice images shown in Figure 2.3.6. The fairly constant artefact in the μ_a images suggests calibration errors since it is non-symmetric and appears in all layers. The dark region in the centre of the absorption images is similar to that seen in the 2D:3D corrected μ_a image shown in Figure 2.3.4 and may therefore be characteristic of images reconstructed from data that have been 2D 3D corrected. Artefacts in the μ'_s images are higher resolution to those in the absorption images. The white regions

around the edges on the μ'_s images originate from the fact that higher resolution mesh elements are used near the surface to better model the regions close to sources and detectors. Forward-mesh basis was used for these reconstructions which can accentuate features in this high density region. Later images (along with those shown in Figure 2.3.4) use pixel basis which effectively smoothes images and can remove this effect (see section 1.3.1.3.1).

Quantitation is poor in these images. This may be due to the choice of very early iterations (5 and 7), which was necessary due to the amplification of image artefacts with higher iterations. In addition however the contrast of the inclusions is very high, and reconstructions of perfect simulated data can rarely quantify such extreme values (Schweiger *et al*, 1997b). Inherent blurring in images will result in peak values being reduced, this effect is explored in appendix A (2.9.8). There is also cross-talk between absorption and scatter despite three temporal datatypes being used, although the noise tolerances of variance and Laplace are thought to be quite poor (see 2.2.6.2), so good convergence to a well-separated solution would not really be expected. Cross-talk may explain why the inclusion at 50mm does not appear to have half the contrast of the other two inclusions, despite only being $5x\mu_a$ and $5x\mu'_s$ (compared to $10x\mu_a$ or $10x\mu'_s$).

Although there is some uncertainty about the exact vertical location of the inclusions due to manufacturing difficulties, clearly a significant number of photons are sampling regions more than 10mm above and below the acquisition plane. Features can be seen in images of planes around 10 mm either side of the plane where the object is located. This is in agreement with the data plots shown in Figure 2.3.1. Interestingly though the inclusions in the out-of-plane images seem to have maintained their (x,y) positions well. Further, it is very clear which planes actually contain the inclusions (since the 'blobs' are brightest in the 50mm, 75mm and 100mm images). This suggests that if an object is mainly homogenous but for one or two discrete inclusions, that 2D slice imaging may prove a quick and simple alternative to full 3D image reconstruction, particularly for preliminary analysis of data quality. Note that these images were reconstructed from absolute data. Results in section 2.3.1.3 implied that difference data (relative to the homogenous phantom) reconstructed in this 2D slice way would have produced results with less artefact without requiring ad-hoc 2D:3D correction.

2.3.2.2 3D absolute image reconstruction of the multi-level phantom data.

We wish to compare the 2D slice imaging results shown above with those obtainable using the full 3D capabilities of TOAST. Data were acquired in a number of geometries on the 3D Multi-level phantom in order to examine the optimal configuration for 3D image reconstruction. The best 3D results were in fact achieved using a selection of the single-plane data plotted in Figure 2.3.1 namely those acquired at 50mm, 65mm, 75mm, 90mm and

100mm. Additional cross plane data were also acquired by distributing sources and detectors over 3 fibre-holder rings so that each source illuminated detectors both in its plane and those in other planes. However image reconstruction from these cross-plane data was not successful. This is suspected to have been a result of poor repeatability in the alignment of the multiple rings used for cross-plane acquisition.

For 3D image reconstruction, all data is read in at once, and the optical property distribution throughout the 3D volume of the mesh found that minimises the error between the data and model. Errors in data are likely to affect images more here than in those shown previously. Firstly because the number of sources and detectors used is now 32×5 = 160, and due to the lengthy time taken for data acquisition, each error value will change slightly with time following calibration measurements. Secondly, the model has to find more unknowns, since the 3D mesh now has 17157 nodes (and 92160 tetrahaedral elements) instead of 3781 nodes for 2D images. Errors in the position of the ring, and its rotational orientation with respect to its position at other levels will also contribute inconsistencies.

Two types of image reconstruction were tested: absolute, and difference relative to the homogenous phantom. Results are compared with images reconstructed from simulated data, derived from a 3D cylindrical model of the phantom with the correct inclusions (the 3D simulation and 3D reconstructions were performed by M. Schweiger and S. Arridge).

For the absolute images, only mean-time data were used and calibration was performed using the methods described in section 2.1.4.1. A single datatype was used because mean-time calibration was found to be most accurate compared to other datatypes (as shown in section 2.2.6.2). Also image reconstruction is faster for a single datatype. The results are shown in Figure 2.3.7. Each image represents a slice through the mesh from top to bottom.

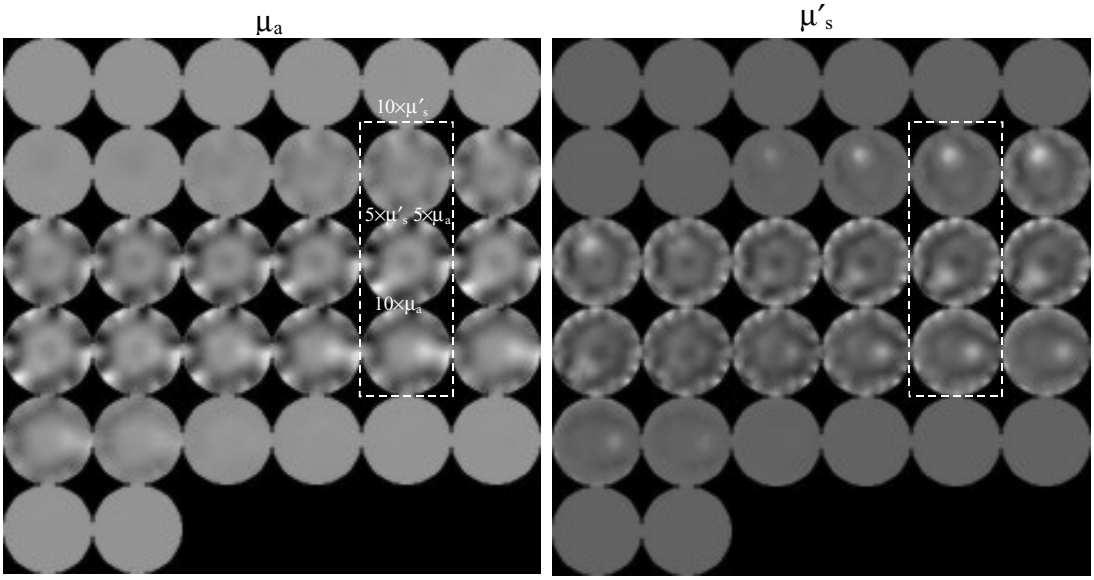

Figure 2.3.7 3D reconstruction from absolute mean-time data. Each image represents a slice through the 3D solution (32 throughout the 140 mm mesh). The image on the top left of each represents the base of the phantom.

2.3.2.2.1 Image summary (3D absolute)

Phantom	3D Multi-level
Mesh	(3D) Cylindrical 140mm tall, 17157 nodes, 92160 linear (tetrahedral) elements
Basis	Forward mesh 3×3×3 median filtered
Starting parameters: μ_a	0.011mm^{-1}
Starting parameters: μ'_s	0.85mm^{-1}
Iteration	16
Sources	32×5
Detectors per source	22
Datatypes	Mean
Calibration	De-conv equivalent
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	15 secs
Wavelength	780nm

Table 2.3.4 Properties of images shown in Figure 2.3.7

2.3.2.2.2 Discussion (absolute 3D)

The results shown in Figure 2.3.7 correspond to the 3D version of the multi-plane images shown in Figure 2.3.6, although here only mean-time data were used. The results using the 3D reconstruction contain more artefact than the 2D planar images. This is likely to be due to the higher sensitivity of the 3D reconstruction to calibration and positional errors and the lower node-density in the 3D mesh compared to the 2D meshes used previously. The μ'_s features can be identified. However, although the μ_a image does show the μ_a inclusion, it is heavily distorted. It is difficult to comment whether significant cross-talk between the μ_a and μ'_s images is present. The properties of the μ'_s image resemble those of the 2D:3D corrected μ'_s image shown in Figure 2.3.4. This may suggest that these features are related to the integrity of the absolute calibrated data.

2.3.2.3 3D difference image reconstruction of the multi-level phantom data.

We can create difference data suitable for 3D image reconstruction by subtracting measurements made in a single plane on a homogenous phantom from each of the multi-level data sets acquired on the 3D multi-level phantom. This requires the assumption that the properties of the two phantoms are fairly similar (and an estimate of the optical properties of the homogenous phantom is available) and that the system stayed stable for the duration of the experiment. In addition it assumes that the upper and lower boundaries of the 3D multi-plane phantom were not affecting data acquired further than 4cm from the top or bottom of the cylinder.

Raw mean-time values were extracted from data acquired on the 3D-Multi-level phantom and the homogenous phantom, (as for the 2D reconstruction in section 2.3.1.3). These data were then reconstructed on the same 3D mesh as the absolute images above, with the same reconstruction parameters. The results are shown in Figure 2.3.8.

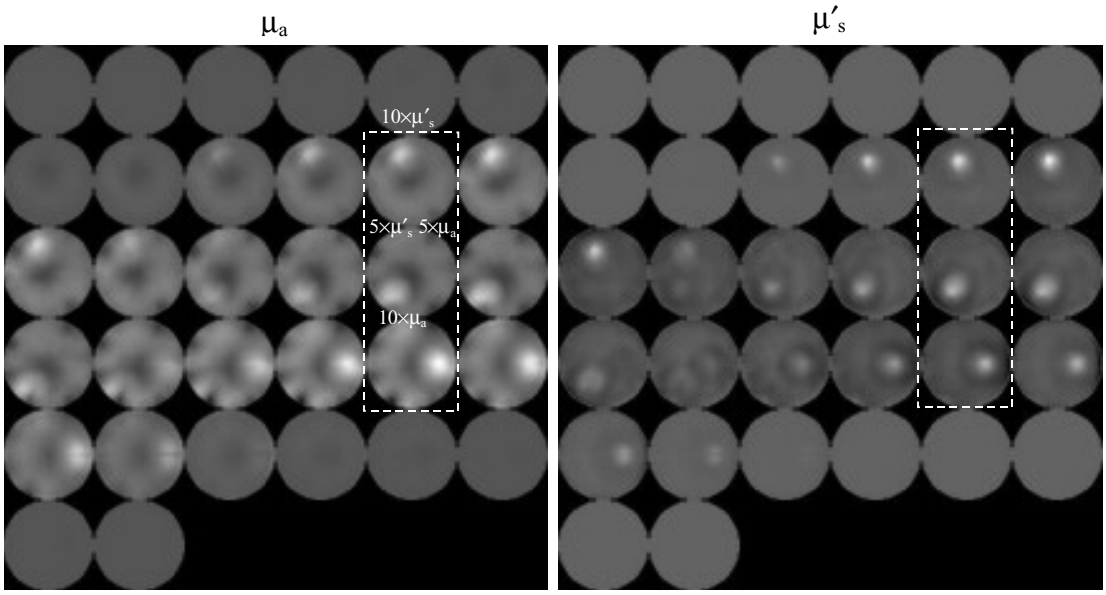

Figure 2.3.8 3D reconstruction from the difference between mean-time data from the 3D Multi-level phantom and the homogenous phantom. Each image represents a slice through the 3D solution (32 throughout the 140 mm mesh). The image on the top left of each set represents the base of the phantom.

2.3.2.3.1 Image summary (3D difference)

Phantom	3D Multi-level - Homogenous
Mesh	(3D) Cylindrical 140mm tall, 17157 nodes, 92160 linear (tetrahedral) elements
Basis	Forward mesh 3×3×3 median filtered
Starting parameters: μ_a	0.011mm^{-1}
Starting parameters: μ'_s	0.85mm^{-1}
Iteration	16
Sources	32×5
Detectors per source	22
Datatypes	(raw) Mean
Calibration	Difference to homogenous
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	15 secs
Wavelength	780nm

Table 2.3.5 Properties of images shown in Figure 2.3.8

2.3.2.3.2 Discussion (3D difference)

These 3D difference images clearly have far less artefact than those reconstructed from absolute data (shown in Figure 2.3.7). There is cross-talk between the two images, although this is likely to be due to the use of a single datatype (see section 1.3.3.4). The reduced amount of artefact compared to the absolute images supports the theory-based suggestion that difference data can compensate for errors much more effectively than absolute calibration methods. This includes both calibration and model errors as detailed in section 2.1.2.3.

2.3.2.4 3D simulation of the multi-level phantom.

A simulation was performed in order to compare data from MONSTIR with what TOAST simulates given the (nominal) geometry and properties of the phantom. Images reconstructed from this simulated data represent the best possible images achievable using

TOAST. This is because if data are simulated using the TOAST forward model on a mesh with a particular optical property distribution, and the image reconstruction uses the same mesh, there is definitely an image solution that corresponds exactly to the input boundary data. Simulations have been traditionally used to optimise reconstruction techniques prior to attempts to reconstruct real experiment data.

Data simulated using TOAST's forward solver are reconstructed using the same reconstruction parameters as the absolute and difference images from experimental data shown above. The same mesh was used for the simulation as for the image reconstruction (known as *inverse crime*). This means that TOAST effectively reconstructs from simulated difference data. The resulting images are shown in Figure 2.3.9. (Note that the orientation of the inclusions in the simulated phantom is reflected in the horizontal axis compared to the real phantom).

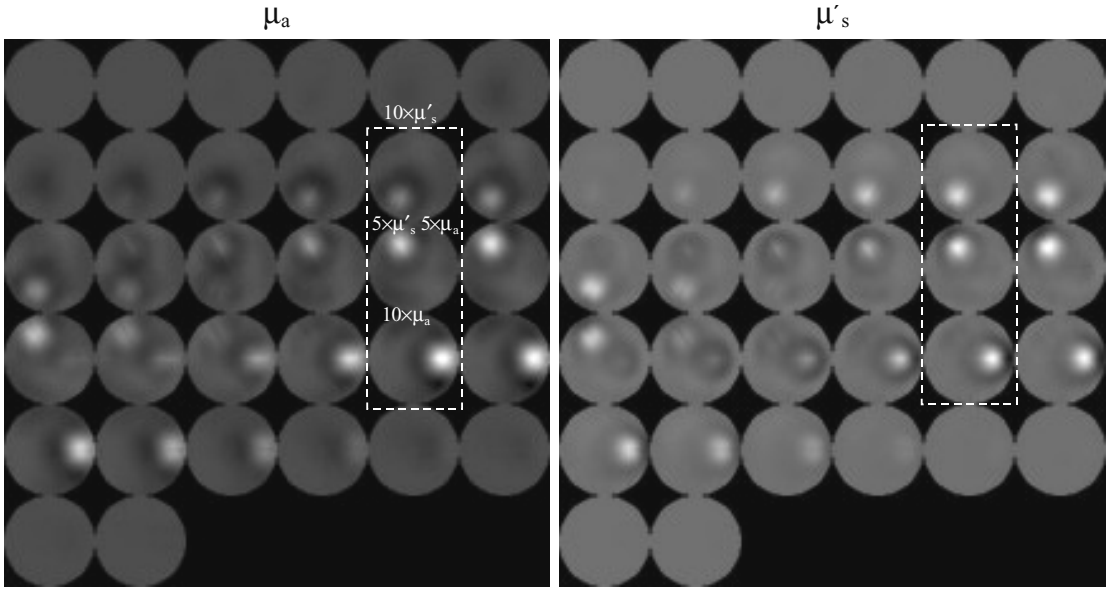

Figure 2.3.9 3D reconstruction from simulated mean-time data. Each image represents a slice through the 3D solution (32 throughout the 140 mm mesh). The image on the top left of each set represents the base of the simulated phantom.

2.3.2.4.1 Image summary (3D simulated)

Phantom	3D Simulation
Mesh	(3D) Cylindrical 140mm tall, 17157 nodes, 92160 linear (tetrahedral) elements
Basis	Forward mesh 3×3×3 median filtered
Starting parameters: μ_a	0.011mm^{-1}
Starting parameters: μ'_s	0.85mm^{-1}
Iteration	16
Sources	32×5
Detectors per source	22
Datatypes	Mean (simulated)
Calibration	Simulated, but inverse crime (difference)
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	-
Wavelength	-

Table 2.3.6 Properties of images shown in Figure 2.3.9

2.3.2.4.2 Discussion (3D simulated)

The images shown in Figure 2.3.9 are very similar to the experimental difference images shown in Figure 2.3.8, although the μ_a simulated images are better quality than the experimental ones. Cross-talk between μ_a and μ'_s is evident, which implies that cross talk in the images from experimental mean-time is probably due to the use of a single datatype and not a result of using experimental data.

Note that even in these 3D image reconstructions, we still see out-of-plane structure. That is that the inclusions appear to extend further vertically that they in fact do. As with the 2D multi-slice images shown in Figure 2.3.6 the true location of the inclusion can be determined by finding the plane where the 'blob' is brightest. Of course this requires the assumption that the inclusion is in fact discrete, which may not be the case in clinical subjects.

The 3D imaging results presented in Figure 2.3.7, Figure 2.3.8 and Figure 2.3.9 were published in (Arridge *et al*, 2000b) (submitted February 2000)

2.3.3 Breast phantom imaging

Section 1.1.3.2 described why the breast was an appropriate target for optical imaging. In order to test the practical aspects of breast data acquisition and image reconstruction, a number of conical resin breast phantoms were manufactured to fit snugly into a specially constructed conical fibre holder (see section 2.7.3.2). The set comprised of two conical phantoms with background optical properties of $\mu_a = 0.007$ mm^{-1} $\pm$ 0.001 mm^{-1} and $\mu'_s = 0.8$ mm^{-1} $\pm$ 0.1 mm^{-1} at 800 nm. One phantom contained three 10 mm high, 10 mm diameter cylinders: one with twice background μ_a, one with twice background μ'_s and one with twice both μ_a and μ'_s. All three objects were placed in the same plane as shown in Figure 2.3.10. The other phantom was identical but was completely homogenous.

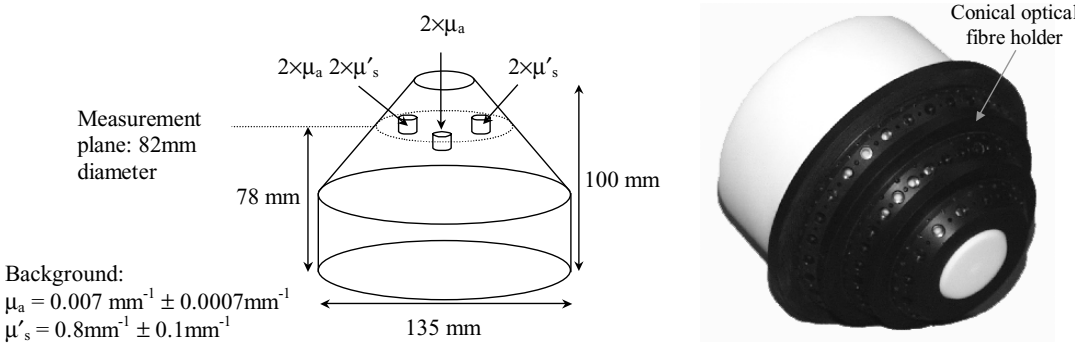

Figure 2.3.10 The breast phantom (containing inclusions) designed to evaluate breast imaging capabilities and test the conical fibre holder shown around the resin phantom (right).

The extension of the cone shape of the phantom into a cylindrical base was a deliberate attempt to consider the effects of the chest wall. Light propagation will not be entirely limited

to the breast tissue alone and if the phantom were manufactured without such as base, data would be affected by unrealistic interactions with this boundary.

Calibration for a conical phantom is complicated by the fact that the sources and detectors are no longer held in a ring that will fit around the absolute calibration tool described in section 2.1.2.1. Note that calibration measurements acquired with the fibre bundles and source fibres in the 70mm diameter ring, using the absolute calibration phantom, *would* be valid for data acquired using the conical fibre holder. However problems arise since the fibre bundles and source fibres are fragile and difficult to move between the two geometries. In the time taken to adapt the set-up for data acquisition following calibration, the system IRFs may not remain stable. Note that the subsequent development of the monstode (see sections 2.1.2.2 and 2.7.2.2.1) allowed sources and detectors to be moved more freely, and also offered the opportunity to perform some in-situ calibration measurements.

However the problems with absolute calibration and the availability of a reference phantom meant that data could be most readily calibrated using difference measurements (see 2.1.4.3). This difference data can be reconstructed in 2D or 3D as shown below.

2.3.3.1 2D breast phantom imaging

As a preliminary experiment, data were acquired using all sources and detectors in a single plane (using the middle ring of the fibre holder, such that the measurement plane coincided with the inclusions). MONSTIR's 32 source and detector fibres were all positioned within the single middle ring. The VOAs were used to shut off 5 detectors each side of an active source. Data were acquired on one phantom, followed by the other. Raw mean-time data were extracted from both sets of measured TPSFs.

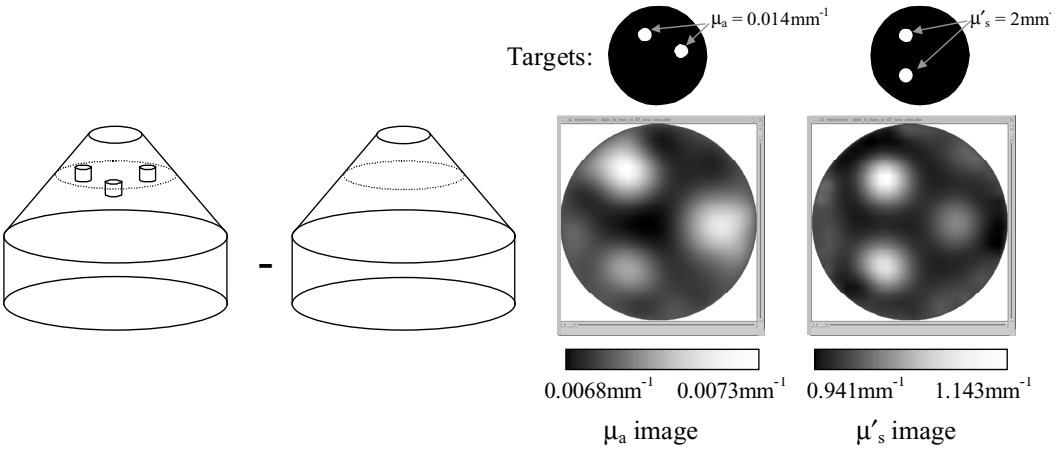

Figure 2.3.11 2D images of the breast phantom at the height of the inclusions, reconstructed for the difference in mean-times between a phantom with structure and an otherwise identical homogenous phantom.

A 2D circular FEM mesh was generated to correspond to the correct size of the sample plane (41mm radius), and difference images were reconstructed. The 10^{th} iteration images are

shown in Figure 2.3.11. 10 iterations took 36 minutes using a 400MHz Sun Ultra-10 workstation.

2.3.3.1.1 Image summary (2D breast phantom)

The images in Figure 2.3.11 were published as part of (Hillman *et al*, 2001c) (submitted January 2001). Note that the breast phantoms were designed, constructed and measured by Hylke Veenstra as detailed in (Hebden *et al*, 2001).

Phantom	Breast phantoms
Mesh	Circular with 3403 nodes, 6615 linear elements, 82mm diameter (2D)
Basis	Pixel 24×24
Starting parameters: μ_a	$0.007mm^{-1}$
Starting parameters: μ'_s	$1mm^{-1}$
Iteration	10
Sources	32
Detectors per source	22
Datatypes	(raw) Mean
Calibration	Difference between Breast - Homog breast
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	30 secs
Wavelength	800nm

Table 2.3.7 Properties of images shown in Figure 2.3.11

2.3.3.1.2 Discussion (2D breast phantom)

The images shown in Figure 2.3.11 are fairly free from artefact, suggesting that good calibration has been achieved through use of the homogenous reference phantom. There is some distortion in the positions of the inclusions, possibly due to the use of a 2D mesh. Quantitatively the images are comparable to the 2D multi-slice images shown previously (Figure 2.3.6Figure 2.3.6). Peak values corresponding to the true inclusion properties are not retrieved due to blurring and the low iteration (10) chosen. Separation of μ_a and μ'_s is not complete, probably due to the fact that only mean-time data were used.

Image quality has not deteriorated notably due to the larger size of the object being imaged (82mm diameter rather than 70mm). Results with this level of parameter separation and resolution could potentially provide sufficient information for a suspicious lesion in the breast to be identified (particularly if multi-wavelength acquisitions were performed, see chapter 2.5). However this data acquisition protocol and image reconstruction method is not optimal for examining a real breast. For example it would be difficult to mount a ring of fibres in the correct plane of a breast to target a suspicious lesion, while also assuming that no out-of-plane structure would affect images and also the cancellation of the 2D:3D mismatch (if difference data were being used).

Further, as was discussed in section 1.1.3.2 the greatest potential of MONSTIR and TOAST is to aid specificity by quantifying the properties of tissue in a particular area in terms of its absorption *and* scatter coefficients. While the images in Figure 2.3.11 have (mostly)

identified the difference between the μ_a and μ'_s inclusions their positions are distorted such that the absorbing *and* scattering inclusion is not coincident in both images. While we know there is a discrete inclusion in the phantom, we can deduce that the same object is responsible for the feature in the two distorted μ_a and μ'_s images. However clinical images may not be so easy to interpret and distortion could cause erroneous diagnosis if a lesion were not well defined.

2.3.3.2 3D breast phantom imaging

In order to evaluate more clinically realistic imaging protocols, another experiment was performed to image the breast phantom in three dimensions. The sources and detector bundles were positioned over the conical surface using the three rings of the fibre holder. New VOA protocols were required to account for the 3D distribution of sources and detectors (see section 1.2.3.3). In addition, since source-detector spacing in this geometry could be up to 110 mm, acquisition was in fact performed in two stages. A single-channel variable attenuator within MONSTIR can adjust the intensity of the source pulses by up to 2.5 ODs via computer control. Data acquisition was performed by using the VOAs to close off all detectors close to the source, and illuminating with a high power to acquire data from detectors far from the source. The source attenuator was then activated such that illumination power was low enough for data to be acquired from detectors close to the source. Some detectors that were very close to sources were not illuminated at all, so 29 or 30 detectors per source were used. Thus the data were not just acquired in a series of planes but 'cross plane' measurements were also acquired so that we are truly sampling the 3D volume.

Data were acquired in exactly the same way on both the breast phantom with inclusions and the homogenous breast phantom. Raw mean-time *and* raw intensity data were extracted from both data sets, and the differences between the data were reconstructed using a 3D mesh.

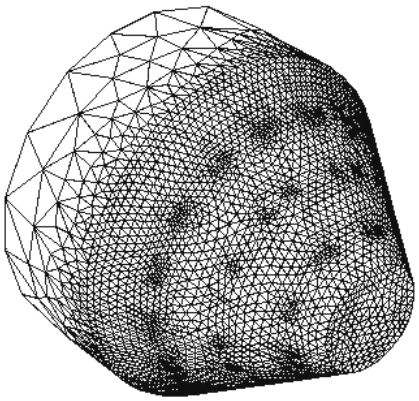

Figure 2.3.12 The 3D mesh generated to reconstruct data acquired on the breast phantoms.

Figure 2.3.12 shows a picture of the surface nodes of the 3D mesh used to reconstruct the data acquired on the breast phantoms. The element size has been modified from uniform to

provide better modelling of photon propagation in areas where photon density is likely to change significantly with distance (as was the case with the 2D mesh as mentioned in section 2.3.2). This is why concentrated areas of nodes are located at the source positions, and large elements make up the 'chest wall' part of the mesh. The mesh has 57766 nodes and 36857 quadratic elements. Image reconstruction (up to the 6th iteration) took 38 hours on a 450 MHz Pentium PC with 1 Gb of RAM.

Figure 2.3.13 shows the images produced using the 3D reconstruction. Each image represents a slice through the phantom from the apex to the base. Each slice is 6 mm apart.

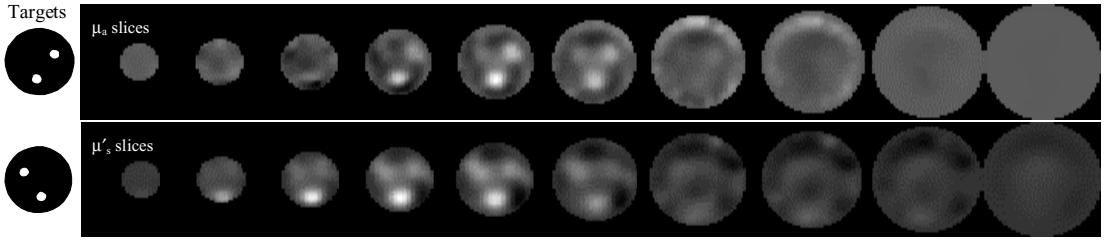

Figure 2.3.13 μ_A and μ'_s images reconstructed from mean-time and intensity differences between the structured and homogenous breast phantoms. Each image shows a slice through the 3D solution, the slices are 6mm apart. The top row of images shows the μ_a solution, and the bottom shown the μ'_s solution. The target images at the palne of the inclusions are shown on the left.

2.3.3.2.1 Image summary (3D breast phantom)

The images shown in Figure 2.3.13 were published in (Hebden *et al*, 2001). Mesh generation and 3D image reconstruction were performed by Hamid Dehghani.

Phantom	Breast phantoms
Mesh	(3D) Conical 57766 nodes and 36857 quadratic elements
Basis	Pixel 32× 32× 18
Starting parameters: μ_a	0.007mm^{-1}
Starting parameters: μ'_s	1mm^{-1}
Iteration	6
Sources	32
Detectors per source	29 or 30
Datatypes	(raw) Mean and intensity
Calibration	Difference between Breast phantom - Homog breast
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	2 × 30 secs
Wavelength	800nm

Table 2.3.8 Properties of images shown in Figure 2.3.13

2.3.3.2.2 Discussion (3D breast phantom)

The 3D images produced using data acquired over the surface of the breast phantom have identified the inclusions in the correct locations. Some cross-talk is visible, the μ_a only feature appears in the μ'_s image and visa versa despite the use of both mean-time and intensity difference data. Also the combined increased μ_a *and* μ'_s feature is the brightest of the three in both images. In fact it should be the same as the μ_a only or μ'_s only feature.

There is more artefact on these 3D breast phantom (difference) images than on the 2D difference images shown in Figure 2.3.11. The image quality also compares poorly to the 3D

multi-level images in Figure 2.3.8. This may be due to the new, and more complicated mesh geometry, along with the larger source-detector separations and the inclusion of cross plane measurements. It may also be due to the use of both mean-time and intensity difference data. If the intensity data were less well calibrated than the mean-time data, convergence of the solution may have been impaired as suggested in section 2.2.6.2.

However, the results have shown that data acquired over the whole volume of a breast-shaped phantom provide sufficient measurement density to identify lesions with 2 × contrast within a phantom which is a realistic size.

2.3.4 Phantoms with axial symmetry

It is often assumed that optical tomography data acquired in a single plane on an object with axial symmetry will represent 2D data ((Hebden *et al*, 1999), (Barbour *et al*, 2001), (Eda *et al*, 1999), (Pogue *et al*, 1995), (Zint *et al*, 2001)). Figure 2.3.14 shows an axially symmetric cylindrical phantom containing three rods (rather than discrete inclusions as in the phantoms explored above). The distribution of optical properties in this phantom can be expressed as $\mathbf{x}(x,y,z)$. The weight function that relates measurements $\mathbf{M}_{n,m}$ to optical properties is the equivalent of the set of PMDFs for the measurements, which is just the Jacobian $J^M{}_{n,m}(x,y,z)$ (see section 1.3.2.1).

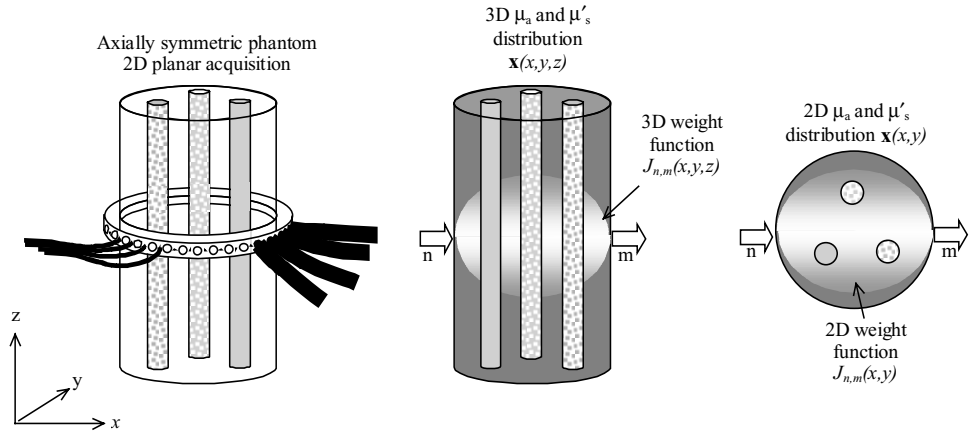

Figure 2.3.14 Definition of an axially symmetric object with a 2D planar measurement geometry. The weight function represents the sensitivity of measurements to changes in optical properties (section 1.3.2.1).

In the simplest, linear form of image reconstruction, it is assumed that a change in measurement $\Delta\mathbf{M}_{n,m}$ corresponds to a change in optical properties $\Delta\mathbf{x}$, the two being linked by a weight function or Jacobian $J^M{}_{n,m}(x,y,z)$ as shown in [2.3.7] (see section 1.3.2.1).

$$\Delta\mathbf{M}^{3D}{}_{n,m} = J^M{}_{n,m}(x,y,z)\Delta\mathbf{x}(x,y,z) \qquad [\ 2.3.7\ ]$$

For experimental measurements, photons travel in three dimensions as demonstrated in Figure 2.3.1. This means that the weight function relating measurements to optical properties

will be three-dimensional. This will be true regardless of the geometry of the detectors (e.g. constraining sources and detectors to a plane does not stop detected photons having migrated out of the plane). If a single source and detector are in the same plane, the intensity field in the object being imaged will be a function of z. Axial symmetry means that:

$$\Delta\mathbf{x}(x,y,z) = \Delta\mathbf{x}(x,y) \qquad [\,2.3.8\,]$$

However we still have:

$$\Delta\mathbf{M}^{3D}{}_{n,m} = J^{M}{}_{n,m}(x,y,z)\Delta\mathbf{x}(x,y) \qquad [\,2.3.9\,]$$

Even if sources and detectors are all in the same z plane, the Jacobian remains a function of z and cannot be replaced with a 2D Jacobian. This was shown in Figure 2.3.2 where both FEM plots and the infinite space Green's functions verify that measurements made on a homogenous (and hence axially symmetric) object with a single source and detector will be different from 2D simulated values.

The only way that '2D data' can be measured on a 3D axially symmetric object is if illumination is via a vertical line-source (or detectors are positioned at multiple heights). This creates an intensity field that is constant over z (since any photon leaving the plane will be replaced by one that has travelled from another plane as shown in Figure 2.3.15).

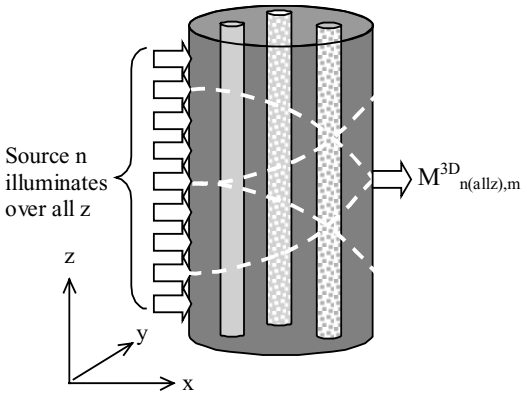

Figure 2.3.15 Only by acquiring data (symmetrically) over the z direction can we effectively measure data in 2D on an axially symmetric phantom

In this limited case, the height of the detector and the weight function are independent of z (and then only if the object is axially symmetric). So we can use a 2D model to solve data acquired in this way, i.e:

$$J^{M}{}_{n(extended),m}(x,y,z) = J^{2D,M}{}_{n,m}(x,y) \qquad [\,2.3.10\,]$$

$$\Delta\mathbf{x}(x,y) = \left[J^{2D,M}{}_{n,m}(x,y)\right]^{-1}\Delta\mathbf{M}^{3D}{}_{n(extended),m} \qquad [\,2.3.11\,]$$

It is unlikely that clinical optical tomography will encounter objects with axial symmetry, so it is not worthwhile to construct a system with linearly extended sources. However since axially symmetric phantoms have already been widely used to acquire pseudo-

2D data with sources and detectors in a plane it is informative to examine how the rods extending beyond the measurement plane will affect reconstructions.

2.3.4.1 2D absolute reconstructions of a phantom with axial symmetry

The first measurements made using MONSTIR were acquired on the 'basic phantom' which has geometry as shown in Figure 2.3.14, has background $\mu_a = 0.01\text{mm}^{-1} \pm 0.002\text{mm}^{-1}$ and $\mu'_s = 1\text{mm}^{-1} \pm 0.2\text{mm}^{-1}$ and contains 3 rods, 8mm in diameter, one with $5x\mu_a$, another with $5x\mu'_s$ and the third with $2x\mu_a$ and $2x\mu'_s$. Data were acquired using the planar fibre holder ring with 5 detectors deactivated either side of the active source using the VOAs (see section 1.2.3.3). The same protocol was also used to acquire data on the homogenous phantom

Although the first ever image produced using MONSTIR was of the axially symmetric 'basic phantom', data were acquired by twisting the phantom within the fibre holder ring containing only 8 detectors (Hebden *et al*, 1999). This resulted in image features quite different to those addressed subsequently, as described in (Hillman, 2001a). Therefore we present images reconstructed from data acquired once the MONSTIR system had been upgraded to 32 detector channels and sources and automated data acquisition was possible.

Calibration measurements were acquired as detailed in section 2.1.2.1. Absolute calibrated mean-time, variance and Laplace (with s= 0.005 ps^{-1}) were extracted from the measured TPSFs using the techniques described in 2.1.4.1. A 2D:3D correction was applied to the absolute data by deriving ad-hoc correction factors from homogenous 2D and 3D models with approximately the same average optical properties as the phantom as described earlier in section 2.3.1.1. Figure 2.3.16 shows the resulting absolute images, which are scaled between their maximum and minimum values.

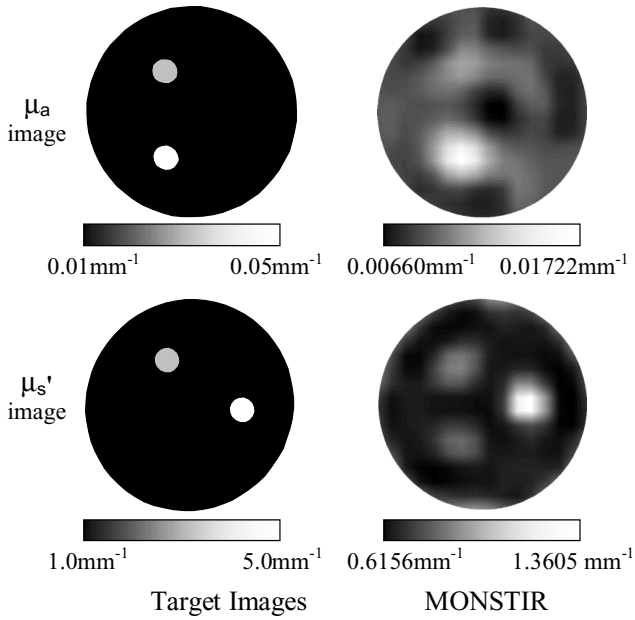

Figure 2.3.16 Images of the 'basic phantom' (Figure 2.3.14.)

It is difficult to compare the 2D absolute images above to results from the non-axially symmetric phantom without examining the potential improvements to the images likely to result from using difference imaging.

2.3.4.2 2D difference reconstructions of phantom with axial symmetry

Raw mean-times were extracted from the TPSFs acquired on both the axially symmetric basic phantom and the homogenous phantoms (using calibration measurements only to guide the calculation window). No 2D:3D correction was applied since results in 2.3.1.3 implied that difference imaging compensates well for intrinsic differences between 2D and 3D data.

The two sets of raw data were read into TOAST and reconstructed on a 2D mesh. The images shown in Figure 2.3.17 are the 5th (μ_a) and 6th (μ'_s) iterations.

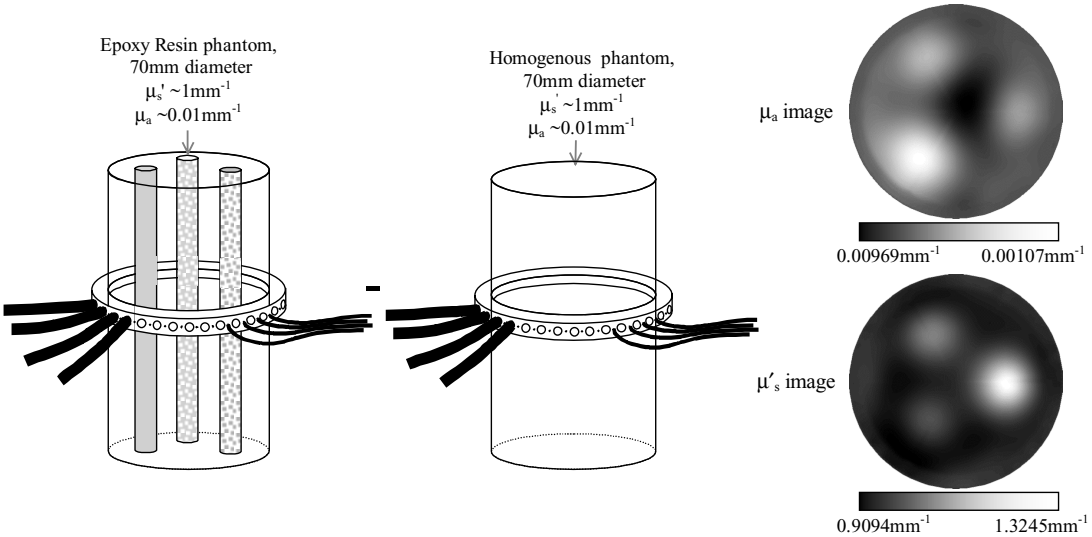

Figure 2.3.17 Difference images of the basic phantom produced using homogenous phantom (mean-time) reference data.

2.3.4.2.1 Discussion (2D images of axially symmetric phantom)

The images of absorption and scatter shown in Figure 2.3.16 have significantly more artefact than those of the same phantom shown in Figure 2.3.17. The use of absolute data in the 2D:3D corrected images is likely to be the main cause of the artefact, in accordance with the results on the multi-plane phantom shown above.

It should be noted that different iterations are shown for the absolute and difference images and different datatypes were used for each reconstruction. The lower iterations were chosen for the difference images since only mean-time data were used and cross-talk worsened in later iterations. Mean-time variance and Laplace were used in the absolute reconstructions in an attempt to improve quantitation and μ_a and μ'_s separation although variance calibration was suspected to have been fairly poor. As suggested in section 1.3.3 the use of three poorly calibrated absolute datatypes may cause more errors due to poor

convergence than using fewer datatypes. Note that Laplace difference reconstruction is not implemented in TOAST. Variance difference data for this phantom were also quite noisy and so were not used.

The main artefact in the absolute images is the blurring between the μ_a features. Despite the differences in the implementation of the two reconstructions, this blurring is unlikely to be due to the presence of rods in the phantom extending beyond the measurement plane since a similar effect is not seen in the difference images of the same phantom.

The distortion of the μ_a rod in the μ_a difference image (causing it to be further out from the centre than the μ'_s rod in the μ'_s image) is very similar to the distortion seen in the 2D difference breast phantom images (Figure 2.3.11), meaning that it is unlikely to be due to the presence of the rods above and below the plane.

In fact there are no discernible features in the absolute and difference images of the basic phantom that seem to relate to the rods extending beyond the plane. However we can ascertain whether there is information within the measurements that corresponds to the out-of-plane presence of the rods by examining a 3D reconstruction of the data as shown below.

2.3.4.2.2 Image summary (absolute 2D basic phantom images)

These results were published in (Schmidt *et al*, 2000a) (submitted July 1999).

Phantom	Basic
Mesh	Circular 3781 nodes, 7392 linear elements (2D)
Basis	Pixel 16×16
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iteration	17
Sources	32
Detectors per source	22
Datatypes	Mean, Cvar, lap (s = 0.005 ps^{-1})
Calibration	De-conv equivalent
2d 3d correction applied?	Yes
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	30 secs
Wavelength	800 nm

Table 2.3.9 Properties of images shown in Figure 2.3.16

2.3.4.2.3 Image summary (difference 2D basic phantom images)

Phantom	Basic - Homogenous
Mesh	Circular 3781 nodes, 7392 linear elements (2D)
Basis	Forward mesh - median filtered (2)
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iteration	$\mu_a = 5$ $\mu'_s = 6$
Sources	32
Detectors per source	22
Datatypes	(raw) Mean
Calibration	Difference to homogenous
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	30 secs
Wavelength	800 nm

Table 2.3.10 Properties of images shown in Figure 2.3.17

2.3.4.3 3D difference reconstructions of phantom with axial symmetry

To examine the effects of having rods in the basic phantom, rather than discrete inclusions, the same data as used for the 2D (mean-time difference) reconstructions above were reconstructed on a 3D cylindrical mesh.

The data were presented to TOAST as if they were collected in two positions, 20 mm apart in the z direction. Due to the axial symmetry of the phantom, approximately the same data would be collected by fibres in a planar ring at any height on the phantom. By providing TOAST with duplicated data in this form we are examining whether TOAST reconstructs continuous rods, or just discrete inclusions at around the height of the rings (as was the case for the 3D Multi-level phantom).

The images shown in Figure 2.3.18 are 1 mm thick slices of the image solution around the region where the fibre holder rings were effectively placed. The 15th iteration is shown for both the μ_a and μ'_s images.

μ_a image slices

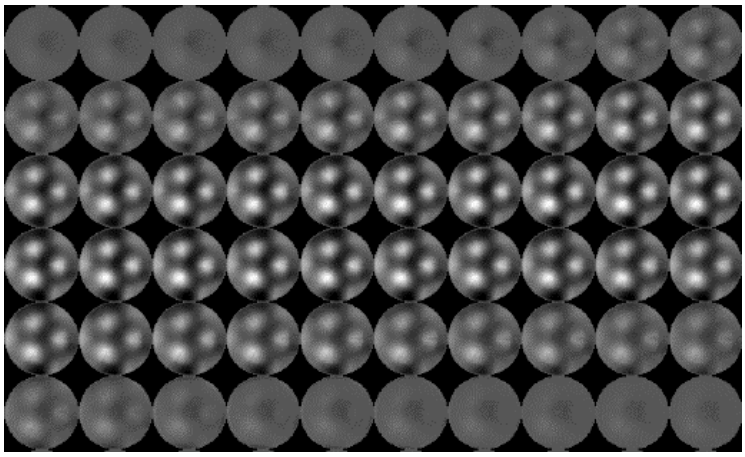

μ'_s image slices

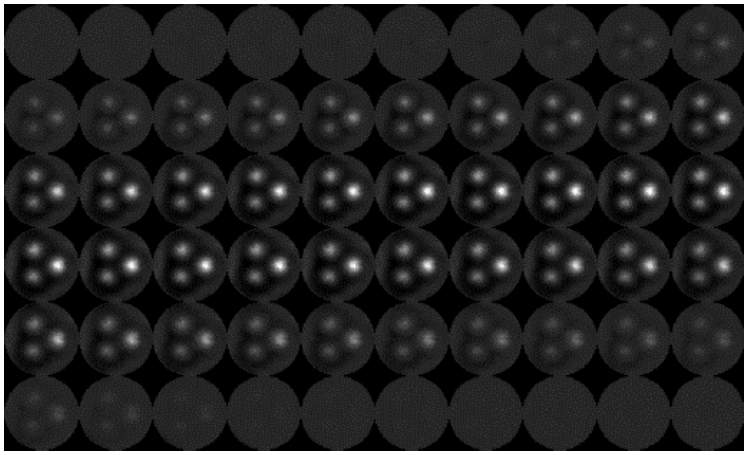

Figure 2.3.18 Images of the basic phantom - homogenous phantom data reconstructed on a 3D mesh, as if two data sets were acquired 20 mm apart vertically. Each slice represents 1 mm of phantom (so the central 60 mm of the solution are shown, 20 mm either side of each acquisition plane).

2.3.4.3.1 Image summary (3D basic - homogenous)

Phantom	Basic - Homogenous
Mesh	(3D) Cylindrical 140 mm tall, 31525 nodes, 21600 quadratic elements. (24 layers, 950 nodes in each)
Basis	Forward mesh - median filtered (2)
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iteration	15
Sources	32×2
Detectors per source	22
Datatypes	(raw) Mean
Calibration	Difference to homogenous
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	30 secs
Wavelength	800 nm

Table 2.3.11 Properties of images shown in Figure 2.3.18

2.3.4.3.2 Discussion (3D basic - homogenous)

The 3D slices show that the data has yielded a reconstruction of continuous rods. Although a later iteration is shown compared to Figure 2.3.17, slight separation of absorption and scatter is also seen. There is also less blurring in the μ_a image and the positions of the both scattering and absorbing inclusion is more consistent between the images. This result suggests that there is more information in data acquired on an axially symmetric phantom in a 2D plane than can be extracted using a simple 2D reconstruction.

2.3.5 Summary and discussion

This chapter has detailed early phantom studies performed using MONSTIR. One of the main issues explored was the effect of using a 2D image reconstruction for measured (and therefore 3D) data. Using 2D reconstructions we found that:

- 2D difference imaging yields better results than applying an ad-hoc 2D:3D correction to absolute data.
- Using 2D absolute or difference imaging requires assumptions about the discrete nature of inclusions and cannot account for the presence of out-of-plane structure.
- Quantitation in 2D (and 3D) images depends on the datatypes used and their respective accuracy and on the (non-linear) iteration chosen. In general, blurring affects peak values and quantitation is poor. Cross talk occurs between μ_a and μ'_s when only a single datatype is used, and also when noise tolerance or consistency between multiple datatypes is poor.
- 2D μ_a and μ'_s images may localise the same object in different positions.
- 2D reconstructions of data acquired on phantoms with axial symmetry are neither improved nor degraded compared to phantoms with discrete inclusions.

Comparison of 2D and 3D image reconstructions demonstrated that

- 3D image reconstruction is ill posed and very slow.
- Quantitation is not dramatically improved by using 3D image reconstructions, in fact very similar results are obtained if difference data are utilised to image a suitable object (Figure 2.3.5). Localisation and co-registration is generally better.
- 3D image reconstruction can be used to determine the μ_a and μ'_s distribution throughout a 3D volume from both data acquired in rings, and data where cross-plane measurements are included.
- Mesh generation for 3D image reconstructions is much more complicated than for 2D reconstructions (even for simple geometric shapes, although in particular for arbitrary geometries).

Improvements in computing speed and capacity, and mesh generation techniques are continually making 3D optical tomographic imaging a more reasonable process. However as a result of its current speed and complexity, many of the following chapters utilise 2D image reconstruction as a much quicker and simpler alternative. Where experimental data is utilised, 2D:3D corrections are applied or difference data are used. Where simulations are used to investigate the intrinsic properties of image reconstruction methods, 2D simulations are used. In this case, since data is simulated on a 2D mesh and subsequently reconstructed in 2D, errors relating to a mismatch do not occur. The results of such simulated experiments can be readily extended to the implications for 3D imaging.

Imaging results using multiple datatypes did not demonstrate many advantages over using just mean-time or mean-time and intensity. Results from chapter 2.2 demonstrated that certain datatypes can be unreliable and susceptible to systematic noise despite calibration. Problems can arise relating to uniqueness when too few datatypes are utilised (as explored in 2.5.6.1), however results are easier to compare and interpret if the same datatypes are utilised for each example. As a result, most reconstructions presented in the following chapters only utilise either mean-time, intensity or both intensity *and* mean-time.

Difference imaging relative to a homogenous phantom was shown to yield impressive results compared to those achievable with absolute calibrated data, both for 2D and 3D image reconstruction. In order to extend difference-imaging techniques to clinical applications however we need to consider the effect of having a poor match between the reference and the object being imaged. Also the effects of acquiring a reference measurement using the object itself (in some different state or at a different wavelength) must be explored.

2.4 Difference imaging with a heterogeneous reference

Chapter 2.1 introduced the concept of difference imaging (section 2.1.2.3). It was shown that calibration can be achieved in a number of ways, but that the use of a reference object is most effective in eliminating systematic errors from measured data. Chapter 2.3 demonstrated that, in limited cases, use of a reference object may cancel model errors including the effect of using a 2D reconstruction for 3D data. However the use of a reference object imposes a number of assumptions on image reconstructions. If a homogenous reference phantom is available, only the optical properties of the (well matched) phantom are required to produce absolute images of absorption and scatter. However practical limitations such as geometry and matching optical properties to clinical subjects mean that manufacture of a homogenous reference phantom is not always likely to be possible.

The other form of difference imaging is where the object being imaged itself is used as the reference phantom. Image data are acquired before and after a change in optical properties within an object that almost invariably has heterogeneous optical properties. This change may be invoked, be a deterioration or change over time, or be the difference in measurement at multiple wavelengths. However the resulting image reconstructions from such difference data can be incorrect. The problems are a direct result of the inherent non-linearity of optical tomographic image reconstruction. This chapter describes the perturbation approximations intrinsic in both linear and non-linear image reconstruction schemes. We then explore how these assumptions affect images of changes in absorption in the presence of a) heterogeneous background scatter, and b) heterogeneous background absorption using simulations. Comparisons of results using non-linear (TOAST) and linear image reconstruction techniques (see section 1.3.2.2), along with experimental studies are included throughout. The origin and dependence of the error are then investigated. We then present a method found to improve results which uses estimates of background structure from baseline images. This is shown to improve quantitation and object localisation in simple images. The significance of this error is discussed in context with similar studies and in relation to successful, reliable clinical imaging as well as implications for other forms of NIR spectroscopy.

Much of the validation of optical tomography has been performed on very simple phantoms with discrete inclusions ((Hebden *et al*, 1999), (Eda *et al*, 1999), (Schmitz *et al*, 2000), (Pogue *et al*, 1995)). Yet many researchers make the assumption that methods found to be effective on simple phantoms can be directly extended to clinical imaging subjects, in particular the use of difference data. The clinical environment imposes multiple constraints on

calibration, data acquisition and image reconstruction techniques. Subject variability and the complex structures involved require consideration of situations where structures do not conform to the assumptions required for imaging methods shown to be successful for simple phantom studies. This chapter demonstrates the limitations of using linear and non-linear image reconstructions, when modelling the exact initial state of the object is not possible. (Hillman *et al*, 2001c).

2.4.1 Linear and non-linear reconstruction schemes.

As introduced in chapter 1.3, image reconstruction in optical tomography is generally achieved by determining a distribution of optical properties within an object that corresponds to a set of measurements made on the object's surface. In our case the diffusion approximation (DA) to the radiative transfer equation is used as the basis of our light propagation model.

As shown in relation to derivation of a PMDF in section 1.3.2.1, the sensitivity of a measurement $y_{n,m}(x(r),\omega)$ made using source *m* and detector *n*, to changes in optical properties $\Delta x(r)$ can be evaluated from a Taylor series expansion:

$$y_{n,m}(x_1(r),\omega) = y_{n,m}(x_0(r),\omega) + \frac{\partial y_{n,m}(x_0(r),\omega)}{\partial x(r)}[x_1(r) - x_0(r)] + ...$$

where $$x(r) = \begin{pmatrix} \mu_a(r) \\ \kappa(r) \end{pmatrix}.$$

[2.4.1]

We can express the first derivative in terms of the Green's function solution to the DA, using the assumption of reciprocity, where $G_m^+(r,\omega)$ and $G_m(r,\omega)$ are the forward and adjoint Green's function solutions (Arridge, 1999):

$$\frac{\partial y_{n,m}(x_0(r),\omega)}{\partial x(r)} = J_{n,m}(x_0(r),\omega) = \begin{pmatrix} G_m^+(r,\omega)G_n(r,\omega) \\ \nabla G_m^+(r,\omega)\cdot\nabla G_n(r,\omega) \end{pmatrix}$$

[2.4.2]

For simple geometries and distributions of optical properties, $G(r,\omega)$ and thus $J_{m,n}(x_0(r),\omega)$ can be derived analytically (Arridge *et al*, 1992), (Arridge, 1995) using, for example the 3D infinite space, the Green's function solution to the DA ([1.3.18]). For more complex geometries, FEM can be used to derive a discrete representation of the Green's operator (Arridge *et al*, 1995a) and thus $J_{n,m}(x_0(r))$ (see section 1.3.1.3.1). Equation [2.4.2] can be generalised for any measurement type($M_{n,m}$), such as integrated intensity or mean-time by applying a suitable operator ($\mathcal{M}$[ ]) to both sides: $M_{n,m}(x_0(r)) = \mathcal{M}[y_{n,m}(x_0(r),t)]$, such that $\mathcal{M}[J_{n,m}(x_0(r),t)] = J_{n,m}^{M}(x_0(r))$ (see section 1.3.3). Evaluating and plotting $J_{n,m}^{M}(x(r))$ onto *r*

for a single source-detector pair, datatype *M* and state *x(r)*, will yield a photon measurement density function (PMDF) as shown in Figure 1.3.5 in section 1.3.2.1.

2.4.1.1 Linear inversion

Converting to datatypes and rearranging [2.4.1] (neglecting higher order terms) we get:

$$\mathbf{M}^{(1)}{}_{n,m} - \mathbf{M}^{(0)}{}_{n,m} = J_{n,m}{}^{M}(x_0(r))[x_1(r) - x_0(r)] \qquad [\ 2.4.3\ ]$$

where $\mathbf{M}$ denotes measured data. Perturbations in the optical properties of an object can be calculated from the differences between experimental measurements $\mathbf{M}_{n,m}{}^{(1)}$ - $\mathbf{M}_{n,m}{}^{(0)}$ using linear inversion of [2.4.3], assuming some given $x_0(r)$ is known to correspond to a reference state $\mathbf{M}_{n,m}{}^{(0)}$ and that $x_1(r)$ - $x_0(r)$ is small.

Figure 2.4.1 shows how [2.4.3] assumes that the change in measurement is linearly related to optical property changes, as long as the changes are small. The importance of calculating $J_{n,m}{}^{M}(x_0(r))$ as a function of the initial state $x_0(r)$ is also clear:

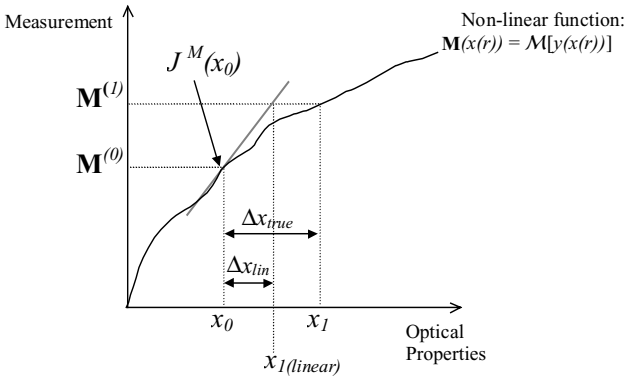

Figure 2.4.1 If the change in optical properties is small, and over this small range we can assume our function is linear, by evaluating the derivative $J_{n,m}{}^{M}(x_0(r))$ at x_0, we will calculate $x_{1(linear)}$ rather than the true x_1

Section 1.3.2.2 details the implementation of linear image reconstruction via Tikhonov regularisation, a one-step linear inversion technique that solves [2.4.3] for $x_1(r)$ - $x_0(r)$ using:

$$\Delta \mathbf{x} = [\mathbf{J}^T\mathbf{J} + \tilde{\lambda}\mathbf{I}]^{-1}\mathbf{J}^T\Delta\mathbf{M}, \qquad [\ 2.4.4\ ]$$

where $\Delta\mathbf{x}$ is the matrix of nodal optical properties sought, $\Delta\mathbf{M}$ is the matrix of changes in measurement, $\mathbf{J}$ is the sensitivity matrix calculated from $J^{M}_{n,m}(x_0(r))$ for a given numerical model, the datatype *M* and initial state $x_0(r)$. In the reconstructions presented in this chapter, $\tilde{\lambda} = \lambda F_{max}$ where F_{max} is the maximum main diagonal element value of the matrix $\mathbf{J}^T\mathbf{J}$ and λ is a regularisation parameter (Dehghani *et al*, 1999a).

2.4.1.2 Non-linear inversion

In cases where the perturbation is large, e.g. an initial distribution $x_0(r)$, close to $x_1(r)$, has not been measured, linear inversion of [2.4.3] will not yield a solution close to $x_1(r)$. Non-linear image reconstruction methods utilise an iterative procedure to allow updates to the

derivative $J^{M}_{n,m}(x(r))$ based on the previously derived perturbations $\Delta x(r)$. So starting from an initial guess of the optical properties $x_g(r)$ and simulated measurement values calculated from that guess $M^{(g)}_{n,m}$, the optical properties $x_1(r)$ represented by a single absolute measurement $\mathbf{M}_{n,m}^{(1)}$ can be found using:

$$\mathbf{M}^{(1)}{}_{n,m} - M^{(g)}{}_{n,m} = J_{n,m}{}^{M}(x_g(r))[x_{g+1}(r) - x_g(r)]$$
$$\mathbf{M}^{(1)}{}_{n,m} - M^{(g+1)}{}_{n,m} = J_{n,m}{}^{M}(x_{g+1}(r))[x_{g+2}(r) - x_{g+1}(r)].....$$

[2.4.5]

until $x_{g+N}(r)$ minimises $\sum_{n,m}\left(\frac{\mathbf{M}^{(1)}{}_{n,m} - M^{(g+N)}{}_{n,m}}{\sigma_{n,m}}\right)^2$,

[2.4.6]

where $\sigma_{n,m}$ is the standard deviation on the measurement (often $\sigma_{n,m}$ is assumed to = $\mathbf{M}^{(1)}{}_{n,m}$ if the standard deviation is not available). Figure 2.4.2 demonstrates that by repeatedly evaluating the derivative for each step, it should be possible to iteratively reach a true solution $x_1(r)$, and that the starting point (or initial guess) does not have to be as close to $x_1(r)$ as is required for linear inversion (Figure 2.4.1).

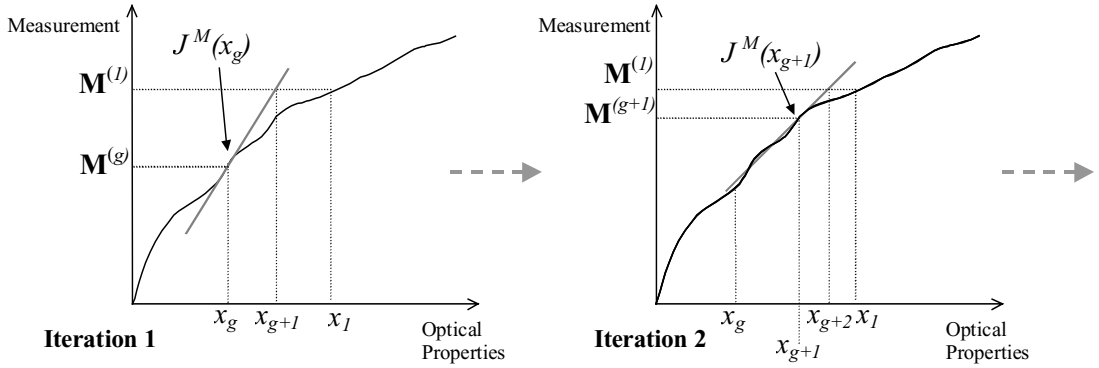

Figure 2.4.2 Starting from an initial guess means that a reference measurement is not required, and by re-evaluating the derivative at each guess of x *it is possible to eventually arrive at a true value of* x_1 *corresponding to* $\boldsymbol{M}^{(1)}$.

If a second data set $\mathbf{M}_{n,m}^{(0)}$ is used (e.g. from a homogenous reference phantom for calibration see section 2.1.4.3) then its properties must be known (as is the requirement for the linear inversion method) such that [2.4.7] is the starting equation used for iterative image reconstruction.

$$\mathbf{M}^{(1)}{}_{n,m} - \mathbf{M}^{(0)}{}_{n,m} + M^{(0)}{}_{n,m} - M^{(g)}{}_{n,m} = J_{n,m}{}^{M}(x_g(r))[x_{g+1}(r) - x_g(r)] .....$$

[2.4.7]

Note that if the model is perfect then $M^{(0)}{}_{n,m} = \mathbf{M}^{(0)}{}_{n,m}$ and [2.4.3] and [2.4.7] are equivalent. Yet unlike [2.4.3], if the model is not perfect, this method will cancel model and systematic errors to some extent..

The non-linear approach allows larger perturbations to be adequately modelled, since it accounts for the effect of inhomogeneities in optical properties on the accuracy of the

sensitivity matrix for the measurements made. This is because with each iteration, $J_{n,m}$ is modified to correspond to the structure indicated by the previous iteration. This also allows 'absolute' images to be derived without the need for a well-characterised reference medium, whose optical properties are close enough to those of the object being imaged to comply with [2.4.3] (assuming that the absolute calibration is sufficiently accurate).

2.4.2 Baseline imaging (homogenous reference measurement)

Imaging relative to a known baseline can be achieved via linear (if the perturbations are small) or non-linear reconstruction techniques. Such images were presented in chapter 2.3, for example Figure 2.3.11 showed a reconstruction of data acquired on two conical resin phantoms. One contained three 10 mm high, 10 mm diameter cylinders: one with twice background μ_a, one with twice background μ'_s and one with twice both μ_a and μ'_s. The other conical phantom was made from the same batch of resin, but was homogeneous.

Since the phantom was manufactured according to a recipe (Firbank *et al*, 1995a), the background optical properties are fairly well known, and the perturbations are discrete and quite small, so in this simple case $x_0(r)$ is approximately known, and $\Delta x(r)$ is small. Both linear and non-linear image reconstruction methods could be expected to provide good results (if adequate scatter - absorption separation is achieved).

Much of the validation of optical tomographic methods has utilised phantoms and simulations of this type, particularly since a further advantage of using experimental reference measurements is that they aid instrument calibration in most cases (see section 2.1.2.3).

2.4.3 Difference imaging (using an uncharacterised reference)

As introduced in section 2.1.4.4 a common approach to imaging is to assess physiological changes in a live subject by initiating or monitoring some change in optical properties. This is both potentially clinically useful, and provides a $\Delta\mathbf{M}_{n,m}$ without the use of an independent reference phantom. A reference phantom for the brain (for example) would have to have optical properties close enough to the subject for linear inversion to be valid. In the clinical environment and given inter-patient variability this is unlikely to be practical (see 2.7.2.7). However, if the head itself is used (at some other state or wavelength) as the equivalent of a reference phantom, no value of $x_0(r)$ will be available. Even for non-linear image production, if *only* $\Delta\mathbf{M}_{n,m} = \mathbf{M}^{(1)}_{n,m} - \mathbf{M}^{(0)}_{n,m}$ is measured, information about $x_0(r)$ is lost. In both cases, some initial guess $x_i(r)$ will have to be used to generate $J^M_{n,m}(x_i(r))$, as shown in [2.4.8].

$$\mathbf{M}^{(1)}{}_{n,m} - \mathbf{M}^{(0)}{}_{n,m} = J^{M}{}_{n,m}(x_i(r))[x_{1\,difference}(r) - x_i(r)]$$

[2.4.8]

Figure 2.4.3 shows why if $x_i(r)$ C $x_0(r)$, a value of $\Delta x(r)$ found using [2.4.8] would not be the solution to [2.4.3] since only knowing $\Delta\mathbf{M}$, the difference $x_{1\ difference}$ - x_i does not necessarily equal x_1 - x_0.

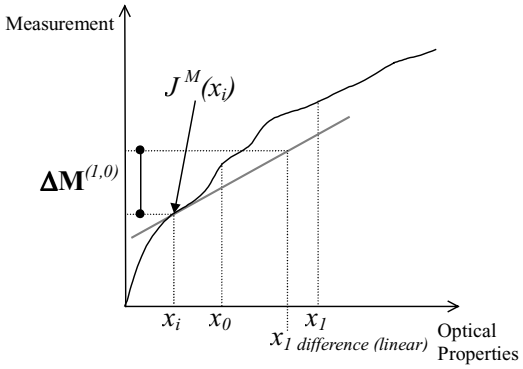

Figure 2.4.3 If only $\Delta\mathbf{M}^{(1,0)}$ (between states x_1 and x_0) is known, and $\mathbf{M}^{(0)}$ and x_0 are not known, an estimate x_i used as the starting condition may result in the solution $\Delta x_{difference} = x_{1\ difference} - x_i$ being far from the true $\Delta x = x_1 - x_0$.

Solving [2.4.8] using a non-linear scheme would involve iteratively solving:

$$\mathbf{M}^{(1)}{}_{n,m} - \mathbf{M}^{(0)}{}_{n,m} + M^{(i)}{}_{n,m} - M^{(g)}{}_{n,m} = J^{M}{}_{n,m}(x_g(r))[x_{g+1}(r) - x_i(r)]$$

[2.4.9]

such that $x_{g+N}(r)$ minimises $$\sum_{n,m}\left(\frac{\mathrm{M}^{(1)}{}_{n,m} - \mathrm{M}^{(0)}{}_{n,m} + M^{(i)}{}_{n,m} - M^{(g+N)}{}_{n,m}}{\sigma_{n,m}}\right)^2 .$$

[2.4.10]

The problem of not knowing the initial conditions means that the sensitivity matrix calculated using $J^{M}{}_{n,m}(x_i(r))$ will not truly represent the data set, and that the perturbation in $x_i(r)$ which minimises [2.4.10] may not represent the true change in optical properties. $x_0(r)$ is contributed to by both background scatter values and background absorption values.

Effectively we are questioning the applicability of the assumption illustrated in Figure 2.4.4. If we measure the change in data between one state and another, can we reconstruct only that change if we have no knowledge of the structure that has remained constant in the object being imaged:

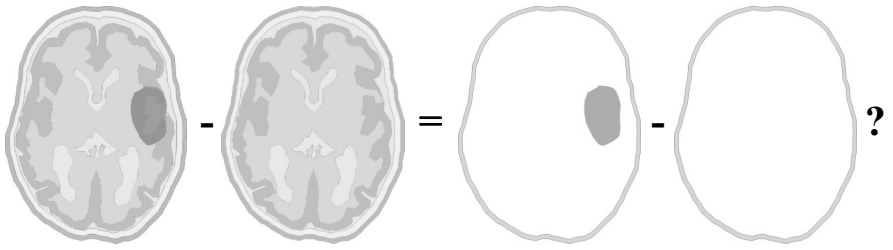

Figure 2.4.4 Perturbation assumptions for linear and non-linear image reconstruction mean that measurements made on an object with changing optical properties may not correspond to images of the change alone reconstructed without prior knowledge of the constant structure within the original object.

Below we investigate the effects of constant background heterogeneous scatter and then heterogeneous absorption on image reconstructions using changes in measurements corresponding to changes in absorption.

2.4.3.1 Background scatter

In cases where only changes in absorption are being considered, it is often assumed that lack of knowledge of the scatter coefficient and its spatial variability is of little consequence. However, clearly the diffusion approximation ([1.3.10]) has a strong dependence on scatter values, which will affect derived absorption values if a solution is sought with an incorrect scatter distribution. This would be the case if absorption changes in the brain were being examined without any knowledge of the structure of the skull, brain, and other gross features that will affect light propagation.

Figure 2.4.5 shows the ratio of the photon density within an object with scattering structure that is 150% of the background μ'_s, compared to the photon density within a homogeneous object. Plots were derived using the TOAST forward model (Schweiger *et al*, 1996), solving for photon density at each node in a 2D mesh in response to a given source position. The mesh had 1459 nodes and 2808 linear elements and was 70 mm in diameter.

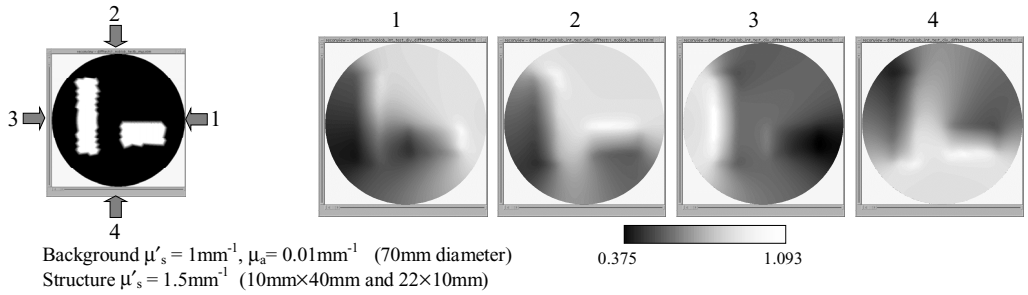

Figure 2.4.5 Photon density maps showing the ratio of photon density within a 2D object with scatter inhomogeneities, relative to the same object without the inhomogeneities, for four different source positions.

Figure 2.4.5 shows that there are regions in the object with scattering structure that are sampled by only a fraction of the photons that would otherwise have contributed to a measurement.

Two cases are considered in Figure 2.4.6: Case 1 (top row) shows the ratio of the photon density in a simulated example of a μ'_s homogeneous 2D model with two μ_a perturbations ($2\times\mu_a$ background, top: 11mm, bottom: 9mm average diameter), relative to the same model without the perturbations. Case 2 (middle row) shows the ratio between two models with a constant scatter structure ($1.5\times\mu'_s$ background as shown in Figure 2.4.5), with and without the two μ_a perturbations. These plots represent the case where we are looking at absorption changes within an object with heterogeneous, constant scatter properties. While the ratio of photon densities for the case where background scatter is homogenous look very similar to

those for heterogeneous scatter, the ratio of the two cases reveals structure corresponding to the constant scatter heterogeneities. This shows that there is > 10% change in the photon density ratio between the cases with inhomogeneous and homogeneous background scatter.

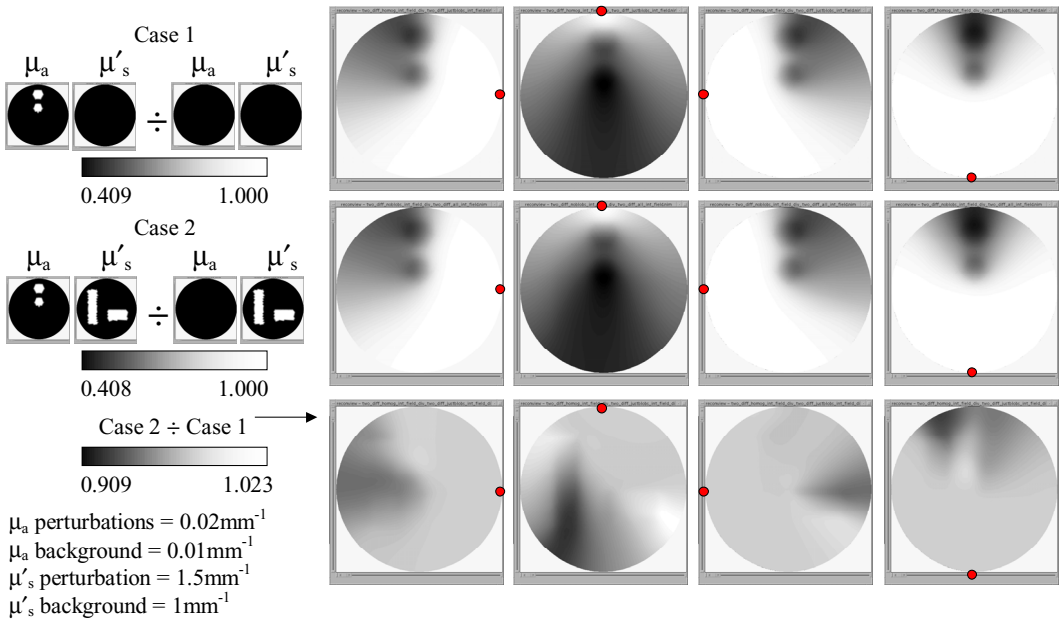

Figure 2.4.6 (Top) Case 1: The ratio of photon densities between an object with two absorbing perturbations relative to homogeneous, (Middle) Case 2: the ratio of photon densities between an object with two absorbing perturbations relative to one without, in the presence of (constant) scattering inhomogeneity, (Bottom) the ratio: (case 1)/(case 2) (The dot shows the source position for each image)

To assess whether lack of knowledge about the scattering structure will affect images reconstructed from the changes in measurement, we can investigate what happens to images if we assume $x_i(r)$ is homogeneous, which is true for case 1, but not for case 2. Figure 2.4.8 shows the 50th iterations of non-linear FEM image reconstructions of (noiseless) simulated mean-time data (top row) and intensity data (bottom row) of the two cases shown in Figure 2.4.6. Both reconstructions assumed that $x_i(r)$ was homogeneous; μ_a = 0.01mm^{-1} and μ'_s = 1mm^{-1}, the scatter solution was held fixed at this value while μ_a was allowed to vary. 704 data values were simulated for a planar tomographic distribution of 32 evenly spaced sources illuminating a fan beam of 22 out of 32 evenly spaced detectors.

Figure 2.4.7 shows corresponding one-step linear image reconstructions, derived by inverting [2.4.3] using λ = 0.01 as the regularisation parameter in [2.4.4]. The sensitivity matrices for mean-time and intensity were evaluated using FEM-derived values of $J^M_{n,m}(x_i(r))$ with $x_i(r)$ homogeneous (as for the non-linear reconstruction). Only changes in μ_a were considered as if μ'_s were constant (which it is), i.e. we solve equations of the form:

$$[\Delta Mean] = [J^{\mu_a}{}_{Mean}][\Delta\mu_a] \qquad [2.4.11]$$

$$[\Delta Intensity] = [J^{\mu_a}{}_{Intensity}][\Delta\mu_a] \qquad [2.4.12]$$

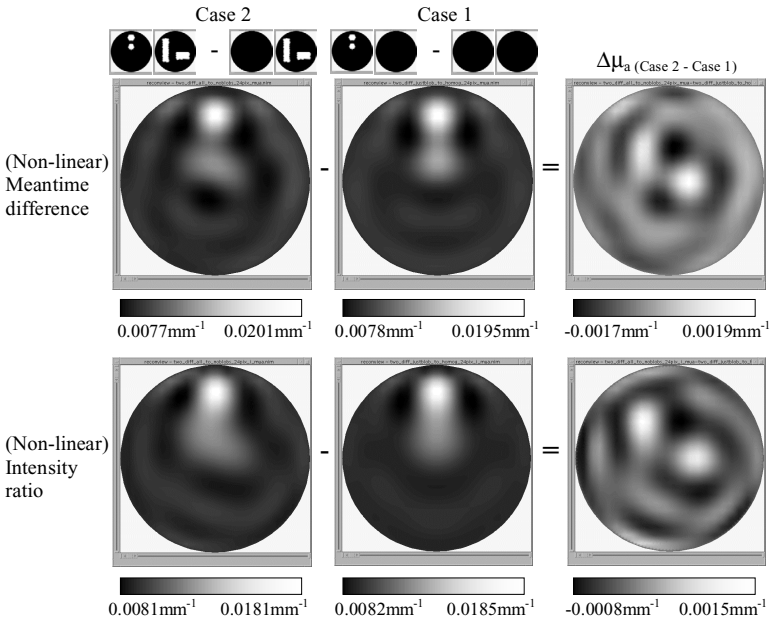

*Figure 2.4.7 Images reconstructed using a **non-linear** solver from Δ<t> (top) and intensity ratio (bottom), for objects with heterogeneous background μ'_s and two μ_a perturbations, relative to the same object without the μ_a perturbations (Case 2), for an object with homogenous μ'_s, with and without the μ_a perturbations (Case 1), and the difference between the case 1 and case 2 images ($\Delta\mu_{a\ (Case\ 2\ -\ Case\ 1)}$). Case 1 and case 2 reconstructions assumed no knowledge of background μ'_s structure, using x_0 (r) = homogeneous μ_a, μ'_s.*

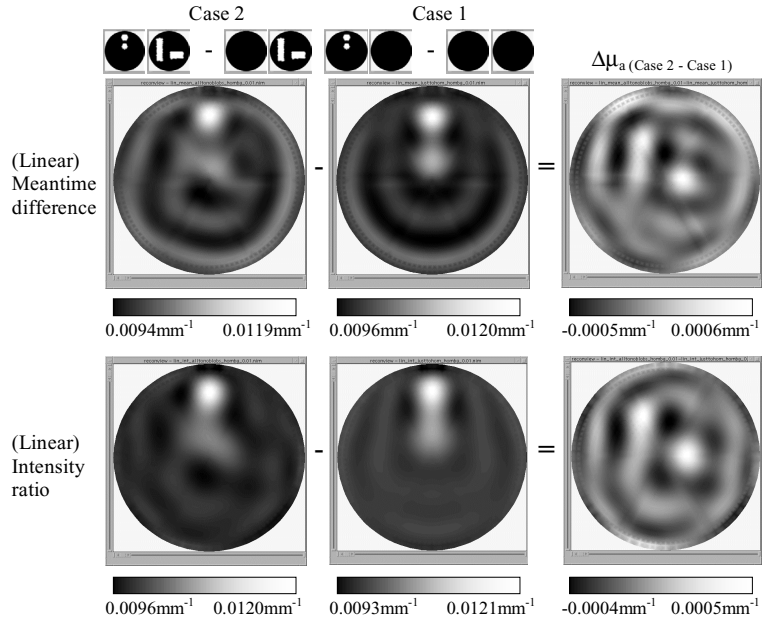

*Figure 2.4.8 Images reconstructed using a **linear** solver from Δ<t> (top) and intensity ratio (bottom), for objects with heterogeneous background μ'_s and two μ_a perturbations, relative to the same object without the μ_a perturbations (Case 2), for an object with homogenous μ'_s, with and without the μ_a perturbations (Case 1), and the difference between the case 1 and case 2 images ($\Delta\mu_{a\ (Case\ 2\ -\ Case\ 1)}$). Case 1 and case 2 reconstructions assumed no knowledge of background μ'_s structure, using x_0 (r) = homogeneous μ_a, μ'_s. (compare with Figure 2.4.7)*

The linear images in Figure 2.4.8 are displayed as absolute μ_a values $x_i(r)+\Delta x(r)$, for comparison with the non-linear results. Clearly, for both types of reconstruction, the images derived from the difference in measurements for case 2 have been distorted by the inapplicability of the approximation that $x_i(r) = x_0(r)$. The differences between the μ_a values

found for cases 1 and 2 are also shown, demonstrating that values vary by up to 36% for mean-time data and 23% for intensity between the two cases for the non-linear reconstruction and 11% and 9% respectively for the linear reconstruction.

2.4.3.2 Background Absorption

Figure 2.4.6 demonstrated that absorption features also affect the areas sampled by photons during a measurement, since they also appear to 'cast shadows'. Therefore background absorption structure that is unknown will affect the ability to reconstruct images from measurements taken before and after a change, or indeed at different wavelengths. In order to test the effects of an unknown background absorbing structure systematically, three cases were simulated as shown in Figure 2.4.9. Case A considers the difference in mean-time and intensity measurements between a homogeneous (μ_a = 0.01mm^{-1} and μ'_s = 1mm^{-1}, 70mm diameter) 2D simulation, and a model with two equal ~10mm inhomogeneities with μ_a = 0.02mm^{-1}. In case B, a region on one half of the mesh is given a μ_a of 0.012mm^{-1}. Data simulated on the mesh with this extra region are compared to a model with the same region of increased μ_a, with the inhomogeneity in that region having a μ_a of 0.022mm^{-1} (such that $\Delta x(r)$ for case A = $\Delta x(r)$ for case B). In case C this background region has μ_a = 0.015mm^{-1}, and a corresponding inhomogeneity with μ_a = 0.025mm^{-1} (again so that $\Delta x(r)$ is the same as for cases A and B).

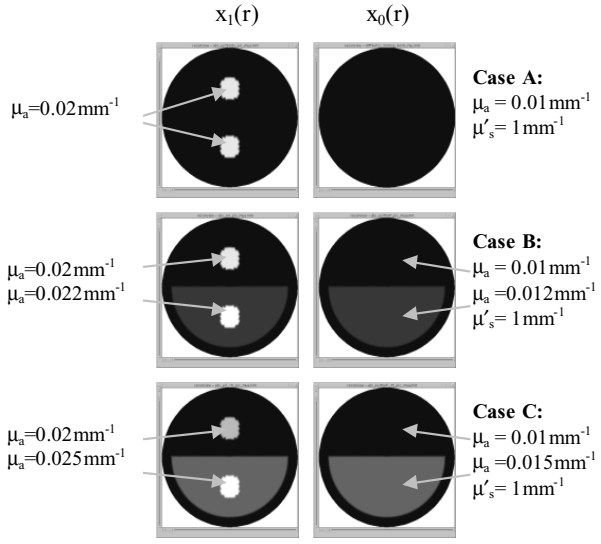

Figure 2.4.9 Three cases simulated to explore the effect of constant background absorption structure on difference images reconstructed from the difference between data simulated for $x_1(r)$-$x_0(r)$.

Figure 2.4.10 shows images reconstructed using a non-linear algorithm assuming no knowledge of the constant background absorption variation.

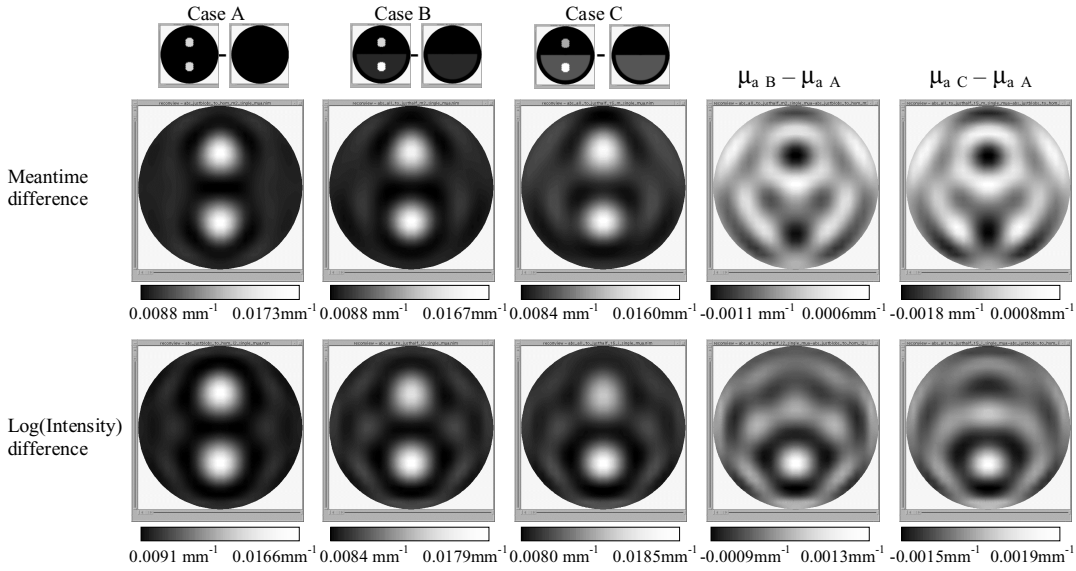

Figure 2.4.10 Non-linear images reconstructed from mean-time (top) and intensity (bottom) data simulated for: Cases A, B and C as described in Figure 2.4.9. Case A represents the correct method, and so residuals of B and C relative to A are also shown (far right).

Figure 2.4.11 shows nodal values of μ_a along the cross section of each image (along the line joining the two small absorbing regions). The effect of the unknown background can be seen to affect the positions and relative magnitude of the two perturbations in the images. Intensity and mean-time are clearly differently affected by the fact that $x_i(r) \neq x_0(r)$ in the image reconstructions.

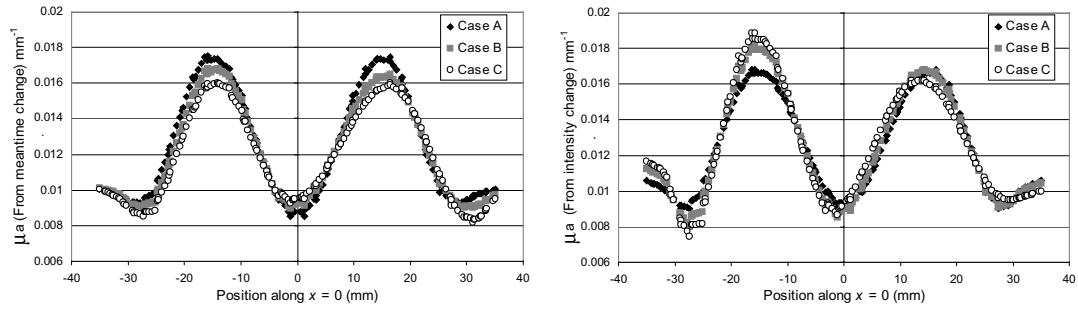

Figure 2.4.11 Cross sectional nodal values of the images in Figure 2.4.10: for mean-time images (left) and intensity images (right).

Figure 2.4.12 shows corresponding linear reconstructions ($\lambda = 0.01$) of cases A to C. Figure 2.4.13 shows corresponding cross sections of the images. The effects on images are quite different between non-linear and linear reconstructions. The FEM mesh used throughout had 3781 nodes and 7392 linear elements, and the same 704 source-detector combinations as described in section 2.4.3.1 were used.

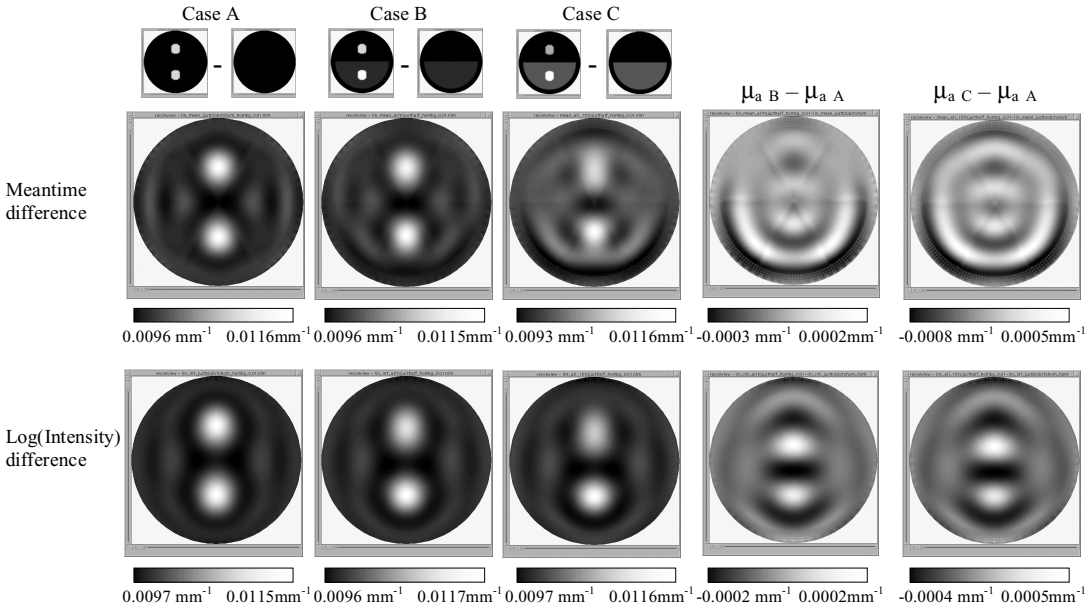

Figure 2.4.12 Linear image reconstructions of cases A to C shown in Figure 2.4.10a, the scales have had 0.01 mm^{-1} added to represent the absolute values indicated by the reconstructed $\Delta x(r)$.

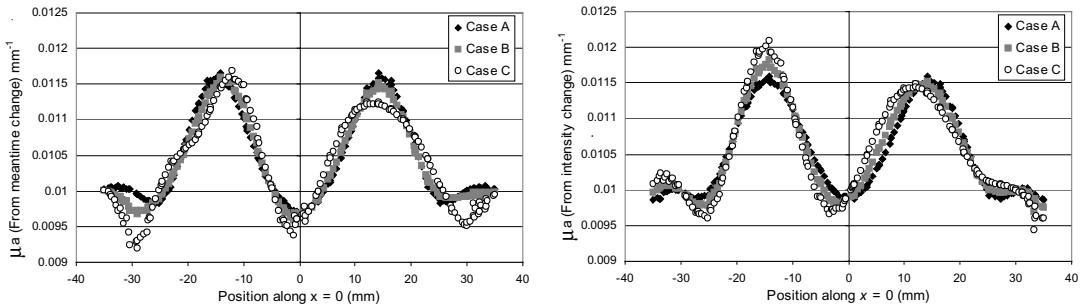

Figure 2.4.13, Cross sectional plots of node values of the images shown in Figure 2.4.12a (+ 0.01mm^{-1}), mean-time images (left) intensity (right).

2.4.4 Experimental investigation

In order to examine the effects of unknown background structure on wavelength-difference reconstructions, two phantoms were constructed that contained wavelength-dependent absorbing features. As with the phantoms described in 2.3 the phantoms were manufactured using epoxy resin. The absorption spectra of the two dyes used in their manufacture used are shown in Figure 2.4.14.

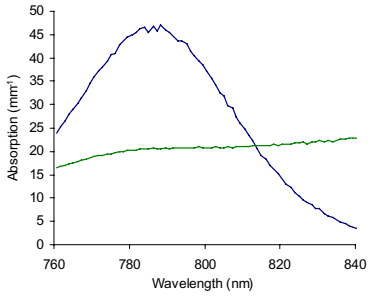

Figure 2.4.14 The absorption spectra of the two dyes used to make wavelength-difference phantoms.

2.4.4.1 Wavelength-difference phantom I

The first wavelength-difference phantom (see Figure 2.4.15) was composed mainly of resin containing the dye with the fairly flat spectrum (such that μ_a = 0.01 mm^{-1} ± 0.002 mm^{-1} at 790nm and 820nm). Its structure consisted of 5 small, discrete (5mm diameter) absorbing rods each containing the dye with the stronger spectral dependence such that they had a ~ 3xμ_a contrast between 820nm and 790nm. The scatter properties of this phantom were μ'_s = 1 mm^{-1} ± 0.2 mm^{-1} throughout and virtually independent of wavelength over this range.

Data were acquired on the phantom with 32 sources and detectors held in a planar ring. 22 detectors per source were illuminated. Data were acquired first with the Ti:sapphire laser tuned to 820 nm. The laser was then re-tuned to 790 nm and data were acquired with the same acquisition protocol as the data at 820nm. Calibration measurements were acquired at both wavelengths using the tools described in 2.1.2.1 although they were only used to determine the window over which to calculate raw mean-time and raw variance from both measured data sets (see 2.1.3.2.3). TOAST (non-linear) was used to reconstruct the difference in the datatypes between the two wavelengths on a 2D mesh. The initial guess was that $\mu_{a,\ initial}$ = 0.01mm^{-1} (homogenous) and μ'_s =1mm^{-1}. Only μ_a was allowed to vary in the reconstruction, μ'_s was assumed to be homogenous, constant and equal to 1 mm^{-1}. The 30th iteration of the resulting $\Delta\mu_{a\ (790nm-820nm)} + \mu_{a,\ initial}$ image is shown in Figure 2.4.15.

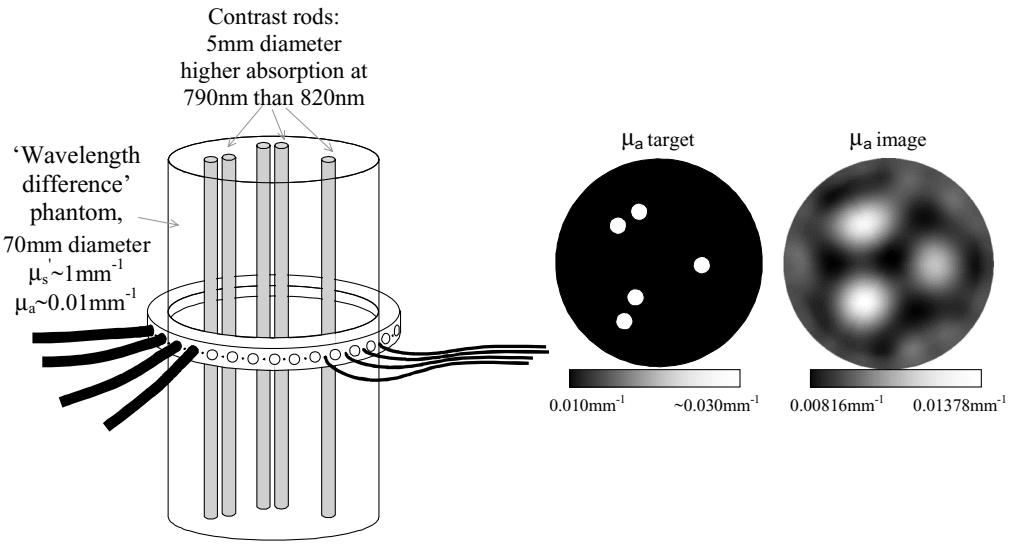

Figure 2.4.15 The wavelength-difference phantom has 5 rods, all with ~ 3× contrast at 790 nm. The μ_a image (right) was reconstructed from the difference in mean-times and variances between acquisitions at 790 nm and 820 nm.

The fact that the phantom contains rods may have affected image quality slightly (see 2.3.4), similarly use of a 2D mesh can cause artefact although less-so with difference imaging (see 2.3.1). The image shown in Figure 2.4.15 is certainly of equivalent quality to previous phantom studies shown in chapter 2.3 where a homogenous reference measurement was used.

However, the 790nm and 800nm states of this phantom are very similar, the approximation that the initial guess is homogenous is clearly sufficient. This is not the case for the more complex phantom described below.

2.4.4.1.1 Image summary (wavelength-difference phantom I)

The images shown in Figure 2.4.15 were published in (Hillman *et al,* 2000b) (submitted November 1999).

Phantom	Wavelength-difference - Homogenous
Mesh	Circular 3781 nodes, 7392 linear elements (2D)
Basis	Forward mesh - median filtered (2)
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iteration	30
Sources	32
Detectors per source	22
Datatypes	(raw) Mean and Cvar
Calibration	Difference between 790 nm - 820 nm data
2d 3d correction applied?	No
Simultaneous	μ_a only
Acquisition time (per source)	30 secs
Wavelength	790nm and 820nm (Ti:Sapphire)

Table 2.4.1 Properties of image shown in Figure 2.4.15

2.4.4.2 Wavelength-difference phantom (II)

The second wavelength-difference phantom was designed and constructed by Felipe Gonzalez and contained much larger absorbing regions, similar to the simulations in 2.4.3.2. Figure 2.4.16 shows the structure of the phantom. The dye with the flat spectrum (Figure 2.4.14) was used for half of the structure and the dye with the sharper spectral dependence used to create regions with μ_a lower than the rest of the phantom at 815nm but higher at 780nm.

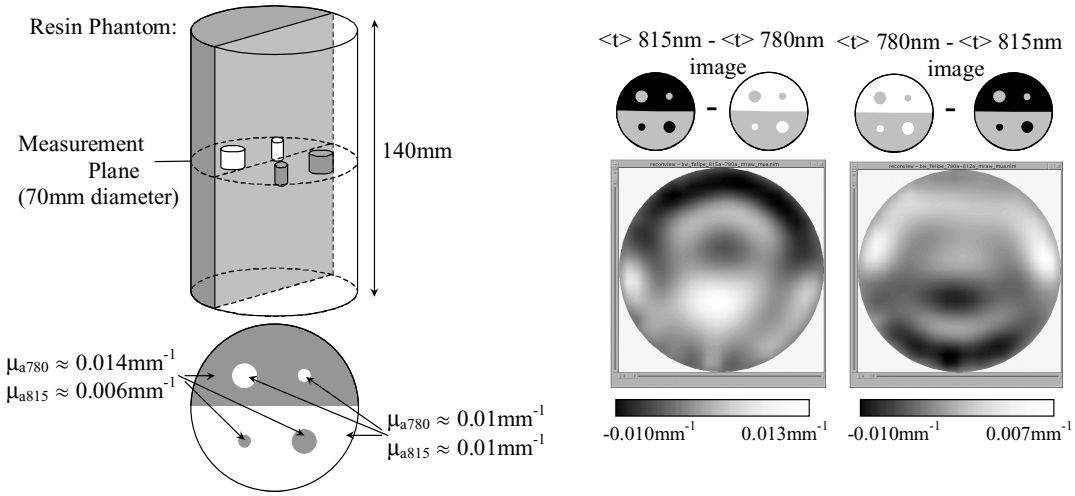

Figure 2.4.16 Images produced from the differences in mean-time data collected at two wavelengths on a phantom with large regions of wavelength dependent absorption (and homogenous scatter). Starting values of homogenous μ_a=0.01 mm^{-1} and μ'_s= 1 mm^{-1} were used, assuming no prior knowledge of the regional absorption parameters of the phantom.

Data acquisition was the same as for the first wavelength-difference phantom; 32×22 measurements were made in the plane indicated with the Ti:Sapphire laser tuned to 815 nm and then 780 nm. TOAST was used to reconstruct the differences in the raw mean-time (only) measurements between the two wavelengths with $\mu_{a,\ initial} = 0.01mm^{-1}$ and $\mu'_s = 1mm^{-1}$ (homogenous) assuming no prior knowledge of the regional absorption structure of the phantom. Figure 2.4.16 shows the resulting $\Delta\mu_a$ images after 30 iterations, μ'_s was kept fixed at $1mm^{-1}$ since the phantom had no internal scatter heterogeneity.

2.4.4.2.1 Image summary

Phantom	Wavelength-difference II
Mesh	Circular 3781 nodes, 7392 linear elements (2D)
Basis	Pixel basis 24 x 24
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iteration	30
Sources	32
Detectors per source	22
Datatypes	(raw) Mean
Calibration	Difference between 780 nm and 815 nm data
2d 3d correction applied?	No
Simultaneous	μ_a only
Acquisition time (per source)	15 secs
Wavelength	780nm and 815nm (Ti:Sapphire)

Table 2.4.2 Properties of images shown in Figure 2.4.16

The images of the second phantom are more distorted than those of the first phantom with simple discrete inclusions. Since use of a 2D mesh in this case may have adversely affected images, since large changes occur out-of-plane (see 2.3.1.4), a 2D simulation was created to determine to origin of the distortion in the wavelength-difference images of phantom II.

2.4.4.3 Simulation of wavelength-difference phantom II

2D simulation based on the geometry and wavelength dependence of the second wavelength-difference phantom was created. Figure 2.4.17 shows the meshes used with TOAST to simulate data. Both intensity (only) changes and mean-time (only) changes were reconstructed (using TOAST e.g. non-linear) using $\mu_{a,\ initial} = 0.01mm^{-1}$ and $\mu'_s = 1mm^{-1}$ as above. The resulting 30^{th} iteration $\Delta\mu_a$ images are shown in Figure 2.4.17. For comparison, the results of subtracting two baseline images (see 2.4.2) reconstructed from simulated absolute mean-time are also shown.

The data is noise-free and genuinely 2D. While good results are produced by subtracting the two absolute images, reconstructions of only changes in data produce distortion of the same form as seen in the experimental results in Figure 2.4.16.

These findings demonstrate that the likely cause of the distortion in the experimental images of the second wavelength-difference phantom are indeed related to the lack of knowledge of the initial absorption structure as suggested in 2.4.3.2.

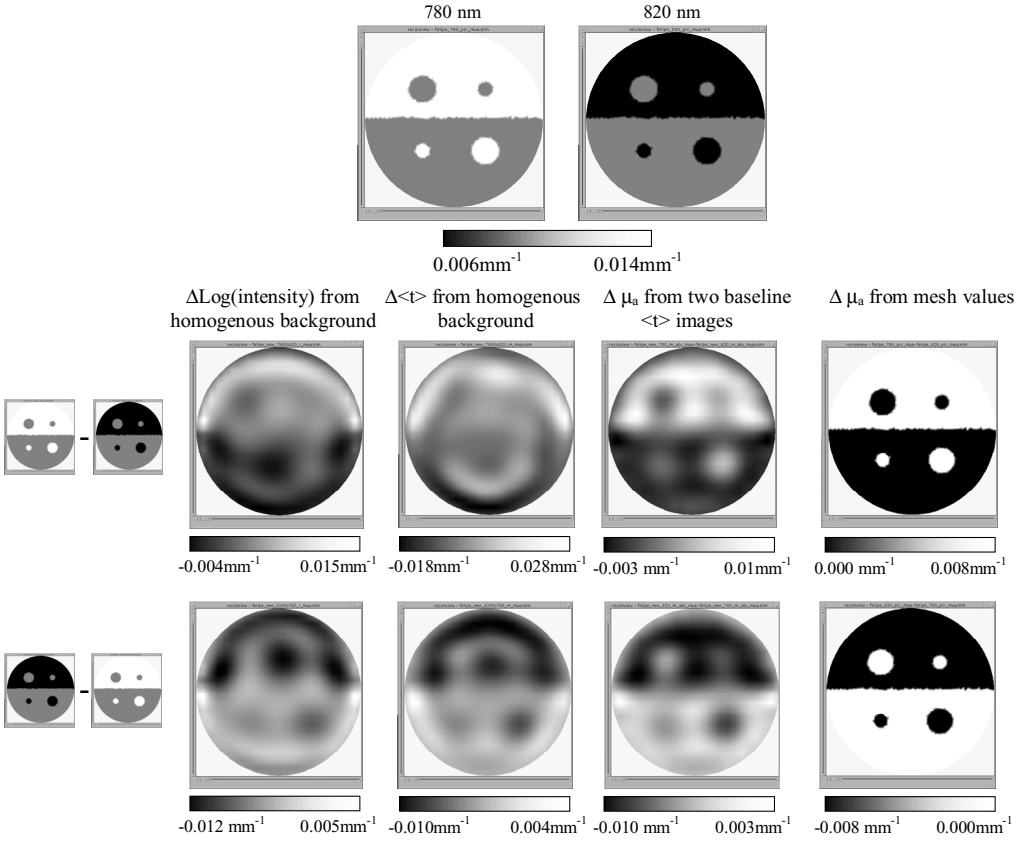

Figure 2.4.17 Image reconstructions from simulated values based on the wavelength dependent phantom shown in Figure 2.4.16. (Top) are the meshes used to generate data, (Far left) are image reconstructed from changes in intensity and (Middle left) changes in mean-time between the two states, starting from homogenous distribution. (Middle right) shows the result of subtracting two baseline μ_a images reconstructed from mean-time at each state, (Far right) shows the mesh values subtracted from each other

2.4.5 The origin of the error

Sections 2.4.3 and 2.4.4 demonstrated that a lack of knowledge about the initial state $x_0(r)$ results in both localisation and quantitation errors when measurements of $\Delta x(r)$ are used to reconstruct difference images. The manifestation of these errors has been shown to depend on the type of reconstruction used (linear or non-linear), the type of data utilised in the reconstruction (mean-time or intensity), and the nature of the unknown structure (scatter, absorption and their magnitude). We can see the origin of these errors by examining the Taylor series expansion of the Jacobian assuming that the initial state chosen $x_i(r)$ is equal to the true initial state $x_0(r)$ plus some error $\delta x(r)$.

$$J^{M}{}_{n,m}(x_i(r)) = J^{M}{}_{n,m}(x_0(r) + \delta x(r))$$
$$= J^{M}{}_{n,m}(x_0(r)) + \frac{\partial J^{M}{}_{n,m}(x_0(r))}{\partial x}[\delta x(r)] + \frac{1}{2}\frac{\partial^2 J^{M}{}_{n,m}(x_0(r))}{\partial x^2}[\delta x(r)^2] + \ldots$$

[2.4.13]

By substituting the first two terms of [2.4.13] into [2.4.8] we can see that the error introduced by estimating the initial optical properties is dependent (to first order) on the rate

of change of the Jacobian (which in turn will depend on the datatype and whether μ_a or μ'_s is being considered) and also on the magnitude of both the error in the estimate $\delta x(r)$ and the size of the change in optical properties being reconstructed ($\Delta x(r)$). Its effect could be considered as an error on the measurement:

$$\mathbf{M}^{(1)}{}_{n,m} - \mathbf{M}^{(0)}{}_{n,m} - \frac{\partial J^{M}{}_{n,m}(x_0(r))}{\partial x}[\delta x(r)][x_1(r) - x_0(r)] = J^{M}{}_{n,m}(x_0(r))[x_1(r) - x_0(r)] \quad [\ 2.4.14\ ]$$

Comparing [2.4.14] with [2.4.8] and [2.4.9] the origin of the differences in the error with linear and non-linear reconstruction schemes is also clear since the extra term will affect the convergence of the non-linear solution.

The first differential of the Jacobian denotes the amount that PMDFs will change with varying μ_a and μ'_s. This can be demonstrated by comparing PMDFs derived on a homogenous object to PMDFs derived for an object with μ_a inclusions as shown in Figure 2.4.18.

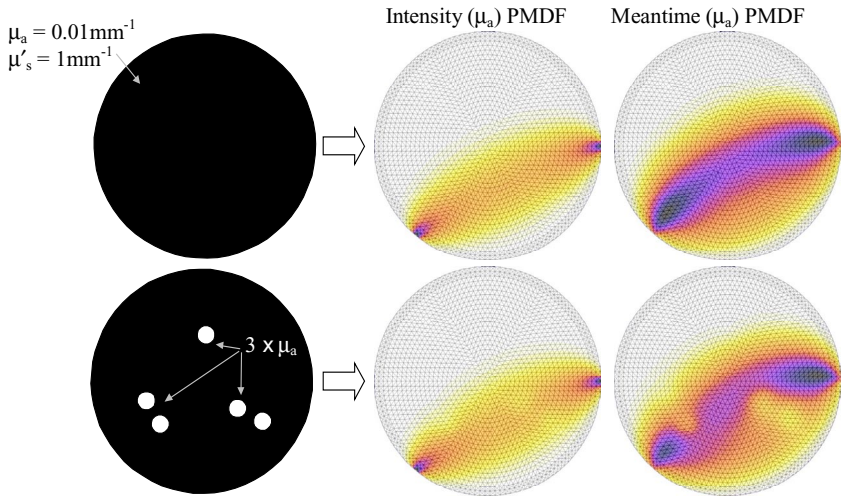

Figure 2.4.18 (Top) PMDFs on a homogenous 2D mesh, (bottom) PMDFs on an object with μ_a structure.

These PMDFs demonstrate similar effects to the intensity field plots shown in Figure 2.4.5 and Figure 2.4.6. The small absorbing features can be seen to affect the PMDFs, in particular the PMDF for mean-time and absorption where the sensitivity of mean-time measurements to changes in the regions obscured by the absorbing structure will be lower (all values are negative so hotter colours = numbers closer to zero). The lower dependence of the intensity μ_a PMDF on the absorbing features illustrates why the extra term in [2.4.14] will affect different reconstructions to different extents (scattering structure and other datatypes will yield different PMDF dependencies).

Thus the perturbation-related limitations demonstrated by the sets of simulations above relate to the fact that the PMDFs used in image reconstruction (for both linear and non-linear reconstructions) will depend on the static optical properties of the object being imaged.

Correct localisation and quantitation of changes in optical properties cannot be fully achieved if the correct PMDF cannot be derived since the background structure is unknown. This is analogous to the 2D:3D mismatch described in section 2.3.1.1 where the wrong (2D) PMDFs were being assumed to described data acquired in 3D.

2.4.6 Overcoming perturbation-related limitations

The reasons for wanting to reconstruct images from changes in measurement on a heterogeneous object were:

1. The ease of calibration offered by looking at *changes* in measurement at a single location such that surface coupling and time-varying instabilities will cancel.
2. The impracticality of creating homogenous or well-characterised reference phantoms that are sufficiently close to the optical properties of the object being imaged to allow linear inversion of [2.4.3] or accurate cancellation between model and measurement in [2.4.9].

In order to overcome limitations imposed by perturbation assumptions we could revert to using only absolute measurements. This would require development of a set of measurements that can be easily calibrated through an independent method (e.g. absolute mean-time only requires a fairly simple calibration to account for the relative delay in each detector channel, although few other datatypes are similarly robust (see section 2.2.6.2)). $\mathbf{M}^{(1)}_{n,m}$ measured with no reference state could then be used to derive absolute $x_1(r)$ using [2.4.5].

Alternatively a reference measurement $\mathbf{M}^{(0)}_{n,m}$ made on a homogenous phantom with known $x_0(r)$ that was quite different to the object of interest $x_1(r)$ could be used for calibration of data to produce baseline images using the non-linear starting condition given by [2.4.7] (since the approximation that $\Delta x(r)$ is small is less important for non-linear reconstruction).

Baseline images could then be used to derive images of optical property changes. For multi-wavelength data, static functional parameters might be derived from a number of single-wavelength baseline images (see chapter 2.5). For changes with time, repeated absolute measurements could be reconstructed as multiple baseline images, allowing changes to be assessed by simply subtracting one image from another. However the requirement for absolute calibration measurements to be effective on data acquired over a longer period of time (or between wavelengths) will result in image artefacts that will (non-linearly) change with time, resulting in potentially erroneous deductions about physiological changes.

However we must consider ways to exploit the calibration advantages of measuring changes in data. Particularly in the clinical environment, equipment will have inherent

instabilities and the need for quick, simple procedures requires that sources and detectors cannot be very easily removed from the subject for baseline calibration. A further advantage of in-situ calibration (via measuring changes) is that datatypes such as intensity or complex amplitude can be better calibrated. In order to baseline calibrate such types of data the coupling coefficients of light to and from the subject must be considered (Boas *et al*, 2001a) since they greatly influence measurements (which is less true for normalised measurements like mean-time and complex phase). The effects of these coupling errors would be greatly reduced by looking at changes in measurements *in situ*. The requirement for multiple combinations of data (e.g. amplitude *and* phase, or intensity and additional temporal moments), in order to separate μ_a and μ'_s parameters is discussed in section 1.3.3.4. In general, additional well calibrated datatypes will improve the accuracy of a reconstruction, so utilising changes in intensity via reference measurements (albeit on a heterogeneous structure) may be advantageous.

A method to reduce the effects of heterogeneous background structure on reconstructions of changes in optical properties from changes in measurement is demonstrated below.

2.4.6.1 Combining baseline and difference imaging

Ideally we need to find a way to reduce the error on our estimate of the background ($\delta x(r)$) since this will reduce the effect of the error term in [2.4.14]. However we also want to be able to take advantage of changes in measurements that are unsuitable for 'absolute' calibration (such as intensity). One possible way to do this would be to combine both techniques, of baseline and difference imaging.

By assuming that at least some baseline data are available, we can reconstruct an estimate of $x_o(r)$ (from mean-time only, for example). Subsequent measurements of changes in either mean-time, or any other parameter (such as intensity) can then be reconstructed using the approximate baseline image as the $x_o(r)$. So having some absolute mean-time measurement of the initial state $\langle \mathbf{t} \rangle^{(0)}_{n,m}$ we iteratively solve:

$$\langle \mathbf{t} \rangle^{(0)}_{n,m} - \langle t \rangle^{(g)}_{n,m} = J_{n,m}{}^{M}(x_g(r))[x_{g+1}(r) - x_g(r)]$$
$$\langle \mathbf{t} \rangle^{(0)}_{n,m} - \langle t \rangle^{(g+1)}_{n,m} = J_{n,m}{}^{M}(x_{g+1}(r))[x_{g+2}(r) - x_{g+1}(r)].....$$

[2.4.15]

yielding the solution $x^{\langle t \rangle}{}_0(r)$. Then, either after some time, assuming the object has remained fairly stable, or using other measurements made simultaneously with $\langle \mathbf{t} \rangle^{(0)}_{n,m}$ we can use the difference between measurements $\mathbf{M}^{(1)}_{n,m}$ - $\mathbf{M}^{(0')}_{n,m}$ (where $\mathbf{M}^{(0')}$ denotes the same or subsequent initial state) in:

$$\mathbf{M}^{(1)}_{n,m} - \mathbf{M}^{(0')}_{n,m} = J_{n,m}{}^{M}(x^{\langle t \rangle}{}_0(r))[x_1(r) - x^{\langle t \rangle}{}_0(r)]$$

[2.4.16]

in the linear inversion scheme, i.e. the <t> baseline image is used to calculate a heterogeneous sensitivity matrix. Alternatively, by incrementing $(x_1(r)-x_0^{<t>}(r))$ following [2.4.5], a non-linear reconstruction can be performed. This produces a solution $x_1(r)$, which is the final state of the object. If the *change* in optical properties is required then $x^{<t>}{}_0(r)$ can be subtracted from the final image. The inclusion of the estimate of the initial state allows better agreement between the sensitivity matrix (PMDFs) used to calculate the optical property variation. The ability to use additional types of data (such as intensity) should allow better assessment of changes in μ_a and μ'_s, assuming that using the estimate $x^{<t>}{}_0(r)$ is more effective than approximating the initial condition as homogeneous.

Figure 2.4.19 shows the images originally reconstructed in section 2.4.3.1, (Figure 2.4.7 and Figure 2.4.8) where constant background scatter structure was shown to affect images of changes in absorption. However here, instead of assuming $x_0(r)$ is homogeneous, reconstructions of the difference data for case 2 have utilised an initial guess of $x^{<t>}{}_0(r)$ from a baseline reconstruction of the absolute <t> simulated values for the initial state of case 2. So $x^{<t>}{}_0(r)$ consisted of the baseline μ'_s image (50th iteration), and homogenous $\mu_a = 0.01\text{mm}^{-1}$. Linear and non-linear results for intensity and mean-time differences are shown.

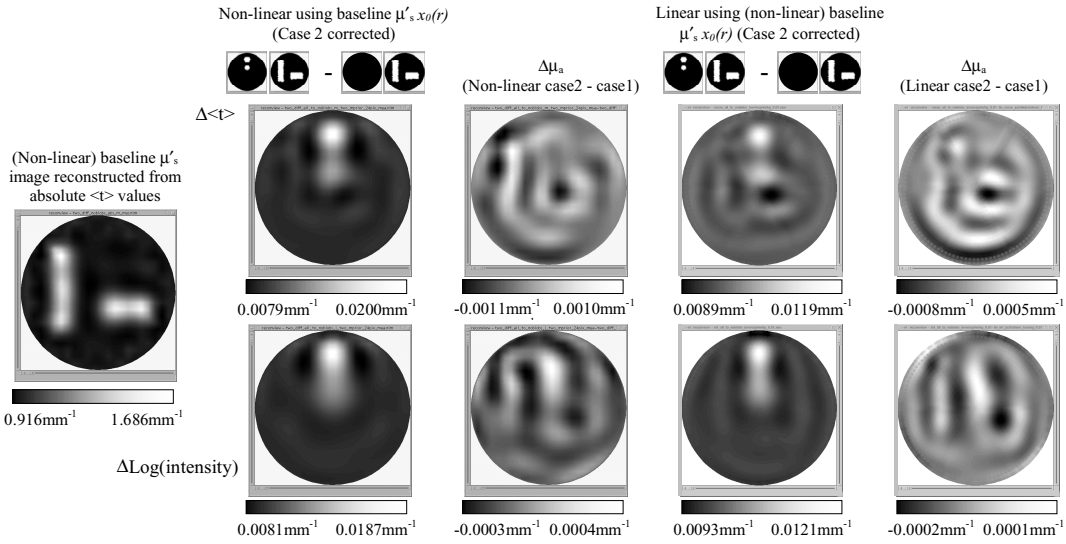

Figure 2.4.19 Images reconstructed using a baseline non-linear scatter image reconstructed from absolute <t> as the initial μ'_s distribution for reconstruction of mean-time (top) and intensity (bottom) difference images, using non-linear (left) and linear (right) methods.

Although the baseline image is not exactly quantitatively correct, its use has improved the difference images reconstructed from both mean-time differences and intensity ratios. The error range has decreased from 0.0036mm^{-1} to 0.0021mm^{-1} for mean-time and from 0.0023mm^{-1} to 0.0007mm^{-1} for intensity, for the non-linear image reconstructions (left). For the linear reconstructions (right) the baseline initial state was used to calculate the sensitivity matrix used in [2.4.4]. Qualitatively the linear images have improved compared to those in

Figure 2.4.8b and for intensity the error range has reduced from 0.0009mm^{-1} to 0.0003mm^{-1}. For mean-time, due to a sharp minimum, the error range has in fact increased from 0.0011mm^{-1} to 0.0013mm^{-1}. Note that although the errors on the linear images are smaller (comparing case 1 to case 2), the signal to noise is better in the non-linear reconstructions.

In the case described in section 2.4.3.2, where background absorption caused errors in images where no $x_0(r)$ was available, again the technique described above can be used to provide a better estimate of the sensitivity matrix, allowing improved images to be reconstructed from changes in measurements. In this case the images shown in Figure 2.4.10 and Figure 2.4.12 were reconstructed again, but utilising a baseline initial μ_a distribution derived from a non-linear image reconstructed using mean-time data (and homogenous, fixed μ'_s = 1mm^{-1}). In the non-linear images shown in Figure 2.4.20, $x^{<t>}_0(r)$ has been subtracted from the resulting images to provide equivalent images to those shown in Figure 2.4.10. The difference between the case A images and the corrected case B and C images are also shown.

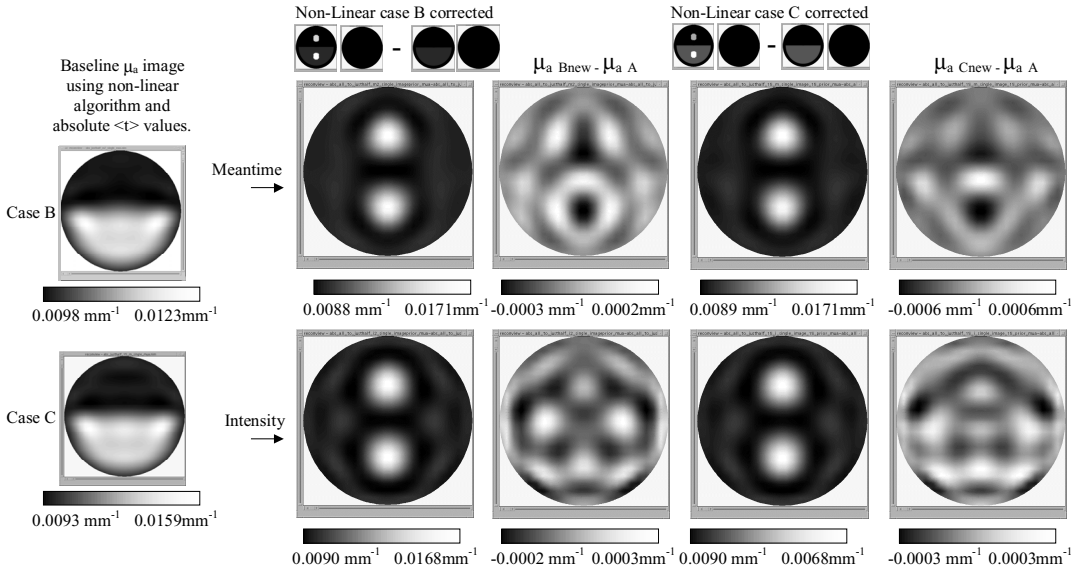

*Figure 2.4.20. **Non-linear** reconstructions using a (non-linear) baseline absorption distribution (far left) as the initial condition for mean-time (top) and intensity (bottom) difference data for cases B and C as shown without correction in Figure 2.4.10. Residuals are the differences between the original case A and corrected B and C images*

Figure 2.4.21 shows corresponding images reconstructed using the linear method. The images in Figure 2.4.21 use the linear method, but utilise the non-linear baseline μ_a distribution to calculate the sensitivity matrices for cases B and C. In both the linear and non-linear images, quantitation and object localisation have been greatly improved by the inclusion of baseline image information.

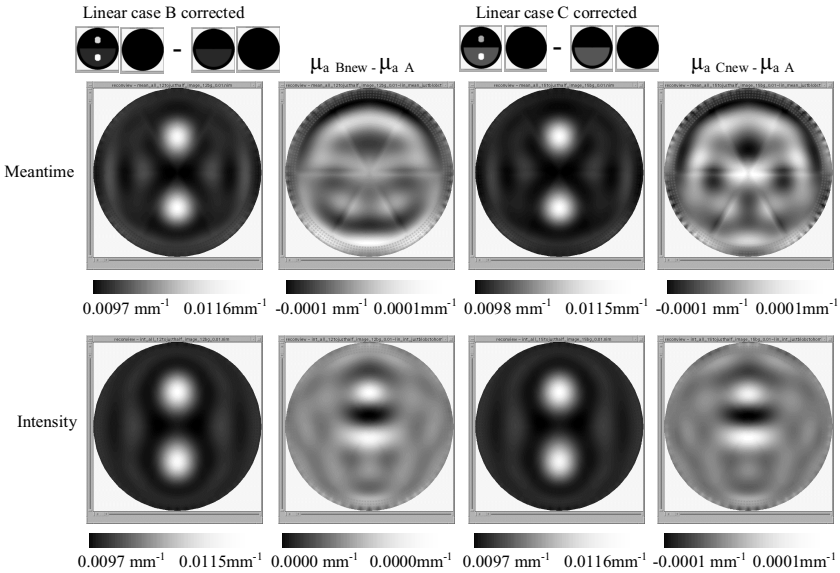

Figure 2.4.21 ***Linear*** *reconstructions using the (non-linear) baseline absorption distribution shown in Figure 2.4.20 as the initial condition for mean-time (top) and intensity (bottom) difference data for cases B and C as shown without correction in Figure 2.4.12 The residuals are the differences between the original case A (see Figure 2.4.12) and corrected linear B and C images. Note that the feint lines in the residuals are artefacts from the 6-sector mesh used in the reconstructions visible since the residuals are so small.*

The images from the simulation based on the wavelength-difference phantom II (Figure 2.4.17) were also reconstructed again utilising baseline initial μ_a distributions derived from the 780nm or 820nm absolute <t> images, as shown in Figure 2.4.22.

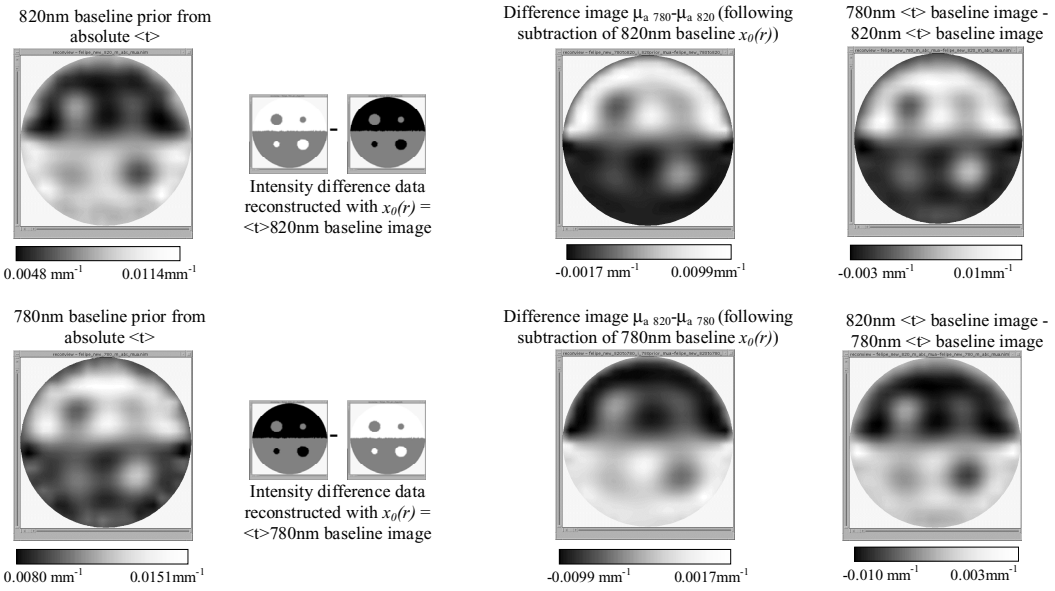

Figure 2.4.22 By using baseline images of absorption, reconstructed from absolute mean-time, as the initial condition for intensity difference images, object localisation and quantitation is improved compared to Figure 2.4.17, where a homogenous background μ_a was assumed.

Only the images produced using changes in intensity are shown, and the initial prior $x^{<t>}{}_1(r)$ has been subtracted from the resulting images to provide equivalent images to those shown in Figure 2.4.17. The difference between two absolute images is shown for comparison. Again, the technique of using an absolute μ_a image prior has reduced the

distortions of the wavelength-difference images compared to those shown in Figure 2.4.16 and Figure 2.4.17.

2.4.7 Summary and discussion

This chapter has demonstrated that the non-linearity of the problem of optical tomography causes images reconstructed from *changes* in measurement on a heterogeneous object to have errors associated with unknown initial conditions.

Simulations were presented demonstrating that changes in absorption in the presence of first background scatter, and then background absorption structure would be affected if no absolute measurement were available of the initial state. Experimental results showed that small wavelength-dependent inclusions in an essentially homogenous phantom could be reconstructed without significant artefact from the difference in data acquired at two wavelengths. A phantom with much larger wavelength-dependent structure was shown to yield images with severe distortion. This was demonstrated to be due, at least in part, to the problem of an unknown initial state by creating a 2D simulation of the phantom that exhibited very similar distortions.

The origin of the error was shown to be quantifiable as an extra term in the equation that is inverted to derive images. This term depends on datatype choice, whether μ_a or μ'_s are considered, the magnitude of the change in μ_a or μ'_s being reconstructed, the error on the estimate of the initial background properties and on the type of imaging (linear or non-linear).

A potential method for overcoming the limitations imposed in difference imaging on a heterogeneous object was then demonstrated. Estimates of the background initial state were reconstructed from well-calibrated <t> only data. Difference images were reconstructed using either subsequent mean-time *or* intensity changes utilising the baseline estimate as a starting condition. This technique was found to provide much improved results in simulations, despite the relative inaccuracy of the initial-state estimate.

The effect identified in this chapter represents a significant potential problem often overlooked by researchers advancing straight from simple phantoms (where little effect is seen) to clinical subjects, where heterogeneities are inevitable and (in the case of the brain) *potentially* overwhelming. The ‘void problem’ explored in chapter 2.6 is an extension of this since CSF between the brain and the skull can cause PMDFs to be significantly distorted.

The way that the photon density (and hence the PMDF) is affected by the presence of absorbing and scattering structure also has wider implications for standard spectroscopy techniques. Regions being probed by the photons will depend on the structure, meaning that

sensitivity of the measurement to a particular region cannot be determined if the structure is unknown. Further, at different wavelengths, the structure will have differing absorption characteristics, and hence differing spatial sensitivities. If the absorption of a region changes it will alter the regional sensitivity of the measurement. This was noted by (Firbank *et al*, 1998), who comments that the tissue volume sampled in single source-detector NIRS would depend on the blood volume in the tissue (or rather its associated μ_a). Hence the manipulation of NIRS data to extract functional parameters may be affected by differing spatial sensitivities of multi-wavelength or successive measurements.

The difference between mean-time and intensity PMDFs as shown in Figure 2.4.18 was noted by (Delpy *et al*, 1997), as a concern for simple spectroscopy techniques that utilise the modified Beer Lambert law, since different measurements would be probing different regions. However, by using an FEM based scheme this effect is accounted for in optical tomography.

Perturbation and differential imaging in optical tomography is explored in (Ntziachristos *et al*, 1999b) in relation to the use of ICG as an intravenous contrast agent (see section 2.7.2.7.5) for (slab geometry) breast imaging. A simulation based on an MRI of the breast is used to explore the problem of background absorption properties of the ICG state being significantly higher than those of the initial state. An analysis concludes that for the chosen geometry and properties, heterogeneity in absorption of the initial-state of the (simulated) breast properties has little effect on resulting images. In addition a correction method is developed where the *average* properties of the ICG state are determined and the large change in background properties accounted for by manipulating the data such that its differences appear to be due only to small perturbations on top of the background change. In this way the small localised changes in absorption corresponding to ICG collecting in vascularised regions can be distinguished from the large background absorption changes using a perturbation model that requires the assumption that Δx is small. These results suggest that for the level of heterogeneity in the adult breast, knowledge of the initial state structure for difference imaging may not be essential. Note however that the geometry of the study described is very different to that utilised by MONSTIR. Note also that ICG is likely to be the only way to invoke changes in the breast, so baseline imaging (relative to a homogenous phantom) will probably be utilised for the breast anyway.

(Cheng *et al*, 1999) describes a simulation study of the effect on topographic linear reconstructions of errors in the *average* background properties. (Pei *et al*, 2001a) presents an analysis of the influence of poor estimates of the average background properties of a reference medium in (fast) dynamic differential linear imaging. The background of the simulation was essentially homogenous. Errors in the reference state average were shown to affect

quantitation and localisation when reconstructing changes in μ_a and κ, although dynamic parameters (e.g. the time-series of relative variations of μ_a from one image to the next) were shown to be well retrieved.

The effects of the layered properties of breast and other tissues on the accuracy of parameter retrieval, given the differing sensitivity to regions with different optical properties, was explored in (Ntziachristos *et al*, 2001), the effects of refractive index mismatches between layers was considered in (Ripoll *et al*, 2001).

Heterogeneity of photon fluence rates in complex structures including a 2D model of the human brain were derived by (Hielsher *et al*, 1998). Transport-based methods were used to account for low scatter in the CSF layer and, in a similar way to plots in Figure 2.4.5 and Figure 2.4.6, revealed that areas sampled by photons in the brain would be affected by static μ_a and μ'_s structure.

The efficacy of the correction method suggested in section 2.4.6.1 will depend on the complexity of the initial state. Retrieval of good estimates of μ_a and μ'_s will require improved absolute datatype calibration and image reconstruction regularisation. (Note however that advantages of this method include the ability to utilise intensity difference data).

It is important to note that the severity of the error identified has only been demonstrated on a very limited set of geometries and conditions. Further work is required to categorically determine whether these effects will degrade images of the neonatal brain beyond clinical usefulness. Further simulations are required, possible using a more clinically realistic model on the neonatal head (preferably in 3D). Genuine conditions such as IVH (and those listed in Table 1.1.2) could be readily simulated, and the optimal method of identifying and localising changes determined.

2.5 Extracting functional parameters from multi-wavelength data.

The ultimate goal of near infrared tomography is to produce 3D images of the functional characteristics deep within volumes of tissue. The applications that we are investigating are imaging of the neonatal brain and the adult breast as described in section 1.1.3 . Maps of the *static* oxygen saturation and blood volume along with the scattering properties of a neonatal brain would allow rapid evaluation of cerebral perfusion. A premature infant, or a term infant that has experienced birth asphyxia may have infarcts, deformities, haemorrhages or general poor control over cerebral haemodynamics that may cause regions of the brain to be starved of oxygen for long periods. Identifying regions of low perfusion would allow treatments to be developed that could prevent resultant permanent disabilities such as cerebral palsy. In breast imaging, the functional and scattering characteristics of lesions will help to distinguish between benign cysts and fibroadenomas and malignant highly vascularised tumours or necrotic regions.

Images of *changes* in saturation and blood volume would also provide useful information. In the case of the neonate, the response to treatment could be monitored, and response to artificially induced changes in cerebral perfusion could be used to evaluate haemodynamic integrity.

Traditionally, optical tomography data has been used to reconstruct images of absorption and scatter (Hebden *et al*, 1999) by either linear or non-linear methods. Chapter 2.4 (section 2.4.4) showed images derived from data acquired at two wavelengths, although only the differences in absorption between the two wavelengths were reconstructed. Images of absolute μ_a and μ'_s, along with absorption changes in the arm due to exercise at both 780 and 820nm were also produced during development of MONSTIR and are described in detail in (Hillman *et al*, 2001b). However, in order to extract the functional characteristics of tissue, a more sophisticated way to combine multi-wavelength data is required. In order to develop robust methods for extraction of functional parameters we need to evaluate how image errors, inherent to optical tomography, will propagate through to derived images of saturation. These include:

- Spatial blurring of inclusions.
- Quantitative inaccuracy.
- Poor or varying convergence of reconstructions.

Measurement-based limitations on the use of multi-wavelength data include the need for separate monochromatic pulsed sources (if time-resolved techniques are being used). Further, multi-wavelength data should be acquired in parallel rather than serially if measurements are to be combined since the properties of tissue will vary over many different timescales. This has only recently become possible thanks to the integration of a new laser into MONSTIR that supplies a pulse train of pulses alternating between 780nm and 815nm (see section 1.2.3.4). Calibration methods and limitations described in chapters 2.1 and 2.2 must apply to the way that real data can be used in deriving saturation e.g. can the data required be accurately extracted from measured data? Further, constraints introduced in chapter 2.4 restrict the suitability of certain types of difference imaging.

From a physiological viewpoint we must consider the fact that the absorption properties of tissue are not solely due to blood. The relations between blood flow, volume and saturation, and variations in haematocrit within vessels of differing sizes also complicate interpretation of results.

We have a number of options of how to use multi-wavelength data. Each with differing advantages and disadvantages according to accuracy of saturation, accuracy of calibration and the type of image required. These can be categorised into four groups: 1) examining just the absorption at different wavelengths without attempting to derive functional parameters (as demonstrated in chapter 2.4 and (Hillman *et al*, 2001b)), 2) deriving absolute saturation from absolute data, 3) deriving absolute saturation data from wavelength-difference data, and 4) deriving *changes* in functional parameters from absolute or difference data.

This section describes studies designed to consider how to combine multi-wavelength data, and ascertain whether functional results obtained from clinical data are likely to be realistic. We start with an introduction to the different methods of combining data. Each method is then explored in turn using simulations and experimental breast and arm data. Some alternative methods of deriving saturation and volume are then described which exploit the intrinsic correlation between data acquired at multiple wavelengths. Finally a summary is presented including a review of the approaches taken by other researchers to extraction of functional parameters from optical data.

2.5.1 Different methods of combining data to derive functional parameters

2.5.1.1 Absolute absorption and changes in absorption at multiple wavelengths

This technique was introduced in chapter 2.4 and (Hillman *et al*, 2001b). Images of the absorption change between two wavelengths can be reconstructed using [2.5.1] or [2.5.2].

$$\left[\mu_{a,\lambda_1} - \mu_{a,\lambda_2}\right] = \left[J_{\mu_a,\lambda_1}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda 1} - \mathbf{M}_{\lambda 2}\right] \quad [\ 2.5.1\ ]$$

$$\mu_{a,\lambda_1} - \mu_{a,\lambda_2} = \left[J_{\mu_a,ref}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda 1} - \mathbf{M}_{ref}\right] - \left[J_{\mu_a,ref}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda 2} - \mathbf{M}_{ref}\right] \quad [\ 2.5.2\ ]$$

[2.5.2] is not susceptible to the errors relating to perturbation assumptions, as demonstrated in chapter 2.4. However, the derivation of single wavelength images, or changes in absorption do not intrinsically represent functional activity. More intuitive values can be derived by combining absolute absorption data to derive saturation and volume as shown below.

2.5.1.2 Deriving absolute saturation and blood volume from absolute data

The problem of deriving functional parameters from optical tomography data and images is somewhat different to the challenge presented by traditional single source-detector near infrared spectroscopy (NIRS) (and most topographic data analysis). Section 1.1.1.2 described that an estimate of the DPF is required, and only changes in measurements can be considered, if the modified Beer-Lambert law is to be of use. In order to extract changes in oxy- and deoxy-haemoglobin multi-wavelength data, or essentially continuous spectra, are often used along with comprehensive compensation for the wide number of other chromophores present in tissue to perform least squares fits to measured spectra over a wide NIR bandwidth. Fitting spectra over a wide range of wavelengths introduces errors associated with the wavelength dependence of pathlength, which can affect convergence of fits to the correct solution. Often it is temporal variations in the derived functional parameters that are of most interest to researchers using NIRS.

The MONSTIR system currently utilises only two wavelengths, 780nm and 815nm. Rather than exact quantitative accuracy, it is the *spatial* contrast of functional parameters that are of diagnostic use, (there is little opportunity to explore the temporal characteristics of variations since the system currently requires around 10 minutes to acquire a single image dataset). The fundamental difference between optical tomography and NRS is that by using numerical or analytical models to reconstruct images, we have the capacity to produce maps of *absolute* absorption, with pathlength compensated for by simultaneously solving for scatter during image reconstruction.

By considering the possible physiologically useful parameters that may be derived from optical tomography data, we define two functional parameters; the fractional oxygen saturation S and the fractional blood volume V of a tissue as:

$$S = \frac{[HbO_2]}{[HbO_2]+[Hb]} \qquad V = \frac{volume\ of\ \ whole\ blood\ in\ tissue}{total\ tissue\ volume} \quad [\ 2.5.3\ ]$$

where [Hb] denotes the concentration of deoxy-haemoglobin, and [HbO_2], the concentration of oxy-haemoglobin. If we are calculating the fraction of whole blood in tissue, we must assume to know the concentration of haemoglobin in blood (the haematocrit). This value is typically ~ 2mM, although it can vary from person to person depending on age, sex, and also on the size of the blood vessel from which the blood was extracted (Cope, 1991). Where $\varepsilon_{HbO2,\lambda1}$ is the molar extinction coefficient of 100% oxygenated whole blood at wavelength λ_1 and $\varepsilon_{Hb,\lambda1}$ is the molar extinction coefficient of 100% deoxygenated whole blood at wavelength λ_1 (for a given haematocrit), we can express the absorption coefficient of a tissue at a particular wavelength in the following way:

$$\mu_{a,\lambda_1} = (1-V)\mu_{a,bg,\lambda_1} + V(\varepsilon_{HbO_2,\lambda_1} S + (1-S)\varepsilon_{Hb,\lambda_1})$$
$$\mu_{a,\lambda_2} = (1-V)\mu_{a,bg,\lambda_2} + V(\varepsilon_{HbO_2,\lambda_2} S + (1-S)\varepsilon_{Hb,\lambda_2})$$

[2.5.4]

where $\mu_{a,bg,\lambda1}$ is the absorption coefficient of the exsanguinated tissue at λ_1. The extinction coefficients of blood can be calculated by multiplying the measured absorption spectrum of pure oxy- and deoxy-haemoglobin (as shown Figure 2.5.1) by the haematocrit.

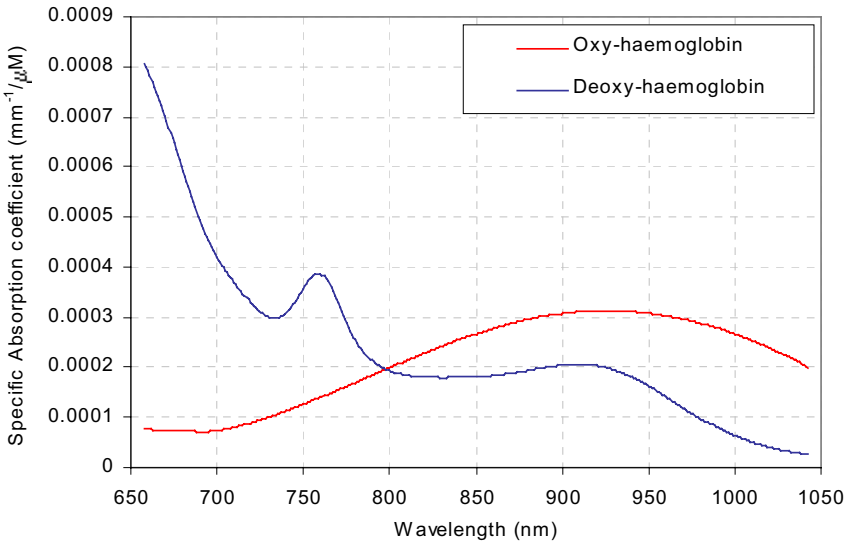

Figure 2.5.1 The absorption spectra of pure oxy- and deoxy-haemoglobin in the near infrared wavelength range.

The value of $\mu_{a,bg,\lambda1}$ includes the absorption of water, and exsanguinated grey and white matter in the brain or adipose and glandular tissue in the breast. Quoted values for the absorption of these tissue over the narrow 780nm to 815nm wavelength range imply that their spectral dependence is likely to be very small compared to the changes associated with blood (Hollis *et al*, 2001), (van-der-Zee, 1992). So, for simplicity we assume that $\mu_{a,bg,\lambda1} \approx \mu_{a,bg,\lambda2}$. Using this assumption, from [2.5.4] and [2.5.3], we can derive expressions for S and V in terms of $\mu_{a,\lambda1}$, $\mu_{a,\lambda2}$ and $\mu_{a,bg}$:

$$S = \frac{(\mu_{a,\lambda_1} - \mu_{a,\lambda_2} + \varepsilon_{Hb,\lambda_2} - \varepsilon_{Hb,\lambda_1})\mu_{a,bg} - \varepsilon_{Hb,\lambda_2}\mu_{a,\lambda_1} + \varepsilon_{Hb,\lambda_1}\mu_{a,\lambda_2}}{(\varepsilon_{HbO_2,\lambda_2} - \varepsilon_{Hb,\lambda_2})(\mu_{a,\lambda_1} - \mu_{a,bg}) - (\varepsilon_{HbO_2,\lambda_1} - \varepsilon_{Hb,\lambda_1})(\mu_{a,\lambda_2} - \mu_{a,bg})}$$

[2.5.5]

$$V = 1 + \frac{\varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda_1} - \varepsilon_{Hb,\lambda_1}\varepsilon_{HbO_2,\lambda_2} + \mu_{a,\lambda_1}(\varepsilon_{HbO_2,\lambda_2} - \varepsilon_{Hb,\lambda_2}) - \mu_{a,\lambda_2}(\varepsilon_{HbO_2,\lambda_1} - \varepsilon_{Hb,\lambda_1})}{\varepsilon_{Hb,\lambda_1}\varepsilon_{HbO_2,\lambda_2} - \varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda_1} + \mu_{a,bg}(\varepsilon_{Hb,\lambda_2} - \varepsilon_{Hb,\lambda_1} + \varepsilon_{HbO_2,\lambda_1} - \varepsilon_{HbO_2,\lambda_2})}$$

[2.5.6]

Note that these equations require an estimate of $\mu_{a,bg}$, along with absolute values for absorption at two wavelengths. In practice these equations can be applied by using each pixel / nodal value in two absorption images to derive images of V and S. The absolute images can be generated from calibrated data, using a non-linear reconstruction. Data calibrated using a reference phantom can also be used to generate baseline absolute images using either linear or non-linear reconstruction as shown in [2.5.7].

$$\mu_{a,\lambda_1} = \mu_{a,\lambda_1} - \mu_{a,ref} + \mu_{a,ref} = \left[J_{\mu_a,ref}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda 1} - \mathbf{M}_{ref}\right] + \mu_{a,ref}$$
$$\mu_{a,\lambda_2} = \mu_{a,\lambda_2} - \mu_{a,ref} + \mu_{a,ref} = \left[J_{\mu_a,ref}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda 2} - \mathbf{M}_{ref}\right] + \mu_{a,ref}$$

[2.5.7]

2.5.1.3 Deriving absolute saturation and blood volume from difference data

While images of the difference in absorption at two wavelengths do not represent any particular functional parameter, utilising the difference between two sets of data with the same geometry provides far superior calibration than trying to derive absolute data from measured TPSFs. By using the dual-wavelength pulse train from two interlaced lasers (as described in 1.2.3.4), error cancellation should be further improved (although currently calibration for wavelength dependent system errors needs to be implemented, see section 2.1.4.5). Note however that improvements to calibration may be outweighed by the effects of unknown initial structure demonstrated in chapter 2.4.

Nevertheless it is interesting to consider whether reconstructed images of ($\mu_{a,\lambda 1}$ - $\mu_{a,\lambda 2}$) from $<t_{\lambda 1}>$-$<t_{\lambda 2}>$ or $\text{Intensity}_{\lambda 1}$/ $\text{Intensity}_{\lambda 2}$ can be used to calculated absolute S and V. Rearranging [2.5.5] and [2.5.6] demonstrates that in fact the two variables cannot be separated if only ($\mu_{a,\lambda 1}$ - $\mu_{a,\lambda 2}$) is available.

$$S = \frac{(\mu_{a,\lambda_1} - \mu_{a,\lambda_2}) + (1 - V)(\varepsilon_{Hb,\lambda_1} - \varepsilon_{Hb,\lambda_2})}{V(\varepsilon_{Hb,\lambda_1} - \varepsilon_{Hb,\lambda_2} - \varepsilon_{HbO_2,\lambda_1} + \varepsilon_{HbO_2,\lambda_2})}$$

[2.5.8]

$$V = \frac{(\mu_{a,\lambda_1} - \mu_{a,\lambda_2})}{(S-1)(\varepsilon_{Hb,\lambda_1} - \varepsilon_{Hb,\lambda_2}) + S(\varepsilon_{HbO_2,\lambda_2} + \varepsilon_{HbO_2,\lambda_1})}$$

[2.5.9]

However, if measurements at 3 wavelengths are available, it should in fact be possible to derive absolute saturation and volume from changes in absorption only as shown in [2.5.10] and [2.5.11].

$$S = \frac{(\mu_{a,\lambda_3} - \mu_{a,\lambda_2})(\varepsilon_{Hb,\lambda_2} - \varepsilon_{Hb,\lambda_1}) + (\mu_{a,\lambda_2} - \mu_{a,\lambda_1})(\varepsilon_{Hb,\lambda_2} - \varepsilon_{Hb,\lambda_3})}{(\mu_{a,\lambda_3} - \mu_{a,\lambda_2})(\varepsilon_{Hb,\lambda_2} - \varepsilon_{Hb,\lambda_1} + \varepsilon_{HbO_2,\lambda_1} - \varepsilon_{HbO_2,\lambda_2}) + (\mu_{a,\lambda_2} - \mu_{a,\lambda_1})(\varepsilon_{Hb,\lambda_2} - \varepsilon_{Hb,\lambda_3} + \varepsilon_{HbO_2,\lambda_3} - \varepsilon_{HbO_2,\lambda_2})}$$

[2.5.10]

$$V = \frac{(\mu_{a,\lambda_3} - \mu_{a,\lambda_2})(\varepsilon_{Hb,\lambda_2} - \varepsilon_{Hb,\lambda_1} + \varepsilon_{HbO_2,\lambda_1} - \varepsilon_{HbO_2,\lambda_2}) + (\mu_{a,\lambda_2} - \mu_{a,\lambda_1})(\varepsilon_{Hb,\lambda_2} - \varepsilon_{Hb,\lambda_3} + \varepsilon_{HbO_2,\lambda_3} - \varepsilon_{HbO_2,\lambda_2})}{\varepsilon_{Hb,\lambda_3}(\varepsilon_{HbO_2,\lambda_1} - \varepsilon_{HbO_2,\lambda_2}) + \varepsilon_{Hb,\lambda_1}(\varepsilon_{HbO_2,\lambda_2} - \varepsilon_{HbO_2,\lambda_3}) + \varepsilon_{Hb,\lambda_2}(\varepsilon_{HbO_2,\lambda_3} - \varepsilon_{HbO_2,\lambda_1})}$$

[2.5.11]

Besides the possibility that data will be far more robust since cancellation of wavelength independent systematic errors will occur, these equations also do not contain the (wavelength independent) $\mu_{a,bg}$ term which appears in [2.5.5] and [2.5.6], potentially removing a significant source of error.

These equations imply that a very good system for deriving absolute saturation and volume would include a 3-wavelength interlaced laser pulse train. Although it must be demonstrated that the lack of knowledge of the background optical properties (including μ'_s – which won't be calculable from $<t_{\lambda 1}>-<t_{\lambda 2}>$ or $\text{Intensity}_{\lambda 1}/\ \text{Intensity}_{\lambda 2}$) are not going to overwhelmingly affect the quantitative accuracy of $(\mu_{a,\lambda 1} - \mu_{a,\lambda 2})$ images, see section 2.4.

Hence, the $\Delta\mu_a$ values for equations [2.5.10] and [2.5.11] can be derived either from reconstructing *changes* in data, or simply by reconstructing three absolute absorption images and then calculating $\Delta\mu_a$ (which will be prone to calibration errors but will not be affected by the unknown initial state, will benefit from not requiring a $\mu_{a,bg}$ estimate and will yield μ'_s).

$$\left[\mu_{a,\lambda_2} - \mu_{a,\lambda_1}\right] = -\left[J_{\mu_a,\lambda_2}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda_1} - \mathbf{M}_{\lambda_2}\right]$$
$$\left[\mu_{a,\lambda_3} - \mu_{a,\lambda_2}\right] = \left[J_{\mu_a,\lambda_2}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda_3} - \mathbf{M}_{\lambda_2}\right]$$

[2.5.12]

$$\mu_{a,\lambda_2} - \mu_{a,\lambda_1} = \left[J_{\mu_a,ref}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda_2} - \mathbf{M}_{ref}\right] - \left[J_{\mu_a,ref}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda_1} - \mathbf{M}_{ref}\right]$$
$$\mu_{a,\lambda_3} - \mu_{a,\lambda_2} = \left[J_{\mu_a,ref}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda_3} - \mathbf{M}_{ref}\right] - \left[J_{\mu_a,ref}{}^{M}\right]^{-1}\left[\mathbf{M}_{\lambda_2} - \mathbf{M}_{ref}\right]$$

[2.5.13]

2.5.1.4 Deriving changes in saturation and blood volume from difference data

For on-going monitoring of the brain of an infant, images showing *changes* in blood volume and oxygen saturation would provide valuable feedback about the effectiveness of treatment. Changes in cerebral haemodynamics could also be induced, for example by varying the levels of inspired oxygen or carbon dioxide. Measuring the changes in saturation and blood volume resulting from these actions may prove a useful diagnostic procedure (see 2.7.2.7.5).

Consider then that we have four measurements, where λ_1 and λ_2 are acquired simultaneously, but state B occurs at some time after state A:

$$\mu^A a,\lambda_1 = (1-V^A)\mu^A a,bg,\lambda_1 + V^A(\varepsilon_{HbO_2},\lambda_1 S^A + (1-S^A)\varepsilon_{Hb,\lambda_1})$$
$$\mu^A a,\lambda_2 = (1-V^A)\mu^A a,bg,\lambda_2 + V^A(\varepsilon_{HbO_2},\lambda_2 S^A + (1-S^A)\varepsilon_{Hb,\lambda_2})$$
$$\mu^B a,\lambda_1 = (1-V^B)\mu^B a,bg,\lambda_1 + V^B(\varepsilon_{HbO_2},\lambda_1 S^B + (1-S^B)\varepsilon_{Hb,\lambda_1})$$
$$\mu^B a,\lambda_2 = (1-V^B)\mu^B a,bg,\lambda_2 + V^B(\varepsilon_{HbO_2},\lambda_2 S^B + (1-S^B)\varepsilon_{Hb,\lambda_2})$$

[2.5.14]

Obviously changes in absorption and hence functional parameters can be calculated from one state to the next from successive absolute saturation and volume images derived using the methods described in sections 2.5.1.2 and 2.5.1.3. However, it is desirable to try and calculate the change in volume ΔV using the difference in absorption between state A and B (so that the method can utilise difference data).

$$\Delta V_{(A-B)} = \frac{(\mu^A a,\lambda_2 - \mu^B a,\lambda_2)(\varepsilon_{Hb,\lambda1} - \varepsilon_{HbO_2,\lambda1}) - (\mu^A a,\lambda_1 - \mu^B a,\lambda_1)(\varepsilon_{Hb,\lambda2} - \varepsilon_{HbO_2,\lambda2})}{\varepsilon_{Hb,\lambda2}\varepsilon_{HbO_2,\lambda1} - \varepsilon_{Hb,\lambda1}\varepsilon_{HbO_2,\lambda2} + \mu_{a,bg}(\varepsilon_{Hb,\lambda1} - \varepsilon_{HbO_2,\lambda1} - \varepsilon_{Hb,\lambda2} + \varepsilon_{HbO_2,\lambda2})}$$

[2.5.15]

Note that ΔV requires knowledge of $\mu_{a,bg}$. We can also derive an expression for the change in oxygen saturation from difference data, however in this case knowledge of the absolute volume at state A and B are required.

$$\Delta S_{(A-B)} = \frac{1}{(\varepsilon_{Hb,\lambda1} - \varepsilon_{HbO_2,\lambda1} - \varepsilon_{Hb,\lambda2} + \varepsilon_{HbO_2,\lambda2})}\left(\frac{(\mu^A a,\lambda_2 - \mu^A a,\lambda_1)}{V^B} - \frac{(\mu^B a,\lambda_2 - \mu^B a,\lambda_1)}{V^A}\right)$$

[2.5.16]

This demonstrates that it is not in fact possible to derive changes in saturation directly from changes in data from one state to the next. It implies that if changes in saturation are required, they can be calculated from the absolute saturation values derived from absolute measurements (section 2.5.1.2) or (if 3 wavelengths are available), from the difference between different wavelength data (section 2.5.1.3). Other researchers who do not have access to absolute data, but rather only have changes in data from one state to the next, will therefore be unable to derive changes in saturation. In this case, dual-wavelength data is commonly used to derive changes in oxy- haemoglobin (Δ[HbO_2]) and deoxy-haemoglobin (Δ[Hb]) concentration only. In this case, $\mu_{a,\lambda1}$ is expressed as:

$$\mu^A a,\lambda_1 = (1-[Hbo_2]^A - [Hb]^A)\mu^A a,bg,\lambda_1 + ([HbO_2]^A\varepsilon_{HbO_2,\lambda_1} + [Hb]^A\varepsilon_{Hb,\lambda_1})$$

[2.5.17]

Using the same notation as in [2.5.4], we get:

$$\Delta[HbO_2] = \frac{(\mu^A a,\lambda_2 - \mu^B a,\lambda_2)\varepsilon_{Hb,\lambda1} - (\mu^A a,\lambda_1 - \mu^B a,\lambda_1)\varepsilon_{Hb,\lambda2} + \mu_{a,bg}(\mu^A a,\lambda_1 - \mu^B a,\lambda_1 - \mu^A a,\lambda_2 + \mu^B a,\lambda_2)}{\varepsilon_{Hb,\lambda1}\varepsilon_{HbO_2,\lambda_2} - \varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda1} - \mu_{a,bg}(\varepsilon_{Hb,\lambda1} - \varepsilon_{HbO_2,\lambda1} - \varepsilon_{Hb,\lambda2} + \varepsilon_{HbO_2,\lambda2})}$$

[2.5.18]

$$\Delta[Hb] = \frac{(\mu^A a,\lambda_1 - \mu^B a,\lambda_1)\varepsilon_{HbO_2,\lambda_2} - (\mu^A a,\lambda_2 - \mu^B a,\lambda_2)\varepsilon_{HbO_2,\lambda1} - \mu_{a,bg}(\mu^A a,\lambda_1 - \mu^B a,\lambda_1 - \mu^A a,\lambda_2 + \mu^B a,\lambda_2)}{\varepsilon_{Hb,\lambda1}\varepsilon_{HbO_2,\lambda_2} - \varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda1} - \mu_{a,bg}(\varepsilon_{Hb,\lambda1} - \varepsilon_{HbO_2,\lambda1} - \varepsilon_{Hb,\lambda2} + \varepsilon_{HbO_2,\lambda2})}$$

[2.5.19]

While these expressions are well known (Hueber *et al*, 2001), many researchers choose to neglect the $\mu_{a,bg}$ term entirely ((Boas *et al*, 2001b), (Bluestone *et al*, 2001)), in which case the expressions reduce to:

$$\Delta[HbO_2] = \frac{(\mu^A{}_{a,\lambda_2} - \mu^B{}_{a,\lambda_2})\varepsilon_{Hb,\lambda 1} - (\mu^A{}_{a,\lambda_1} - \mu^B{}_{a,\lambda_1})\varepsilon_{Hb,\lambda 2}}{\varepsilon_{Hb,\lambda 1}\varepsilon_{HbO_2,\lambda_2} - \varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda 1}}$$

[2.5.20]

$$\Delta[Hb] = \frac{(\mu^A{}_{a,\lambda_1} - \mu^B{}_{a,\lambda_1})\varepsilon_{HbO_2,\lambda_2} - (\mu^A{}_{a,\lambda_2} - \mu^B{}_{a,\lambda_2})\varepsilon_{HbO_2,\lambda 1}}{\varepsilon_{Hb,\lambda 1}\varepsilon_{HbO_2,\lambda_2} - \varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda 1}}$$

[2.5.21]

Note that although we can calculate $\Delta[HbO_2]$ and $\Delta[Hb]$ from available state-difference data, it is not possible to convert these values into saturation changes. While a change in volume can be expressed in terms of changes in oxy- and deoxy-haemoglobin:

$$\Delta V = [HbO_2]^A + [Hb]^A - ([HbO_2]^B + [Hb]^B) = \Delta HbO_2 + \Delta Hb,$$

[2.5.22]

change in saturation is not a function of $\Delta[HbO_2]$ and $\Delta[Hb]$:

$$\Delta S = \frac{[HbO_2]^A}{[HbO_2]^A + [Hb]^A} - \frac{[HbO_2]^B}{[HbO_2]^B + [Hb]^B}$$

[2.5.23]

The changes in absorption used in the above equations would be derived (using linear reconstruction notation), by using either the differences in data [2.5.24] (which is susceptible to perturbartion approximation errors – see section 2.4), or from the changes in absolute absorption images [2.5.25] (Although these absolute absorption values could be similarly used to calculate absolute saturation and volume (and hence subsequent changes) using [2.5.5] and [2.5.6]). Arm 'state-difference' images in (Hillman *et al*, 2001b) utilised [2.5.24]:

$$\left[\mu^A{}_{a,\lambda_1} - \mu^B{}_{a,\lambda_1}\right] = \left[J_{\mu_a,A,\lambda_1}{}^M\right]^{-1}\left[\mathbf{M}^A{}_{\lambda_1} - \mathbf{M}^B{}_{\lambda_1}\right]$$
$$\left[\mu^A{}_{a,\lambda_2} - \mu^B{}_{a,\lambda_2}\right] = \left[J_{\mu_a,A,\lambda_2}{}^M\right]^{-1}\left[\mathbf{M}^A{}_{\lambda_2} - \mathbf{M}^B{}_{\lambda_2}\right]$$

[2.5.24]

for absolute or baseline data:

$$\mu^A{}_{a,\lambda_1} - \mu^B{}_{a,\lambda_1} = \left[J_{\mu_a,ref}{}^M\right]^{-1}\left[\mathbf{M}^A{}_{\lambda_1} - \mathbf{M}_{ref}\right] - \left[J_{\mu_a,ref}{}^M\right]^{-1}\left[\mathbf{M}^B{}_{\lambda_1} - \mathbf{M}_{ref}\right]$$
$$\mu^A{}_{a,\lambda_2} - \mu^B{}_{a,\lambda_2} = \left[J_{\mu_a,ref}{}^M\right]^{-1}\left[\mathbf{M}^A{}_{\lambda_2} - \mathbf{M}_{ref}\right] - \left[J_{\mu_a,ref}{}^M\right]^{-1}\left[\mathbf{M}^B{}_{\lambda_2} - \mathbf{M}_{ref}\right]$$

[2.5.25]

2.5.2 Testing derivation of saturation and volume from absolute images.

A simulation was performed to test the accuracy of absolute saturation and fractional blood volume parameters derived from dual-wavelength data. The optical properties of the simulation were based on the molar extinction coefficients of whole blood at 780nm and 815nm (assuming that the haemoglobin concentration of whole blood = 2mM (section

1.1.2.2)) and calculated using [2.5.4] with $\mu_{a,bg} = 0.01mm^{-1}$. The simulation is 2D and has a number of inclusions with different saturation to the background, different volume to the background, or different saturation *and* volume to the background.

Forward data (mean-time and intensity) were simulated using TOAST on a 2D mesh with 2945 nodes, and with 16 coincident evenly spaced sources and detectors (assuming monstodes were being used, see sections 2.1.2.2 and 2.7.2.2.1). Images were reconstructed using intensity only, mean-time only, and both intensity and mean-time together. Mean-time images tended to have high spatial frequency artefacts, intensity only images exhibit cross talk and low resolution, therefore only the dual datatype (*mean and intensity*) images are presented in this section.

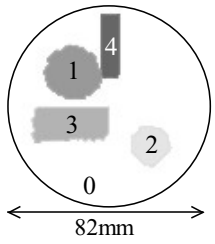

Region	SO_2 (%)	Vol (%)	μ_a (mm^{-1}) (780nm)	μ_a (mm^{-1}) (815nm)	μ'_s (mm^{-1})
0	76	4	0.024786	0.026527	1
1	98	4	0.023290	0.027222	1
2	50	3	0.022414	0.021779	1
3	76	6	0.032178	0.034790	1
4	76	4	0.024786	0.026527	1.5

$\varepsilon_{HbO2,780}$	0.33886
$\varepsilon_{Hb,780}$	0.50874
$\varepsilon_{HbO2,815}$	0.44212
$\varepsilon_{Hb,815}$	0.36316

(mm^{-1})

$\mu_{a,bg} = 0.01mm^{-1}$

Figure 2.5.2 Optical properties of simulation to test derivation of saturation and volume. (right) extinction coefficients of whole blood used to calculate suitable optical properties (from CCD spectrometer measurements).

2.5.2.1 Linear reconstructions of simulation

A homogenous reference state was assumed to have the same properties as region 0 for each wavelength. TOAST was used to simulate reference data and Jacobians (PMDFs) based on each reference. TOAST generates diffusion coefficient (κ) PMDFs, so for simplicity we reconstruct for κ rather than μ'_s ($\kappa = 1/(3(\mu'_s + \mu_a))$). Eight Jacobians in total were required for μ_a, κ, intensity and mean-time for both wavelengths. Images of μ_a and κ were reconstructed using a simple linear reconstruction method given by:

$$\begin{bmatrix} \Delta I_{815nm-815nm(\mathrm{hom})} & \Delta I_{780nm-780nm(\mathrm{hom})} & \Delta < t >_{815nm-815nm(\mathrm{hom})} & \Delta < t >_{780nm-780nm(\mathrm{hom})} \end{bmatrix}$$

$$= \begin{bmatrix} J_{\mu_a,815nm(\mathrm{hom})}^{I} & 0 & J_{\kappa,815nm(\mathrm{hom})}^{I} & 0 \\ 0 & J_{\mu_a,780nm(\mathrm{hom})}^{I} & 0 & J_{\kappa,780nm(\mathrm{hom})}^{I} \\ J_{\mu_a,815nm(\mathrm{hom})}^{<t>} & 0 & J_{\kappa,815nm(\mathrm{hom})}^{<t>} & 0 \\ 0 & J_{\mu_a,780nm(\mathrm{hom})}^{<t>} & 0 & J_{\kappa,780nm(\mathrm{hom})}^{<t>} \end{bmatrix} \begin{bmatrix} \mu_{a,815nm} - \mu_{a,815nm(\mathrm{hom})} \\ \mu_{a,780mn} - \mu_{a,780nm(\mathrm{hom})} \\ \kappa_{815nm} - \kappa_{815nm(\mathrm{hom})} \\ \kappa_{780nm} - \kappa_{780nm(\mathrm{hom})} \end{bmatrix}$$

[2.5.26]

While equations will be shown in this form throughout this section, in fact data were normalised before being solved in this way, as detailed in section 1.3.3.5. Tikhonov regularisation was used to invert the linear equation shown in [2.5.26], as described in section 1.3.2.2. The nodal values of the $\mu_{a,780nm}$-$\mu_{a,780nm(hom)}$ and $\mu_{a,815nm}$-$\mu_{a,815nm(hom)}$ images were converted to absolute by adding the reference state values ($\mu_{a,780nm(hom)}$ and $\mu_{a,815nm(hom)}$) and are shown in Figure 2.5.3. The mean kappa image is shown at the bottom of Figure 2.5.4.

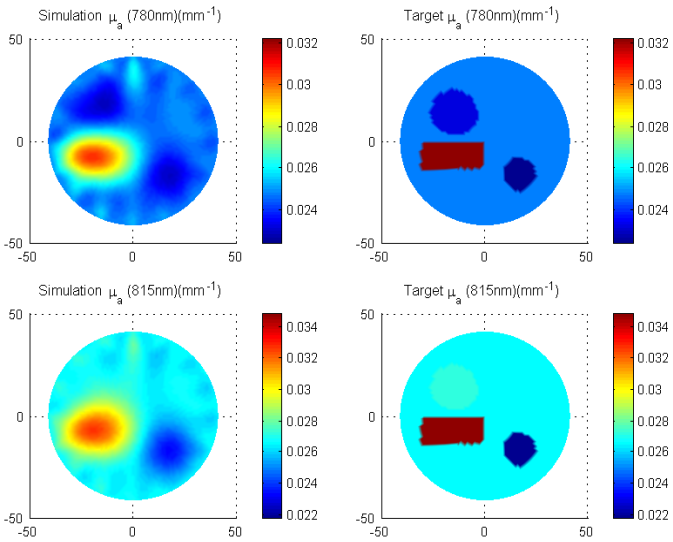

Figure 2.5.3 (left) Simple linear reconstructions of simulation absolute 780nm and 815nm absorption (plotted on the same colour scale as their target images on the right)

The nodal values of the absolute absorption images were then manipulated using [2.5.5] and [2.5.6] to produce the saturation and volume images shown in Figure 2.5.4.

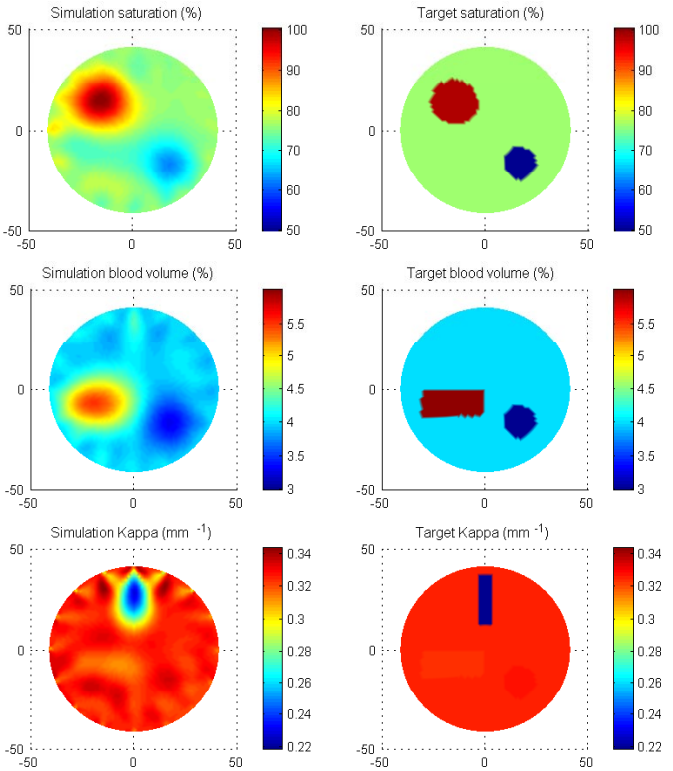

Figure 2.5.4 (left) saturation, and volume values derived from absolute absorption images, and (bottom) the mean absolute kappa image, (right) the target (original mesh values) for saturation, volume and kappa. Target and image have the same colour scale.

Figure 2.5.4 demonstrates that the three parameters (S, V and κ) can be separated. The image values are slightly less extreme than the target values (the peaks are lower and the troughs are more shallow), but otherwise the retrieved parameters are fairly close to the targets. Interestingly, although the original μ_a images are blurred and not quantitativley correct, the saturation image is closer to the target values than the volume image.

As can be seen from [2.5.5] and [2.5.6], absolute values of saturation will be affected by errors in $\mu_{a,bg}$, $\mu_{a,780nm}$ and $\mu_{a,815nm}$ and errors in the estimates of the extinction coefficients $\varepsilon_{*,\lambda*}$. Figure 2.5.5 shows plots that demonstrate the effect on derived saturation and volume of errors in $\mu_{a,bg}$, and (identical) additive and multiplicative errors in both $\mu_{a,780nm}$ and $\mu_{a,815nm}$. The plots were generated using the properties of region 0, with both μ_a values multiplied by 0.8, 1, 1.2 or 1.4, or with -0.0001, 0, 0.0001 or 0.001 mm^{-1} added.

The absolute functional parameters obtained are clearly very sensitive to knowledge of the background absorption, as well as to scaling errors in $\mu_{a,780nm}$ and $\mu_{a,815nm}$. Volume can become negative if too high a value of $\mu_{a,bg}$ is used, and saturation will rise until a discontinuity causes it to drop to a negative value.

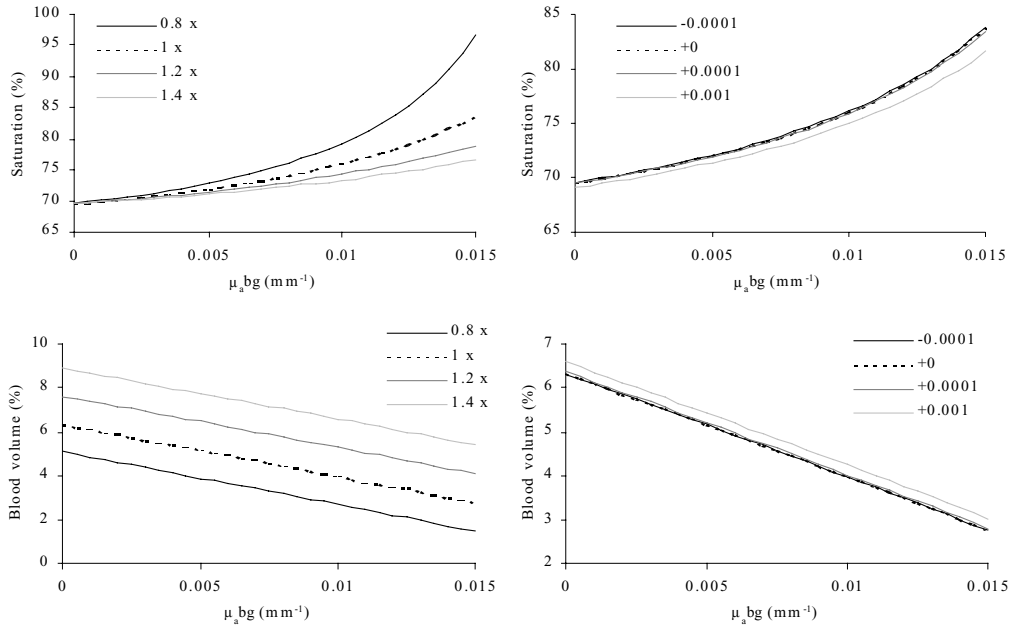

Figure 2.5.5 Plots showing the variation of calculated saturation (top) and blood volume (bottom) values when background μ_a is incorrect, or identical multiplicative (left) and additional (right) errors affect the absolute μ_a multi-wavelength images. (using 76% saturation and 4% blood volume values and μ_{abg}= 0.01mm^{-1})

2.5.2.2 Non-linear reconstructions of simulation

In order to test whether the improved quantitation generally achievable using non-linear image reconstruction algorithms provides superior accuracy when deriving saturation and fractional blood volume, TOAST was use to reconstruct the simulated data described above. The reference state mesh was used as the starting conditions for each wavelength's reconstruction (equivalent to [2.5.26]). The same meshes were used as for the linear reconstructions (note that the coarseness of the FEM mesh was found to have a significant affect on quantitation, see (Hillman *et al*, 2002)). Forward-mesh basis was used to allow direct comparison with linear results. μ_a and μ'_s were reconstructed simultaneously for each wavelength.

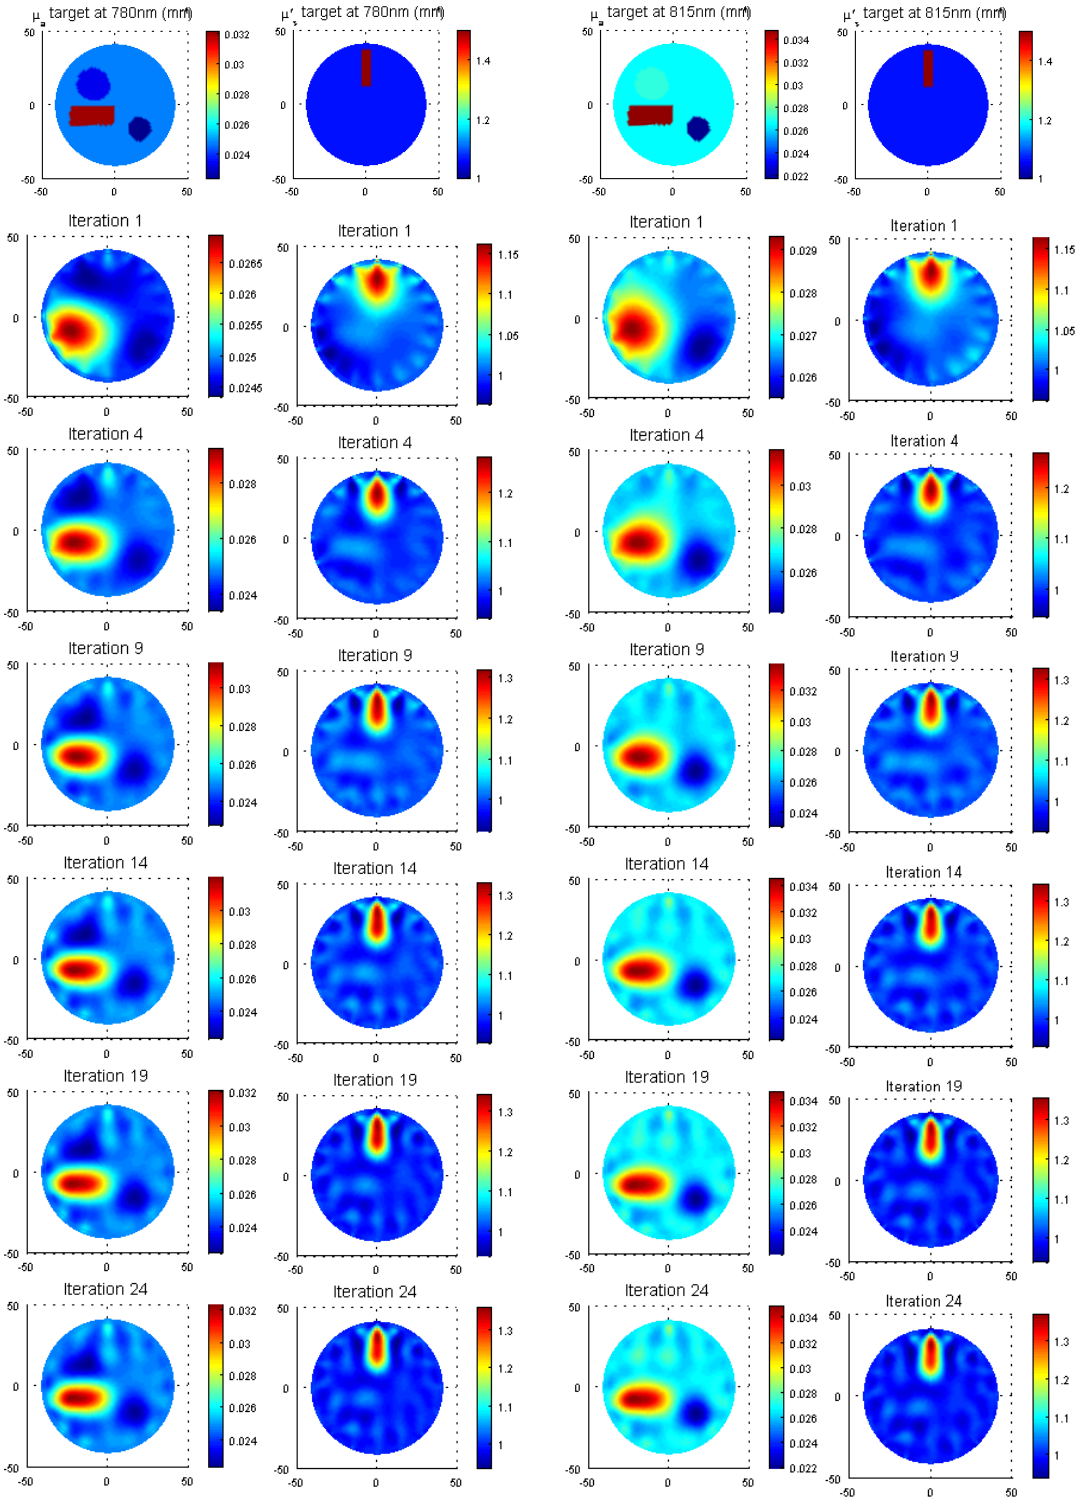

Figure 2.5.6 Absolute images of $\mu_{a,780nm}$ and $\mu_{a,815nm}$ derived from noiseless simulated mean-time and intensity. Iterations 1,4,9,14,19 and 24 are shown. Each has its own colour-scale whose limits are given to the right of each image.

For non-linear image reconstruction, convergence occurs via a series of iterations. A definite stopping criterion has yet to be developed, and as such it is not possible to estimate *which* iteration is the optimal for derivation of saturation and volume. A series of iterations is therefore shown to demonstrate the changes in images with increasing convergence. Note that reconstruction basis, and many other reconstruction parameters will subtlety affect the way in

which image solutions converge (see section 1.3.1.3.1). Each iteration in Figure 2.5.6 is scaled to its maximum and minimum values to emphasise the changes in shape of the features with changing iteration number. Figure 2.5.7 shows saturation and blood volume calculated from the nodal values (using [2.5.5] and [2.5.6]) of the absolute absorption images shown in Figure 2.5.6.

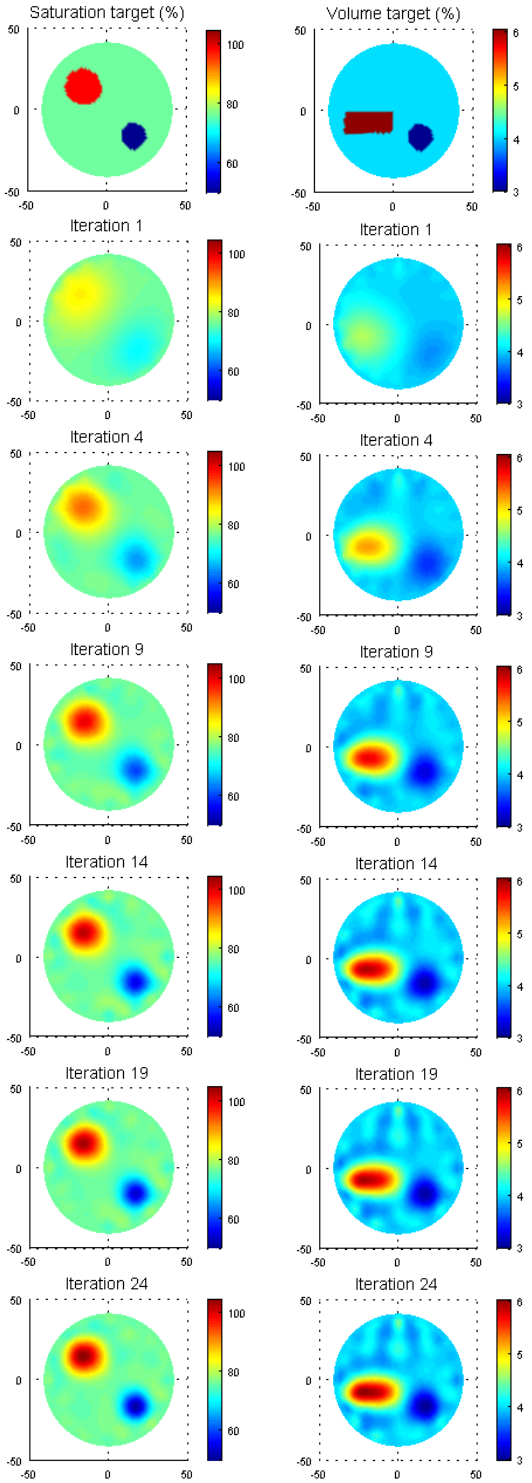

Figure 2.5.7 Absolute saturation (left) and volume (right) images of derived from the non-linear absolute images in Figure 2.5.6. Each, including the target images at the top, are scaled to a common colour-scale.

Each 'iteration' of saturation and volume has been calculated using $\mu_{a,780nm}$ (iteration 1) and $\mu_{a,815nm}$ (iteration 1), and then $\mu_{a,780nm}$ (iteration 2) and $\mu_{a,815nm}$ (iteration 2) etc. The images shown in Figure 2.5.7 are scaled to the maximum and minimum values within all the iterations and target to demonstrate how quantitation varies with iteration.

The resolution and quantitation of the higher iteration non-linear absorption, scatter, saturation and blood volume images is better in the non-linear case (although the values found using linear reconstruction are comparable to ~iteration 5.). This is to be expected since the process of iterating aims to find the optimal distribution of absorption and scatter that corresponds to the data and is not bound by the limitations of a perturbation approximation. In fact, with higher iterations, the peak values actually overshoot the target values. This is counteracted by the blurring of the features, such that regions around each peak may be lower than the true value, and further, regions around each inclusion boundary will have apparently higher values than they should. Without rigorous constraints, e.g. knowledge of the boundary of each region ((Kolehmainen *et al*, 2000), (Schweiger *et al*, 1999a)) image reconstruction of measurements of diffuse light will always suffer from distortion by blurring.

The above results considered only simulated data with no noise. The data were created and reconstructed on the same mesh, reducing possible meshing errors, and the exact background and reference starting parameters were known and used in the reconstruction. In reality, measured data will have noise, background values need to be estimated, and meshing artefacts will not be compensated for. Below, the methods described above are repeated but with noise added to the simulated data.

2.5.2.3 Linear reconstructions of simulated data with noise added

In order to explore the effect of noise on the quantitative accuracy attainable in the derivation of saturation and blood volume, noise was added to the simulated data used in sections 2.5.2.1 and 2.5.2.2 above. The same linear and non-linear reconstructions as above were then performed (still providing the correct background and reference properties and using the same mesh).

The noise was obtained by processing two dual-wavelength homogenous phantom data-sets (hom1 and hom2) acquired in the same geometry as the simulation using MONSTIR during the void experiments detailed in chapter 2.6. The mean-time noise was extracted by finding the average mean-time per source-detector separation. This average was then subtracted from the original homogenous data set, and the difference (representing system noise) added onto the clean simulated mean-time values (using hom1_w1 and hom1_w2 to generate noise for simulated data at w1 ≡ 780nm and w2 ≡ 815nm). For the simulated

reference data sets (wavelengths 1 and 2) noise was extracted from hom2_w1 and hom2_w2. Using notation from section 2.2.6.2:

$$Average_{\,d}[\mathcal{M}[y(t)]] = \frac{1}{32}\sum_{n=1}^{32}\mathcal{M}[y^{d,n}(t)] = \frac{1}{32}\sum_{n=1}^{32}\mathcal{M}[D^{d,n}(t)] + \frac{1}{32}\sum_{n=1}^{32}\mathcal{M}[I^{d,n}(t)]$$

[2.5.27]

where $y^{d,n}(t)$ is one of the 32 TPSFs acquired with a source-detector pair that is a distance d apart (assuming a circular planar arrangement). D is the correct TPSF and I is the error on D in the measurement y. Since the phantom is cylindrical and homogeneous:

$$\frac{1}{32}\sum_{n=1}^{32}\mathcal{M}[D^{d,n}(t)] = \mathcal{M}[D^{d}(t)],$$

[2.5.28]

and the average of all $I(t)$ terms is just the average correction for all sources and detectors. So assuming that the average of the system noise is the same for hom1 and hom2, we are adding:

$$\begin{aligned}
&\mathcal{M}[y^{(\mathrm{hom}\,2)^{d,n}}(t)] - Average_{\,d}[\mathcal{M}[y^{(\mathrm{hom}\,2)}(t)]] - \mathcal{M}[y^{(\mathrm{hom}\,1)^{d,n}}(t)] - Average_{\,d}[\mathcal{M}[y^{(\mathrm{hom}\,1)}(t)]] \\
&= \mathcal{M}[I^{(\mathrm{hom}\,2)^{d,n}}(t)] - \mathcal{M}[I^{(\mathrm{hom}\,1)^{d,n}}(t)] - \frac{1}{32}\sum_{n=1}^{32}\mathcal{M}[I^{(\mathrm{hom}\,2)^{d,n}}(t)] + \frac{1}{32}\sum_{n=1}^{32}\mathcal{M}[I^{(\mathrm{hom}\,1)^{d,n}}(t)] \\
&\approx \mathcal{M}[I^{(\mathrm{hom}\,2)^{d,n}}(t)] - \mathcal{M}[I^{(\mathrm{hom}\,1)^{d,n}}(t)]
\end{aligned}$$

[2.5.29]

to the simulated difference data. These noise values are quite small compared to the absolute data. The difference between the noisy and original data sets can be seen if the (data – reference) mean-times are plotted as shown in Figure 2.5.8.

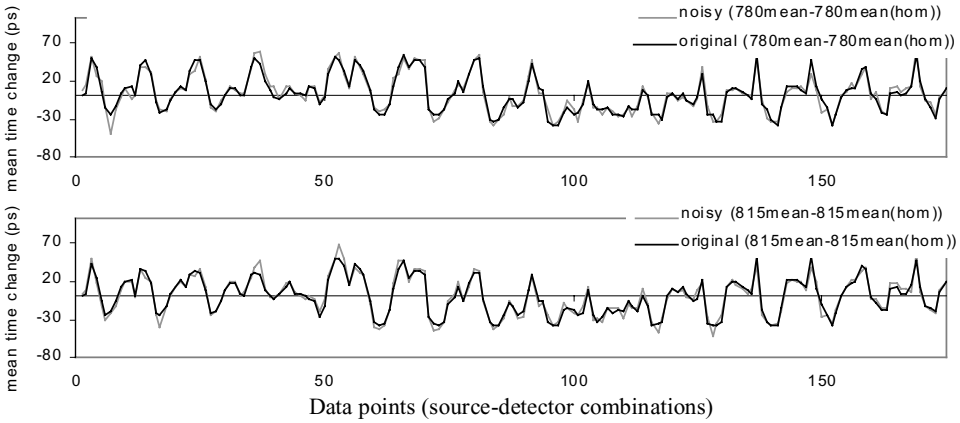

Figure 2.5.8 The simulated mean-time (data – reference), original and artificially noisy values for 780nm and 815nm.

For the intensities, the ratio of the intensities of hom1_w1/hom2_w1 and hom1_w2/ hom2_w2 were calculated. The ratio for each source was normalised by its mean due to the secondary fibre switch problem described in section 2.1.1.3, and the simulated data were then multiplied by the normalised intensity ratios. The simulated reference intensity data were not adjusted since when TOAST calculates log(data)-log(ref) the fact that the data is multiplied

by the ratio is equivalent to having multiplied each data set by the corresponding noise. The ratios of data / reference for intensities are shown for original and noisy data in Figure 2.5.9.

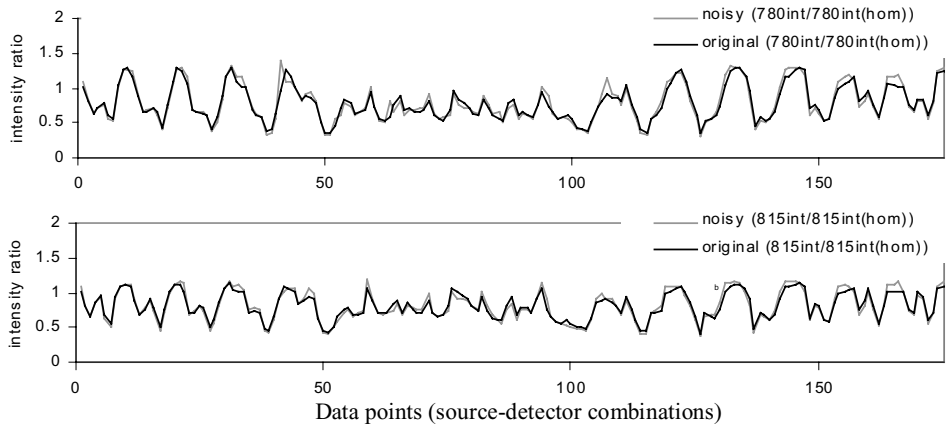

Figure 2.5.9 The simulated intensity (data / reference), original and artificially noisy values for 780nm and 815nm

The level of noise on the resulting simulated noisy data is intended to represent the level equivalent to a 'difference imaging' experiment, where data from a structured phantom (or the breast) and the data from a homogenous reference phantom are used together to eliminate common convolved errors. This approach was chosen in preference to simulating statistical noise on the data, since a very complex noise model would be required to relate every source of system noise (see section 2.1.1), to effects on extracted mean-time and intensity errors. Continuation of this work could include development of a more thorough noise model to assess its affect on results.

Figure 2.5.10 shows the resulting linear absorption images reconstructed from the noisy data.

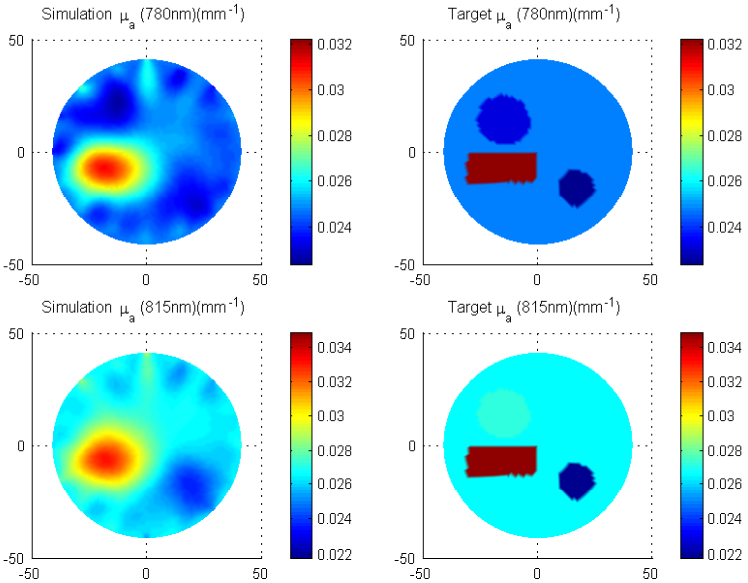

Figure 2.5.10 Simple linear reconstructions of simulation (with noise added to the raw data) absolute 780nm and 815nm absorption (plotted on the same colour scale) (Compare with Figure 2.5.3)

Figure 2.5.11 shows the saturation and fractional blood volume images derived from the images in Figure 2.5.10, along with the (average) kappa image derived from the noisy data. The target images are also shown, each scaled to a common scale with its corresponding image.

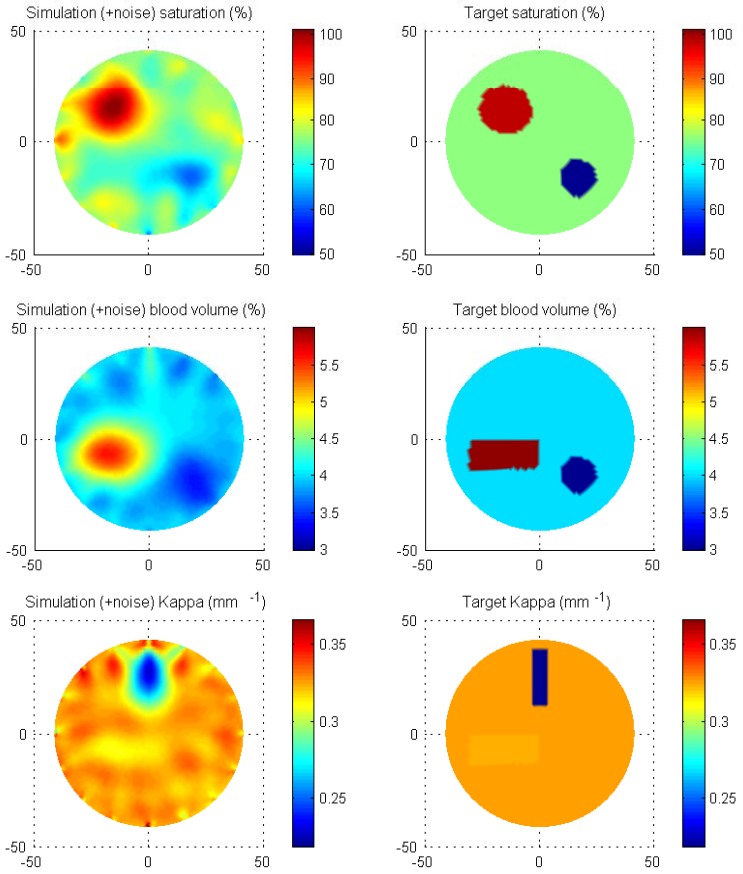

Figure 2.5.11(left) saturation, and volume values derived from the (noisy) absolute absorption images shown in Figure 2.5.10, and the mean absolute kappa reconstruction (from noisy data), (right) the target (original mesh values) for saturation, volume and kappa. Each target and image are shown with the same colour scale. (Compare with Figure 2.5.4).

While the saturation and volume results from the noisy data have clearly been affected by the presence of noise, the level of quantitative accuracy remains similar to that of the noise-free data (see Figure 2.5.3 and Figure 2.5.4). Besides the distortion of the low saturation inclusion in the saturation image, the spatial accuracy of the images has been well conserved.

2.5.2.4 Non-linear reconstructions of simulated data with noise added

The simulated data with noise added, described in section 2.5.2.3 were also reconstructed using TOAST, to compare the effect of a non-linear reconstruction on saturation and volume derivation from noisy data. Figure 2.5.12 shows absolute absorption and scatter non-linear reconstructions (for 780nm and 815nm) of the simulated data + noise. A series of iterations is shown to demonstrate the way in which the solution converges. Each image is scaled to its

own maximum and minimum to accentuate the changing shapes of the features with increasing iteration number.

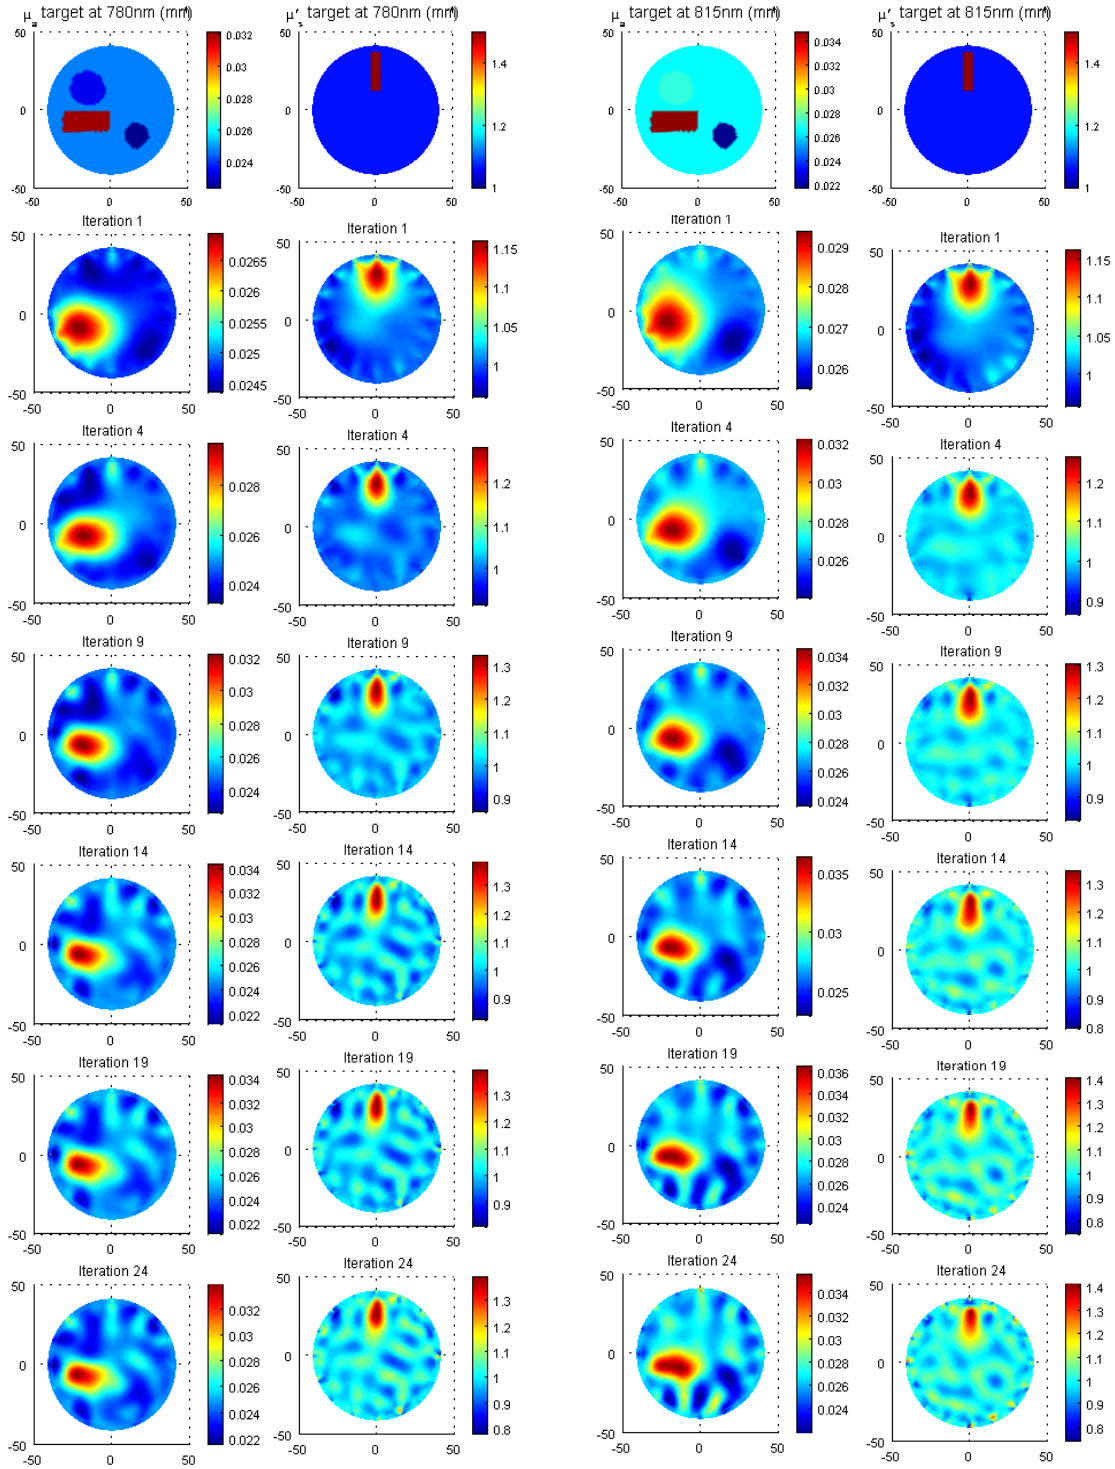

Figure 2.5.12 absolute μ_a and μ'_s at 780nm (left) and 820nm (right), derived from simulated data + noise. Each image is scaled to its maximum and minimum values (shown on the colour-bar to the right of each image)

Saturation and volume were calculated from each iteration μ_a image pair. The results are shown in Figure 2.5.13. Each image (including its target) is scaled to a common scale of the maximum and minimum within all the images.

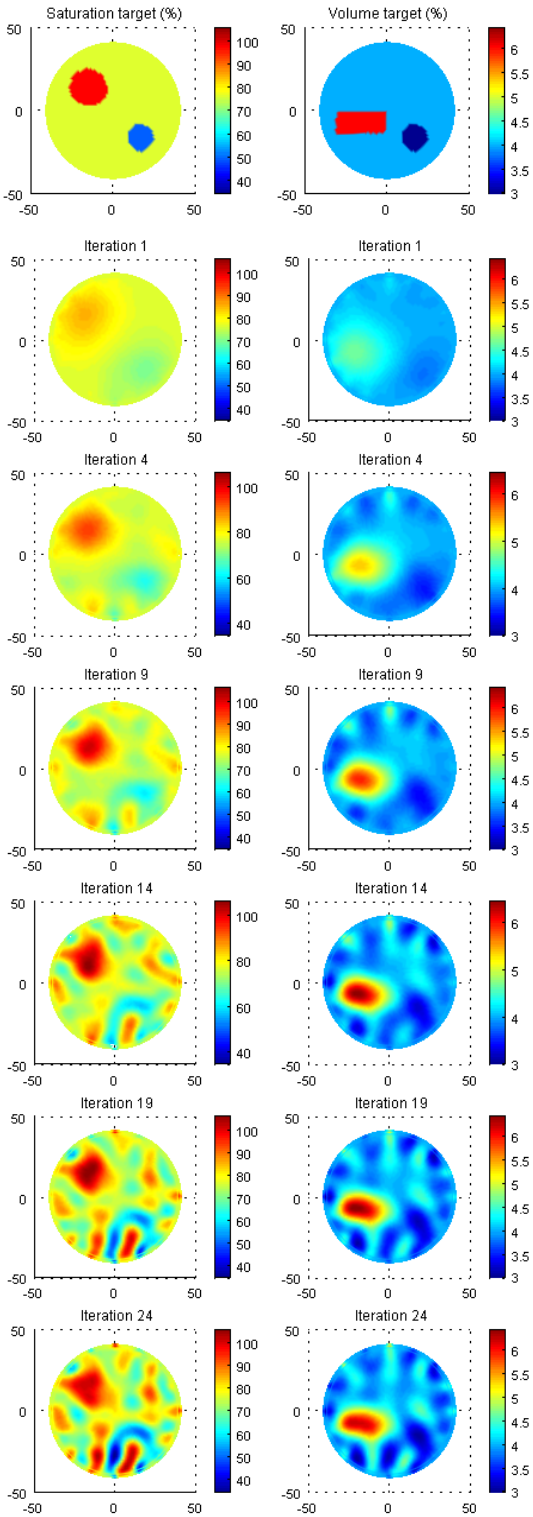

Figure 2.5.13 Absolute saturation (left) and volume (right) images of derived from the (noisy) non-linear absolute images in Figure 2.5.12. Each, including the target images at the top, are scaled to a common colour-scale.

In contrast to the linear and non-linear results using noise-free data, the results shown above highlight a major potential problem with using un-constrained non-linear images to derive absolute saturation and volume. While in the noise-free case (Figure 2.5.7), higher non-linear iterations provided progressively improved quantitation and resolution, in the presence of noise we see that saturation and volume images gradually become more distorted

and inaccurate with higher iterations. The addition of noise to the linear reconstructions appears to have had relatively little effect in comparison, although the linear result is still similar to iteration 4 or 5. The reason for the distortion in the non-linear reconstructions is likely to be a subtle propagation of errors through the non-linear reconstruction of the two absolute absorption images. Because the non-linear reconstruction re-calculates the Jacobian for each iteration (see 1.3.2.3), the presence of noise can cause a change in the apparent position of an inclusion. The way in which the two absolute images converge becomes different, and they can no longer be properly compared node-by node.

This implies that, in the case of real data, manipulation of unconstrained non-linear absorption images (particularly higher iterations) to derive saturation and volume can create misleading images, with apparently high resolution structure, due only to mismatches in the spatial localisation between the two μ_a images. Methods of constraining absorption solutions in non-linear reconstructions, in mind of the requirements of saturation and volume derivation are described in section 2.5.6 .

The methods described above for derivation of absolute saturation and volume from absolute μ_a are demonstrated and evaluated on data acquired on the adult breast and exercising forearm later in sections 2.5.4 and 2.5.5 . Below, we use a similar simulation to the above, but consider the effectiveness of having *three* wavelengths and reconstructing the differences between them.

2.5.3 Testing derivation of saturation and volume from 3-wavelength data

Section 2.5.1.3 implied that it is not possible to derive absolute saturation or volume from images showing *changes* in absorption between only two wavelengths. Yet it was found that absolute saturation and volume *can* be derived from only the differences between data acquired at 3 wavelengths ([2.5.10] and [2.5.11]). A further advantage of the 3-wavelength absolute saturation and volume equations is that the $\mu_{a,bg}$ term (representing the absorption of all other chromphores besides blood as in [2.5.4]) is eliminated if assumed to be wavelength independent. Calculating absolute saturation and volume from dual-wavelength absolute μ_a images required an estimate of $\mu_{a,bg}$ which was shown to have a potentially significant effect on derived saturation and volume if incorrect (see Figure 2.5.5).

Superior calibration can be achieved if differences in data are reconstructed, although chapter 2.4 demonstrated that erroneous images could result in cases where initial state heterogeneity is unknown.

The effectiveness of reconstruction for changes in absorption at multiple wavelengths is explored below using the same simulation as before, but with additional data simulated at 800nm. No noise is added to the data, assuming the best-case that data is perfectly calibrated. The values in the simulation are shown in Figure 2.5.14.

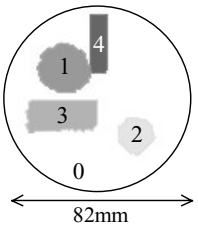

	SO_2 (%)	Vol (%)	μ_a (mm^{-1}) (780nm)	μ_a (mm^{-1}) (800nm)	μ_a (mm^{-1}) (815nm)	μ'_s (mm^{-1})
0	76	4	0.024786	0.025423	0.026527	1
1	98	4	0.023290	0.025526	0.027222	1
2	50	3	0.022414	0.021476	0.021779	1
3	76	6	0.032178	0.033134	0.034790	1
4	76	4	0.024786	0.025423	0.026527	1.5

$\varepsilon_{HbO2,\ 780}$	0.33886
$\varepsilon_{Hb,\ 780}$	0.50874
$\varepsilon_{HbO2,\ 800}$	0.39838
$\varepsilon_{Hb,\ 800}$	0.38668
$\varepsilon_{HbO2,\ 815}$	0.44212
$\varepsilon_{Hb,\ 815}$	0.36316

$\mu_{a,bg}$ =0.01mm^{-1}

Figure 2.5.14 Parameters for three-wavelength derivation of saturation and volume.

We can approach the reconstruction of 3-wavelength data in two ways. We can derive 3 absolute μ_a (an μ'_s) images and then, node-by-node, calculate saturation and volume using [2.5.10] and [2.5.11], eradicating the need to estimate $\mu_{a,bg}$ (assuming the background absorption is wavelength independent), however this method does not improve the ability to calibrate data. The second approach is to reconstruct *only* the changes in absorption using only changes in data. This offers improved calibration, but potentially introduces problems associated with perturbation assumptions (unknown initial state), and also does not yield μ'_s images. The linear version of this method is shown in [2.5.30]. By using changes relative to 800nm, we only need to calculate the Jacobians based on the 800nm state. However these Jacobians should, in theory, contain prior information about the heterogeneous structure in the simulation at 800nm although in practice it is not realistic to assume that such information will be available.

$$\begin{bmatrix}\Delta I_{780-800} & \Delta I_{815-800} & \Delta < t >_{780-800} & \Delta < t >_{815-800}\end{bmatrix}$$
$$= \begin{bmatrix} J_{\mu_a,800}{}^{I} & 0 & J_{\kappa,800}{}^{I} & 0 \\ 0 & J_{\mu_a,800}{}^{I} & 0 & J_{\kappa,800}{}^{I} \\ J_{\mu_a,800}{}^{<t>} & 0 & J_{\kappa,800}{}^{<t>} & 0 \\ 0 & J_{\mu_a,800}{}^{<t>} & 0 & J_{\kappa,800}{}^{<t>} \end{bmatrix} \begin{bmatrix} \mu_{a,780} - \mu_{a,800} \\ \mu_{a,815} - \mu_{a,800} \\ \kappa_{780} - \kappa_{800} \\ \kappa_{815} - \kappa_{800} \end{bmatrix}$$

[2.5.30]

Initially, we explore the effectiveness of using three absolute images to derive absolute saturation and fractional blood volume, using linear and then non-linear methods. The use of wavelength-difference data is then demonstrated.

2.5.3.1 Using 3 linear absolute μ_a images to derive saturation and volume

Three absolute μ_a (and κ) images were reconstructed using the simulated mean-time and intensities by adding two extra columns to the data and Jacobians in [2.5.26]. The images

and their targets are shown in Figure 2.5.15. Each target and image pair is scaled to the maximum and minimum value within the pair. κ images for each wavelength are not shown, but are very similar to each other and to those shown in Figure 2.5.4.

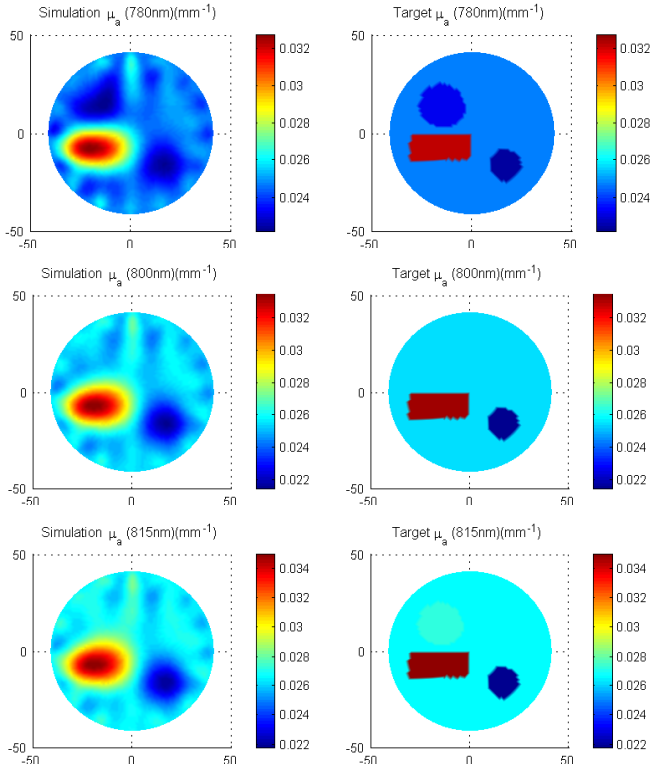

Figure 2.5.15 Three linear absolute absorption images for the three wavelengths: 780nm, 800nm and 815nm.

Figure 2.5.16 shows images representing the differences between the absolute images (800-780nm and 815-800nm). These demonstrate the information being supplied to the absolute saturation and volume calculations. Their targets are also shown on a common scale.

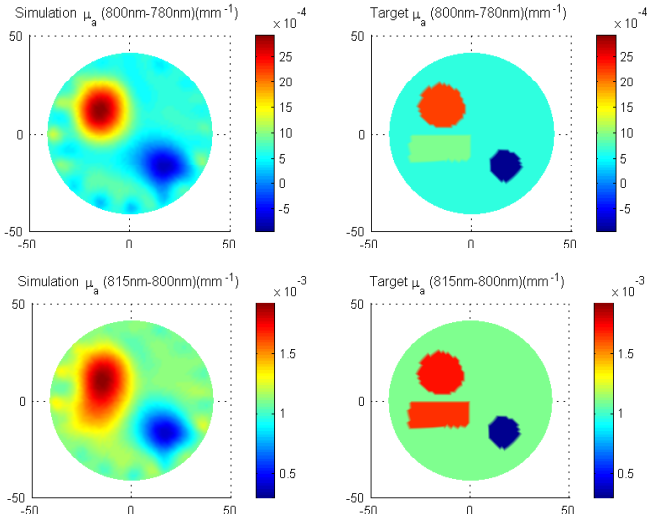

Figure 2.5.16 The difference between linear absolute absorption images (800-780 nm and 815-800nm, representing the information supplied to the absolute saturation and volume calculations.

Figure 2.5.17 shows the absolute saturation and blood volume images calculated from the three absolute μ_a images in Figure 2.5.15, using equations [2.5.10] and [2.5.11] along with their target images.

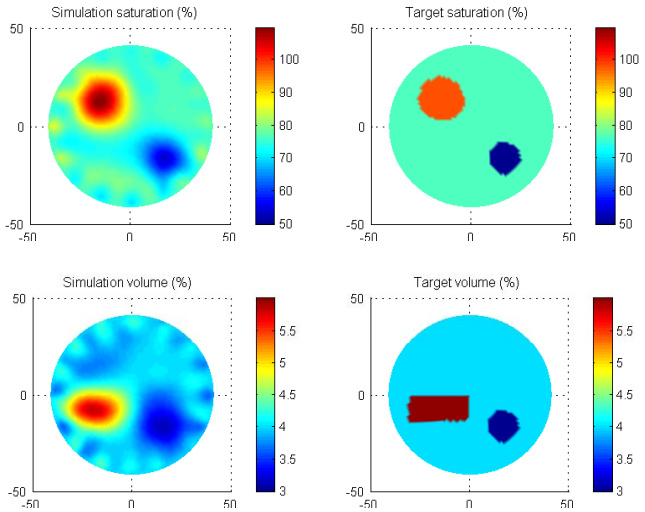

Figure 2.5.17 Absolute saturation and volume images calculated from three absolute linear images (representing 3-wavelength acquisition)

As we would expect the results above are very similar to those shown in Figure 2.5.4, given that no noise was added to the data, and no perturbation conditions were violated. The quantitation is better than the dual-wavelength absolute images, which might be expected given the additional information provided by the extra wavelength.

2.5.3.2 Using 3 non-linear absolute μ_a images to derive saturation and volume

The three-wavelength data were also used to reconstruct three non-linear absorption (and κ) images for comparison with the linear results. The series of iterations for the three wavelengths, along with the target images, are shown in Figure 2.5.18 (the left 3 columns). Each iteration is scaled to its maximum and minimum values. κ images are not shown, but are very similar to each other and to those shown in Figure 2.5.6. Figure 2.5.18 also shows the difference between the non-linear absolute μ_a images (800-780nm and 815-800nm, two right-hand columns), representing the information supplied to the calculation of absolute saturation and volume. Each image showing the difference between the absolute images is shown on a common scale to all the other iterations, and the target image.

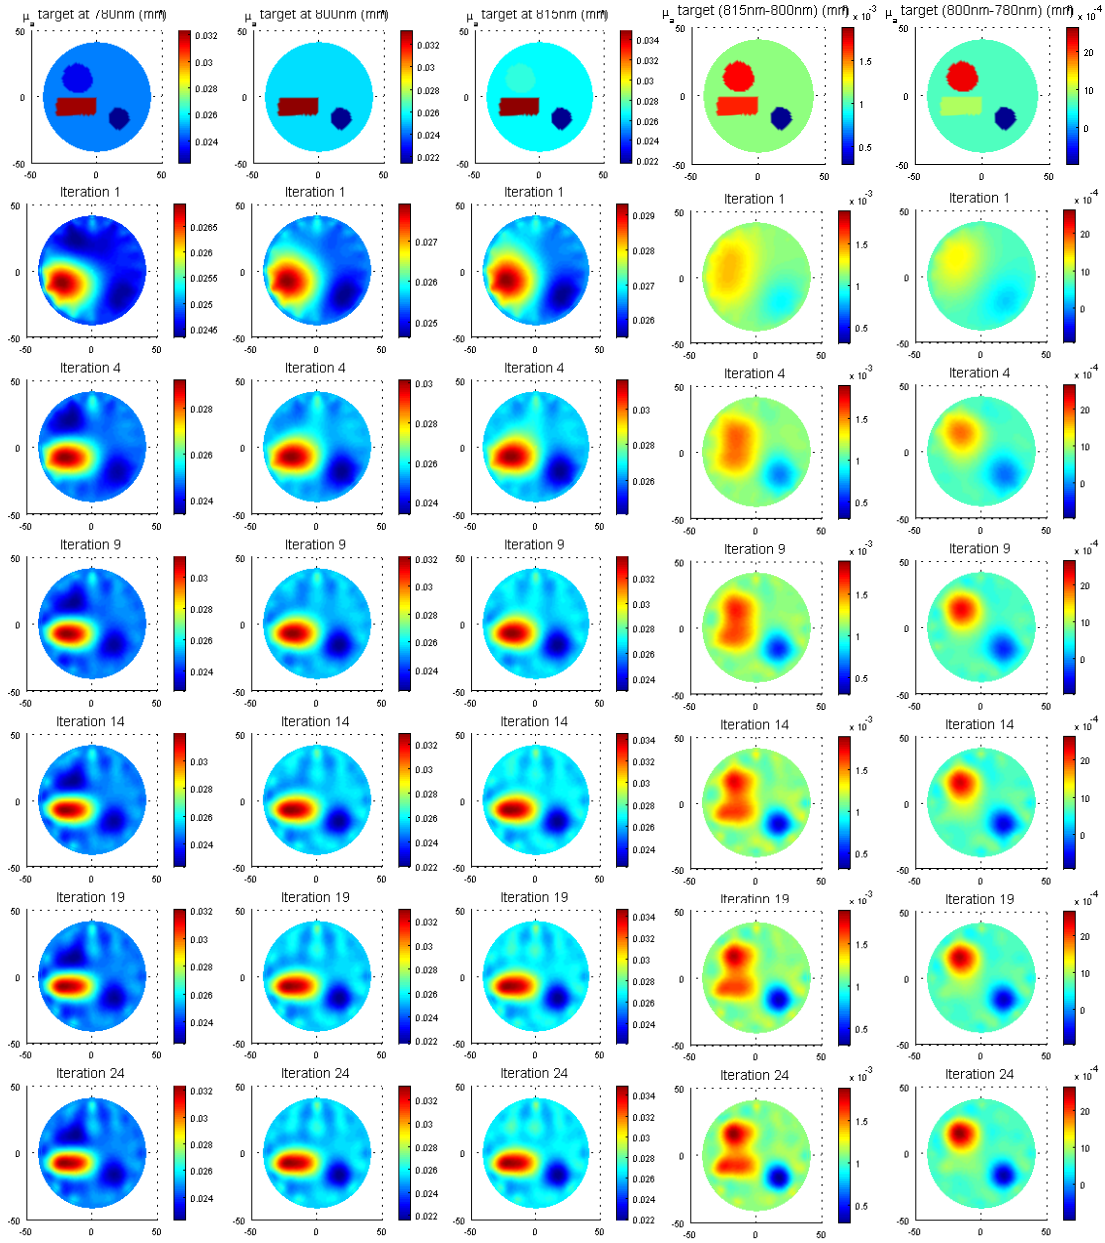

Figure 2.5.18 Non-linear absolute absorption images for 780nm, 800nm and 815nm simulated values. (right had two columns) show the difference between the absolute images (800-780nm and 815-800nm).

Absolute saturation and volume images calculated for each iteration from the set of 3 non-linear absolute absorption images shown in Figure 2.5.18 are shown in Figure 2.5.19. Each iteration is scaled to a common scale along with the target image.

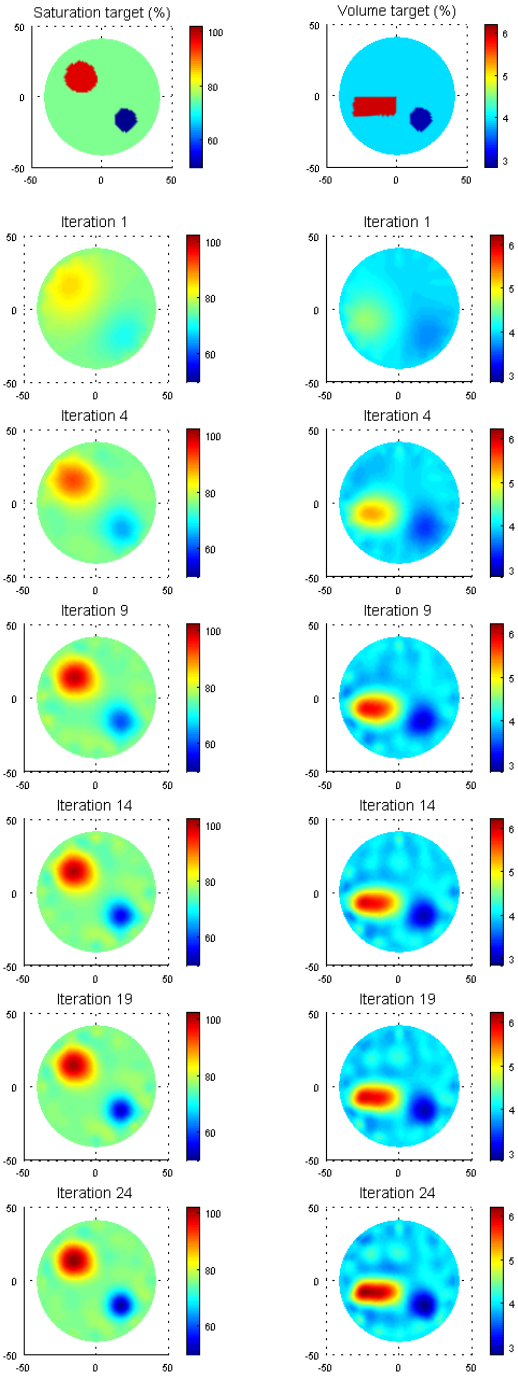

Figure 2.5.19 Absolute saturation and volume images calculated from three absolute non-linear images (780nm , 800nm and 815nm). Each image is scaled to a common scale.

As with the linear case, the three absolute wavelength results are very similar to those calculated using only two wavelengths (Figure 2.5.7) although quantitation and localisation is slightly improved. This has demonstrated that combining 3-wavelength absolute images reconstructed from noise-free data can result in good quality images, *without* knowledge of $\mu_{a,bg}$. However, the promise of the three wavelength equations for absolute saturation and volume lie in their ability to utilise wavelength-difference images only. The effectiveness of using images reconstructed from changes in wavelength is explored below.

2.5.3.3 Using 2 linear λ-difference μ_a images to derive saturation and volume

Linear images of the difference in wavelength between 800-780nm and 815-800nm were reconstructed using the form shown in [2.5.30]. By using 800nm as the starting point for each difference reconstruction, the same Jacobian could be used to generate each wavelength-difference image. The Jacobian was generated based on a homogenous distribution, assuming that the initial state of the difference data is not known see chapter 2.4). Solutions for μ_a and κ were sought, although there is no change in μ'_s between the three wavelengths. Chapter 2.4 demonstrated that difference data would contain information about static background structure, although it is unlikely to be possible to reconstruct anything from the very small signal that it contributes. Therefore the κ results are not shown since they showed little structure and are not used to derive saturation and volume.

The difference images and their targets are shown in Figure 2.5.20. The lambda factor for the Tikhonov regularisation (see section 1.3.2.2) had to be reduced to improve resolution of these images since the changes in data were smaller than for generation of absolute images.

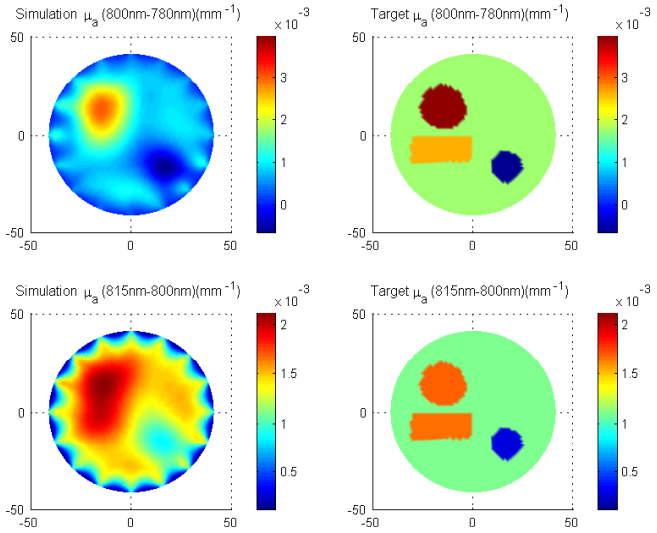

Figure 2.5.20 Linear wavelength-difference images derived from the change in intensity and mean-time between 800 and 780nm, and 815 and 800nm alongside their targets. (Compare with Figure 2.5.16).

Figure 2.5.21 shows absolute saturation and volume images calculated using nodal values of the images shown in Figure 2.5.20 as values for (μ_{a2}-μ_{a1}) and (μ_{a3}-μ_{a2}) in equations [2.5.10] and [2.5.11].

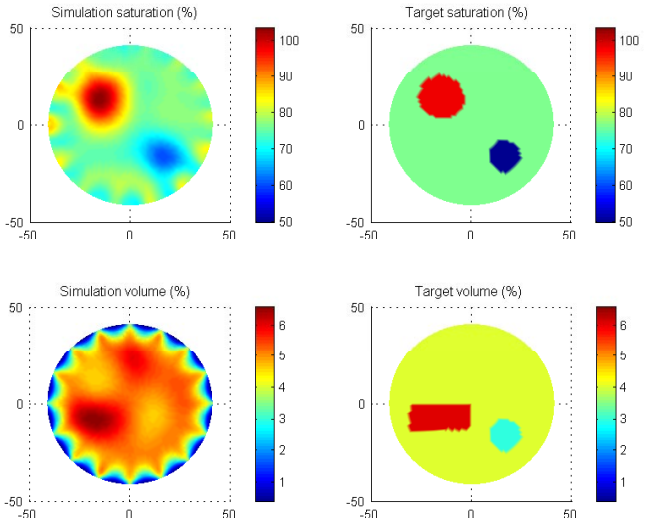

Figure 2.5.21 Absolute saturation and volume images derived from the linear wavelength-difference images shown in Figure 2.5.20 (Compare with Figure 2.5.17).

The saturation and volume images derived from wavelength-difference images are very distorted compared to those derived from either 2 or 3 wavelength absolute images. Two factors are likely to be affecting the images. The first is the fact that the starting condition for the wavelength-difference image is not correct. The input data represented a perturbation from one structured state to another, whereas a homogenous background was used to generate the Jacobian. Secondly, the star-shaped artefact around the edges of the volume image is likely to be due to the inability of the chosen reconstruction method to account for net changes in the background absorption between two states (since region 0 is wavelength dependent). The artefact can also be seen around the difference images in Figure 2.5.20. Since there were only 16 sources and detectors, regions very close to the edge of the mesh are barely sampled, and hence the background change in μ_a is not accounted for in these regions. This is probably worsened by the fact that the Jacobians were generated using 'forward-mesh' basis, causing the solution to be more spiky than it would have been had basis such as 'pixel' been used (see section 1.3.1.3.1). The artefact is not so apparent in the saturation image because equation [2.5.10] has $(\mu_{a2}-\mu_{a1})$ and $(\mu_{a3}-\mu_{a2})$ as multipliers within both numerator and denominator. A constant multiplicative error on both would therefore cancel. This implies that while the values around the edges of the difference images are incorrect, their ratio is still fairly correct. For volume, equation [2.5.11] has $(\mu_{a2}-\mu_{a1})$ and $(\mu_{a3}-\mu_{a2})$ only in the numerator, and therefore the errors remain in the volume image.

2.5.3.4 Using 2 non-linear λ-difference μ_a images to derive saturation and volume

The accuracy of saturation and volume images derived from wavelength-difference images generated using a non-linear reconstruction can also be explored. Figure 2.5.22 shows images reconstructed from the difference in data between 800-780nm and 815-800nm.

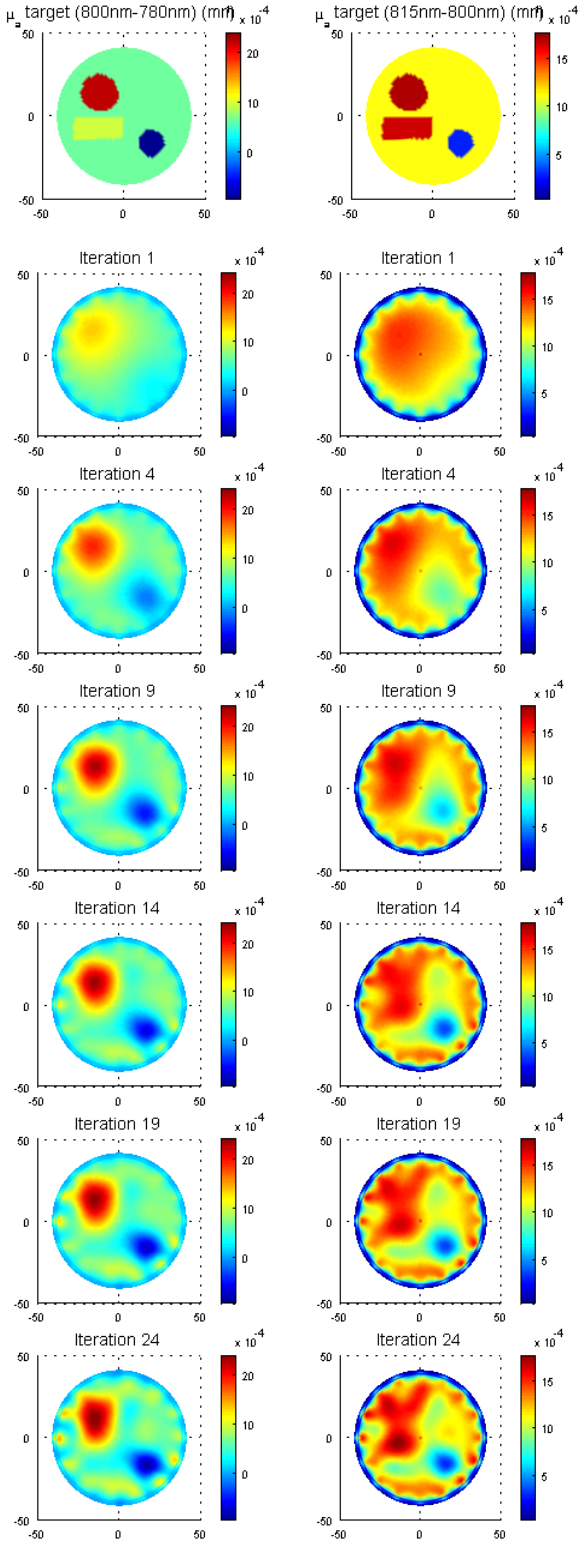

Figure 2.5.22 Non-linear wavelength-difference reconstructions from a homogenous starting condition (Compare with Figure 2.5.18).

The homogenous properties of the background of the 800nm simulation were used as the starting conditions for the reconstruction. The images shown have had this initial μ_a subtracted from them so that they are comparable to the linear images. Each iteration is normalised to the maximum and minimum within the series of iterations and the target. Figure

2.5.23 shows absolute saturation and volume images derived from the wavelength-difference images in Figure 2.5.22.

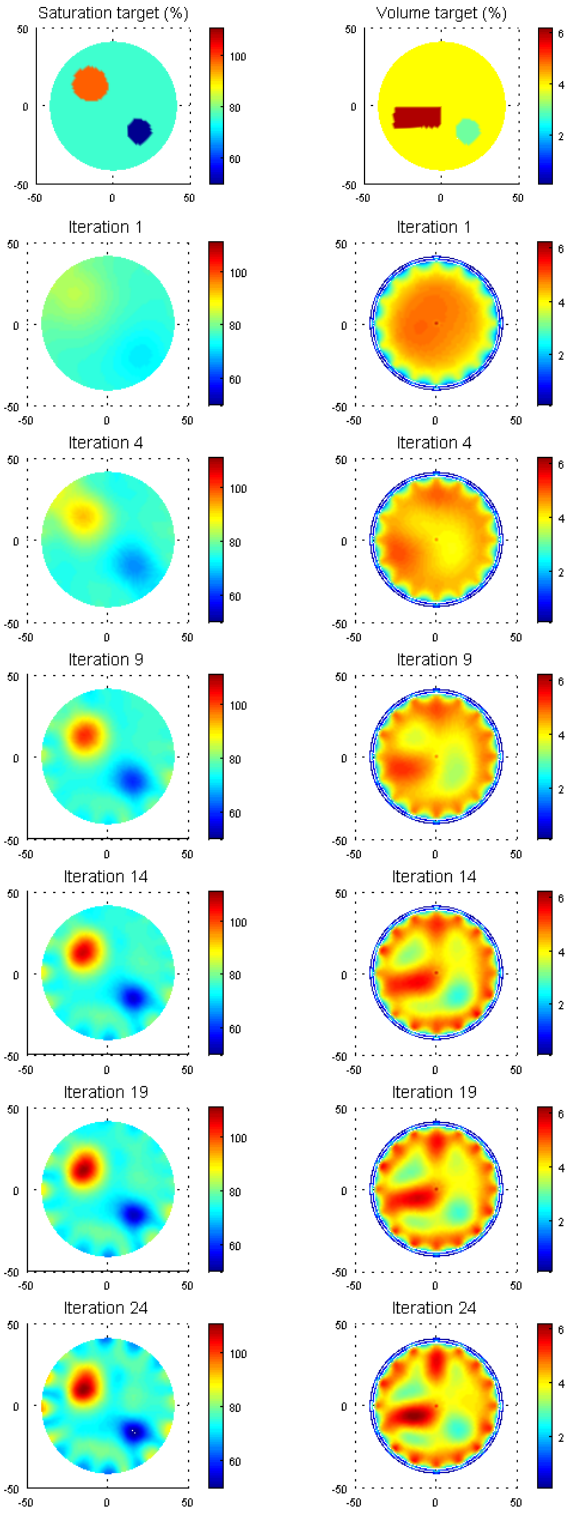

Figure 2.5.23 Absolute saturation and volume images derived from non-linear wavelength-difference images shown in Figure 2.5.22

In the same way as in the linear wavelength-difference results, edge artefacts have caused the non-linear volume images to have a bad ring artefact. The presence of the

scattering object is evident in the volume image too, implying that its presence is not accounted for by using a homogenous Jacobian (see chapter 2.4).

The saturation image becomes quite distorted with higher iterations, but is both quantitatively and qualitatively reasonable. The edge artefacts that have caused distortion in the volume images may be reduced both for linear and non-linear reconstructions by applying specific regularisation to the reconstruction. Pixel-basis for example will provide a much smoother solution for (μ_{a2}-μ_{a1}) and (μ_{a3}-μ_{a2}), although it is possible that applying a basis will change the way in which the two wavelength-difference images converge and cause distortion in the resulting saturation and volume images. A continuation of this investigation would be to apply better regularisation to the solutions to improve the absolute volume results. Further analysis is also required to determine the magnitude of the effect of heterogeneous background structure. Preliminary results (not shown) implied that while compensating for the unknown background structure did improve images, the star-shaped artefact remained.

2.5.3.5 Sensitivity of three-wavelength equations to errors in μ_a

Figure 2.5.24 shows plots of saturation and volume derived from three μ_a values with either additive or multiplicative errors on all or just one of the three μ_a values using [2.5.10] and [2.5.11]. The values are based on those of region 0 (see Figure 2.5.14).

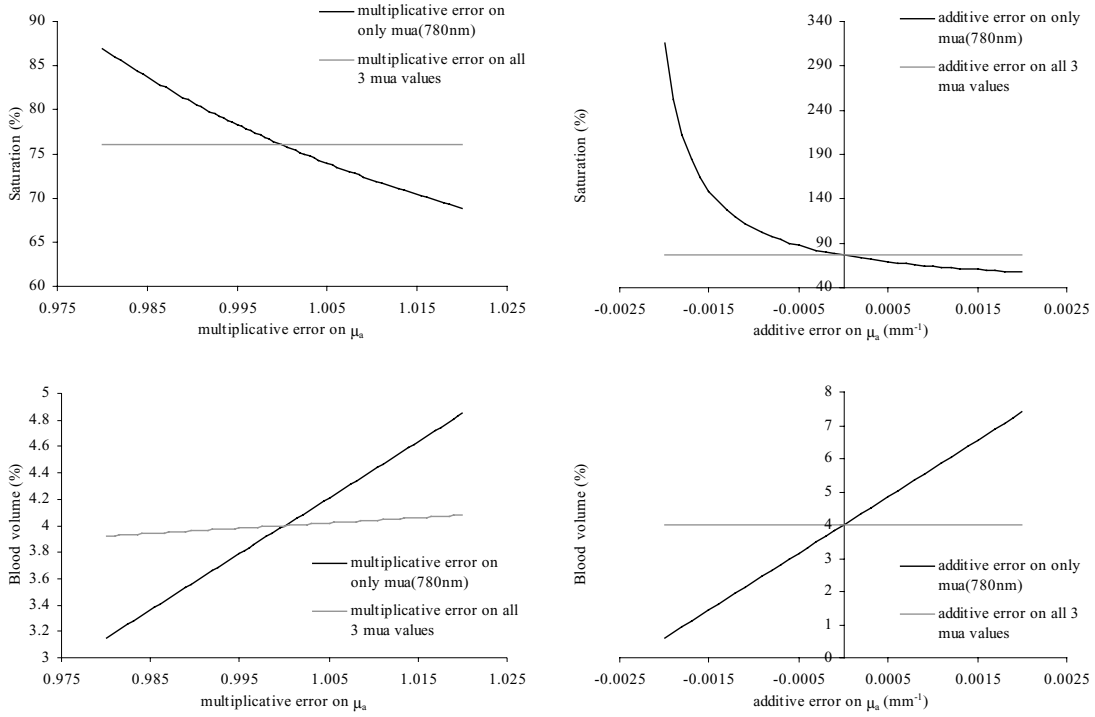

Figure 2.5.24 Plots of saturation and volume derived from erroneous three wavelength μ_a values (corresponding to region 0).

The saturation and volume are shown to vary with errors on the μ_a values. Saturation is shown to be independent of both additive and multiplicative errors, so long as the same errors affect all three μ_a values (or their differences). This was found to be true for edge artefacts in sections 2.5.3.3 and 2.5.3.4. However, if different errors occur in each wavelength, the saturation values derived can be badly affected. This is equivalent to the sort of errors likely to result from reconstruction of noisy data as demonstrated in sections 2.5.2.3 and 2.5.2.4. Volume values are shown to be affected by identical multiplicative errors in all three wavelengths, which accounts for the star-shaped artefacts in the volume images shown above. Volume values are more badly affected by multiplicative errors that are different between the three wavelengths. Additive errors affecting all μ_a values don't affect volume (since differences in μ_a are all that matter), however when different additive errors affect each μ_a the derived volume values will be wrong. Of course, unlike the two-wavelength equations ([2.5.5] and [2.5.6]) examined in section 2.5.2 , the three wavelength equations do not contain a $\mu_{a,\,bg}$ term, which eliminates a large potential source of error.

Using three-wavelength-difference data to calculate absolute saturation data has been shown to be effective. Volume images are affected by edge artefacts that could potentially be improved by smoothing and regularisation. Wavelength-difference images themselves may contain errors due to unknown background conditions, although the severity of the error is likely to depend on the specific background distribution, which will depend on for example whether brain or breast is being imaged. The effect of this error could be quantified with the development of more clinically oriented simulations. Methods of potentially correcting these errors were demonstrated in section 2.4.6 .

As was shown in section 2.5.1.4, state-change difference data cannot be used to generate saturation changes. If a three wavelength system were available, and the problems described above can be overcome, it may eventually be possible to derive changes in saturation and volume, along with static distributions using purely wavelength-difference data (although μ'_s would not be available). If problems with wavelength-difference imaging are insurmountable, this section has demonstrated that three *absolute* images can be used to generate a better estimate of absolute saturation and volume that only two absolute images. An added advantage is that the three wavelength equations do not require an estimate of $\mu_{a,bg}$ (assuming it is wavelength independent), suggesting that a three wavelength device would offer a number of advantages over a two wavelength system.

2.5.4 Reconstructing saturation and volume from breast data

Simulations in sections 2.5.2 and 2.5.3 have demonstrated that the retrieval of functional parameters, using linear and non-linear reconstructions, is not perfect even for data generated from μ_a values corresponding exactly to the saturation and volume values sought. Further, the addition of very little noise, compared to absolute measured data, onto data causes later iterations of non-linear absolute μ_{a1} and μ_{a2} images to distort such that they infer erroneous saturation and volume distributions when combined. While the results above imply that better ways of deriving saturation and volume may exist (see section 2.5.6), they demonstrate that we cannot expect to extract very accurate and reliable saturation and volume images from clinical data. Further aspects of experimental imaging that are likely to degrade absolute image quantitation and localisation (as introduced in preceding chapters) include:

- Calibration (monstodes or absolute) errors see section 2.1.2 .
- Calibration (reference phantom or state change) errors and unknown reference state see section 2.1.2.3.
- Inadequate datatypes (<t> only, otherwise need to use reference), see section 1.3.3 .
- 2D – 3D correction section 2.3.1.1.
- Mesh resolution, reconstruction basis, other TOAST settings, iteration chosen and general model accuracy.

Since we have access to a number of sets of data from clinical subjects, dual-wavelength data is used below to derive images of saturation and volume, using the processing methods introduced in section 2.5.1 . The purpose of this section is to explore whether physiologically reasonable values for saturation and volume are retrieved from measured data. However, much more work will be required to properly evaluate whether the functional images presented are truly accurate. This may include future investigations for example using blood-based phantoms. In addition, results could be compared with positron emission tomography images (PET - where a radioactive oxygen isotope is used to trace the functional activity in the body), functional magnetic resonance imaging (FMRI - where the magnetic properties of deoxy-haemoglobin are exploited to investigate regions of changing activation) or even invasive techniques such as tissue biopsies.

2.5.4.1 Breast saturation and volume images from linear absolute images

Data were acquired using MONSTIR on the breast of a (40 year old) healthy adult volunteer using the technique described in section 2.7.3 . Figure 2.5.25 shows images of absolute $\mu_{a,780nm}$ and $\mu_{a,815nm}$ derived using a linear reconstruction of the real data. Reference data were acquired on a homogenous phantom immediately after the breast data were acquired. The Jacobians for the linear reconstruction were generated based on the phantom's

properties, using a 2D mesh with 5890 nodes. Since difference data is being used, the 2D:3D mismatch is likely to have only a small effect (see section 2.3.1.4) Figure 2.5.26 shows the corresponding κ images of the breast. Raw intensity and raw mean were extracted from measured TPSFs, and the difference between the breast and homogenous phantom data read-in to the linear reconstruction.

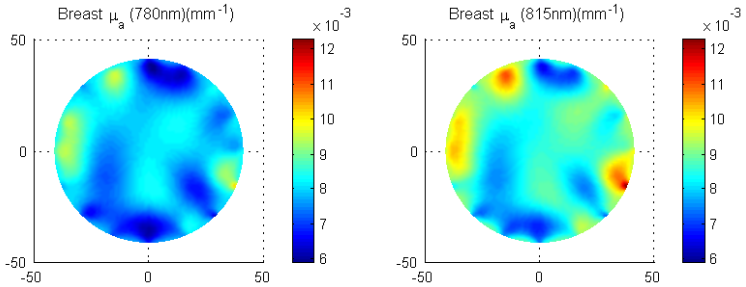

Figure 2.5.25 Absolute 780nm and 815nm absorption images of the adult breast, derived using a simple linear reconstruction

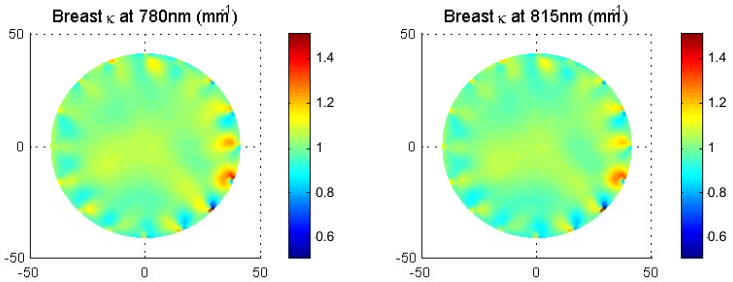

Figure 2.5.26 Absolute linear kappa image of the breast at 780nm and 815nm

Figure 2.5.27 shows the corresponding oxygen saturation and blood volume images derived from the absolute μ_a images in Figure 2.5.25 and the estimate that $\mu_{a,bg} = 0.003\text{mm}^{-1}$.

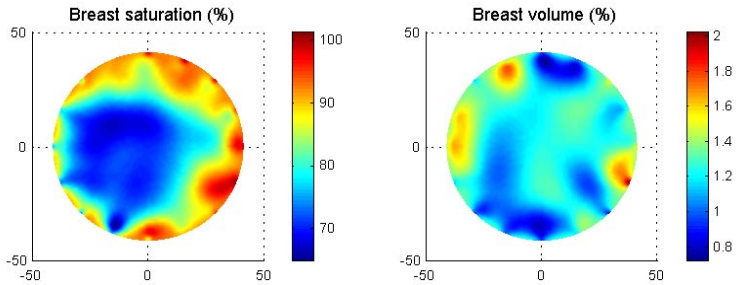

Figure 2.5.27 (left) Oxygen saturation image (right) fractional blood volume image of the adult breast derived using absolute absorption values derived using a simple linear reconstruction.

Structure can be seen in the saturation, volume and kappa images. The saturation image implies that the central area of the breast has a low saturation compared to the areas closer to the surface. The volume image may have revealed the presence of veins running under the surface of the skin, although the features may also be attributable to artefacts representing noisy sources or detectors. This may be a similar effect to that seen in three-wavelength derivations where, as discussed in section 2.5.3.5, the fact that μ_a does not appear in the denominator of the volume equation, multiplicative errors in two absolute images are less

likely to cancel than in saturation images. The kappa images also have edge artefacts, although they do reveal an area of lower scatter that appears to coincide with the region of low perfusion and low blood volume.

While the simulations in sections 2.5.2.1 and 2.5.2.3 demonstrated that these results are not likely to be entirely quantitatively accurate, this method has revealed some spatial heterogeneities in functional parameters within the breast. The saturation and volume values derived are similar to the average values found by (McBride *et al*, 1999) and (Cerussi *et al*, 2001).

Early reconstructions of these data assumed the homogenous phantom to have nominal values of μ_a = 0.007mm^{-1} and μ'_s = 1mm^{-1}. However, detailed analysis of the raw data revealed that the homogenous phantom itself had fairly significant wavelength dependence. Without correcting for this, the two absolute images would be incorrect, not only since the Jacobians would be wrong, but one wavelength would be offset relative to the other. By using the method described in appendix A (2.9.7) (monstode) calibrated mean-times extracted from the homogenous phantom data were plotted v/s distance, and assuming that μ'_s is wavelength independent and ~ 1mm^{-1} the phantom was found to have ~ μ_a = 0.0077mm^{-1} at 780nm and μ_a = 0.0088mm^{-1} at 815nm. The Jacobians were therefore calculated based on these values, and these values were added onto the resulting difference images to convert them into absolute μ_a before calculation of saturation. For comparison, Figure 2.5.28 shows the saturation and volume images derived from these data, but without accounting for the wavelength dependence of the reference phantom. Clearly the wavelength dependence of the reference phantom has affected quantitation quite badly. In all cases a Tichonov regularisation factor of 0.2 was used.

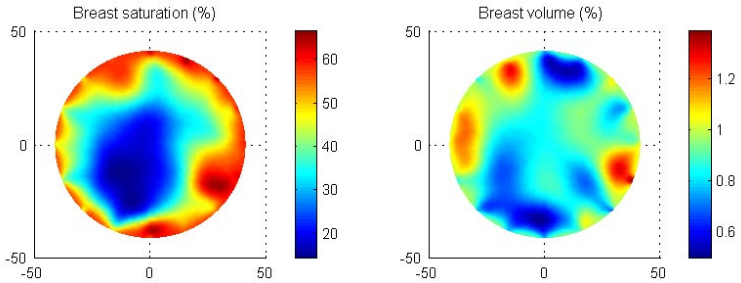

Figure 2.5.28 Saturation and fractional blood volume images of the adult breast without correction for the wavelength dependence of the homogenous reference phantom.

Figure 2.5.29 shows saturation and volume images derived from data acquired on a second adult volunteer. Again a region of lower saturation is seen nearer the centre of the breast, although a small region of high volume and saturation is visible near the very centre. The saturation values that are > 100% imply that the absolute images have been affected by a bad source or detector on the right hand side of the image.

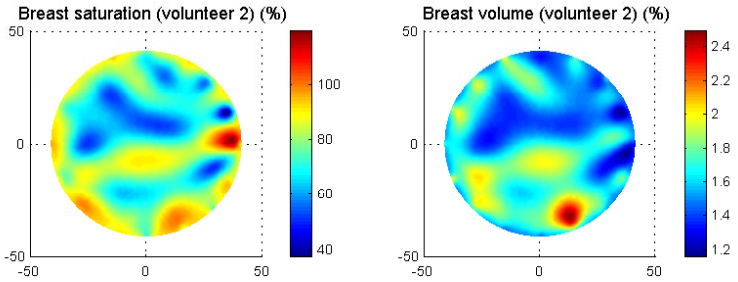

Figure 2.5.29 Volunteer 2: (left) Oxygen saturation image (right) fractional blood volume image of the adult breast derived using absolute absorption values derived using a simple linear reconstruction.

The values of the extinction coefficients of whole blood used to calculate saturation and volume were the same as those used in the simulations, and assume a constant haematocrit of 2 mML^{-1}. The estimate that $\mu_{a,bg} = 0.003mm^{-1}$, used to calculate saturation and volume from μ_a images, is based on the fact that water absorption between 780nm and 815nm is $0.0022mm^{-1} \pm 0.0004mm^{-1}$ and that the absorption of lipids and other constituents of the tissue may be slightly higher than this. The value for $\mu_{a,bg}$ could also not be higher than the minimum μ_a value in the absolute μ_a images (~$0.006mm^{-1}$). The wavelength dependence of this background is assumed to be negligible in the equations used. This is a reasonable assumption over this small wavelength range (780-815nm) but would not hold if there was a much greater spread in the wavelengths used.

Figure 2.5.30 shows plots of the values of saturation and volume that would result from erroneous values of $\mu_{a,bg}$ and identical additive and multiplicative errors in the two absolute μ_a images.

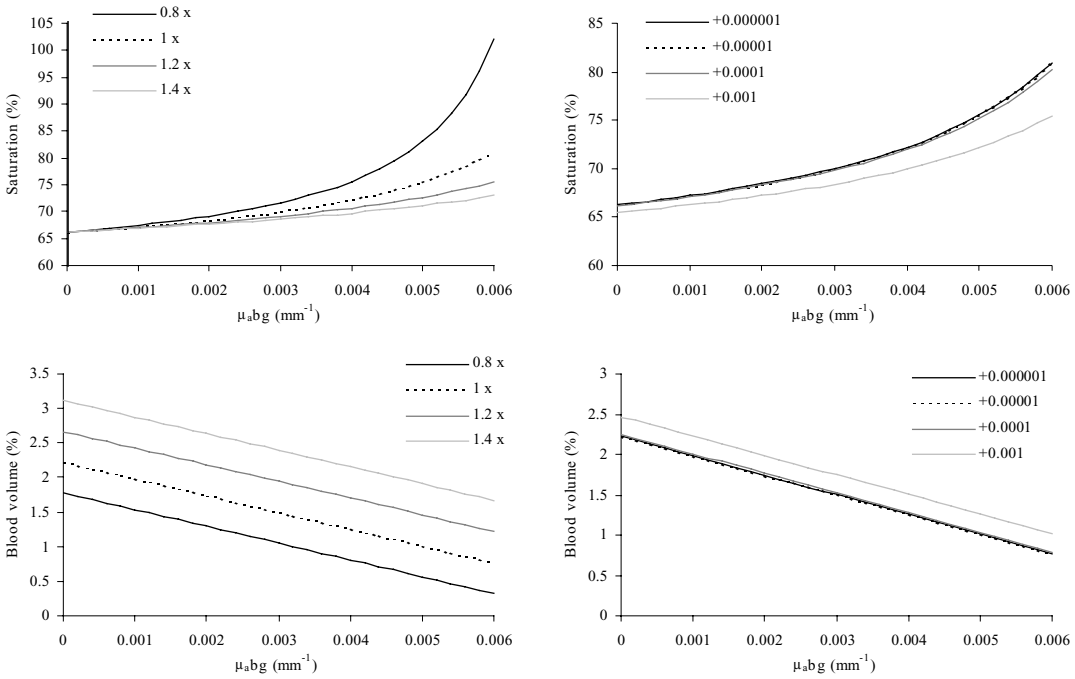

Figure 2.5.30 Plots showing the variation of calculated saturation (top) and blood volume (bottom) values when background μ_a is incorrect, or identical additional (left) or multiplicative (right) errors affect the absolute μ_a multi-wavelength images. (using 70% saturation 1.5% volume, $\mu_{a,bg} = 0.003mm^{-1}$)

The plots in Figure 2.5.30 are based on values corresponding to 70% saturation, 1.5 % volume and $\mu_{a,bg}$ = 0.003mm^{-1}, ($\mu_{a,\ 780}$ = 0.0088mm^{-1} and $\mu_{a,815}$ = 0.0092mm^{-1}). The plots demonstrate that if the estimate of $\mu_{a,bg}$ is too low, that saturation will be too low, and volume will be too high. Additive errors will affect saturation and volume less than multiplicative errors. Note however that as for the three wavelength saturation and volume equations described in section 2.5.3.5, *different* errors in each absolute image are likely to cause erroneous saturation and volume images.

2.5.4.1.1 Image summary (Linear absolute breast images)

Subject	Breast: healthy volunteers 1 and 2
Mesh	41mm radius, circular 2945 nodes 5704 linear elements (2D)
Basis (used to generate PMDFs)	Forward mesh
Starting parameters: μ_a	0.0077mm^{-1} @ 780nm and 0.0088mm^{-1} @ 815nm
Starting parameters: μ'_s	1 mm^{-1}
Iterations	Linear λ=0.2
Sources	16
Detectors per source	11
Datatypes	Raw mean + raw intensity
Calibration	Difference to homogenous phantom
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	15 secs
Wavelength	780 nm + 815nm (interlaced)
μ_a background estimate	0.003mm^{-1}

Table 2.5.1 Properties of images shown in section 2.5.4.1

2.5.4.2 Breast saturation and volume images from non-linear absolute images

The same data acquired on the breast of the 1st volunteer were also reconstructed using TOAST to allow comparison between linear and non-linear reconstructions. Two sets of absolute absorption and scatter images were reconstructed separately, starting from initial estimates based on the wavelength dependent values of the reference phantom. Intensity and mean-time (raw) difference data were used, so a 2D-3D correction was not applied (see section 2.3.1.1). Successive iterations of the absolute μ_a and μ'_s images are shown in Figure 2.5.31. Each image is scaled to its maximum and minimum values as shown by the colour-bar to the right of each iteration.

As with the simulated images, the linear results resemble the third or fourth iteration of the non-linear results. As the iterations increase, a point at 4 o'clock starts to dominate the image, and is likely to be due to bad data from the source or detector at that position. This is similar to the effects we see in simulated images with noise added as shown in Figure 2.5.12. Note that clearly erroneous data points were removed from measured data during pre-processing, but some subtler errors will remain.

The μ'_s images also look similar to the linear result. The edge artefacts may imply that modelling of the distributed source (due to the use of monstodes) is not suitably accurate, or

that the initial background scatter properties are quite far away from $\mu'_s = 1mm^{-1}$, or that there is little genuine heterogeneity in the scatter in the breast of this subject. The use of a smoothing reconstruction basis (e.g. pixel basis) may also improve these edge artefacts (see section 1.3.1.3.1).

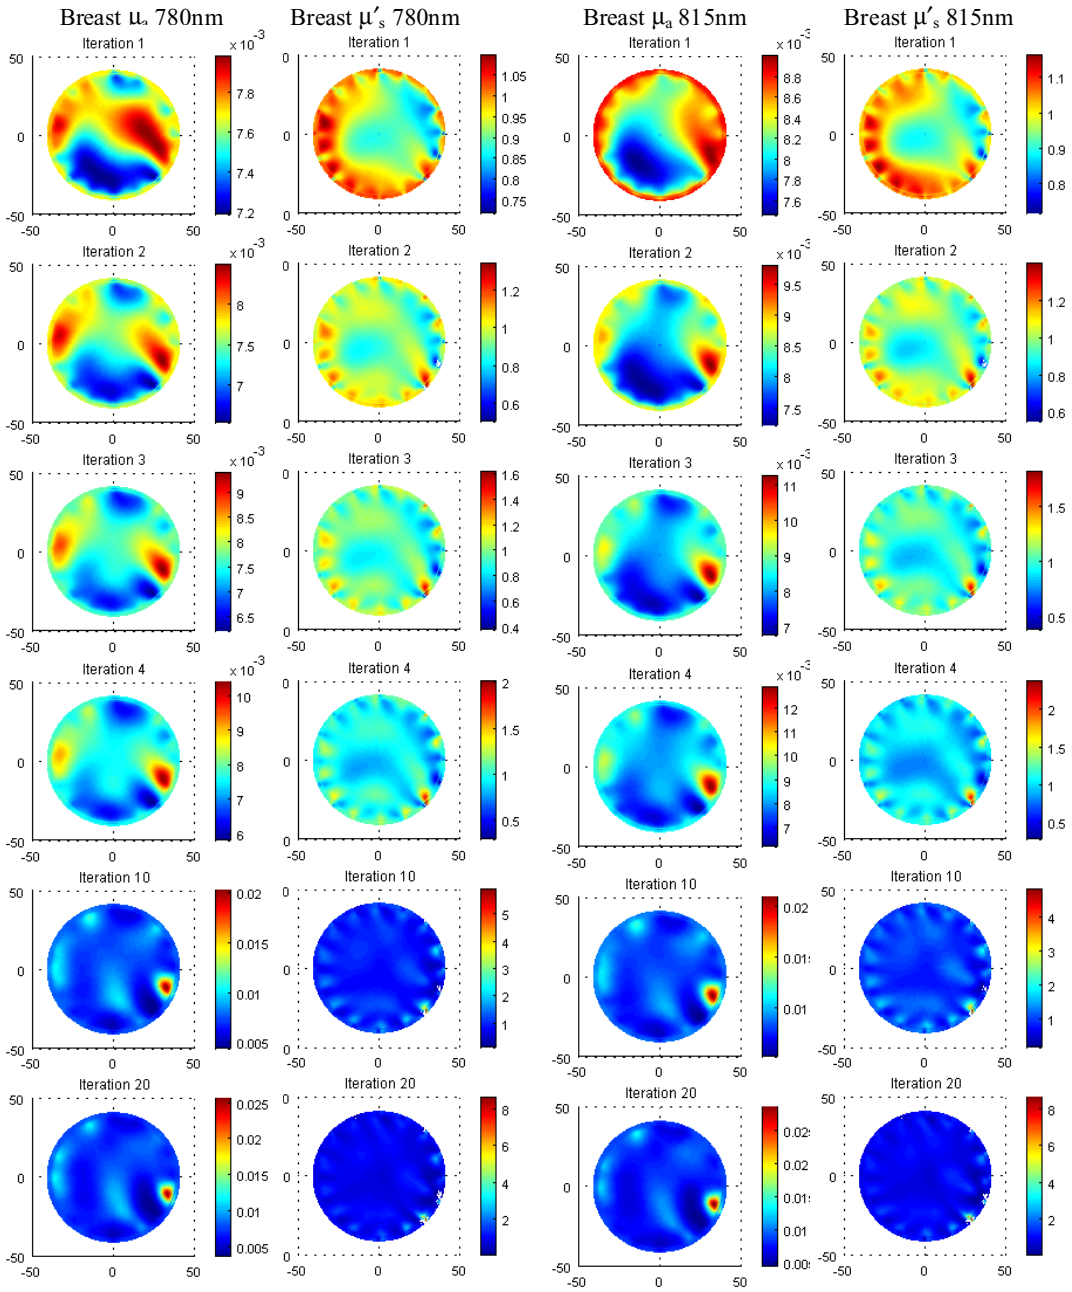

Figure 2.5.31 Non-linear breast images: absolute μ_a and μ'_s at 780 (left) and absolute μ_a and μ'_s at 815nm (right). Each successive iteration is scaled between its maximum and minimum values.

The absolute saturation and volume images derived from the absolute μ_a images in Figure 2.5.31 are shown in Figure 2.5.32. As with the simulations, each iteration pair was used to derive a corresponding saturation or volume 'iteration'. Again the non-linear results resemble the 3rd or 4th iteration of the non-linear images. Quantitation is very similar between

the linear images and these iterations of the non-linear results. As the feature due to the bad source or detector begins to dominate, saturation values increase to greater than 100%.

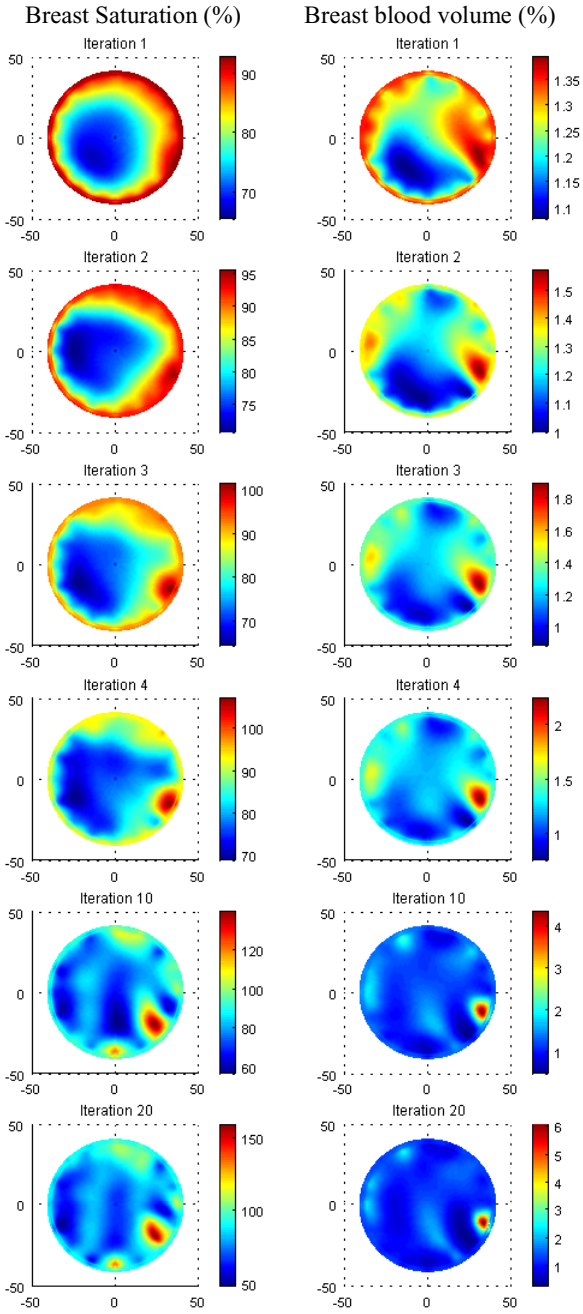

Figure 2.5.32 Non-linear absolute saturation and volume images derived from absolute absorption images shown in Figure 2.5.31.

2.5.4.2.1 Image summary (Non-linear absolute breast images)

The reconstruction parameters for the non-linear images were identical to those shown in the linear summary Table 2.5.1 except iterations 1 to 20 are shown rather than linear results with $\lambda = 0.2$. Forward mesh basis was used.

2.5.5 Reconstructing saturation and volume from arm data

Preliminary studies of the arm using MONSTIR were published in (Hillman *et al*, 2000b) and (Hillman *et al*, 2001b). The Ti:Sapphire laser was used, and in the latter publication, state-difference images at two wavelengths were presented, although extraction of functional parameters was not attempted. Below, reconstructions are performed using a set of data acquired on the arm using the new dual-wavelength pulse fibre-laser. The subject was 'subject B' referred to in (Hillman *et al*, 2001b).

The experimental protocol consisted of acquiring resting data (with simultaneous 2-wavelength acquisition) followed by an active set of data where the subject used the fingers to grip a tension meter. A monstode calibration measurement was then made followed by another resting and then a further active data set. Finally a homogenous phantom data set was acquired. The 32 monstode were positioned around the arm in two rings of 16. The rings were 23mm apart. For this preliminary analysis, image reconstruction was performed using only data from the first 16 sources and detectors (within one planar ring) since evaluating performance on a 2D mesh is much faster and provides images that can be directly compared with earlier results. An automated system was used to calculate the optimal VOA setting for each source detector pair, so each source illuminated between 11 and 14 detectors around the ring of 16 (see section 1.2.3.3).

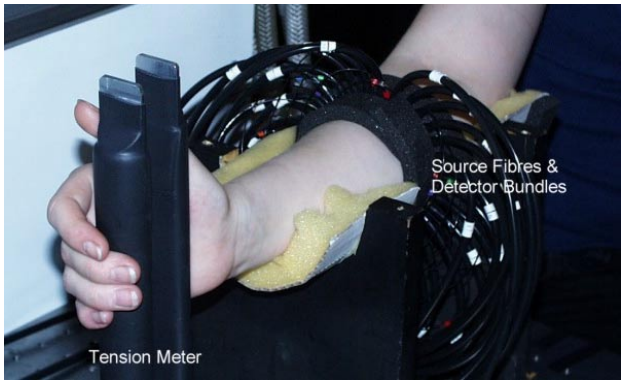

Figure 2.5.33 A photograph of the experimental set-up used for arm imaging experiments

Originally it had been anticipated that the homogenous phantom data could be used as a reference data set to allow linear reconstruction of the data (and use of intensity as a datatype for absolute imaging). However the optical properties of the arm were found to be so very different from those of the phantom that the data was unsuitable for reconstruction using the simple perturbation approximation.

Below we present images of the arm that have been processed to explore the methods of reconstructing dual-wavelength data described above, and to examine the capacity of these techniques to detect activation-related changes in functional parameters. The results also provide additional validation to the results published in (Hillman *et al*, 2001b).

2.5.5.1 Saturation and volume arm images from absolute absorption images

Absorption and scatter images were reconstructed using TOAST from (monstode) calibrated mean-time data. As with the previous arm images, an ad-hoc 2D-3D correction was applied by using a simulated 3D data-set based on the homogenous background properties of the arm as a reference data set, and starting the 2D reconstruction from the same background properties (see section 2.3.1.4). Starting values of $\mu_{a,780nm}$ = 0.03296mm^{-1} and $\mu_{a,820nm}$ = 0.03591mm^{-1} were estimated as described in appendix A (2.9.7 . Figure 2.5.34 shows the resulting μ_a and μ'_s images (iteration 3 throughout) for each of the two active and resting experiments at the two wavelengths.

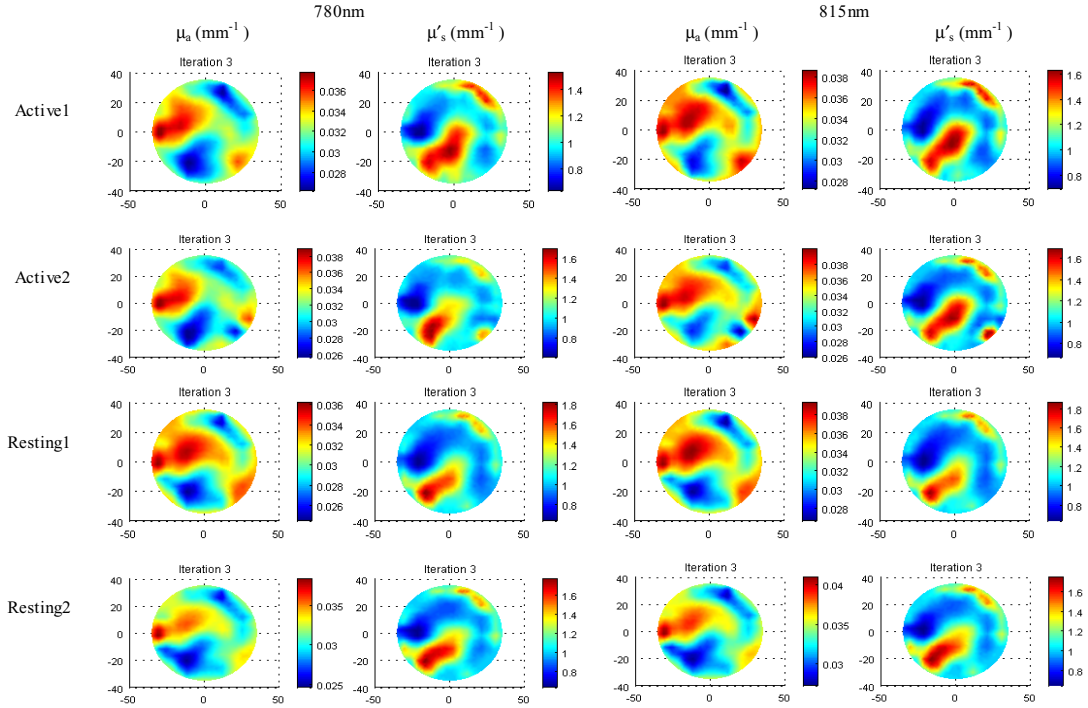

Figure 2.5.34 Absolute absorption and scatter images (iteartion3) mean-time only for all wavelengths and states.

Figure 2.5.35 shows the anatomy and an MRI of subject B's arm (as described in (Hillman *et al*, 2001b)) rotated to correspond with the images from the new data. The pink regions on the MRI indicate the likely position of muscles involved in the finger flexing exercise performed.

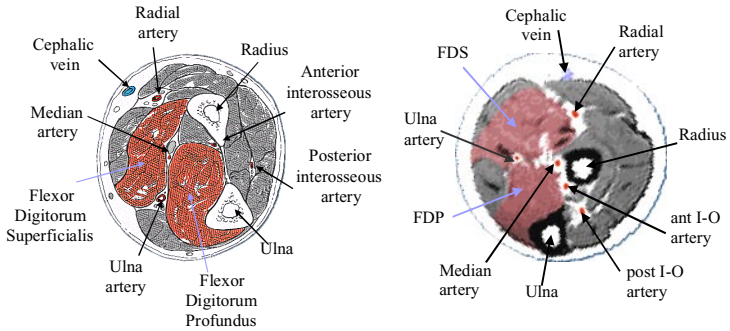

Figure 2.5.35(Left) the basic anatomy of the forearm (modified from (Hall-Craggs, 1995)), the Flexor digitorum profundus and superficialis are highlighted (right) an MRI of the arm of subject B rotated to correspond to the geometry of the latest images.

The red regions on the absolute scatter images represent the bones, as visible in the MRI. The images are similar to those shown in (Hillman *et al*, 2001b), although the lower resolution of these new images may relate to the reduced number of sources and detectors used (16 coincident sources and detectors compared to 32 interspersed sources and detectors). As in (Hillman *et al*, 2001b), it is quite possible that the images resulting from the use of just mean-time data will suffer from cross talk between absorption and scatter. Indeed since absorption causes a decrease in mean-time but scatter causes an increase, we would expect μ_a and μ'_s images to be the inversion of each-other if affected by cross talk. The images in Figure 2.5.34 do appear to have some correlation between low absorption areas and high scatter areas and visa versa. However it is difficult to ascertain whether this is because the areas of high scatter, e.g. corresponding to the bone, actually *are* less absorbing than surrounding tissue. The series of iterations of the absolute images shown in Figure 2.5.34 were used to derive images of saturation and volume for each resting1, active1, resting2 and active2 state. The absolute saturation images are shown in Figure 2.5.36. Each active and resting pair is shown on the same colour-scale.

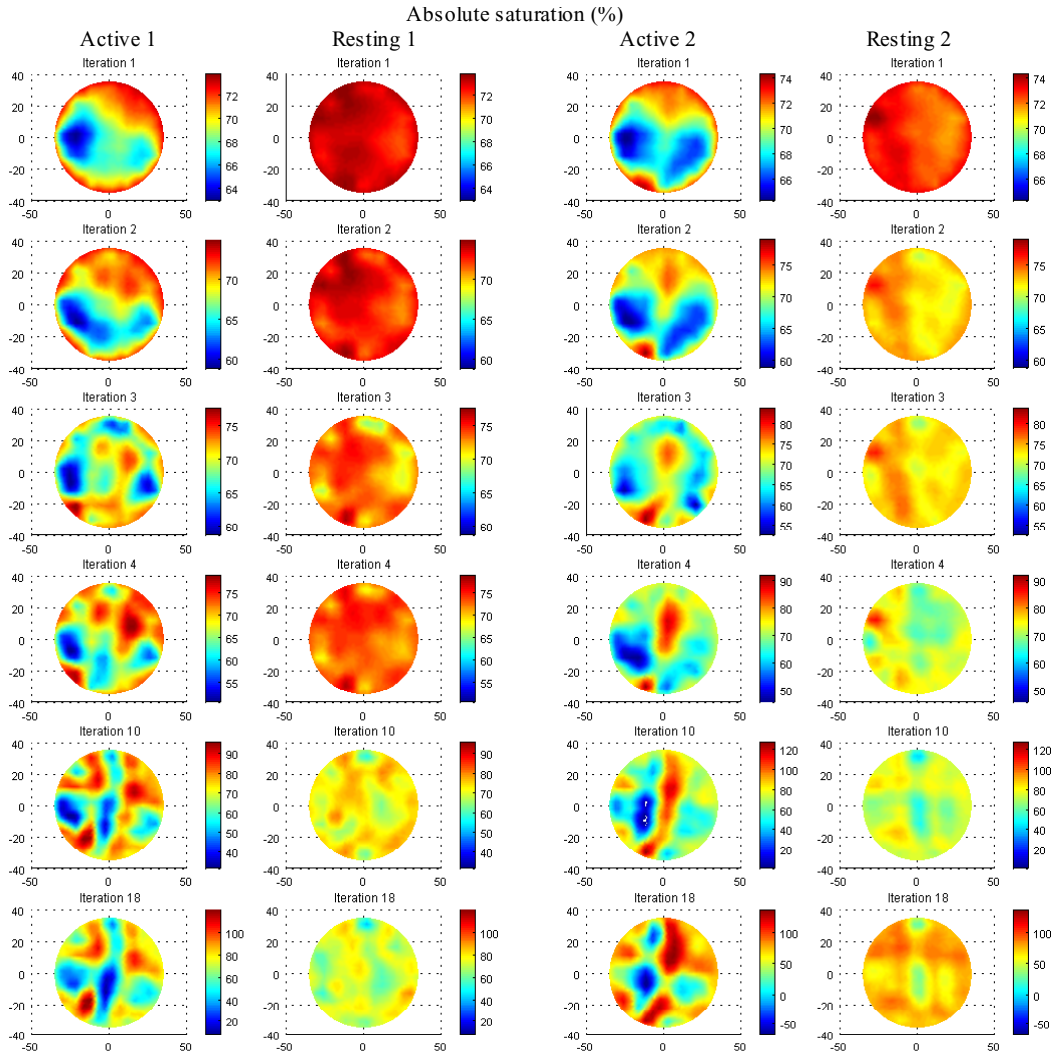

Figure 2.5.36 Absolute saturation images of the arm, active and resting, fromabsolute μ_a images (Figure 2.5.34.)

Figure 2.5.37 shows the corresponding fractional blood volume images derived from the absolute μ_a images in Figure 2.5.34. Again each active- resting pair have been scaled to the same colour range. As with the breast images, a value of $\mu_{a.bg} = 0.003mm^{-1}$ was used, along with the extinction coefficients based on adult blood at 2mM.

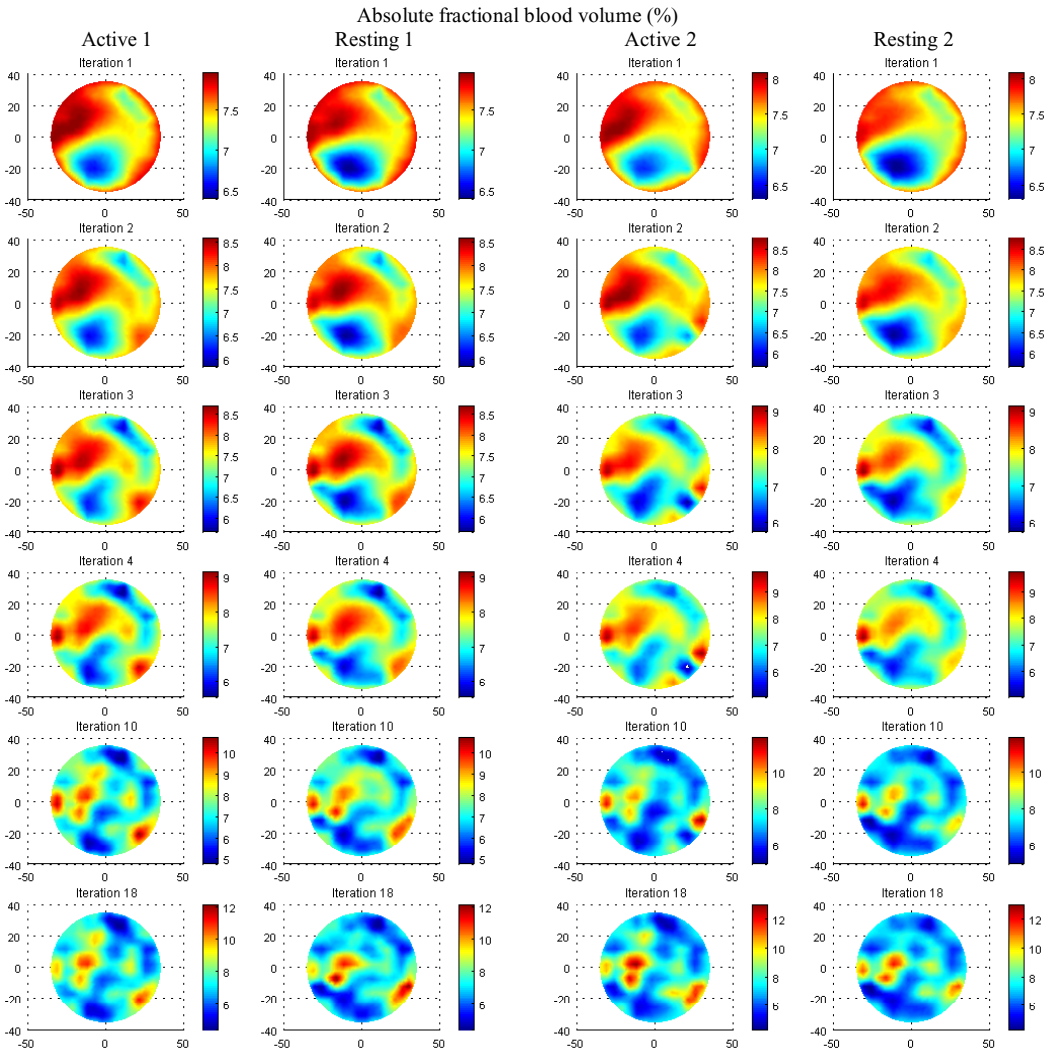

Figure 2.5.37 Absolute fractional blood volume images of the arm while active and resting derived from the absolute μ_a images in Figure 2.5.34.

These results show a definite apparent change in oxygen saturation with exercise in the forearm. The volume results imply that volume is very consistent between data sets, although a slight change can be seen between active and resting sets. As the iterations increase we see an increase in high spatial frequency structure, likely to be due to the distortion of the absolute μ_a images due to noisy data (as demonstrated by the simulation in Figure 2.5.13).

As with the breast images however we must assume that experimental errors will propagate through the derivation of saturation and volume. In this case we are using 2D-3D corrected mean-time only to reconstruct functional parameters, which is unlikely to produce quantitatively reliable results. However it is undeniable that a change in saturation is visible

between active and resting states. It is not possible to eliminate the possibility that, through slight movement of the arm during exercise, the coupling between skin and the monstodes changed. Or that slight movement of the bones has caused a distortion in the absorption images that appears as a saturation change. However, the relative uniformity of both resting images, the second of which was reconstructed from data acquired *after* the first active set, imply that changes in position or coupling are not solely responsible for the apparent saturation change. Further, mean-time should be more robust to changes in coupling than other datatypes, and since only mean-time is being used, coupling changes are unlikely to have contributed significantly to the images.

2.5.5.1.1 Image summary (Non-linear absolute arm images)

Subject	Arm : healthy volunteer (subject B)
Mesh	Circular 1201nodes 2304 linear elements (2D)
Basis (used to generate PMDFs)	Forward mesh
Starting parameters: μ_a	0.03296mm^{-1} @ 780nm and 0.03591mm^{-1} @ 815nm
Starting parameters: μ'_s	1 mm^{-1}
Iterations	3 (Figure 2.5.34) and 1 to 18
Sources	16
Detectors per source	11
Datatypes	Mean
Calibration	Deconv equivalent (monstode)
2d 3d correction applied?	Yes
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	10 secs
Wavelength	780 nm + 815nm (interlaced)
μ_a background estimate	0.003mm^{-1}

Table 2.5.2 Properties of images shown in section 2.5.4.1

2.5.5.2 Non-linear wavelength-difference arm images

Section 2.5.1.3 described how it is possible, in theory, to derive images of absolute saturation and volume from only *changes* in absorption between three wavelengths. Section 2.5.3 investigated the potential benefits of using this method, which include the fact that an estimate of $\mu_{a,bg}$ is not required and that quantitation improves when absolute μ_a images at three wavelengths are combined (compared to dual wavelength results). However the most attractive possibility of using the three-wavelength equations is that they only need to use *changes* in absorption, which may be possible to reconstruct from only changes in data. Simulations showed that wavelength-difference images (where only changes in data are reconstructed) can suffer from distortions due to unknown background properties (see chapter 2.4) and the reconstruction's inability to account for global perturbations in background properties. Where difference images were calculated by subtracting two absolute images, the saturation and volume results were very good.

Although the MONSTIR system only has two wavelengths available at present, it is informative to compare wavelength-difference images reconstructed from clinical data with the difference between absolute μ_a images (at each wavelength) to evaluate the likely success

of a 3-wavelength system. Figure 2.5.38 shows wavelength-difference images reconstructed using TOAST for the active and resting data sets. Since the changes are now within a range where the perturbation approximation might hold (unlike for absolute - homogenous phantom data as described above), we can use both intensity and mean-time changes to reconstruct images of the change in absorption. Note that the changes in scatter were also reconstructed simultaneously, but are not shown. Starting values for the reconstruction corresponded to the estimate of μ_a at 780nm (0.033mm^{-1}) since we are reconstructing 815nm - 780nm data.

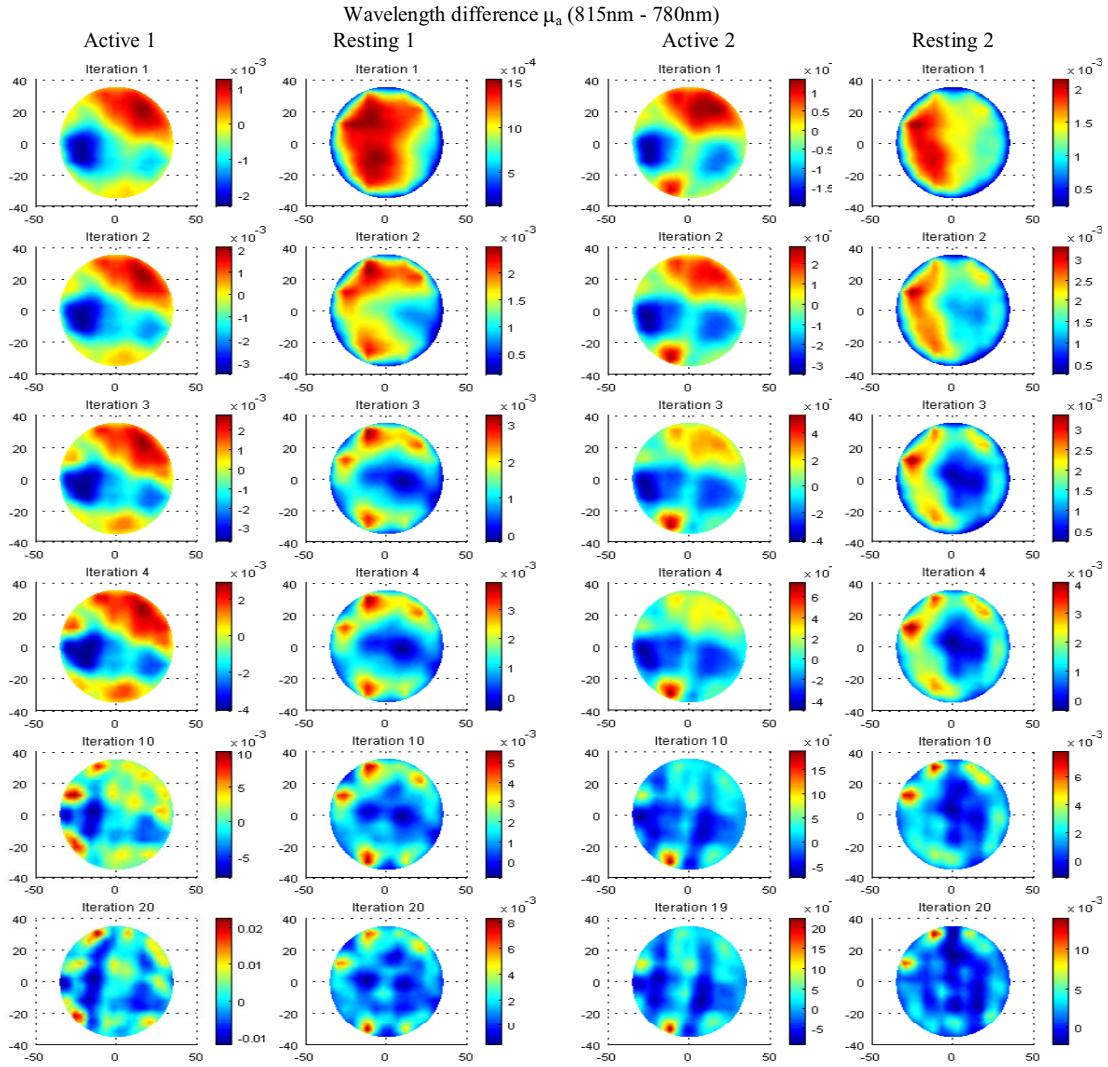

Figure 2.5.38 Wavelength-difference non-linear μ_a images of the arm, reconstructed from the change in mean-time and intensity between wavelengths for two active and resting states.

As was described in section 2.1.4.5 the wavelength dependence of the MONSTIR system has yet to be fully characterised. In the future simple calibration measurements could potentially account for wavelength dependent features, allowing subsequent acquisition of only wavelength-difference data. However currently there are errors associated with the unknown delay between the 780nm and 815nm and the possible difference in the 780nm and 815nm pulse intensity (and their differing propagation through the system). As a consequence, we must in fact calibrate the mean-time data to be used for wavelength-

difference imaging using monstode measurements to eliminate wavelength dependent delays in the system. It is not possible to use monstode measurements to calibrate intensity, so the differences between raw values were used, which means that wavelength related intensity errors will not have been compensated for. No 2D-3D correction was applied because difference data is being used (see section 2.3.1.1).

Corresponding absolute 815nm – absolute 780nm images are also shown in Figure 2.5.39 derived from all the iterations of the absolute images shown in Figure 2.5.34.

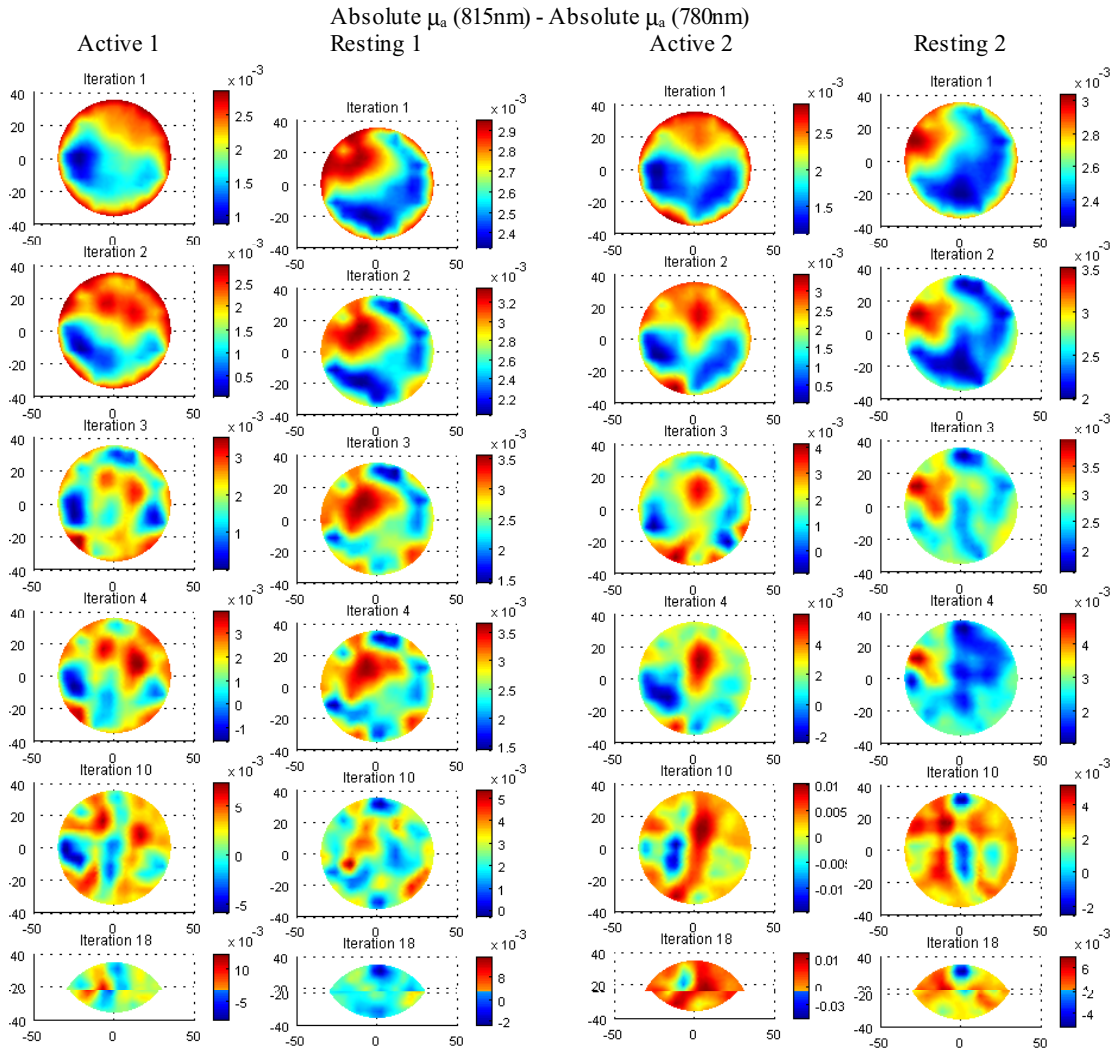

Figure 2.5.39 Images showing the differences between absolute μ_a non-linear images at two wavelengths (absolute815nm – absolute780nm). For active1, resting1, active2 and resting2.

The results for wavelength-difference images (using mean and intensity) agree fairly well with the difference in absorption between two absolute (mean-time only, 2D-3D corrected) μ_a images. Quantitative agreement between the images is within the correct order of magnitude, although the early active difference data imply negative values, whereas the absolute differences do not. This could be explained by the starting values used in the absolute reconstructions (see Table 2.5.2) where the 815nm starting value is higher than the

780nm one. The difference images used only the 780nm μ_a (homogenous) value as their starting point and this was subtracted from the resulting images to give differences. This means that the perturbation will be centred about zero for the wavelength-difference images, but about the positive 815nm-780nm starting values for the absolute-absolute images.

The fact that the wavelength-difference (active) images have the same form as the saturation images is encouraging since a contribution from intensity data has not affected the evident change in that region. This also implies that absolute saturation values *could* be derived from such wavelength-difference images (if a further wavelength were available).

However similar analysis of breast data (not shown) did not reveal a good correlation between wavelength-difference and absolute(wavelength1)-absolute(wavelength2) images. These results imply that more imaging tests are required and the accuracy of wavelength-difference calibration (with regard to quantitative wavelength-difference imaging) would need to be evaluated more thoroughly if a three-wavelength system were to be considered.

2.5.5.2.1 Image summary (Non-linear wavelength-difference arm images)

Subject	Arm : healthy volunteer (subject B)
Mesh	Circular 1201nodes 2304 linear elements (2D)
Basis (used to generate PMDFs)	Forward mesh
Starting parameters: μ_a	0.03296mm^{-1} (780nm)
Starting parameters: μ'_s	1 mm^{-1}
Iterations	1 to 20
Sources	16
Detectors per source	11
Datatypes	Mean + raw intensity
Calibration	Monstode and Difference (815nm relative to 780nm)
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s (μ'_s not shown)
Acquisition time (per source)	10 secs
Wavelength	780 nm + 815nm (interlaced)

Table 2.5.3 Properties of images shown in Figure 2.5.38

2.5.5.3 Non-linear state-change arm images

The arm data can also be used to explore state-difference imaging. As was discussed in section 2.5.1.4, it is not possible to derive *changes* in saturation from changes in data acquired when some state change happens or is induced. Below we compare the results of subtracting absolute saturation and volume images (derived from absolute absorption images), to produce images showing the change in saturation and volume, with μ_a state-difference images at two wavelengths, and the difference in absolute absorption between two states, at two wavelengths. Figure 2.5.40 shows the change in saturation and volume due to exercise derived from the absolute saturation and volume images in Figure 2.5.36.

The images reveal a repeatable reduction in saturation of around 10% during exercise, in a region which agrees well with the position of the flexor digitorum profundus and superficalis as shown in Figure 2.5.35. There also seems to be a corresponding slight blood

volume increase in the region of the active muscle and a slight decrease in blood volume in other regions of the arm (although for volume(1), distortion implies a bad source or detector is affecting image quality).

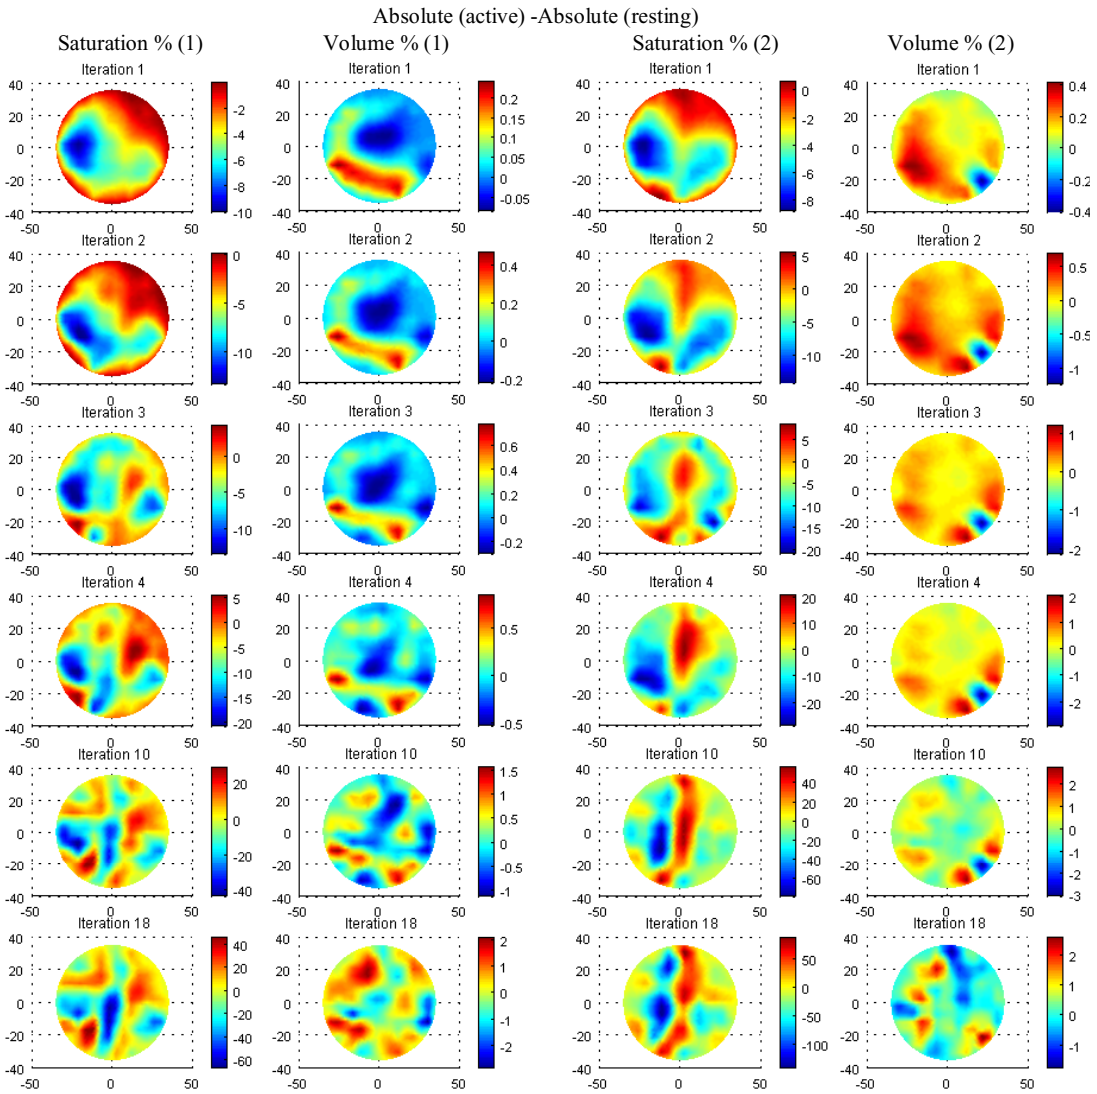

Figure 2.5.40 Changes in saturation and fractional blood volume in the arm during exercise, (active – resting) derived by subtracting absolute saturation images derived from absolute absorption images.

Figure 2.5.41 shows state-difference images derived from changes in both mean-time and intensity between the active and resting data sets. No 2D:3D correction was applied since difference data is being used (see 2.3.1.1). These images are comparable to those derived from early data sets acquired on the arm as shown in (Hillman *et al*, 2001b). As with wavelength-difference images we would expect these images to be affected by the unknown initial state, also there may be poor error cancellation due to the time between acquisition of the data sets that are being compared. However, the state-difference images in Figure 2.5.41 reveal increases in absorption at 780nm (in agreement with a deoxygenation) and further, a very similar increase in absorption at 780nm was seen in the same subject in the original arm imaging experiment (see (Hillman *et al*, 2001b)).

As described in section 2.5.1.4, equations [2.5.18] and [2.5.19] could be used to convert state-change images into images of change in deoxy-haemoglobin (ΔHb) and change in oxy-haemoglobin (ΔHbO_2), and hence changes in volume, although not in saturation.

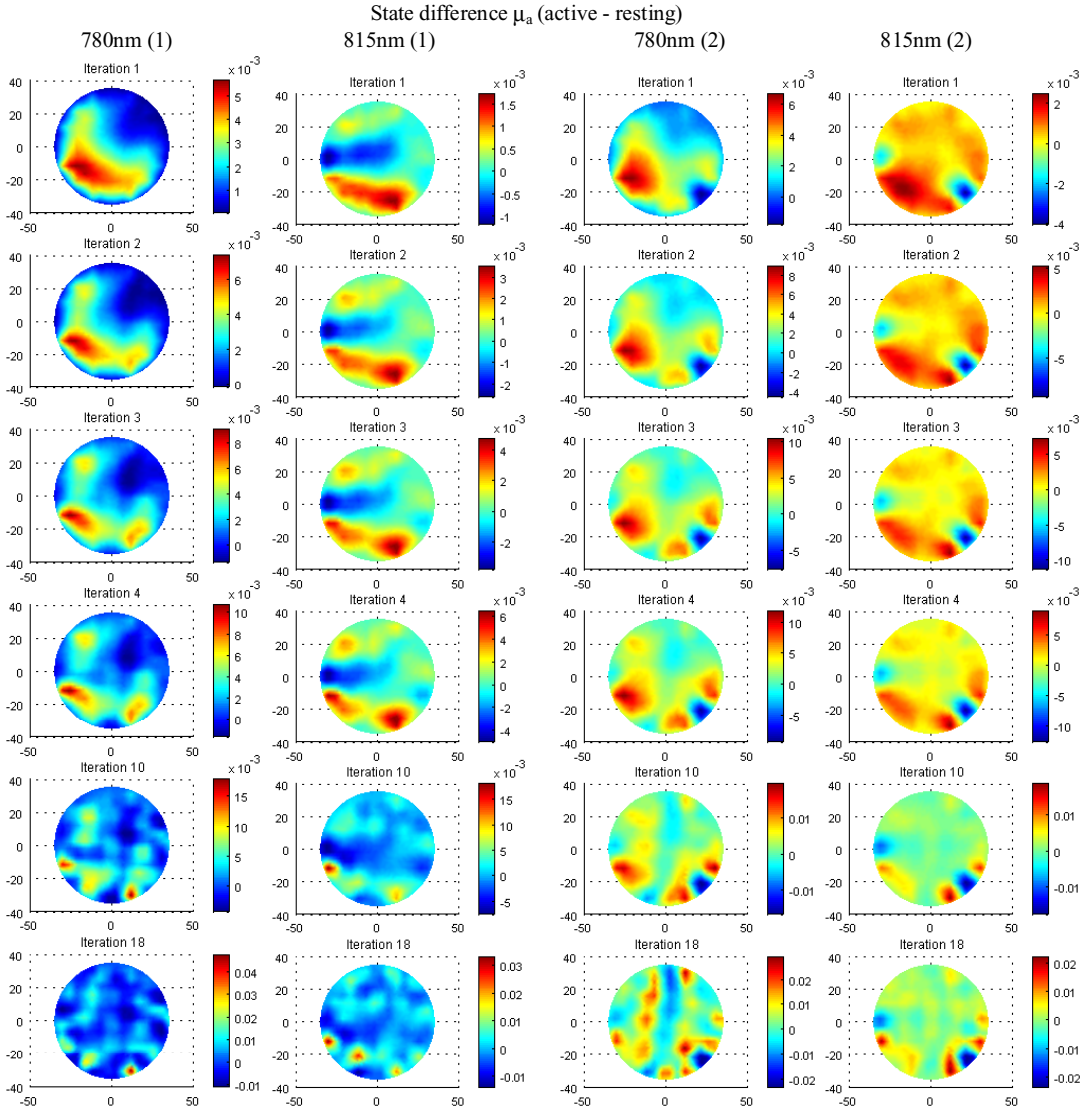

Figure 2.5.41 Non-linear state-difference arm images derived from changes in mean and intensity between an active and resting state.

2.5.5.3.1 Image summary (Non-linear state-difference arm images)

Subject	Arm : healthy volunteer (subject B)
Mesh	Circular 1201nodes 2304 linear elements (2D)
Basis (used to generate PMDFs)	Forward mesh
Starting parameters: μ_a	0.03296mm^{-1} @ 780nm and 0.03591mm^{-1} @ 815nm
Starting parameters: μ'_s	1 mm^{-1}
Iterations	1 to 18
Sources	16
Detectors per source	11
Datatypes	Raw mean + raw intensity
Calibration	Difference (Active relative to resting state)
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s (μ'_s not shown)
Acquisition time (per source)	10 secs
Wavelength	780 nm + 815nm (interlaced)

Table 2.5.4 Properties of images shown in Figure 2.5.41

We can compare the state-difference images to the difference between the absolute active and resting images at the two wavelengths as shown in Figure 2.5.42.

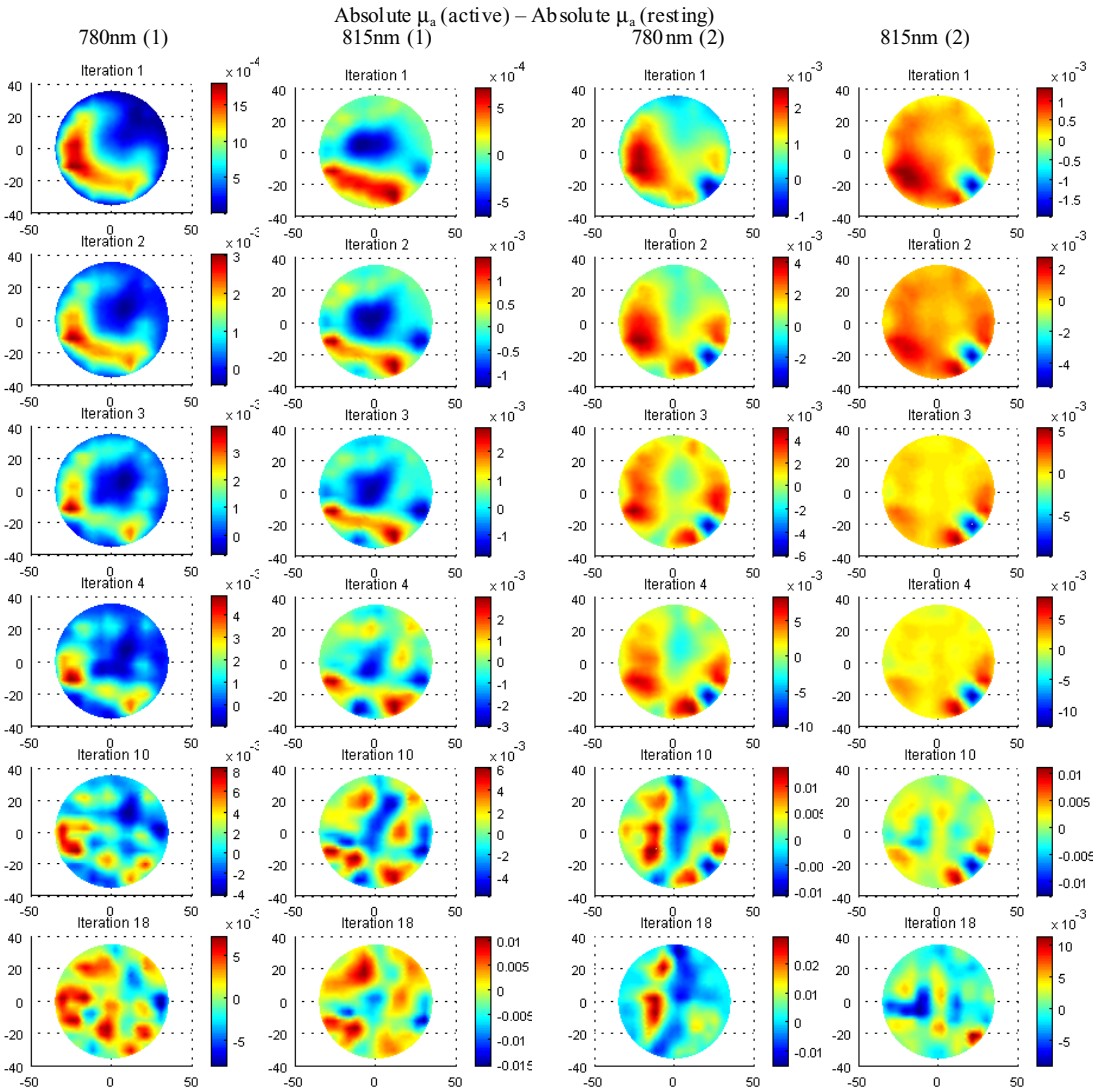

Figure 2.5.42 Non-linear absolute (active) – absolute (resting) arm images

As with the wavelength-difference images, the magnitude of the absorption changes in the absolute – absolute images is higher than in the difference images. However, again the general form of the images is very similar, suggesting the state-difference images *may* be suitable for reconstruction of ΔHb and ΔHbO_2 images. This would require a more rigorously constrained difference reconstruction and would be suitable in cases where calibration is impractical and only changes between successive images are required. The fairly good agreement between the (intensity *and* mean-time) state-difference and absolute – absolute images also implies that the use of mean-time only data in reconstructing absolute images, while resulting in some cross talk, appears to conserve sufficient information about absorption to agree with mean-time and intensity difference reconstructions.

2.5.6 Constraining reconstructions for derivation of saturation and volume

The simplistic analysis of absolute images presented above treats data acquired at each wavelength as an independent measurement. Linear reconstructions were shown to be somewhat more tolerant of noise than non-linear reconstructions for calculating absolute saturation and volume (see sections 2.5.2.3 and 2.5.2.4), since the non-linear absolute μ_a images did not converge in the same way. Regularisation and constraining of non-linear image reconstructions is being investigated (Arridge, 1999). However, an advantage of dual-wavelength data is that there are certain intrinsic correlations between data acquired at one wavelength and another. Physical characteristics of tissue can be exploited to better constrain reconstructions and potentially make derivation of saturation and volume more reliable. Some methods that could be included into reconstructions (linear or non-linear) to improve results from multi-wavelength data would include:

- Assuming that μ'_s is the same at both wavelengths (over a fairly small bandwidth)
- Assuming saturation and volume must be 7 0 and > 100%.
- Assuming known values for certain components e.g. skull or skin
- Assuming certain fractions of different tissues are present
- Correlating μ_a solutions with known spectra
- Solving in terms of saturation and volume rather than μ_a

The possibility of assuming that scatter is the same at both wavelengths is briefly explored below. In addition, a possible method for implementation of a reconstruction that reconstructs for saturation and volume is described. Modifications to TOAST's non-linear solver could allow these and the other constraints to be applied to dual-wavelength reconstructions. For initial investigations, simple linear reconstructions are presented.

2.5.6.1 Assuming μ'_s is the same at both wavelengths

A very simple 2D simulation was devised to test the potential improvements to dual-wavelength image reconstruction achievable by assuming scatter is the same for both wavelengths. We can compare the results for μ_a solutions resulting from the solution of an equation of the form of [2.5.26] e.g:

$$\begin{bmatrix} \Delta I_{815nm-815nm(\mathrm{hom})} & \Delta I_{780nm-780nm(\mathrm{hom})} & \Delta < t >_{815nm-815nm(\mathrm{hom})} & \Delta < t >_{780nm-780nm(\mathrm{hom})} \end{bmatrix}$$
$$= \begin{bmatrix} J_{\mu_a,815nm(\mathrm{hom})}{}^{I} & 0 & J_{\mu'_s,815nm(\mathrm{hom})}{}^{I} & 0 \\ 0 & J_{\mu_a,780nm(\mathrm{hom})}{}^{I} & 0 & J_{\mu'_s,780nm(\mathrm{hom})}{}^{I} \\ J_{\mu_a,815nm(\mathrm{hom})}{}^{<t>} & 0 & J_{\mu'_s,815nm(\mathrm{hom})}{}^{<t>} & 0 \\ 0 & J_{\mu_a,780nm(\mathrm{hom})}{}^{<t>} & 0 & J_{\mu'_s,780nm(\mathrm{hom})}{}^{<t>} \end{bmatrix} \begin{bmatrix} \mu_{a,815nm} - \mu_{a,815nm(\mathrm{hom})} \\ \mu_{a,780mn} - \mu_{a,780nm(\mathrm{hom})} \\ \mu'_{s,815nm} - \mu'_{s,815nm(\mathrm{hom})} \\ \mu'_{s,780nm} - \mu'_{s,780nm(\mathrm{hom})} \end{bmatrix}$$

[2.5.31]

with the results of solving an equation of the form:

$$\begin{bmatrix}\Delta I_{815nm-815nm(\text{hom})} & \Delta I_{780nm-780nm(\text{hom})} & \Delta < t >_{815nm-815nm(\text{hom})} & \Delta < t >_{780nm-780nm(\text{hom})}\end{bmatrix}$$

$$= \begin{bmatrix} J_{\mu_a,815nm(\text{hom})}^{\ I} & 0 & J_{\mu_s',815nm(\text{hom})}^{\ I} \\ 0 & J_{\mu_a,780nm(\text{hom})}^{\ I} & J_{\mu_s',780nm(\text{hom})}^{\ I} \\ J_{\mu_a,815nm(\text{hom})}^{<t>} & 0 & J_{\mu_s',815nm(\text{hom})}^{<t>} \\ 0 & J_{\mu_a,780nm(\text{hom})}^{<t>} & J_{\mu_s',780nm(\text{hom})}^{<t>} \end{bmatrix} \begin{bmatrix} \mu_{a,815nm} - \mu_{a,815nm(\text{hom})} \\ \mu_{a,780mn} - \mu_{a,780nm(\text{hom})} \\ \mu'_{s,both} - \mu'_{s,both(\text{hom})} \end{bmatrix}$$

[2.5.32]

The simulation has two absorbing regions with different absorptions at 'two wavelengths', and a region with higher and identical scatter at both wavelengths as shown in Figure 2.5.43. The mesh had 1459 nodes and was 70mm in diameter with refractive index = 1.4. Mean-time and intensity data were simulated using TOAST for the 780nm and 815nm distributions, along with a homogenous reference state assumed to be wavelength independent with the same properties as region 0. PMDFs were also generated using TOAST based on the background reference properties. So in this simple case $J_{*,815nm}^{\ *} = J_{*,780nm}^{\ *}$.

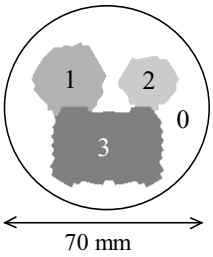

Region	μ_a (mm^{-1}) (780nm)	μ_a (mm^{-1}) (815nm)	μ'_s (mm^{-1})
0	0.01	0.01	1
1	0.005	0.02	1
2	0.02	0.005	1
3	0.01	0.01	2

Figure 2.5.43 Simulation geometry and properties used to test constraining dual-wavelength reconstructions by assuming scatter is wavelength independent.

κ PMDFs generated by TOAST were converted into μ'_s PMDFs via the chain rule since κ will not be wavelength invariant. Where $dM/d\mu'_s$ represents the μ'_s Jacobian:

$$\frac{dM}{d\mu'_s} = \frac{dM}{d\kappa}\frac{d\kappa}{d\mu'_s} + \frac{dM}{d\mu_a}\frac{d\mu_a}{d\mu'} \quad \text{but} \quad \frac{d\mu_a}{d\mu'} = 0$$

$$\text{since } \kappa = \frac{1}{3(\mu_a + \mu'_s)}$$

$$\frac{dM}{d\mu'_s} = -3\frac{dM}{d\kappa}\kappa^2$$

[2.5.33]

Figure 2.5.44 shows linear images reconstructed using the unconstrained equation [2.5.31] (note that Jacobians are in fact normalised as detailed in section 1.3.3.5). Images of $\mu_{a,780nm}$, $\mu_{a,\ 820nm}$ and $\mu'_{s,780nm}$ and $\mu'_{s,815nm}$, reconstructed using intensity only, mean-time only and both mean and intensity are shown. All images are displayed on the same colour-scale as their corresponding target. Figure 2.5.45 shows images reconstructed from the same data, but using an equation of the form of [2.5.32], yielding $\mu_{a,780nm}$, $\mu_{a,\ 820nm}$ and $\mu'_{s,both}$ for intensity data, mean-time data and mean *and* intensity data. (lambda = 0.5 for both sets)

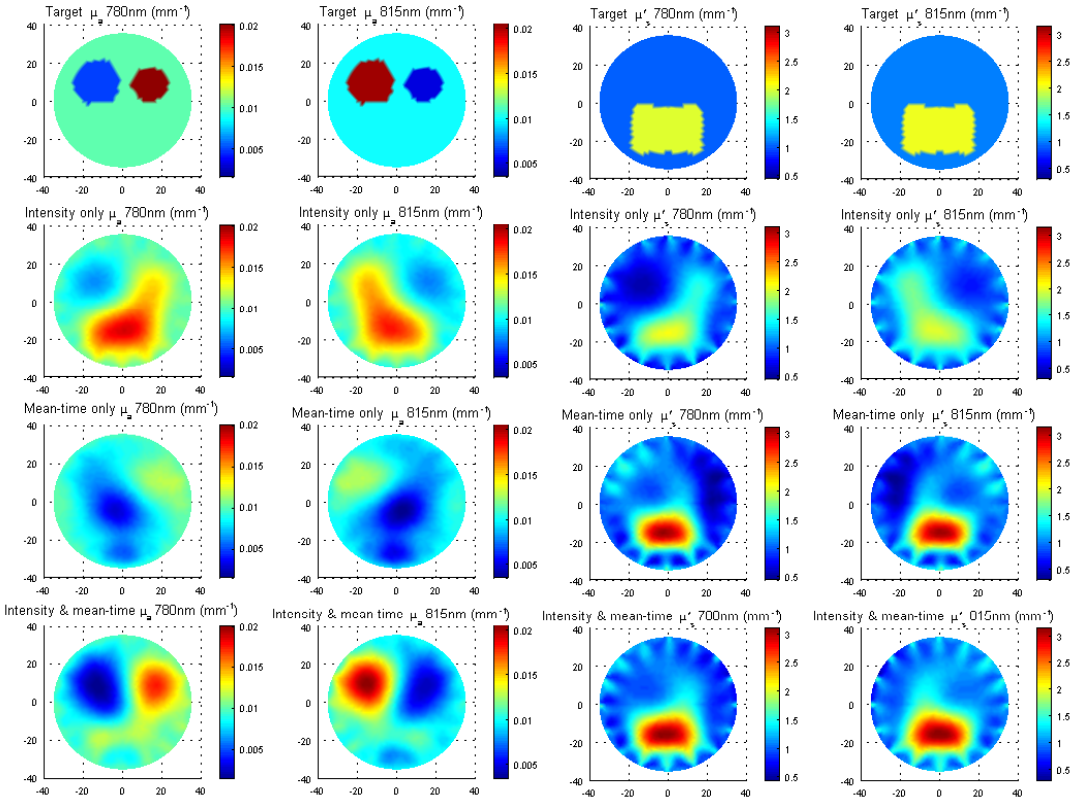

*Figure 2.5.44 Absolute absorption and scatter target images (very top) and images derived from only intensity (top), only mean-time (middle) and both intensity and mean-time (bottom), allowing **scatter to be wavelength dependent**.*

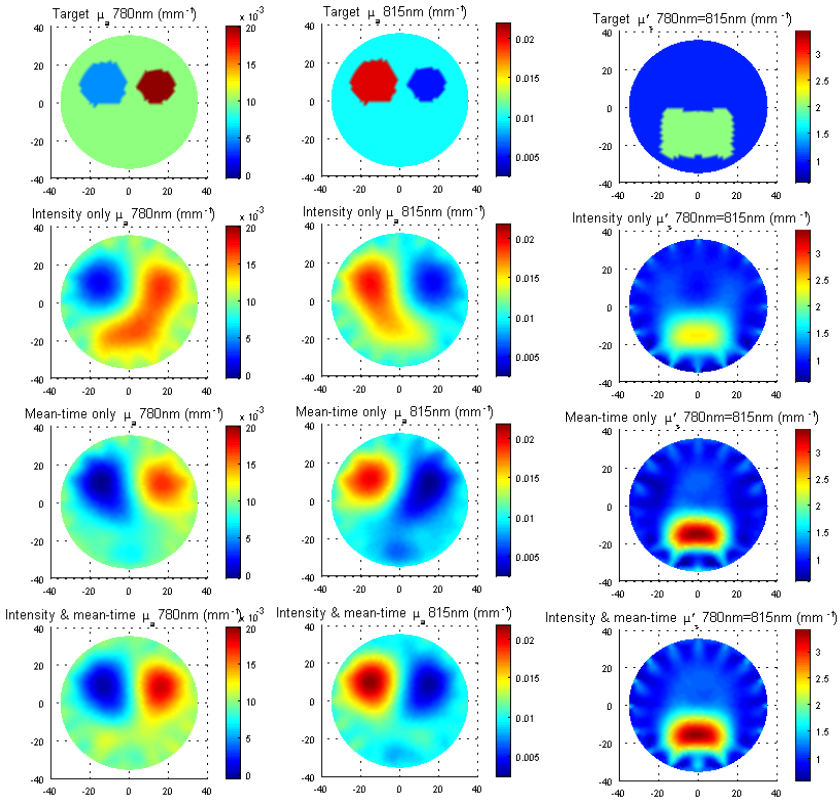

*Figure 2.5.45 Absolute absorption and scatter target images (very top) and images derived from only intensity (top), only mean-time (middle) and both intensity and mean-time (bottom), forcing the **scatter solution to be the same** for both wavelengths.*

There is clearly a great improvement in the images resulting from constraining the scatter solution at the two wavelengths, particularly for the mean-time-only and intensity-only results. While separation isn't complete in the constrained intensity images, a significant amount of separation has been achieved. (Arridge *et al*, 1998c) proves that it is *not* possible to separate absorption and scatter from measurements of intensity at a single wavelength. This new technique implies that it may be possible to separate scatter and absorption from intensity alone if sufficient wavelengths are available. The mean-time data are also much improved with the use of the constraint.

As with the mean-time-only absolute arm images shown in (Hillman *et al*, 2001b) and section 2.5.5.1, the unconstrained images display evidence of cross talk since the absorption images resemble the inverse of the scatter images. This technique of constraining cannot currently be applied to the arm data since the data is not suitable for linear reconstruction, and a non-linear reconstruction scheme with this constraint applied has yet to be implemented. While the above simulations were repeated with data where noise had been added (as in section 2.5.2.3), results looked very similar to the constrained and unconstrained images and so are not shown. Once TOAST is able to apply this scatter constraint, its effectiveness on noisy non-linear reconstructions should be evaluated.

The arm results in section 2.5.5.2 did however imply that absolute(815nm) – absolute(780nm) images from mean-time only data were very similar to mean and intensity wavelength-difference images. This suggests that the effect of cross talk on μ_a images may cancel when two wavelengths are subtracted if we can assume that μ'_s is the same for each wavelength. Figure 2.5.46 shows images generated by subtracting the 815nm absorption images from 780nm images for both un-constrained (Figure 2.5.44) and constrained scatter (Figure 2.5.45). Linear wavelength-difference images are also provided for comparison (reconstructed using the changes in data between one wavelength and the other using PMDFs based on a homogenous background of $\mu_a = 0.01\text{mm}^{-1}$ and $\mu'_s = 1\text{mm}^{-1}$).

Figure 2.5.46 supports the results from the arm data. In the unconstrained images, while quantitation isn't correct for intensity-only and mean-time-only results, the region of cross talk and the distortion due to the scattering object appears to have cancelled out. However the results also show that a dramatic improvement in the accuracy of the difference results from constraining the scatter solution to be wavelength independent. The wavelength-difference images show characteristic errors due to the unknown background scattering structure as demonstrated in chapter 2.4. Arm wavelength-difference images were clearly not affected as badly as this since the higher scattering bones in the arm and the level of intrinsic μ_a heterogeneity are small compared to the regions in these simulations. These results suggest

that quantitation in the functional arm images in section 2.5.5.1 are unlikely to be as good as they would be had intensity also been available, and furthermore if a scatter constraint could have been used. However the *shape* of changes in absorption will apparently have been conserved, implying that the saturation and volume images have some qualitative accuracy.

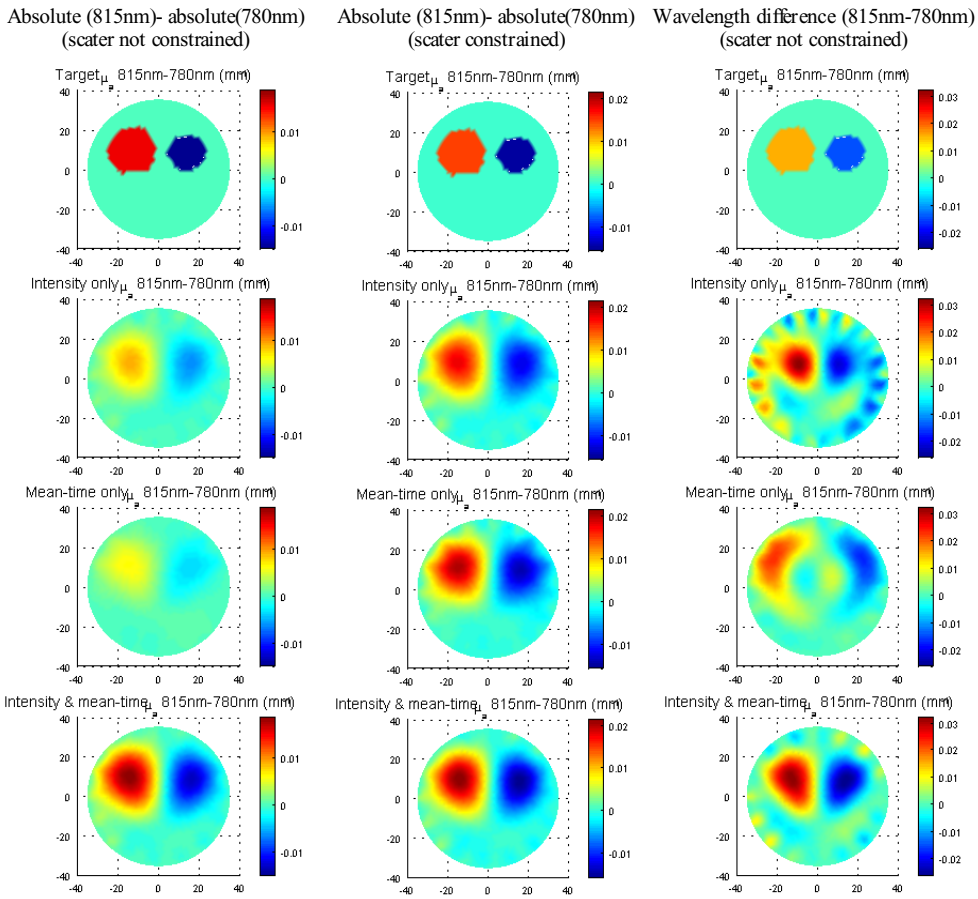

Figure 2.5.46 Images showing the apparent difference in absorption between 815nm and 780nm derive from (left) subtracting the absolute (unconstrained) μ_a images from Figure 2.5.44, (middle) subtracting the absolute (constrained) μ'_s images from Figure 2.5.45, and (right) wavelength-difference images reconstructed from the change in data between the two wavelengths.

It was found that even better μ_a and μ'_s separation could be obtained for scatter constrained intensity-only and mean-only images if the reconstruction was weighted to get a better solution for the μ'_s images than the μ_a images, by multiplying the absorption Jacobians in [2.5.32] by 0.1 (and then dividing the resulting μ_a image by 0.1). However this also reduced the quantitation of the absorption images to ~ 0.1 of their true values.

This section has demonstrated that, if it can be assumed that scatter is independent of wavelength over the range being used, significant improvements in images can be achieved by forcing the reconstruction to solve for only 3 rather than 4 images (or 4 rather than 6 if three wavelengths are used). It would be valuable to implement this kind of constraint into TOAST to explore whether improved quantitative clinical images, which are more robust to noise, can be produced.

2.5.6.2 Solving in terms of saturation and volume rather than μ_a

Section 2.5.6.1 above demonstrated that dual-wavelength images could be improved by putting a constraint on the reconstruction that μ'_s is the same at two wavelengths. Section 2.5.2.4 demonstrated that non-linear absolute absorption images can yield erroneous values in saturation and volume if are noisy. A more direct way to constrain image reconstructions of saturation and volume that will not require two μ_a images to converge in a similar way would be to directly reconstruct for saturation, volume and scatter instead of $\mu_{a,,780nm,}$ $\mu_{a,815nm}$ and μ'_s. If we want to reconstruct a change in a given property ΔF from a change in measurement ΔM, we need to solve:

$$\Delta M = J_F{}^M \Delta F$$

[2.5.34]

Where $J_F{}^M$ represents the Jacobian relating a change in parameter F to a change in measurement M. If we assume that F is a functional parameter e.g. saturation or volume, which will affect both measurements at both 780nm and 815nm, we can express a change in F as:

$$\Delta F = \frac{dF}{dM_{780(\text{hom})}} \Delta M_{780nm-780(\text{hom})} + \frac{dF}{dM_{815(\text{hom})}} \Delta M_{815nm-815(\text{hom})}$$

[2.5.35]

Using the chain rule we get:

$$\frac{dF}{dM_{780(\text{hom})}} = \frac{\partial F}{\partial \mu_{a,780(\text{hom})}} \frac{\partial \mu_{a,780(\text{hom})}}{\partial M_{780(\text{hom})}}$$

[2.5.36]

From [2.5.36] we can deduce that a Jacobian in [2.5.34] for F = saturation S will be:

$$J_{S,780nm(\text{hom})}{}^M = \frac{dM_{780(\text{hom})}}{dS} = \left(\frac{\partial S}{\partial \mu_{a,780(\text{hom})}} \right)^{-1} J_{\mu_a,780(\text{hom})}{}^M$$

[2.5.37]

where $J_{\mu a}{}^M = \partial M / \partial \mu_a$, and if F = volume V:

$$J_{V,780nm(\text{hom})}{}^M \frac{dM_{780(\text{hom})}}{dV} = \left(\frac{\partial V}{\partial \mu_{a,780(\text{hom})}} \right)^{-1} J_{\mu_a,780(\text{hom})}{}^M$$

[2.5.38]

The total equation the would need to be solved would then have the form:

$$\begin{bmatrix} \Delta I_{815nm-815nm(\text{hom})} & \Delta I_{780nm-780nm(\text{hom})} & \Delta < t >_{815nm-815nm(\text{hom})} & \Delta < t >_{780nm-780nm(\text{hom})} \end{bmatrix}$$

$$= \begin{bmatrix} J_{S,815nm(\text{hom})}{}^I & J_{V,815nm(\text{hom})}{}^I & J_{\mu'_s,815nm(\text{hom})}{}^I \\ J_{S,780nm(\text{hom})}{}^I & J_{V,780nm(\text{hom})}{}^I & J_{\mu'_s,780nm(\text{hom})}{}^I \\ J_{S,815nm(\text{hom})}{}^{<t>} & J_{V,815nm(\text{hom})}{}^{<t>} & J_{\mu'_s,815nm(\text{hom})}{}^{<t>} \\ J_{S,780nm(\text{hom})}{}^{<t>} & J_{V,780nm(\text{hom})}{}^{<t>} & J_{\mu'_s,780nm(\text{hom})}{}^{<t>} \end{bmatrix} \begin{bmatrix} S - S_{(\text{hom})} \\ V - V_{(\text{hom})} \\ \mu'_s - \mu'_{s(\text{hom})} \end{bmatrix}$$

[2.5.39]

assuming that we also want to apply the constraint that scatter is equal at both wavelengths. Using [2.5.37] and [2.5.38] we can see that [2.5.39] can also be written as:

$$\begin{bmatrix}\Delta I_{815nm-815nm(\hom)} & \Delta I_{780nm-780nm(\hom)} & \Delta<t>_{815nm-815nm(\hom)} & \Delta<t>_{780nm-780nm(\hom)}\end{bmatrix}$$

$$= \begin{bmatrix} J_{\mu_a,815nm(\hom)}^{I} & 0 & J_{\mu_s',815nm(\hom)}^{I} \\ 0 & J_{\mu_a,780nm(\hom)}^{I} & J_{\mu_s',780nm(\hom)}^{I} \\ J_{\mu_a,815nm(\hom)}^{<t>} & 0 & J_{\mu_s',815nm(\hom)}^{<t>} \\ 0 & J_{\mu_a,780nm(\hom)}^{<t>} & J_{\mu_s',780nm(\hom)}^{<t>} \end{bmatrix} \begin{bmatrix} \frac{\partial \mu_{a,815(\hom)}}{\partial S} & \frac{\partial \mu_{a,815(\hom)}}{\partial V} & 0 \\ \frac{\partial \mu_{a,780(\hom)}}{\partial S} & \frac{\partial \mu_{a,780(\hom)}}{\partial V} & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} S-S_{(\hom)} \\ V-V_{(\hom)} \\ \mu_s'-\mu_{s(\hom)}' \end{bmatrix}$$

[2.5.40]

The differential terms in [2.5.40] can be derived from the absolute saturation and volume equations ([2.5.5] and [2.5.6]) in section 2.5.1.2 such that we get:

$$\frac{\partial V}{\partial \mu_{a,\lambda_1}} = \frac{-(\varepsilon_{Hb,\lambda_2}-\varepsilon_{HbO_2,\lambda_2})}{\varepsilon_{Hb,\lambda_1}\varepsilon_{HbO_2,\lambda_2}-\varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda_1}+\mu_{a,bg}(\varepsilon_{Hb,\lambda_2}-\varepsilon_{Hb,\lambda_1}+\varepsilon_{HbO_2,\lambda_1}-\varepsilon_{HbO_2,\lambda_2})}$$

[2.5.41]

$$\frac{\partial V}{\partial \mu_{a,\lambda_2}} = \frac{(\varepsilon_{Hb,\lambda_1}-\varepsilon_{HbO_2,\lambda_1})}{\varepsilon_{Hb,\lambda_1}\varepsilon_{HbO_2,\lambda_2}-\varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda_1}+\mu_{a,bg}(\varepsilon_{Hb,\lambda_2}-\varepsilon_{Hb,\lambda_1}+\varepsilon_{HbO_2,\lambda_1}-\varepsilon_{HbO_2,\lambda_2})}$$

[2.5.42]

$$\frac{\partial S}{\partial \mu_{a,\lambda_1}} = \frac{(\mu_{a,\lambda_2}-\mu_{a,bg})(\varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda_1}-\varepsilon_{Hb,\lambda_1}\varepsilon_{HbO_2,\lambda_2}+\mu_{a,bg}(\varepsilon_{HbO_2,\lambda_2}-\varepsilon_{HbO_2,\lambda_1}+\varepsilon_{Hb,\lambda_1}-\varepsilon_{Hb,\lambda_2}))}{\left((\varepsilon_{Hb,\lambda_2}-\varepsilon_{HbO_2,\lambda_2})(\mu_{a,\lambda_1}-\mu_{a,bg})-(\varepsilon_{Hb,\lambda_1}-\varepsilon_{HbO_2,\lambda_1})(\mu_{a,\lambda_2}-\mu_{a,bg})\right)^2}$$

[2.5.43]

$$\frac{\partial S}{\partial \mu_{a,\lambda_2}} = \frac{(\mu_{a,\lambda_1}-\mu_{a,bg})(\varepsilon_{Hb,\lambda_1}\varepsilon_{HbO_2,\lambda_2}-\varepsilon_{Hb,\lambda_2}\varepsilon_{HbO_2,\lambda_1}+\mu_{a,bg}(\varepsilon_{HbO_2,\lambda_1}-\varepsilon_{HbO_2,\lambda_2}+\varepsilon_{Hb,\lambda_2}-\varepsilon_{Hb,\lambda_1}))}{\left((\varepsilon_{Hb,\lambda_2}-\varepsilon_{HbO_2,\lambda_2})(\mu_{a,\lambda_1}-\mu_{a,bg})-(\varepsilon_{Hb,\lambda_1}-\varepsilon_{HbO_2,\lambda_1})(\mu_{a,\lambda_2}-\mu_{a,bg})\right)^2}$$

[2.5.44]

The volume differentials are independent of $\mu_{a,\ 780nm}$ and $\mu_{a,\ 820nm}$. This means that the volume Jacobians for 780nm ([2.5.38]) and 815nm will reduce to just absorption Jacobians multiplied by a constant. However, the saturation differentials are clearly functions of $\mu_{a,780nm}$ and $\mu_{a,\ 815nm}$. This means that large changes in saturation will not be properly modelled by a linear saturation reconstruction.

[2.5.40] can be implemented in a simple linear form to examine the results of reconstructing data from the dual-wavelength simulation detailed in Figure 2.5.2, section 2.5.2 . μ_a and μ'_s Jacobians were converted to saturation and volume Jacobians by multiplying by the appropriate differential. The values of the differentials were calculated based on initial estimates of $\mu_{a,780nm}$ and $\mu_{a,815nm}$ (the background or region 0 properties), and $\mu_{a,bg} = 0.01mm^{-1}$ (the correct value) was used. This resulted in values of:

$$\frac{\partial \mu_{a,780nm}}{\partial S} = -10394.20 \qquad \frac{\partial \mu_{a,780nm}}{\partial V} = 79.456$$

$$\frac{\partial \mu_{a,815nm}}{\partial S} = 9298.63 \qquad \frac{\partial \mu_{a,815nm}}{\partial V} = 170.948$$

The Jacobians were normalised such that the reconstruction had the form:

$$\left[\frac{\Delta M}{mean|\Delta M|}\right] = \left[\frac{F_o \times \frac{\partial x_\lambda}{\partial F} \times J_{x,\lambda}{}^{M}(x_o)}{mean|\Delta M|}\right]\left[\frac{\Delta F}{F_o}\right]$$

[2.5.45]

where F_o is an initial estimate of the saturation or volume (76% and 4% respectively were used). x is μ_a or μ'_s . Note that the μ'_s Jacobians were not multiplied by a differential factor, but were converted from κ PMDFs using [2.5.33]and multiplied by an initial estimate of μ'_s ($1mm^{-1}$) and divided by the mean change in data. Linear images were reconstructed using intensity and mean-time data together. Each resulting image had the initial values of saturation, volume or μ'_s added on to convert them from difference to absolute images. The results are shown in Figure 2.5.47.

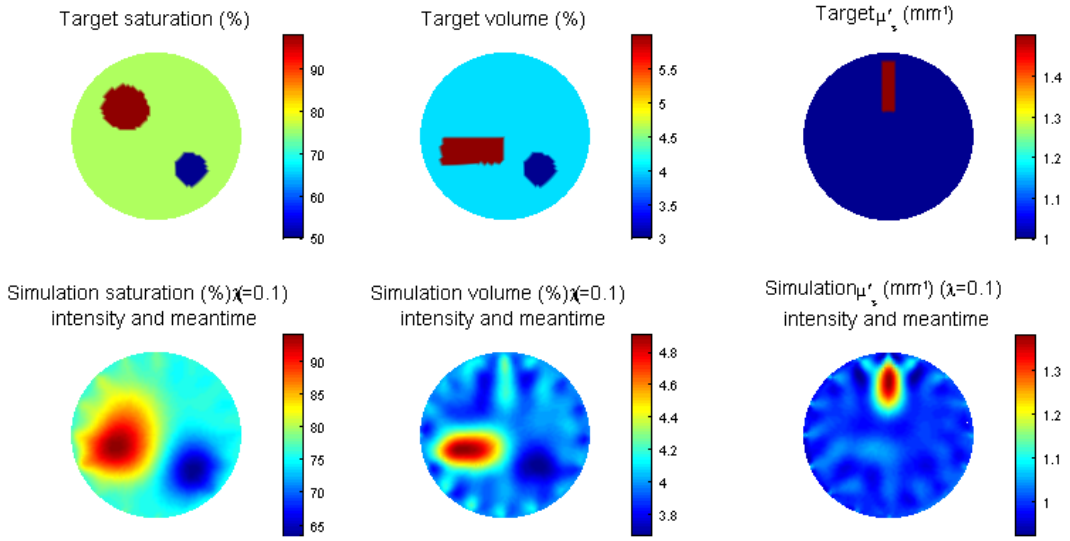

Figure 2.5.47 Absolute saturation (left), volume (middle) and μ'_s (right) images derived directly from dual-wavelength data using intensity and mean-time data.

A lambda factor of $0.1mm^{-1}$ was used to reconstruct the images in Figure 2.5.47. Higher values of lambda reduced quantitation. Lower values of lambda worsened cross talk. Cross talk between saturation and volume, can be seen in the saturation image, light cross talk between scatter and volume is also evident. The high saturation peak has been distorted by the highly absorbing high-volume region, which in fact should have the same saturation as the background. This may be due to the fact that the saturation Jacobians depend on μ_a, making them sensitive to regions with large changes in μ_a. It may also be an intrinsic problem with this method. However, if a non-linear implementation were possible, updates to the saturation Jacobians with successive iterations may overcome this problem.

2.5.7 Summary and discussion

This chapter has presented investigations into a number of approaches to extraction of functional parameters from multi-wavelength optical tomography data. Three methods of extracting functional data were proposed:

- Absolute μ_a images reconstructed separately for two or more wavelengths could be combined node-by-node to yield absolute saturation and fractional blood volume images.
- The difference between absorption images at three or more wavelengths (or 2 or more 'wavelength-difference' images) could be combined to produce images of absolute saturation and fractional blood volume.
- Changes in data with changes in functional state of the object at two or more wavelengths could be used to determine changes in oxy- and deoxy- haemoglobin concentrations.

The equations used are valid only in the limited case where $\mu_{a,bg}$ due to water and other tissue constituents is wavelength independent, e.g. when we are using two wavelengths that are quite close to each-other (780nm and 815nm).

The first two of these methods were evaluated using a simple simulation which had optical properties derived to represent particular saturation and volume properties. In general, the blurring and poor quantitation of diffuse image reconstruction meant that the exact saturation and volume values could never be correctly retrieved. Blurring causes the peak value of μ_a inclusions to be lower than it should be, and a region surrounding the inclusion will have a higher apparent absorption than it should. In general this propagates through into saturation and volume values retrieved. This implies that it will be saturation *contrast* (or volume contrast) that may allow clinical interpretation of results, while comparing peak values in images may be misleading. It is difficult to determine whether inaccurate nodal values will detract from the diagnostic usefulness of images that have clearly revealed discrete regions of high or low blood volume or saturation with good contrast relative to the background.

The first simulation demonstrated deriving absolute saturation and volume from absolute μ_a images using both linear and non-linear reconstructions. Linear reconstructions produced images with improved quantitation with higher iterations. However the use of unconstrained non-linear images reconstructed from data with systematic-type noise added was shown to result in images that progressively deteriorated with increasing iterations. Linear results were affected to a lesser degree by the presence of noise.

The wavelength-difference method was investigated in two forms. By deriving images from three absolute μ_a images (at each of three wavelengths) it is possible to reconstruct absolute saturation and volume without knowledge of $\mu_{a.bg}$ (assuming it is wavelength independent). Noise-free simulation results implied that this method held promise for reconstructing images with *better* resolution and quantitation than the two wavelength results, for both linear and non-linear reconstructions, also yielding μ'_s images.

If three wavelengths are available, absolute saturation and volume can also be derived from only the differences in measured data between three wavelengths. As was discussed in chapter 2.4 wavelength-difference imaging will suffer from perturbation-related errors due to unknown static parameters. In addition this technique would not yield absolute scatter images. However difference data has been shown to be well calibrated (see section 2.1.2.3) and measurements of only the changes in intensity and mean-time between two wavelengths would be easy to acquire in the clinical setting requiring virtually no independent calibration measurements. This approach was explored with simulated data. Wavelength-difference images reconstructed from noise-free simulated data were found to have a number of artefacts. Some of these related to the effects of reconstructing a global background perturbation, which could be improved with smoothing and regularisation of the reconstruction algorithm. In general saturation results for this method were satisfactory, while volume results were quite badly affected by artefacts.

The methods found for extracting absolute saturation and volume from absolute μ_a images were then applied to measured clinical data. Since the MONSTIR system only has two wavelength acquisition, it was not possible to experimentally validate the use of wavelength-difference images to derive absolute saturation and volume. However wavelength-difference and absolute-absolute differences for the two wavelengths available can be compared to explore the consistency and quality of such images from experimental data.

Data acquired on the breast of an adult volunteer were reconstructed (in 2D but relative to a homogenous phantom, see section 2.3.3.1). Linear absolute saturation and volume images using two absolute (baseline) images (one for each wavelength) were shown to have physiologically reasonable values despite the many assumptions made in reconstructing the data. Saturation varied spatially between 65 and 100% and fractional blood volume (assuming 2mM haematocrit) was found to be between 0.8 and 2%. These values are in fairly good agreement with (McBride *et al*, 1999). The accuracy of the distribution of scatter, saturation and volume is not known and would require further studies to validate. Non-linear reconstruction of the same data yielded similar results, although the images were found to

become progressively more distorted with higher iterations, implying that noise was affecting the data (as demonstrated in section 2.5.2.4).

Data acquired on the arm using the new multi-wavelength laser in MONSTIR were then reconstructed. The arm imaging protocol involved performing a gripping exercise during two of the acquisitions, and resting during the other two. Monstode-calibrated mean-time only data were used (with a 2D:3D correction applied) to generate (non-linear) absolute μ_a and μ'_s images of the four data sets at 780nm and 815nm. These images were then used to derive absolute saturation and volume images. There was a marked difference between the absolute saturation at the active and resting states for both repetitions of the protocol. The relatively homogenous ~75% saturation of the resting arm showed an apparent oxygenation decrease (to ~64%) in the region corresponding to the muscles thought to be responsible for the gripping exercise. Values found for the blood volume ranged from ~6.5% to 8% (assuming 2mM haematocrit) and remained fairly constant between the active and resting states although a slight volume increase of ~ 0.5% was seen in both repetitions in the region where the saturation images implied an activation-related deoxygenation. As with the preceding (added noise) simulated results (section 2.5.2.4), the saturation and volume images became more distorted with higher iterations.

Wavelength-difference images of the arm were compared when reconstructed from changes in data (including intensity), and when derived from two absolute images. The results compared well qualitatively (especially since intensity and mean-time were used for the wavelength-difference images, whereas the absolute - absolute images used absolute mean-time only). The locations of the perturbations in both cases also agreed well with the saturation images. Quantitative comparison of the two images suggest that further work is required to understand the relation between wavelength-difference images (and related calibration issues - section 2.1.4.5) if construction of a 3-wavelength system is to be considered.

State-difference arm images were then reconstructed and compared with the difference between absolute images for active and resting. Saturation and volume *changes* with activation were also derived from the absolute saturation and volume images for comparison. The state-difference images agree surprisingly well with the absolute(active)-absolute(resting) images. Suggesting that the heterogeneity of the background is not affecting the difference images significantly (see chapter 2.4). Further, the location and magnitude of the perturbation at 780nm in the state-difference image corresponds remarkably well to the state-change image of the same subject shown in (Hillman *et al*, 2001b), acquired some 22 months previously using a different laser and fibre arrangement. The magnitude of the change at 815nm is also comparable, although the shape is quite distorted.

The distribution of the saturation and volume *changes* is fairly different from the 780nm and 815nm μ_a changes seen in the state-difference images. Even so, the 780nm increase in absorption roughly corresponds to the region where deoxygenation is seen. Volume increases seem to correspond more to regions where absorption at 820nm increases.

The results on the arm are very positive. Beyond the possibility that changes in data with exercise are due to changes in coupling with movement of the arm (which is unlikely since similar results are obtained for mean-time and intensity data), we have measured a repeatable, and physiologically reasonable, apparent change in saturation and blood volume in the arm with exercise. The results from simulations in this and preceding chapters indicate that the absolute images produced *will* have errors relating to 2D:3D mismatch, the effects of noise on non-linear reconstructions and inherent blurring. Where difference data are used, errors relating to perturbation assumptions and (for intensity results) multi-wavelength calibration will also affect images. These results need to be validated with further work, for example using 3D reconstructions, phantoms containing haemoglobin and improved calibration procedures (creating a more suitable homogenous reference phantom).

The final study investigated the possibility of constraining image reconstructions by utilising the fact that data acquired at two wavelengths will be correlated. It was shown that by constraining a linear reconstruction of dual-wavelength data, such that the scatter solution is forced to be the same at both wavelengths, that image quality could be improved. Further, this technique provided much better quantitation and separation of absorption and scatter in images reconstructed from only one datatype (intensity or mean-time). The differences *between* the absolute μ_a images at each wavelength were also investigated for each datatype combination with and without the μ'_s constraint. The constrained data yielded excellent quantitatively accurate results. However in addition, the difference in the unconstrained images, while not quantitatively accurate, readily revealed the location of the absorption perturbations. The μ_a / μ'_s cross-talk related features in the absolute images had cancelled, again through the assumption that scatter is wavelength independent (which is reasonable over this wavelength range).

Finally, a method for reconstructing saturation, volume and scatter directly from multi-wavelength data was outlined. An example result from a linear implementation of this method was presented, although theory suggests that a non-linear solver would be more suitable. Further work is required to determine whether this approach holds advantages for extraction of functional parameters.

To date, exploration of extraction of functional parameters from optical tomography images has been limited. (McBride *et al*, 1999) describes experiments using a frequency

domain tomographic system with a tuneable Ti:Sapphire laser source. Clear plastic beakers were filled with intralipid (see 2.2.1) and human blood. Nitrogen and oxygen were bubbled through the mixture to change its oxygenation (from fully deoxygenated to fully oxygenated) and single-source to multiple detector measurements were made with the source tuned to 750nm, 800nm and then 830nm. Average background oxygenation and volume values were extracted by calculating μ_a from amplitude and phase measurements using an iterative FEM algorithm (although an ad-hoc correction was required to account for the effects of the clear walls of the beaker holding the liquid). The spectra for water and lipid were subtracted from the retrieved μ_a values, based on an estimate of the water and lipid concentration of the mix. A least-squares fit to the three wavelengths was then performed using spectra for oxy- and deoxy-haemoglobin, to determine their relative concentrations. A co-oximeter probe was used to measure partial pressure of oxygen in the mix, indicating values within 15% of the optical tomography results. A similar method was used to determine the average saturation and volume within the breasts of two healthy subjects. Values found for a 30 year old woman were 70% saturation and 63μM total Hb (so fractional blood volume =100x63/2000=3.15% assuming 2mM), and 67% saturation and 22μM total Hb (1.1%) for a 60 year old woman. Heterogeneous phantom studies were then performed using an intralipid mix, with one or two thin walled 25mm diameter plastic tubes containing contrasting mixes placed within the (84mm diameter) beaker. Absolute saturation and total Hb images were produced from absolute μ_a images at two wavelengths (750nm and 802nm). No attempt was made to derive absolute saturation and blood volume from clinical images.

By fitting spectra to multi-wavelength absolute images, the effect of spurious data points could potentially be reduced. However the wider wavelength rage used by (McBride *et al*, 1999) necessitates compensation of water and lipid spectra, such that the advantage of using 3 wavelengths (where the wavelength independent $\mu_{a,bg}$ cancels) cannot be exploited.

Studies of the accuracy of functional parameter extraction via NIRS (using the modified Beer-Lambert law) have been presented by (Delpy *et al*, 1997), (Boas *et al*, 2001b) and (Firbank *et al*, 1998). In particular, (Firbank *et al*, 1998) uses a similar technique to that described in this chapter when comparing NIRS results with a simple simulation based on equations of the form of [2.5.4]. The effect of $\mu_{a,bg}$ inaccuracy on fractional Hb concentration was examined (in a similar way to Figure 2.5.5). The effects of vessel diameter were discussed, noting that the apparent μ_a derived using NIRS for a given concentration of blood in tissue will decrease if the blood is contained in larger vessels (although this should be better accounted for in optical tomography since spatial heterogeneity can be accounted for).

Three-dimensional CW optical tomography images of changes in oxy- and deoxy-haemoglobin in the frontal lobe during a Valsalva manoeuvre are reported in (Bluestone *et al*, 2001). Data were acquired over the forehead at two wavelengths (760nm and 830nm, simultaneously via low frequency modulation see section 1.2.2.1), and a 3D FEM based image reconstruction was used to extract $\Delta\mu_{a,760nm}$ and $\Delta\mu_{a,830nm}$. Equations of the form of [2.5.20] and [2.5.21] were utilised to extract Δ[Hb] and $\Delta[HbO_2]$, assuming $\mu_{a,bg} = 0$. During the manoeuvre, the measured signal changed by up to 40%. Reconstructed images indicate regional changes in oxy-haemoglobin concentration that are noted to be much larger than those observed by other researchers. Use of a simplistic homogenous model to reconstruct difference data acquired on the (necessarily heterogeneous) head may have adversely affected the accuracy of these results (as chapter 2.4 would suggest). The authors state that it is also possible that signal changes are originating from elsewhere such as brain movement or cutaneous circulation fluctuations but that work is currently underway to validate the presented results.

2.6 Effects of non-scattering regions (voids)

The modelling and image reconstruction techniques described so far have allowed us to describe light propagation through diffusive media via the diffusion approximation to the RTE (see 1.3.1.1). However, the major assumption of the diffusion approximation is that $\mu'_s >> \mu_a$ in all regions. The problem of accounting for regions that do not comply with this assumption has been the topic of much research in recent years (Arridge *et al*, 2000a), (Hielsher *et al*, 1998) particularly in reference to brain imaging. If the effects of having such regions prove to be beyond that capabilities of a diffusion approximation – based image reconstruction algorithm they may represent a major obstacle in the development of a successful clinical optical tomography (brain) imaging system.

This chapter introduces the origin of the 'void problem' in brain imaging, and describes the theory behind a hybrid radiosity-diffusion version of TOAST that can model light propagation through regions with $\mu'_s \approx \mu_a$. An experiment is then described that aimed, for the first time, to validate the results from the 3D void-TOAST code. Detailed analysis of the measured data is then presented, investigating the match with simulated values, the effects of the refractive index mismatches, and the ability of MONSTIR to measure signal originating from inclusions within a diffusing region beyond a layer where $\mu'_s \approx \mu_a$. Preliminary image reconstructions of the simulated and measured data using a diffusion-only algorithm are presented. Implications of the results for MONSTIR, void-TOAST and optical tomography of the brain in general are then discussed.

2.6.1 The clinical origin of the 'void problem'

For optical brain imaging, we must consider the structures within the head that may violate the diffusion approximation. Figure 2.6.1 shows the layers that light must pass through to penetrate the brain. The brain is wrapped in a number of membranes including the dura and the arachnoid membrane. The space between the arachnoid membrane and the brain (the subarachnoid space) is filled with cerebrospinal fluid (CSF). Small trabeculae and blood vessels also weave through this gap. CSF is a clear ($\mu'_s \approx 0$, $\mu_a \approx 0$) straw-coloured liquid made up mostly of water, it cushions the brain, which is effectively floating, and distributes nutrients and chemicals. CSF is continually circulating, being produced and reabsorbed.

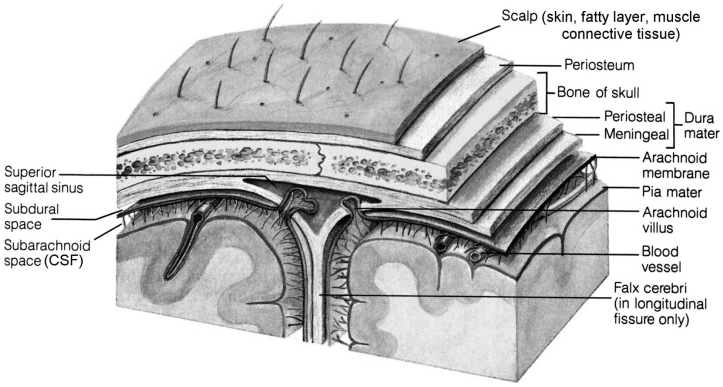

Figure 2.6.1 The layered structure of the adult head

CSF also fills any sulci (areas where the brain folds in on itself) and the ventricles, which are located in the centre of the brain as shown in Figure 2.6.2 (and Figure 1.1.10 in section 1.1.3.1). Intraventricular haemorrhage can affect the peri-ventricular white matter (see Table 1.1.2). Ideally optical tomography would be able to resolve the functional integrity of the tissue surrounding the ventricles. Further, neonatal hydrocephalus can result from congenital malformations and intraventricular haemorrhage, causing the CSF to build up in the brain, enlarging the ventricles and potentially restricting periventricular perfusion (Figure 2.6.2).

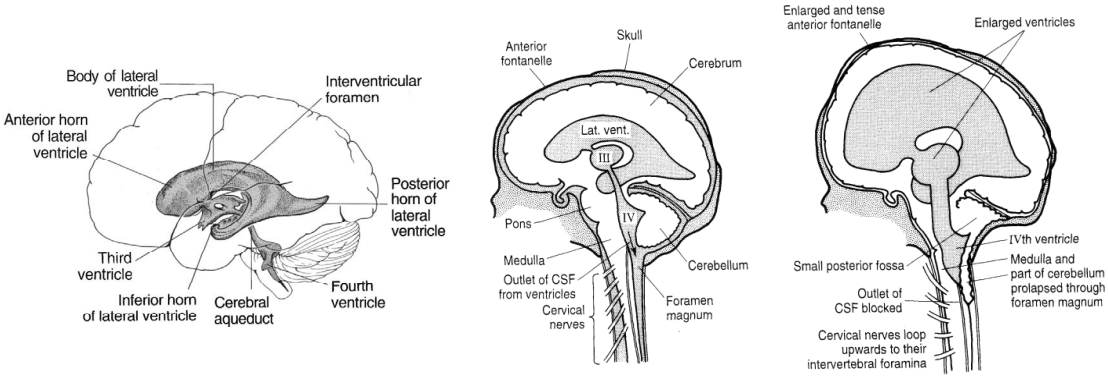

Figure 2.6.2 (left) The ventricles of the adult brain (reproduced from (Crossman et al, 1995) (p8))(right) Development of neonatal hydrocephalus following obstruction of the circulation by the Arnold-Chiari malformation (the protrusion of part of the cerebellum through the foramen magnum obstructs the downward flow of CSF) (from (Kelnar *et al*, 1995) *(p228))*

It has been suggested that the presence of CSF in the head will affect the propagation of light into the brain during imaging and NIR spectroscopy (Okada *et al*, 1997), (Dehghani *et al*, 2000a). Further, since the diffusion approximation is not valid for tissues with $\mu_a \approx \mu'_s$, CSF may affect the ability of diffusion equation based FEM to model light propagation through these 'voids', and hence the ability of algorithms such as TOAST (see section 1.3.2.3) to reconstruct images of the brain.

Early investigation of light propagation through 'clear layers' implied that 'light piping' would occur if the layers were smooth and light was able to bounce freely over large distances before being obstructed. This was shown to affect sampling depth and volume of the diffusing region beyond a clear layer in reflection (topographic) geometries (Okada *et al*, 1997),

(Firbank *et al*, 1995b), and for tomographic data (Ripoll, 2000a), (Dehghani *et al*, 2000c). The relevance of this result to brain imaging and spectroscopy was questioned however since the surface of the brain is not entirely smooth, and so light will be unlikely to find may paths where it can propagate freely and unobstructed through the subarachnoid space (Takahasi *et al*, 1998). However, in the case of the premature infant, the surface of the brain has very few convolutions, as shown in Figure 2.6.3.

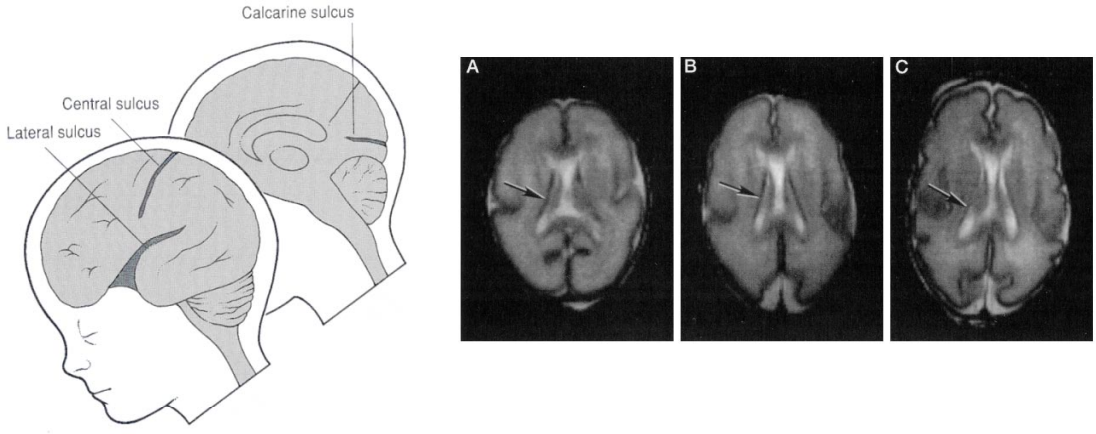

Figure 2.6.3 Only the major folds are present in the cortex of a 28 week old neonate (reproduced from (FitzGerald, 1996) *(p6)).(right) T2-weighted MRI images of the brain of a 732-g infant born at 27 weeks gestation and first imaged at 2 days of age(A), then at 28 (B) and 30 (C) weeks. Note that at 27 weeks the cortex is almost smooth (Website, 3)*

The scattering nature of the CSF layer surrounding the brain has also been questioned since small trabeculae *may* increase the net scattering of light in the region (Okada *et al,* 2000).

The void problem is difficult to study in vivo as it is not easy to judge whether signals have been truly affected by light piping and the propagation of light through the ventricles. Ways to compensate for the effects of voids require a method of simulating the propagation of light through both highly scattering and low scattering regions. Further, possible image reconstruction algorithms that truly include void regions require such a model (see section 1.3.2).

2.6.2 Modelling void regions

As described in section 1.3.1.4.2, numerical methods that can tolerate tissues that do not comply with the constraints of the diffusion approximation include Monte Carlo (Okada *et al*, 1997) and full radiative transport equation models (Hielsher *et al*, 1998), (Dorn, 2000). Monte Carlo models take a long time to perform and are not suitable for use in an iterative image reconstruction scheme. RTE models are very difficult to solve, although progress is being made in implementing a transport-based approach to the void problem, (Dorn, 2001). The

hybrid FEM model described below is described in (Arridge *et al*, 2000a), is potentially a versatile and fast method for simulating and reconstructing void data.

2.6.2.1 Hybrid Diffusion / Radiosity algorithm for modelling voids

Modelling the propagation of light through non-scattering regions can be readily achieved via Radiosity (used widely in computer graphics to model lighting and reflections from diffusing surfaces). However, if we have a diffusing object that contains non-scattering regions (or voids), neither a diffuse model nor a radiosity model is sufficient. Hence, a version of TOAST called "Void TOAST" has been developed that incorporates both models (diffusion and radiosity). A similar CW only model is described by (Takahasi *et al*, 1998).

Modelling the propagation of light using a specific model for a specific region requires the boundaries of the regions to be known. In order to create a model, the visibility of all points on the boundaries of the voids must be established. Diffusion approximation FEM can then be used to model the field at the internal boundaries, while radiosity is simultaneously used to establish how photons will propagate through the void. The radiosity solution for a non-scattering void region Ω, with $\mu_a = \mu_a(\Omega)$, within a diffusing medium (as shown in Figure 2.6.4) yields an expression for the isotropic source of photons at *a*:

$$q_0(a) = \int_{\partial\Omega} h_{a,b} \frac{\hat{s}_{a,b} \cdot \hat{v}(b) \hat{s}_{a,b} \cdot \hat{v}(a)}{\pi d^2_{a.b}} e^{-\mu_a(\Omega) d_{a,b}} \mathcal{B}\Phi(b) db \,,$$

[2.6.1]

where $h_{a,b}$ is equal to 1 when points *a* and *b* are mutually visible, but equal to zero if not (e.g. *a* and *c* in Figure 2.6.4). The distance between points *a* and *b* is given by $d_{a,b}$, $\mathcal{B}\Phi(b)$ is the boundary flux ($\Gamma(b)$) where $\Phi(b)$ is the (diffusion region) field at *b*. The inward normals at points *a* and *b*, along with the direction of the line joining *a* and *b* are defined in Figure 2.6.4.

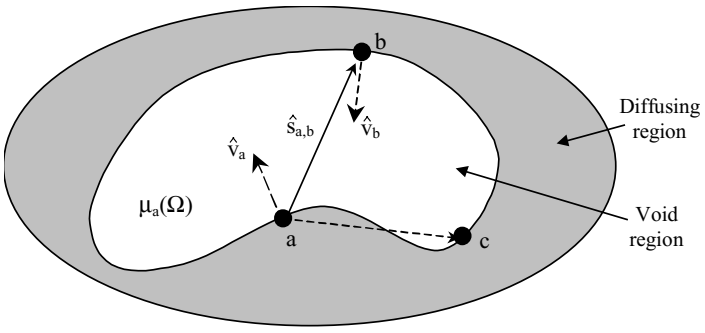

Figure 2.6.4 Radiosity requires the visibility of two points a, b to be considered, points a and c are not mutually visible.

The radiosity source term in [2.6.1] can be used to modify the finite element operator derived in section 1.3.1.3.1, such that an extra matrix **E** describes coupling across all voids *k*. The visibility is now only considered between the nodes on the internal boundaries and refractive index is assumed to be uniform throughout. The coupling matrix is defined as:

$$\mathbf{E} = \sum_k E^{(k)}{}_{i,j} = \sum_k \mathcal{B} \int_{A_i} u_i(a) \int_{A_j} u_j(b) h_{a,b} \frac{\hat{s}_{a,b} \cdot \hat{v}(b) \hat{s}_{a,b} \cdot \hat{v}(a)}{\pi d^2{}_{a.b}} e^{-\mu_a(\bar{\Omega}) d_{a,b}} d^{n-1} b d^{n-1} a$$

[2.6.2]

such that the FEM operator in [1.3.31] becomes:

$$\left[\mathbf{K}(\kappa) + \mathbf{C}(\mu_a c) + \mathcal{B}A + i\omega\mathbf{B} - \mathbf{E}\right]^{-1}$$

[2.6.3]

2D Void-TOAST has been used to model the changes in PMDF that would results from 2D distributions of voids (e.g. from an MRI scan of a premature infant, see Figure 2.6.5 below). It can also be used to estimate the amount of light penetration achieved in topographic measurements, if a 2D mesh were created from an adult head MRI (Dehghani *et al,* 2000a). Recently, 3D FEM modelling integrating voids has been developed. Preliminary results are given in (Riley *et al,* 2000) however complex geometries are yet to be explored. As described in section 2.3.1.4, a 3D model is required since quantitative analysis of results from 2D models will not necessarily agree with results from a 3D distribution of voids.

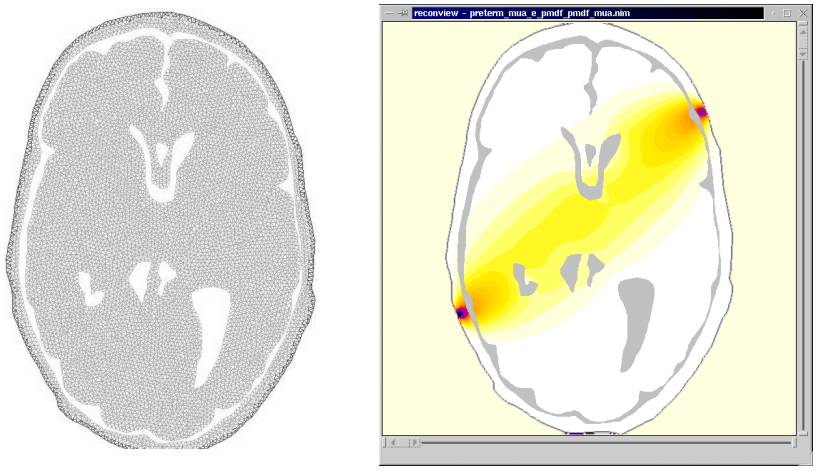

Figure 2.6.5 Left: a complex 2D mesh based on an MRI of a pre-term baby head. . Right: PMDF for intensity and μ_a for the 2D mesh, assuming that all regions besides the voids are homogenous. (simulation by Hamid Dehghani)

Image reconstruction can be achieved using void-TOAST. 2D studies (Dehghani *et al,* 2000b) and (Dehghani *et al,* 2000c) have shown that:

- Forward (absolute) data simulated using void TOAST cannot be reconstructed successfully using a homogenous mesh and diffusion-only-TOAST.
- If the void mesh (with the same boundaries) and void TOAST are used to reconstruct the void TOAST forward data, images are much improved compared to using diffusion only.
- If data are simulated using void TOAST, and reconstructed on a similar (but not identical) void mesh, or noise is added to the forward data, reconstruction using void TOAST is badly affected.

This means that in the case of data acquired on an infant's head, if the CSF regions are sufficiently sizeable and non-scattering that data are affected, absolute image reconstruction

using void-TOAST is only possible if the exact (3D) boundaries of the void regions are known. This is an impractical constraint since, even though MRI images can reveal the geometry of the clear regions, accuracy will depend on the algorithm used to extract the 3D boundaries from a set of images. Further the head of a premature infant will change shape in the space of hours, and grow day by day, altering the accuracy of MRI-derived information on the clear regions. Methods to find the internal boundaries of void regions from NIR data, as part of the inverse problem, are being investigated (Dorn, 2001), (Kolehmainen *et al,* 2000). Also, studies to establish the limits of accuracy required to use void-TOAST for measured data are currently being performed, although in general, improved models of true clinical void distributions are required to evaluate the magnitude of the void problem.

Although the 2D void model has been used to demonstrate a number of aspects of light propagation through clear layers and reservoirs only Monte Carlo has been used to demonstrate that the hybrid model is correct (Arridge *et al,* 2000a). An early, 2D version of the Void-TOAST code was compared with Monte Carlo simulations and reflection – geometry experimental measurements using a streak camera in (Okada *et al*, 1997). The 3D version of Void-TOAST was compared to Monte Carlo in (Riley *et al*, 2000) but until now, has not been compared to experimental data. An experiment using MONSTIR (see 1.2.3) was devised to examine the accuracy of the 3D void model, and also to examine the practical constraints on measurements and image reconstructions in the presence of a void region. Prior to development of a 3D void model, it was not possible to perform experimental validation, since a 2D simulation cannot be expected to agree with an experiment where photons are free to migrate in all directions (see section 2.3.1).

2.6.3 Experimental validation of the 3D void model.

2.6.3.1 Experimental set-up for the void phantom

Since the majority of experimental void studies have only considered reflection-geometry fibre positioning, these studies were designed to use fibres in a true tomographic arrangement. The 3D version of void-TOAST required a mesh that was composed of simple geometric shapes. Further, since light is likely to propagate for large distances within a void region, using a cylindrical void is not a suitable choice. The simplest 3D geometry to model is a spherical void, either as a full non-scattering region, or with a concentric diffusing region inside. A spherical phantom would be very difficult to make and measure using existing MONSTIR fibre holders. The following phantom was therefore designed (see Figure 2.6.6).

Epoxy resin, mixed with TiO_2 particles (white paint) and dye were mixed, cured, and machined into a cylindrical shape (70mm in diameter) using techniques and the recipe

described in (Firbank *et al*, 1995a). The centre of the cylinder was the machined away to leave a beaker-shaped shell 5mm thick. A glass round-bottomed flask, 53mm in diameter with a 49mm inner diameter, was then positioned inside the beaker, centred on a small indentation in the centre of the base of the beaker, and held at the very top by a polystyrene cuff to prevent it from moving (see Figure 2.6.8). A black Delrin circular fibre holder was then positioned around the beaker, level with the widest part of the flask, and 16 source and detector fibres (in monstodes – see section 2.7.2.2.1) were attached.

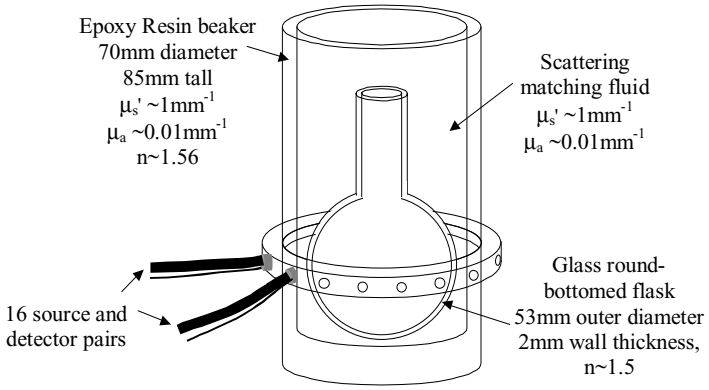

Figure 2.6.6 The phantom used to acquire 'void data'. The image on the right demonstrates that the flask can be removed to provide a 'homogenous' phantom for comparison with data collected in the presence of the void.

As shown in Figure 2.6.7, the phantom could then be filled with:

1) Scattering liquid (μ_a = 0.01mm^{-1} μ'_s = 1mm^{-1}) in the beaker and within the round-bottomed flask, to create a ~2mm thick spherical void region with diffusing regions either side.

2) Scattering liquid around the outside of the flask, but clear (μ'_s = 0) inside the flask, to create a large spherical non-scattering region.

3) Having removed the flask, the beaker could be filled with scattering liquid, creating a homogenous, completely diffusing phantom, for comparison with data acquired in the presence of a void.

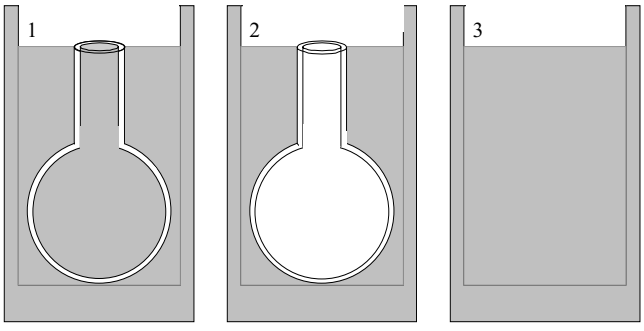

Figure 2.6.7 The three possible configurations of the void phantom. 1) clear layer, 2) clear region, 3) homogenous.

In addition to the effects of voids on light propagation, the phantom was designed so that small inclusions with optical contrast could be suspended, off axis inside the flask. Two

cylindrical inclusions were used, one 10 mm in diameter and 10mm tall and the other 6mm in diameter and 7mm tall. Each could be positioned 15mm and 17mm from the central axis of the phantom. Their optical properties were nominally μ_a = 0.05mm^{-1} and μ'_s = 1mm^{-1}. Thus data could be acquired with and without inclusions, for both the clear layer and homogenous configurations.

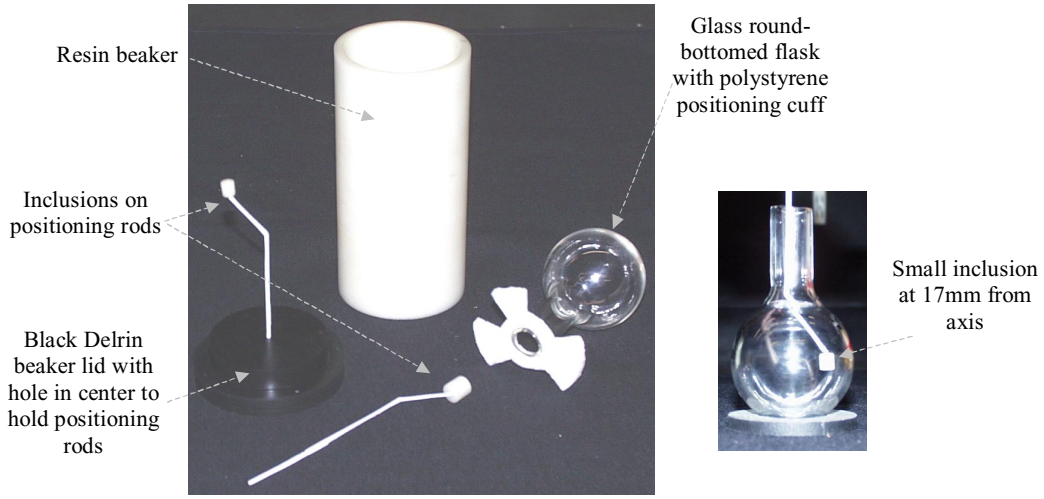

Figure 2.6.8 Photographs of the elements of the void phantom showing how the inclusions were positioned inside the glass flask.

2.6.3.2 Simulations of the void phantom

Simulations of the phantom were created by Jason Riley, who wrote the 3D void-TOAST code (Riley *et al*, 2000). Figure 2.6.9 shows the 3D finite element mesh created to model the round-bottomed flask within the cylindrical beaker.

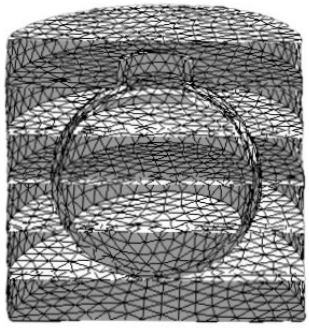

Figure 2.6.9 The finite element mesh used to simulate void-data using the hybrid void-TOAST model.

Simulations were performed for cases 1, 2 and 3 as defined in Figure 2.6.7. The scattering regions were assumed to have μ_a=0.01mm^{-1} and μ'_s = 1mm^{-1}. The non-scattering regions had μ_a = 0.002mm^{-1} (based on water) and μ'_s=0. Refractive index was assumed to be 1.56 (that of resin) throughout. Two simulations were also created that contained an inclusion for cases 1 and 3 representing the large inclusion at a radial distance of 17mm. The refractive index was assumed to be 1.56 throughout. The scattering regions had μ_a = 0.01mm^{-1} and μ'_s = 1mm^{-1}, non-scattering regions had μ_a = 0.002mm^{-1} and μ'_s = 0 (based on water, see 1.1.2.1). The simulated mean-time results (no inclusion) are shown in Figure 2.6.10 for cases 1 to 3.

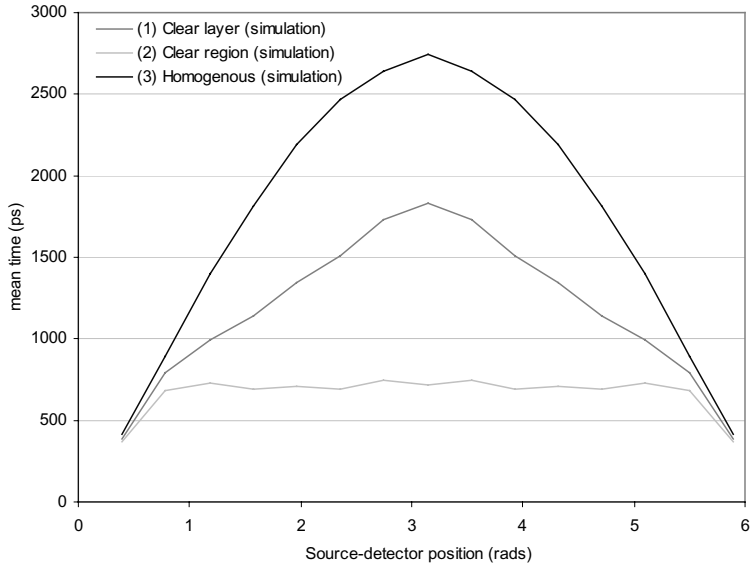

Figure 2.6.10 Simulated mean-time values for the three cases illustrated in Figure 2.6.7.

Points plotted represent the 15 detectors around the fibre-holder ring opposite the source as shown in Figure 2.6.11 (this is the same format as the data plots in section 2.2.6.2)

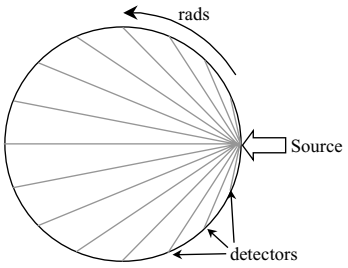

Figure 2.6.11 The source-detector geometry used for the simulated data.

The intensity simulations are plotted as the ratio of case1/case3 and case2/case3. This is because it is not possible to accurately calibrate absolute intensity measured by MONSTIR (see section 2.1.2.4), but taking that ratio between two intensity measurements, acquired with the same VOA positions gives a good representation of the relative intensity of signals on one object to those on another.

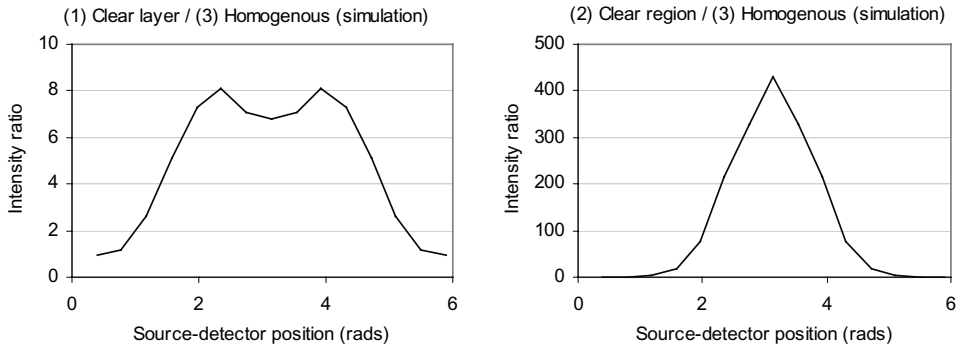

Figure 2.6.12 The ratio of simulated intensity measurements for cases 1 and 2 (relative to the homogenous case 3).

It is also possible to use void-TOAST to simulate the intensity field Φ within each of the phantoms. The images shown in Figure 2.6.13 show the simulated field due to a single source

(incident from the right) plotted across the x-y plane and the x-z plane. The images are scaled to a common (logarithmic) colour scale.

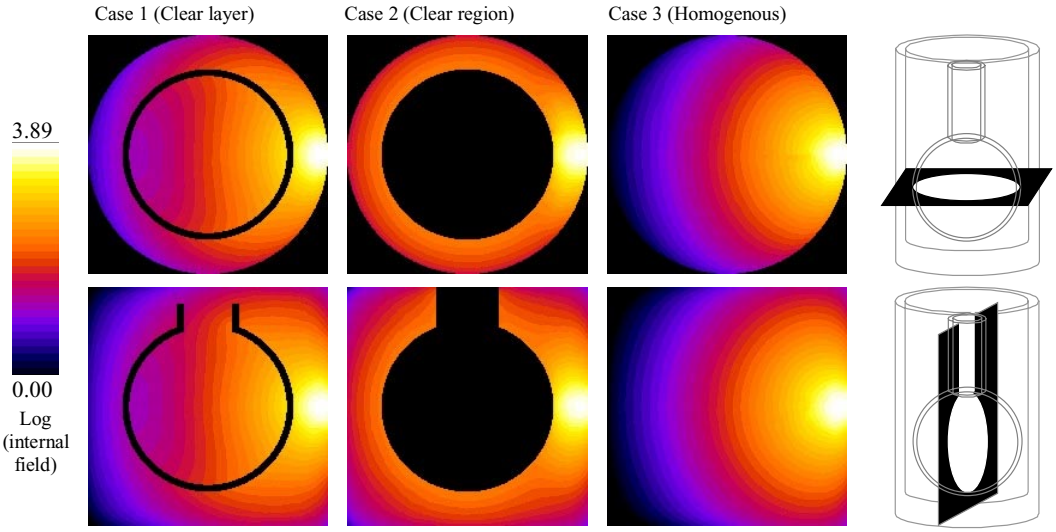

Figure 2.6.13 Plots of the internal log(intensity) fields due to a single source on each of the three phantoms. (x-y plane and x-z plane). Common (logarithmic) colour scale ranges from 0 to 3.89.

2.6.3.3 Experiment 1 (intralipid)

Intralipid is an emulsion of water and soy-bean oil, used as an intravenous nutrient and commonly used as a scattering liquid in NIR phantoms as described and used earlier in chapter 2.2. In the first void experiment, a mixture of intralipid and dye was used as the scattering liquid as shown in Figure 2.6.7. A volume of intralipid was diluted and mixed with NIR absorbing dye to have an approximate $\mu_a = 0.01mm^{-1}$ and $\mu'_s = 1mm^{-1}$ using the methods described in section 2.2.1 . Data were acquired for cases 1, 2 and 3 (with distilled water used to fill the flask for case 2). Data were also acquired with the large inclusion suspended inside the flask (with the flask filled with intralipid as for case 1) 17mm and 15mm from the axis.

Acquisition definition files (used to pre-define the positions of the VOAs for each source position, see section 1.2.3.5) were generated to ensure that the maximum number of sources and detectors were used. The three cases described generated very different intensity variations around the phantom (as suggested by the simulated fields in Figure 2.6.13). For cases 1 and 3 (and corresponding data with inclusions present) it was possible to use the same ADF (thus allowing the intensity data to be directly ratio'd). Additional data were acquired using 'near' and 'far' acquisitions; where the detectors close to the source are shut-off, and a high intensity illumination is used to collect data for distant detectors. Close detectors are then opened, and a neutral density (ND) filter placed between the laser output and the delivery fibre to reduce the illumination intensity to avoid damaging close detectors (a similar technique was used to measure the 3D breast phantom described in section 2.3.3.2). For case 2, a larger ND filter and a slightly different ADF file were needed, although the 'near' and

case 2 ADF files were the same for the opposite 11 detectors, so intensity is only compared for these combinations.

Mean-time data were extracted from the measured TPSFs and calibrated using measurements made from back-scattered light within the monstodes and a separate source calibration measurement (see section 2.1.4.2). The data were then processed by assuming that the phantom was entirely axially symmetric. The absolute mean-time data for each equidistant source-detector pair were averaged (using the same techniques as for data shown in section 2.2.6.2). The resulting profiles are plotted in the same way as the simulated data, and are shown in Figure 2.6.14. The standard deviation of each point within the profile is shown by the error bars. One detector (number 13) was eliminated from the averaging since its TPSFs were distorted due to a damaged CFD.

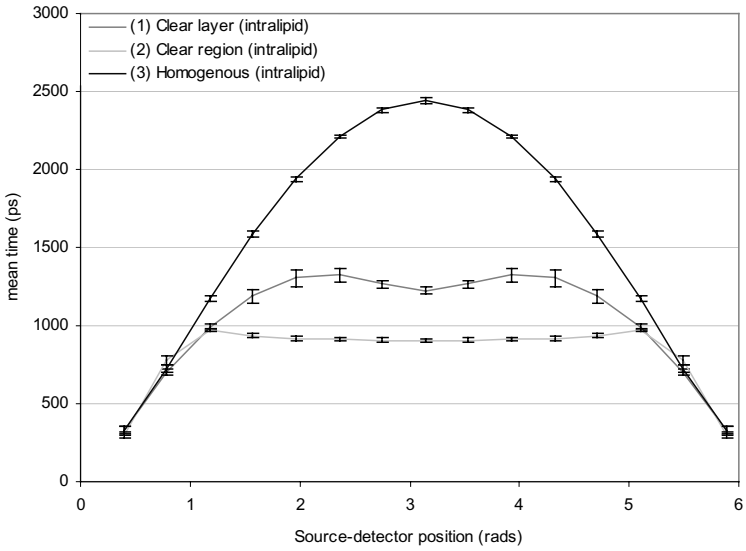

Figure 2.6.14 Averaged mean-times for experiment 1 (using intralipid as the scattering liquid) for all three cases.

Clearly, the mean-time results of this experiment do not agree well with the simulated plots shown in Figure 2.6.10. Especially for case 2 (the clear layer), where the distinct peak in the centre of the simulated plot is a dip in the experimental result. The ratio of the intensity measured on case 1, relative to case 3, and case 2 to case 3 are shown in Figure 2.6.15.

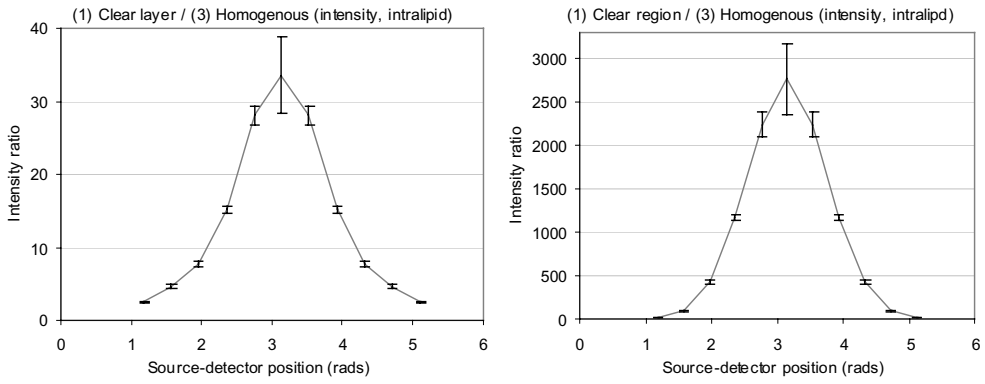

Figure 2.6.15 The ratio of intensity measurements for cases 1 and 2 (relative to case 3) for experiment 1.

As described in section 2.1.2.4, it is only possible to determine features of intensity data by using ratios of MONSTIR measurements. The data shown in Figure 2.6.15 were processed by calculating the raw intensity from each TPSF (following background subtraction). These raw data were then ratio'd. All data from detector 13 were removed (since it was faulty), and any spikes 4 or more times larger than their neighbouring data points were removed. The intensity ratio profile for each source was then normalised. This was necessary because of a problem with the secondary fibre switch (described in section 2.1.1.3). At the time of this experiment, a fault in the unit meant that the time taken for a shutter to open might be different to the time taken by others. Further, this time was variable for a particular switch. This meant that the exposure time for a given source may be different from one data acquisition to the next, scaling the ratio for a particular source according to whether the switch in the first data collection had taker more or less time to open in the second. The normalised ratio data were then averaged in the same way as the mean-time data, and the standard deviations calculated. The final average profile and standard deviations were then scaled by the mean normalisation factor, assuming that the exposure time in the first acquisition was just as likely to be less than that in the second as it was likely to be more. The ratios were also scaled by an additional factor due to the differing ND filters required in some acquisitions.

Again, the intensity ratio results from this experiment do not agree well with the simulated values shown in Figure 2.6.12. It was suspected that the difference between the simulation and the experimental data might be due to the refractive index of the intralipid solution. The flask is made of glass, which has a nominal refractive index of 1.5. The intralipid, a mixture of oil and water, is likely to have a refractive index of around 1.4. This means that total internal reflection can occur within the glass walls of the flask. The resin beaker was also likely to have a refractive index of 1.56, causing a further mismatch. Neither of these refractive index mismatches is accounted for in the void-TOAST simulation. A second experiment was designed to test this, as described below.

2.6.3.4 Experiment 2 (resin)

In order to find a scattering liquid that would be a better refractive index match to the glass flask, initially adding sugar to the intralipid was investigated. However, changing the refractive index mismatch between the small oil droplets and the water in the intralipid changes the scattering properties of the intralipid in an unpredictable way. Instead it was decided to try using un-cured resin as the scattering liquid. In order to make a solid phantom, such as the beaker used in this experiment, epoxy resin is mixed with scattering TiO_2 particles and dye. 3 parts resin to 1 part hardener are mixed, and the resulting liquid cured in an oven. For this experiment, a volume of resin was mixed to the same recipe as the solid beaker, but

with the 1/4 hardener replaced with additional resin. This liquid's scattering and absorbing properties (including the anisotropy factor etc.) are likely to match the surrounding beaker far better than the intralipid solution used in experiment 1. The refractive index of this liquid was found to be 1.59 ± 0.04 by using MONSTIR to measure the difference in the temporal delay of light passing through an empty cylinder compared to the cylinder filled with the (clear) liquid resin (n=ct/d). This value agrees well with the value of 1.56 measured for cured resin.

The protocol for experiment 1 was repeated using the un-cured resin as the scattering liquid instead of intralipid (and pure un-cured resin instead of water for the clear-region case). New ADF files were required in some cases. In addition to cases 1, 2 and 3, data were also acquired with the large inclusion positioned at 15mm and 17mm from the central axis, both with the flask in place (with the inclusion within the flask as shown in Figure 2.6.8), and with the inclusion simply suspended in the homogenous phantom. The data were processed in the same way as for experiment 1. Mean-time results are shown in Figure 2.6.16.

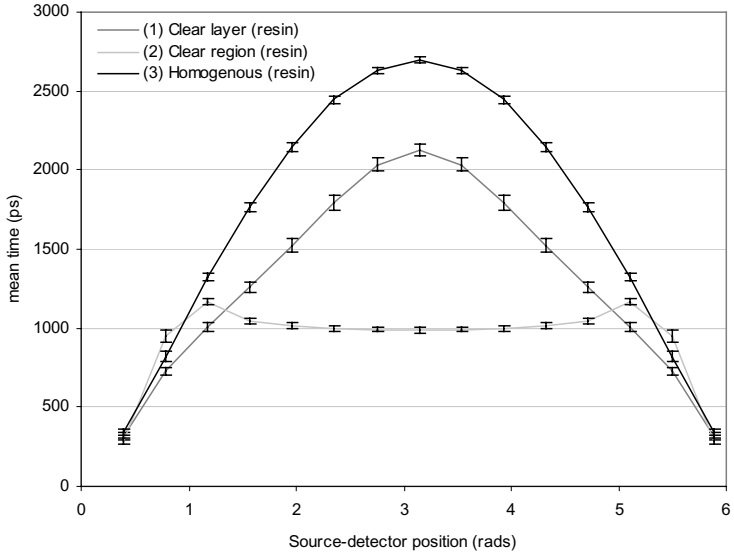

Figure 2.6.16 Averaged mean-time results for experiment 2 (with resin as the scattering liquid) for all three cases.

These data look much more like the simulated data shown in Figure 2.6.10 than the results from experiment 1 shown in Figure 2.6.14. The resin experiment intensity ratios are shown in Figure 2.6.17. These data also look more like the simulated plots in Figure 2.6.12.

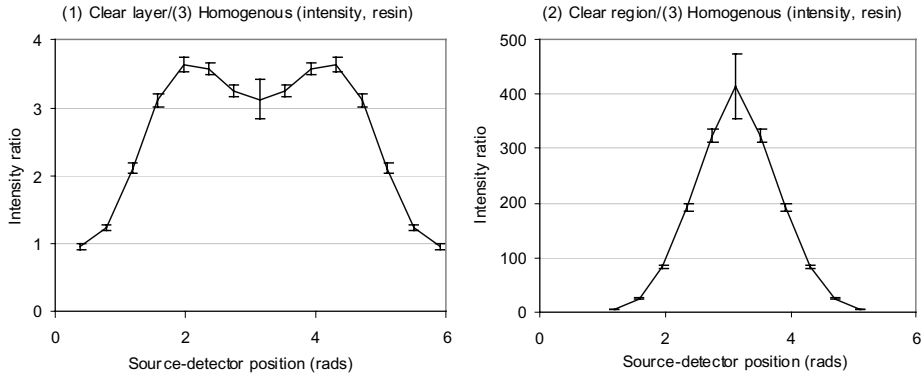

Figure 2.6.17 The ratio of intensity measurements for cases 1 and 2 (relative to case 3) for experiment 2 (resin).

2.6.4 Analysis of non-inclusion results

While the results from experiment 2, using resin as the scattering fluid, match the simulated results better than those from experiment 1, where intralipid was used, the agreement is not exact. Below, we investigate each 'case' in turn. Analysis of the data acquired with and without the inclusion is presented in section 2.6.5 .

2.6.4.1 Case 3 (homogenous)

Figure 2.6.18 shows the homogenous data for the simulation, experiment 1 and experiment 2. We can see that the resin and simulated cases match very well, indicating that the optical properties used in the model are a good match to the properties of the experiment. The intralipid curve is slightly lower than the others, this may be due to the refractive index mismatch between the beaker and the intralipid, and also due to the absorbing and scattering properties of the intralipid being slightly different to those of the resin.

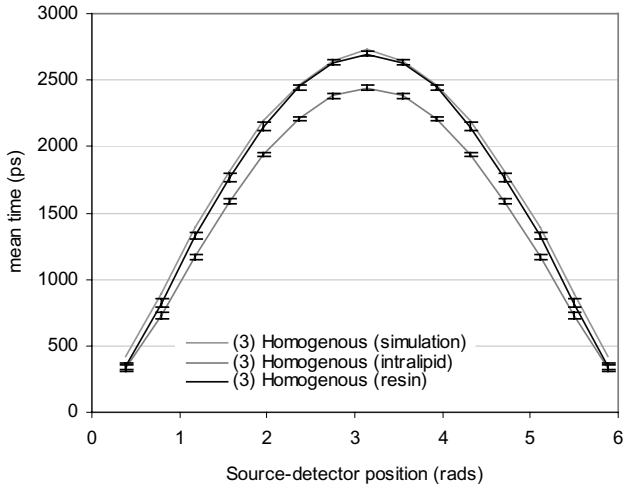

Figure 2.6.18 The homogenous (case 3) mean-time results of experiments 1 (intralipid), 2 (resin) and the simulation.

2.6.4.2 Case 1 (Clear layer)

Figure 2.6.19 shows the mean-time results for the clear layer (case 1, left). The plot on the right shows the intensity ratios of the clear-layer data to the homogenous case data for experiments 1 and 2 and the simulation.

These clear layer results suggest that the mismatch between the intralipid results and the simulation were indeed due to refractive index mismatch. With intralipid either side of the glass shell, the refractive index mismatch is such that very efficient specular reflections can occur within the shell. This means that much of the light that enters the glass walls of the flask is readily propagated round to the other side of the phantom, without passing through any more diffusing material than the small amount between the source and the flask, and the

flask and the detector. Since the light that has passed through the diffusing region (e.g. through the body of the flask) will have been significantly attenuated, for the intralipid experiment, the majority of photons arriving at the far detector will be those that have passed through the glass, relatively unattenuated. Since they took the route through the non-scattering layer, they will take less time to reach the other side, which is why the mean-times for the furthest detector (source-detector separation = π), are lower than those for the simulated case. This may explain the distinct dip in the mean-time data seen only in the intralipid results.

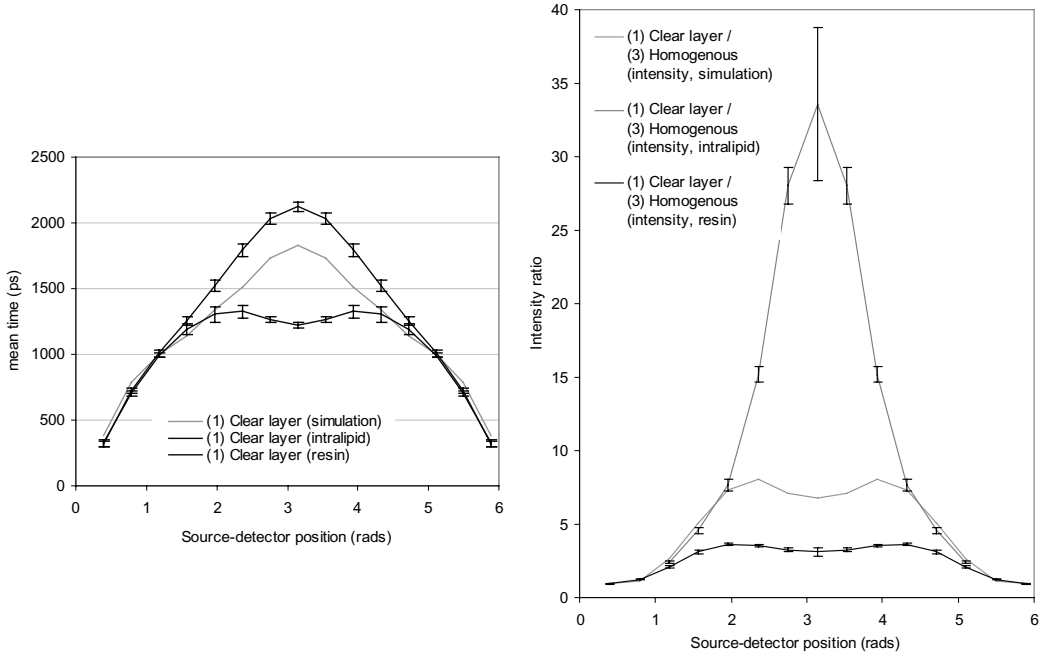

Figure 2.6.19 (Left) The clear layer (case 1) mean-time results of experiments 1 (intralipid), 2 (resin) and the simulation, (right) the clear layer:homogenous ratios for experiments 1 and 2, and the simulation.

For the simulated case, effectively there is no refractive index mismatch and hence no efficient specular reflections occurring at the flask / liquid interface (although photons *will* be back-scattered into the void region). This explains why the mean-times for the simulation are longer than for the intralipid results, since transport in the clear layer is less efficient. In addition, the intensity is much higher on the opposite side of the flask for the intralipid experiments compared to the simulation. This agrees with the hypothesis that refractive index mismatch between the intralipid and the glass caused most of the light detected to preferentially travel through the clear layer.

If refractive index match is in opposite direction, light is less likely to enter the void at all. This is the case of the resin: glass interface, where reflections can occur on the *outer* surface of the void (since the refractive index of the resin ($n \approx 1.56$) is higher than that of the glass ($n=1.5$)). This means that light is more likely to stay in the diffusing areas than in the intralipid case, or the simulated case, and may explain why the mean-times for the resin experiment are somewhat higher than those for the simulation. This also agrees with the

intensity results, where the intensity detected on the opposite side of the phantom for the resin data is lower than for the simulation and intralipid data, since many of the photons have passed through diffusing and attenuating material before reaching the other side. Note also that the increased refractive index of the resin compared to intralipid may account, in part, for the faster propagation of light through the intralipid and hence the lower corresponding mean-times.

In general the shape of the resin and simulated results demonstrate that a clear-layer represents a 'short-cut' for light. If a large fraction of the light detected has arrived at the detector having primarily travelling within the void, comparatively fewer photons will have probed the diffuse regions. This is illustrated well by the simulated field plots in Figure 2.6.13 where a bright region can be seen all the way round the boundaries of the void implying that light is back-scattering and remaining within the void rather than passing through the central diffusing region. This effect is explored in section 2.6.5 later where the ability of MONSTIR to resolve signals relating to perturbations within the diffusing central region is compared for each case.

2.6.4.3 Case 2 (clear region)

Figure 2.6.20 shows the mean-time results for the clear region (case 2, left). The plot on the right shows the intensity ratios of the clear-region data to the homogenous case data for experiments 1 and 2 and the simulation.

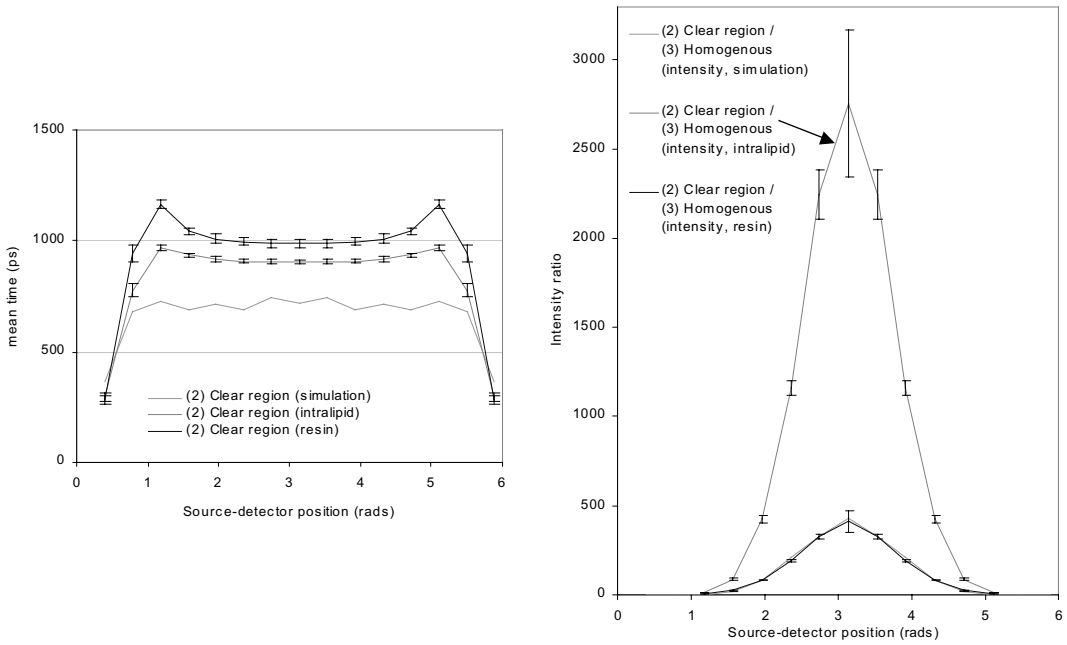

Figure 2.6.20(Left) The clear region (case 2) mean-time results of experiments 1 (intralipid), 2 (resin) and the simulation, (right) the clear region:homogenous ratios for experiments 1 and 2, and the simulation (simulation overlaps resin line).

The plateau in the mean-time results resembles the results shown in (Okada *et al*, 1997) where a reflection-geometry measurement was made on a slab-like phantom with layers representing scalp, skull CSF and brain. A streak camera was used to measure mean-time with differing source-detector spacings. A plateau was reached after a certain source-detector separation since the majority of light was passing directly through the void and then through just the small amount of scatterer near to the source and the detector. The plateaux occurred once the spacing was greater than twice the depth of scatterer between the source and the clear layer. An angle of 1.5 rad, which is where our data reaches a plateau, corresponds to an approximate separation of ~ 2cm which agrees well with the fact that the sources and detectors are ~ 9mm from the void.

In general, intensity ratios are *much* bigger than for the clear-layer data (Figure 2.6.19) in all cases. This is supported by the intensity field plots in Figure 2.6.13 which imply that the majority of light passes straight across the void region. In addition the mean-time values are all much lower, again implying that the majority of light is taking a short-cut through the void region.

Refractive index mismatch in this case clearly has a large effect. The intensity of the intralipid case is much larger than for the simulated and resin cases, the intensity for the resin and simulated cases are very similar. However it is difficult to explain why the results of the two experiments and the simulation results are so different. It is suspected that the absorption properties of the void region are also affecting results. The μ_a value chosen for all of the clear regions (including the glass) in the simulation ($0.002mm^{-1}$) was based on water, although the absorption of glass is likely to be lower than this. This could explain why, in reality the intralipid case transmitted light more efficiently than the simulation would imply. The resin experiment would have specular reflections occurring within the very centre of the void, although perhaps the absorption of the resin, or the relative inefficiency of the specular reflections cause the intensity results to resemble the simulation.

The differences in the experimental mean-time plots are attributable to the refractive index mismatch, since light will travel ~ 1.2 x faster through the low refractive index water than through the high refractive index resin. It appears that since specular reflection can occur for both of the experimental results, but not for the simulated case, that the experimental mean-times are longer since light has bounced repeatedly within the clear region. Note however that the difference between each mean-time in this case is only small compared to cases 1 and 2.

The analysis of the three cases detailed above imply that if a void layer is of higher refractive index than the surrounding regions, that there will be very little probing of the

diffusing regions by the detected light (relative to the light that has passed through the clear layer). We can investigate this hypothesis by examining the data acquired with and without inclusions present within the central diffusing region.

2.6.5 Analysis of inclusion data

We can analyse the data acquired with the inclusions present to further investigate the ability of the model to predict measured data. We can also use the data to explore the hypothesis that the lower the clear layer / surrounding tissue refractive index ratio, the better the penetration of light into the central diffusing region (e.g. the brain). While image reconstruction using the 3D void model was not attempted for these experimental data, we can compare the results obtained using a 2D diffusion-only model (TOAST) to reconstruct data from the simulated, resin and intralipid data.

Data were simulated for the source-detector geometry used in experiments 1 and 2, with a region of higher absorption and scatter corresponding to the properties and 17mm radial position of the largest of the inclusions. It is difficult to see evidence of the inclusion in absolute data without calculating the difference between data acquired at a reference state where the inclusion is not present. We therefore examine the inclusion data by looking at the difference between the homogenous case with and without the inclusion, and between the clear-layer case, with and without the inclusion. The approximate positions of the inclusion for the homogenous and clear layer cases are shown in Figure 2.6.21.

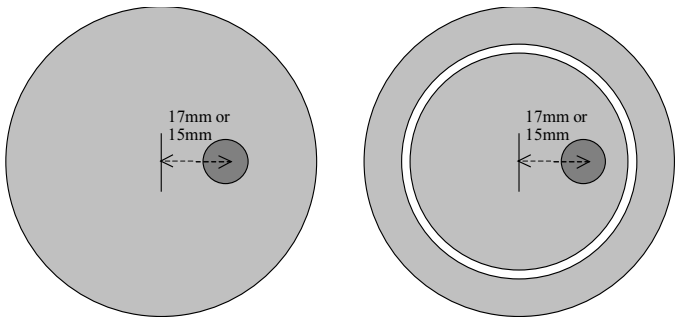

Figure 2.6.21 The position of the large inclusion in a cross section of the homogenous (case 3) and clear-layer (case 1) geometries. The inclusion may be positioned either 15mm from the central axis, or 17mm.

In order to get a more general quantification of the signal due to the blob, the data were analysed by calculating the average detected signal per source (for all detectors), or the average signal for each detector (for all sources). This method, in effect, looks for the shadow of a heterogeneity, since a source positioned such that most of the light emerging from it must pass through the inclusion, will yield a lower total detected (intensity) signal than a source positioned further from the inclusion. Similarly, a detector positioned such that much of the light reaching it has passed through the inclusion will have a lower average detected

(intensity) signal for each source. These averages would be flat lines if the phantom were entirely axially symmetric. Mean-time data would also reveal a characteristic perturbation when averaged in this way. This method was chosen because the raw difference data were found to be quite noisy, and averaging data for each source-detector spacing as utilised earlier would not reveal the inclusion.

Since the phantoms, and simulations were configured such that the inclusions were positioned close to source and detector 1, and therefore the signal that we wish to identify will be level with source 1, the data are presented twice on the x-axis, equivalent of casting the shadow round the phantom twice.

While the average of all detectors per source generally contains a better signal, for intensity data we must use the average of all sources per detector because of the problem with the exposure time consistency due to the secondary source fibre switch described above. Data were extracted by taking the ratio (for intensities) and the difference (for mean-times) between two sets, removing outliers and calculating the mean and standard deviation (not shown for clarity). Intensities were normalised for each source and then re-scaled, as described in section 2.6.3.3.

2.6.5.1 Homogenous (no void) inclusion mean-time differences

Figure 2.6.22 shows the (average) difference in mean-time between the homogenous (i.e. no void) phantom with and without the inclusion for the simulation and experiment 2 (resin). The change in data for the resin experiment is shown for the inclusion being both in the 17mm radial position and the 15mm radial position (see Figure 2.6.21). Note that data in Figure 2.6.22 (and Figure 2.6.23, Figure 2.6.24 and Figure 2.6.25) is repeated as if the source were scanned round twice from 0 to 4π to accentuate signal corresponding to the inclusion.

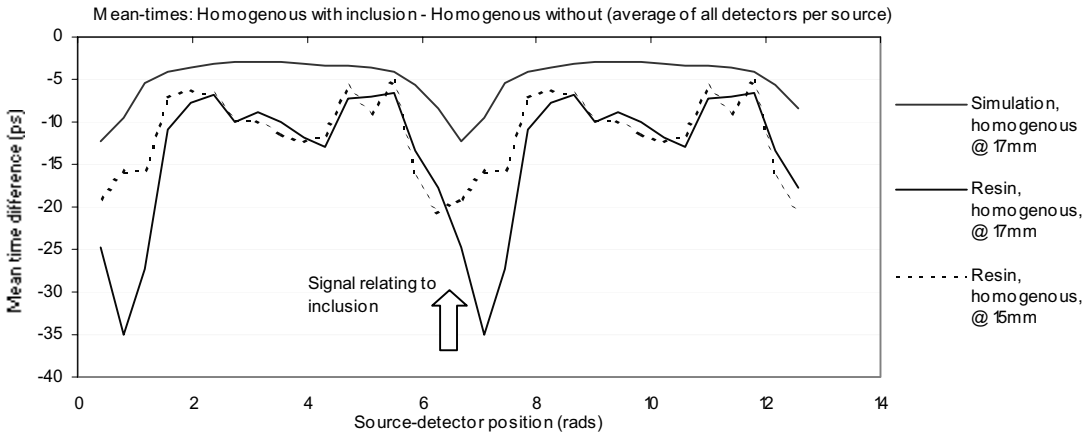

Figure 2.6.22 Mean-time difference averages for all detectors per source for the homogenous (case 3) simulation (with and without the large inclusion at 17mm from the axis) and the homogenous experiment 2 (resin) data (with and without the large inclusion at both 17 and 15mm axial displacement). The inclusion is at ~ 0 rad and 2π rad as indicated by the arrow (data is repeated as if the source were scanned round twice from 0 to 4π).

This plot implies that, while the data is noisy, the presence of the inhomogeneity has affected the measured mean-time in the homogenous case. The data for the inclusion axially displaced by 15mm show a lower signal than for the 17mm position, as we would expect. The fact that the simulated value is smaller than both of the measured data sets implies that the inclusion in the model may not be as large, or may not be very accurately modelled. The difference in the apparent location of the inclusion is probably due to the position of the inclusion changing when swapped from its 15mm to 17mm position (see Figure 2.6.8).

2.6.5.2 Clear-layer inclusion mean-time differences

Figure 2.6.23 shows the change in mean-time due to the inclusion when the clear-layer flask is in place, and the light must pass through the clear layer before reaching the inclusion. Simulated data and data for experiments 1 and 2 are shown (for intralipid and resin). Note that these data are (clear layer + inclusion)-(clear layer without inclusion) rather than the data relative to the homogenous case.

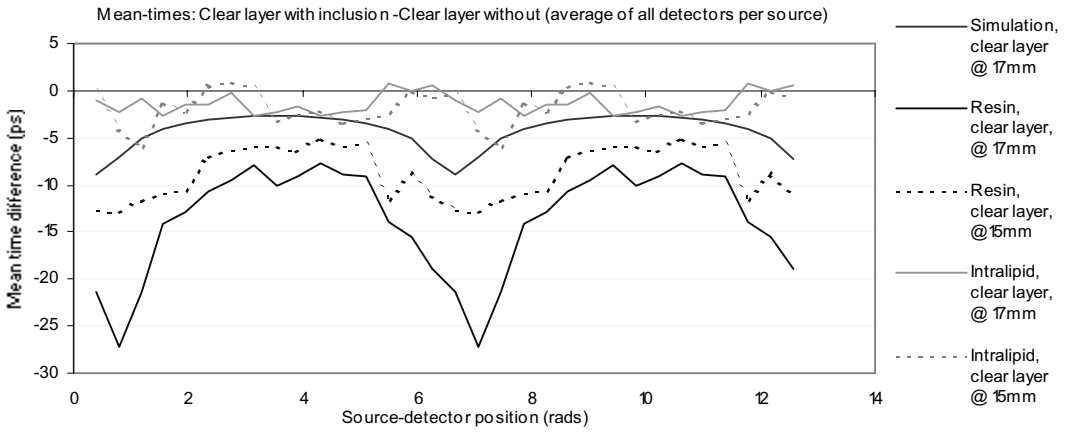

Figure 2.6.23 Mean-time difference averages for all detectors per source for the clear-layer (case 1) simulation (with and without the large inclusion at 17mm from the axis) and the clear-layer experiment 2 (resin) data (with and without the large inclusion at both 17 and 15mm axial displacement), and the clear-layer experiment 1 (intralipid) data (with and without the large inclusion at both 17 and 15mm axial displacement). The inclusion approximately in the same place as in Figure 2.6.23.

The signal due to the inclusion for the resin, 17mm position, clear-layer case is very large, almost as large as the signal for the homogenous case shown in Figure 2.6.22. It is also possible to see the signal from the inclusion at 15mm radial distance for the resin experiment. Again, as with the homogenous case the simulated value is smaller than expected. Note however that the data for the 1st experiment, using intralipid, shows no discernible signal. This supports the refractive-index related deductions described above, that very little light penetrates into the central diffusing regions in the case where the void has a higher refractive index than the surrounding material.

2.6.5.3 Homogenous (no void) inclusion intensity ratios

Figure 2.6.24 shows averaged intensity ratio data for the homogenous case, for the resin experiment and from the simulation. These experimental intensity ratios are more noisy than the mean-time difference plots (note that detector 5 has been omitted completely since it exhibited fluctuations between the homogenous and inclusion data acquisitions).

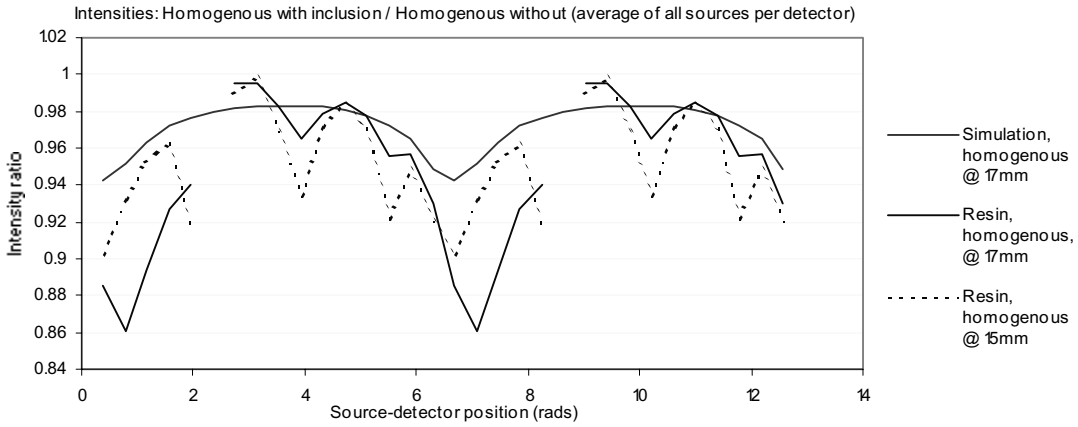

Figure 2.6.24 Intensity ratio averages for all sources per detector for the homogenous (case 3) simulation (with and without the large inclusion at 17mm from the axis) and the homogenous experiment 2 (resin) data (with and without the large inclusion at both 17 and 15mm axial displacement), the inclusion is at ~ 0 rad and 2π rad

Despite the noise, a signal corresponding to the inclusion is visible for both the 17mm and 15mm cases. (the slightly different position of the inclusion related peak corresponds to the slight differences in the positioning of the inclusion during the experiment).

2.6.5.4 Clear-layer inclusion intensity ratios

Figure 2.6.25 shows the averaged intensity ratio plots for the clear-layer case, with and without the inclusion, for experiments 1 and 2 and the simulation. (note that missing points correspond to the faulty detector 13). Note also the these data represent (clear layer + inclusion)/(clear layer without inclusion) rather than being relative to the homogenous case.

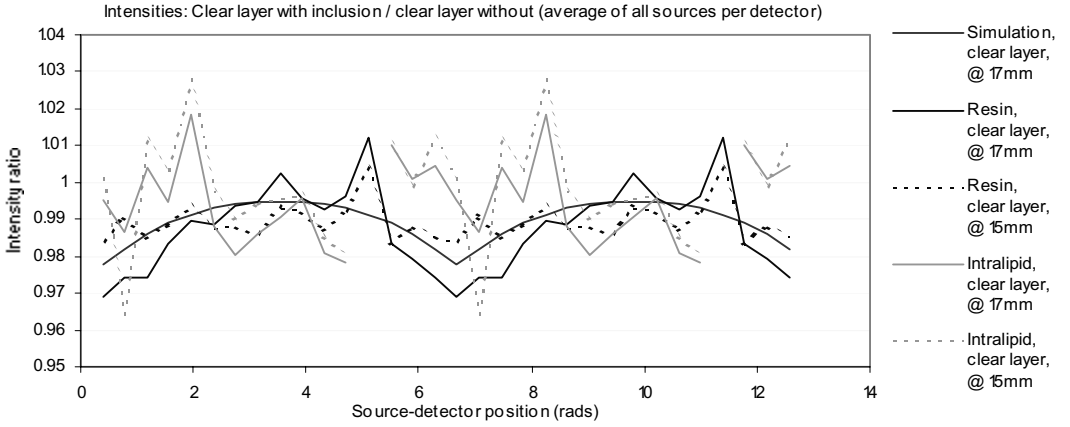

Figure 2.6.25 Intensity ratio averages for all sources per detector for the clear-layer (case 1) simulation (with and without the large inclusion at 17mm from the axis) and the clear-layer experiment 2 (resin) data (with and without the large inclusion at both 17 and 15mm axial displacement), and the clear-layer experiment 1 (intralipid) data (with and without the large inclusion at 17 and 15mm axial displacement).The inclusion is at ~ 0 rad and 2π rad.

The intensity ratio averaged data for the clear-layer is very noisy, but again, we can see a small signal that seems to correspond to the inclusion in the resin data, more clearly at 17mm but slightly at 15mm. Again, as with the mean-time results the experiment 1 intralipid data is very noisy and doesn't appear to correspond to the inclusion.

The purpose of this analysis was to determine whether there is any information in the measured mean-times and intensities that relates to the presence of an inclusion. The data were processed quite significantly (particularly the intensity data) to remove spurious data points so the signal-to-noise of the raw data is worse than that shown. However 'bad' data points can often be attributed to particular sources or detectors, allowing us to attempt image reconstruction.

2.6.6 Image reconstruction of inclusion data

While the data are quite noisy, we can subjectively remove 'bad' detectors, and correct for intensity shifts due to the fibre switch. This allows us to reconstruct the inclusion data described above to evaluate the effects of the clear layer on the ability to detect a perturbations beyond the non-scattering layer. We can also compare the experimental results with those corresponding to the simulation to validate the model.

The 3D void-TOAST model was not used to reconstruct experimental data since implementation of the techniques required was not complete. However we *can* validate results published in (Dehghani *et al*, 2000b) and (Dehghani *et al*, 1999b) which relate to the efficacy of diffusion-model-only reconstructions of simulated void data. It must be remembered however that the smooth spherical geometry of the void is a very limited case, and that certain properties of images reconstructions may be worsened or improved as a result of the geometry. Much more work will be required to extend these results to a level where conclusive comment can be made on the effect of CSF on optical brain monitoring.

In addition, the reconstructions presented were performed on a 2D mesh for speed and simplicity, despite the data being necessarily 3D. The implications of using a 2D reconstruction are described at length in chapter 2.3. In general however, difference data is relatively unaffected by a 2D:3D mismatch as long as the object being imaged is simple and there is no significant out-of-plane structure. The effects of the presence of the 3D void are not known. When absolute data are reconstructed without a 2D:3D correction, there is likely to be significant distortion.

2.6.6.1 (Homogenous + inclusion) – (homogenous without inclusion)

Figure 2.6.26 shows μ_a reconstructions of the simulated (left) and measured (resin experiment, right) intensity ratio (top) and mean-time (bottom) between the homogenous phantom with and without the large inclusion at 17mm. Figure 2.6.27 shows the corresponding μ'_s images. The data used in these reconstructions are equivalent to those shown in Figure 2.6.24 and Figure 2.6.22. Contributions from noisy detectors 4, 6 and 13 were removed from the 1 to 16 detector set prior to image reconstruction. The images are all the 4th iterations.

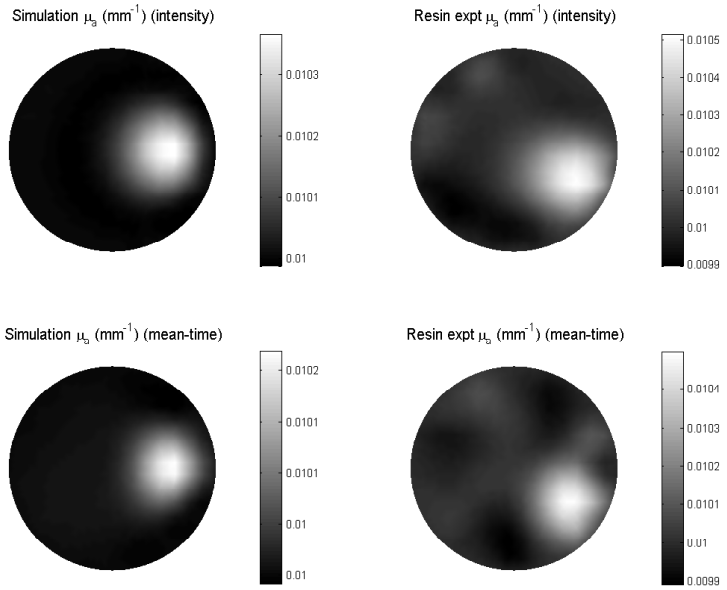

Figure 2.6.26 ***Absorption****: (top) Intensity difference images (bottom) mean-time difference images of the* ***homogenous*** *phantom with and without the large inclusion at ~17mm (Left) shows simulated data (right) shows measured (resin) data all iterations 4.*

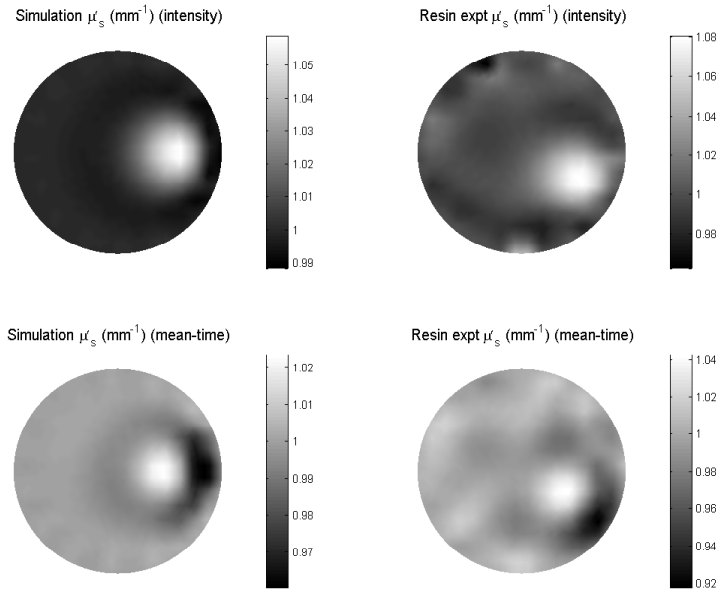

Figure 2.6.27 ***Scatter*** *(μ'_s): (top) Intensity difference images (bottom) mean-time difference images of the* ***homogenous*** *phantom with and without the large inclusion at ~17mm (Left) shows simulated data (right) shows measured (resin) data all iterations 4.*

The results for the measured and simulated data agree well, although a diffusion-only model was used to simulate the forward data in this case, so the quality of the results is likely to be good. Quantitatively the μ_a peak-values values are low due to blurring and the low iteration chosen, however they are higher for the reconstructions of the measured than for the simulated data. This supports the suggestion above that the simulated properties of the inclusion are smaller (or it is physically smaller) than those of the inclusion used in the experiment, as was implied by the magnitude of the perturbation signal in Figure 2.6.24 and Figure 2.6.22. Although the inclusions should be the same scatter as the background resin, a dual parameter reconstruction was performed, and a small scatter signal can be seen in the μ'_s images. Again the agreement between the simulation and the experiment is qualitatively good. The slight shift in the position of the inclusion implied by the raw data agrees with the position of the inclusion in the experimental data images.

2.6.6.1.1 Image summary (homogenous difference images)

Phantom	Case 3 (homogenous) void phantom
Mesh	Circular 1201 nodes 2304 linear elements (2D)
Basis	Pixel 16 × 16
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iterations	4
Sources	16
Detectors per source	11
Datatypes	Intensity (top) raw mean (bottom)
Calibration	Difference (to homogenous)
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	10 secs
Wavelength	780 nm

Table 2.6.1 Properties of images shown in Figure 2.6.26 and Figure 2.6.27

2.6.6.2 Absolute (and baseline) images of void data with the inclusion

We have a number of choices of how to reconstruct the void data. As introduced in preceding chapters we may use:

- Absolute data
- The difference between the object of interest and a homogenous phantom
- The difference between data acquired on the same object in two different states.

Figure 2.6.28 shows μ_a (top) and μ'_s (bottom) images reconstructed from (far left) the simulated inclusion (large @ 17mm) absolute mean-time data (with a clear-layer), (left) the absolute mean-time reconstructions for the resin (case 1) with clear layer and inclusion data. The images on the right are 'baseline' images where the difference between the mean-time (clear layer + inclusion) – (homogenous case 3) is reconstructed from simulated and (far right) resin measured data. A 2D:3D mismatch will affect the images on the left, since they utilise absolute data (see 2.3.1.2). The 2D:3D mismatch should cancel to some extent for the

images on the right, although the clear layer will evidently distort the way that the photons propagate out-of-plane. The absolute mean-time data were calibrated using a monstode calibration measurement set acquired during the resin experiment (see section 2.1.4.2).

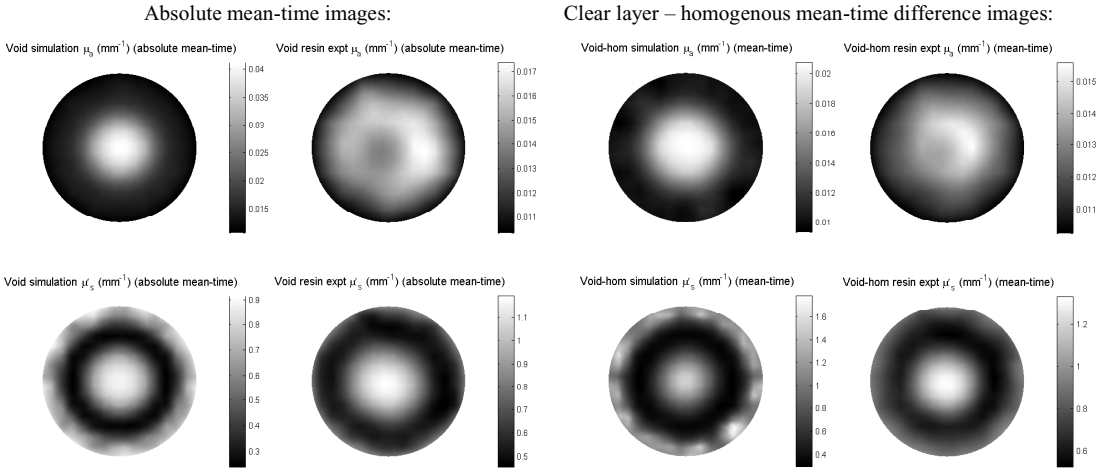

Figure 2.6.28 (far left) Absolute mean-time reconstructions of simulated and (left) data acquired in the resin experiment in the presence of the clear-layer (with the large inclusion at 17mm) (right) mean-time difference reconstructions of clear-layer data – homogenous (case 3) data for simulation and (far right) measured resin data. (All are iterations 6, μ_a top, μ_s' bottom)

The simulated results agree with the findings of (Dehghani *et al*, 2000c), although a dual parameter reconstruction has been used here rather than simply solving for absorption. There is little evidence of the inclusion in the absolute, or baseline (difference to homogenous) images shown in Figure 2.6.28. The structure seen is likely to relate to the presence of the void since TOAST has tried to find a distribution that would yield data of the shape shown in earlier data plots for clear-layer data (Figure 2.6.10). Without doing a 3D reconstruction it is not possible to categorically eliminate the possibility that the 2D:3D related artefact has overwhelmed the reconstruction. While the contrast of the central bright regions is very high compared to the images in Figure 2.6.26 (although these images are the 6th rather than the 4th iteration), it is likely that the majority of the contrast is due to the way that TOAST has tried to compensate for the observed data. In fact if we reconstruct the simulated data without the inclusion present in exactly the same way (not shown), we see an almost identical region in the centre of the resulting absorption and scatter images.

Images reconstructed from measured data indicate the presence of the inclusion, particularly in the absorption images where the central bright region is shifted to the right, closer to the location of the inclusion (see Figure 2.6.28). The contrast is high but is again likely to be due to the distortion of the data by the void rather than the presence of the inclusion .

Clearly images of this quality would not be suitable for subsequent calculation of saturation and blood volume. Further, the data has implied the presence of an inclusion in the

centre, which appears even in images where there was not inclusion. The void-TOAST algorithm (along with 3D modelling) aims ultimately to be able to reconstruct images from data such as these. The model would need prior knowledge of the boundaries of the clear region as defined in [2.6.2], (although techniques to retrieve these from optical data are being developed (Dorn, 2001), (Kolehmainen *et al*, 2000)). This would mean that absolute data could be acquired on the head, and the effects of the CSF could be eliminated by its inclusion into the forward model.

However, we can also try reconstructing the difference between data acquired with and without the inclusion present, while assuming that the initial estimate for the diffusion reconstruction is homogenous.

2.6.6.2.1 Image summary (absolute and baseline clear-layer images)

Phantom	Case 1 (clear layer) void phantom
Mesh	Circular 1201 nodes 2304 linear elements (2D)
Basis	Pixel 16 × 16
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iterations	6
Sources	16
Detectors per source	11
Datatypes	Mean / raw mean
Calibration	Monstode calibrated / difference to homogenous
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	10 secs
Wavelength	780 nm

Table 2.6.2 Properties of images shown in Figure 2.6.28

2.6.6.3 (Clear layer + inclusion) – (clear layer without inclusion)

The raw intensity ratios and raw mean-time differences between data acquired with and without the inclusion within the clear layer were used to reconstruct the images shown in Figure 2.6.29 and Figure 2.6.30. Simulated data, the resin experiment data and the intralipid experiment data were all used to explore whether the data shown in Figure 2.6.23 and Figure 2.6.25 could yield images showing the presence of the inclusion. The initial conditions for the reconstruction were a homogenous (2D) region with μ_a = 0.01mm^{-1} and μ'_s= 1mm^{-1} to correspond to the nominal properties of the rein and intralipid. The 4^{th} iterations are shown.

Data for these reconstructions were processed to remove detector 13 only, although the intensity data were manipulated to reduce the effects of the variable delay in the fibre-switch opening as described above in 2.6.3.3 and in 2.1.1.3.

The presence of the inclusion has been clearly identified in the simulation and resin experiment difference-images shown in Figure 2.6.29 and Figure 2.6.30. These are a dramatic improvement to the absolute and baseline imaging results shown in Figure 2.6.28. The intralipid experiment has not yielded results that indicate the presence of an inclusion.

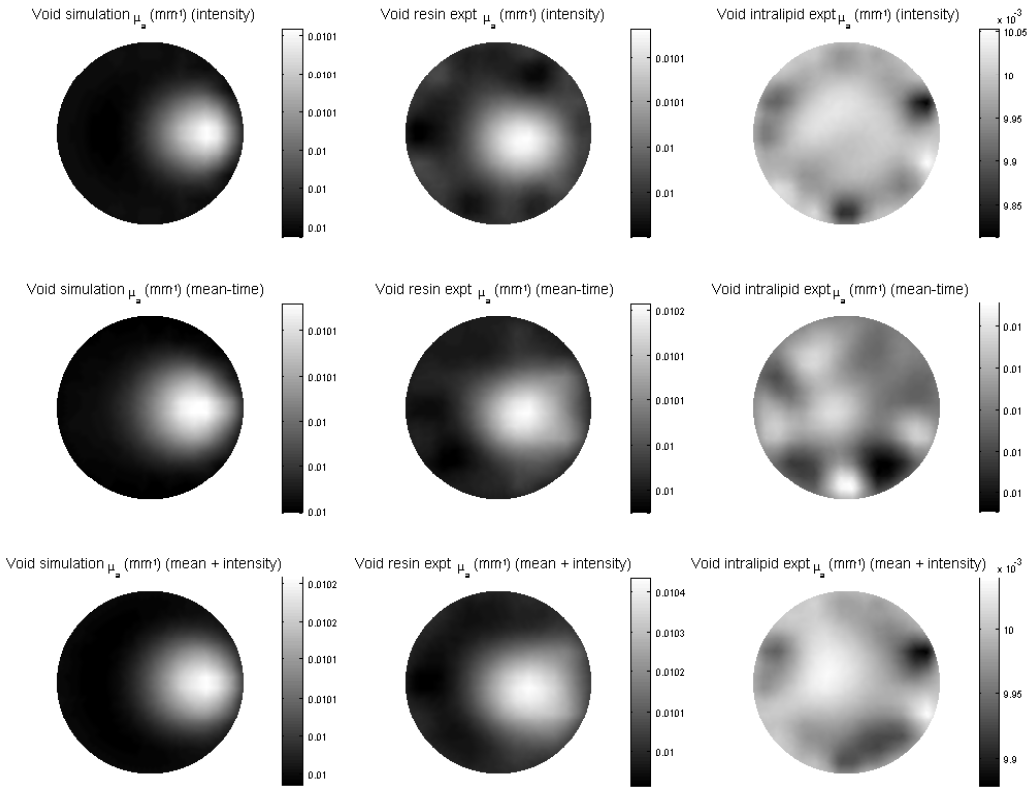

Figure 2.6.29 ***Absorption****: Intensity (top) mean-time (middle) and intensity and mean-time (bottom) difference images of data with the inclusion within the* ***clear-layer*** *and without for simulated (left), resin experiment (middle) and intralipid experiment (right).*

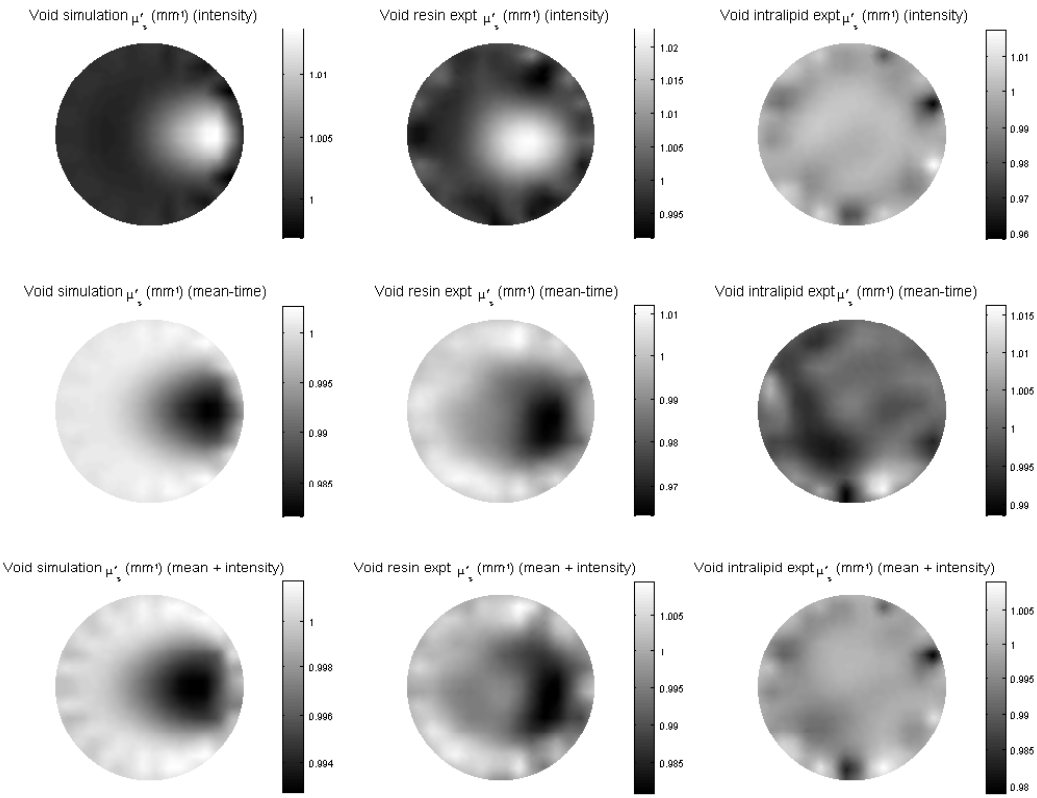

Figure 2.6.30 ***Scatter*** *(μ′s): Intensity (top) mean-time (middle) and intensity and mean-time (bottom) difference images of data with the inclusion within the* ***clear-layer*** *and without for simulated (left), resin experiment (middle) and intralipid experiment (right).*

An interesting effect is seen in the mean-time scatter results. The homogenous difference images shown in Figure 2.6.26 identify the inclusion as absorbing, although slight cross talk was seen into the μ'_s image, with a slight dipole pattern in the mean-time only images. In the clear-layer difference case, mean-time and mean and intensity images imply a region of low scatter at the position of the inclusion for both the simulated an measured data. The origin of this effect is not known, but should be explored in more detail. (Note that intensity reconstructions, cross-talk for high μ_a will imply high μ'_s as illustrated in appendix A (Figure 2.9.3).

Published results based on (2D) void models have mainly considered the challenge of reconstructing absolute data acquired on an object with a clear layer (starting from results like those in Figure 2.6.28), (Dehghani *et al*, 2000b), (Dehghani *et al*, 2000c). (Pei *et al*, 2001b) show images reconstructed using changes in data acquired on a cylindrical object with a clear cylindrical void layer with moving μ_a inclusions within the centre. They demonstrate that the presence of three moving objects can be identified but that their positions and quantitation are quite distorted. With higher iterations the simulated difference images become slightly distorted at the edge, although the resin experimental results localise the inclusion (correctly) further towards the centre.

The intralipid results did not identify the inclusion. This may have been due to the lack of penetration of light through the central diffusing region due to the lower refractive index of the intralipid compared to the glass clear layer, as hypothesised above from the raw data. This would suggest that signal relating to the perturbation may be very small. In addition, the refractive index of the intralipid relative the resin inclusion itself may have distorted the amount of light that could pass through the inclusion and therefore contain signal to be detected. A further possibility is that since data were acquired on a different day to those acquired on the resin phantom, the noise-properties of the MONSTIR system, or some other aspect of system performance may have affected results. (Note though that the standard deviations on the homogenous mean-time data are no larger than for the resin data, see Figure 2.6.14 and Figure 2.6.16)

Another interesting observation is that while the averaged data plotted in Figure 2.6.25 seemed to imply that signal-to-noise on the perturbation signal was worse for intensity measurements, the results above have demonstrated that equally good results can be obtained from the intensity measurements as from the mean-times.

2.6.6.3.1 Image summary (homogenous difference images)

Phantom	Case 1 (clear layer) void phantom
Mesh	Circular 1201 nodes 2304 linear elements (2D)
Basis	Pixel 16 × 16
Starting parameters: μ_a	0.01 mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Iterations	4
Sources	16
Detectors per source	11
Datatypes	Raw intensity (top) raw mean (bottom)
Calibration	Difference to Clear-layer without inclusion
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	10 secs
Wavelength	780 nm

Table 2.6.3 Properties of images shown in Figure 2.6.29 and Figure 2.6.30

2.6.7 Summary and discussion

This chapter explored important aspects of light propagation through non-scattering regions, with reference to the effect of CSF in the neonatal brain. An experiment was described that was designed to test the validity of a recently developed 3D hybrid radiosity-diffusion model for simulating objects containing void regions. Three cases were explored: a spherical clear layer, a spherical clear region and a simple homogenous case without non-scattering regions. Two experiments were performed, the first had a matching fluid (intralipid) with a low refractive in contact with the (high refractive index) void region, the second had a matching fluid (resin) with a higher refractive index than the void. Data were also acquired when an absorbing inclusion was placed within the central part of the phantom with and without the clear layer present. A preliminary comparison of (non-inclusion) simulated data and data acquired in the resin experiment is presented in (Riley *et al*, 2001).

Averaged mean-time and intensity data for the non-inclusion case were examined and the differences in the simulation and the two experiments were found to be well explained by the differences in the refractive indices of the components. It was hypothesised that where $n_{void} > n_{tissue}$ (the intralipid experiment), light is rapidly transported to the detectors via the clear layer and fails to probe the diffusing region in the centre of the void. When where $n_{void} < n_{tissue}$ (the resin experiment), light is transmitted in the void layer less efficiently and effectively probes the diffusing regions better than the $n_{void} > n_{tissue}$ case. Analysis also implied that void regions (e.g. ventricles) would affect data in a way that could potentially introduce image errors. Unless the location and properties of the clear regions were known a-priori (allowing a hybrid model to be used), a reconstruction would attempt to account for the anomalous transport of light across such regions with a combination of erroneous absorption and scatter. Investigation of such clear *regions* should continue, to evaluate whether it is possible to get good sensitivity to regions of tissue around the ventricles which are of particular diagnostic interest (see section 1.1.3.1).

Analysis was then presented of the data acquired with and without an absorbing inclusion present. Signal corresponding to the presence of the inclusion was discernable in simulated data and data from the resin experiment for both the homogenous and the clear layer cases. For the intralipid experiment it was not possible to identify features corresponding to the inclusion. The inclusion data were then reconstructed to verify the results of the raw data analysis.

Homogenous (no clear layer) simulated and resin experiment data were reconstructed using *difference imaging* (relative to the state without the inclusion). These identified and localised the inclusion. Absolute images were then reconstructed using clear-layer (resin and simulated) absolute mean-time data (with no 2D:3D correction). Baseline data (clear-layer + inclusion relative to the homogeneous state) were also reconstructed. The resulting images were badly distorted, with no evidence of the inclusion for the simulated data, and only a slight indication of its presence for the resin data.

Difference images were then reconstructed from the *change* in data with the clear-layer in place, with and without the inclusion present for simulated, resin experiment *and* intralipid experiment data. The μ_a results for the simulated and resin experiment data localised the inclusion in approximately the correct position, but with lower quantitation to that achieved without the 2mm void-layer present. It was not possible to identify the inclusion in the images from the intralipid data. μ'_s difference images were also reconstructed, which should have been uniform (since the inclusion was nominally the same μ'_s as the background), however images implied a decrease in μ'_s at the site of the inclusion for mean-time reconstructions, whereas an *increase* was seen in the homogenous (no clear layer) results.

In general the predictions of the 3D void-TOAST model have been shown to agree well with measured data. Most of the differences seen between the model and the experimental data can be attributed to inconsistencies between the model parameters and those of the measured phantoms. In particular the effect of refractive index mismatches on measured data is not currently incorporated into the model.

The experimental study of the effects of refractive index mismatches between a void layer and the surrounding tissue has revealed an important consideration for those analysing the effects of CSF on the penetration of NIR into the brain. The refractive index of brain tissue and the subarachnoid membrane will invariably be higher than that of the CSF. This is the case represented by the resin experiment. The preliminary results of this study implied that penetration of light into the brain will be better than indicated by a simulation that cannot account for refractive index mismatch (and hence specular reflection). To confirm this suggestion, accurate measurements should be made to quantify the mismatchs between the

brain tissue, membranes and CSF. However the availability of recently acquired neonatal brain data (see 2.7.2.3) presents the opportunity to examine the mean-time profile of light emerging from the head to give a very preliminary insight into which of the situations ($n_{void}>n_{tissue}$, $n_{void} = n_{tissue}$ and $n_{void} <n_{tissue}$) correspond best to the baby brain.

Figure 2.6.31 shows the clear-layer data for the simulation, experiment 1 and 2 and the homogenous data from the simulation plotted v / s source-detector separation. The shape of the resulting line at longer distances appears to characterise the case where refractive index of the clear layer is lower, the same or higher than that of the surrounding tissue.

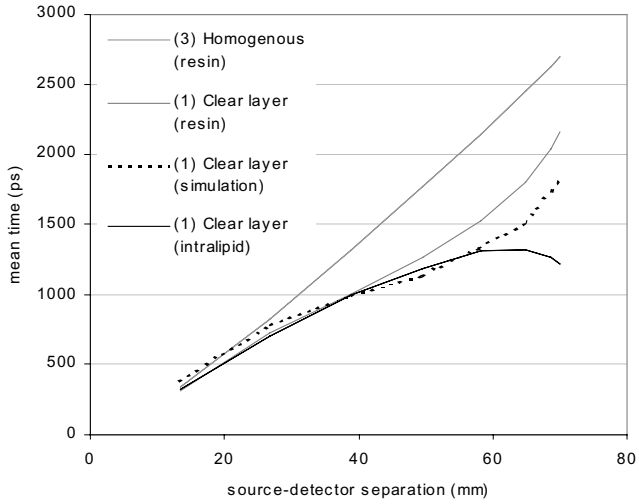

Figure 2.6.31 Plot of mean-time profiles for each clear-layer experiment and the simulation, as well as the homogenous case, plotted v/s source-detector separation

Figure 2.6.32 shows a plot of the calibrated mean-times measured on the head of the second neonate studies. These data were extracted from the full data set as their source-detector positions represented an approximate ring around the head. Data are shown for each wavelength (780nm and 815nm) with the mean and standard deviation of the mean-time for that given source-detector spacing (± 1mm).

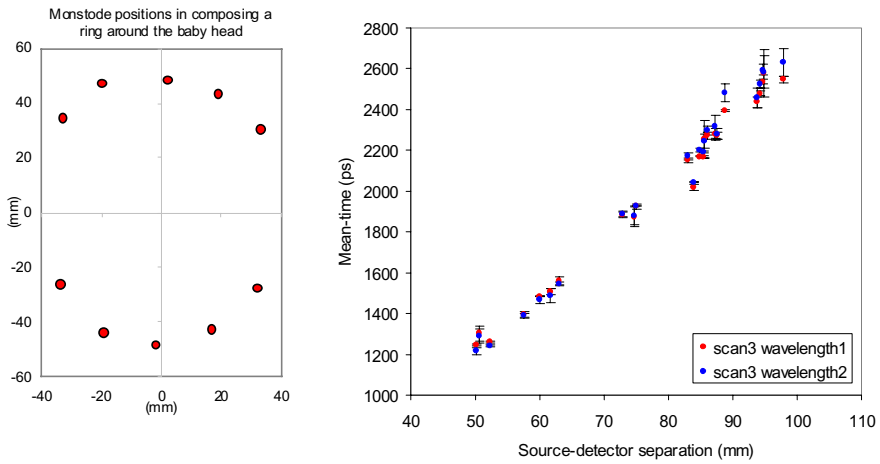

Figure 2.6.32 (right) a plot of calibrated mean-time values extracted from measurements made on the second neonate for source-detector spacings corresponding to the positions shown on the left which represent a plane of fibres positioned around the head (see 2.7.2.3)

The plot of the neonatal data does not reveal any features that imply light is taking a dramatic short cut across the head via the CSF. However the standard deviations are larger at larger distances, so it is difficult to produce a fit that will allow quantitative comparison with the plots in Figure 2.6.31. This method could be extended by looking at the variation of mean-time with distance over all the source-detector arrangements that comprise a ring.

Image reconstruction of the simulated and resin experiment difference data (with and without an inclusion in the presence of a clear layer) was successful in identifying and (roughly) localising the inclusion. This was to be expected from examination of the raw data where a signal from the inclusion could be clearly identified. However it is suspected that the quality of these images is a result of the geometry chosen. In direct analogy to the problem identified in chapter 2.4, we are acquiring difference data on an object with unknown heterogeneous structure. The structure will affect the sensitivity of the measurement compared to that considered in a reconstruction assuming a homogenous initial state. Where the object is rotationally symmetric (i.e. concentric rings, as in case above, and also in (Pei *et al*, 2001b)), the PMDFs (of the measurement) will be the same for any given source-detector spacing (perhaps all probing only a shallow depth of the inner diffusing region). Despite this a perturbation positioned off axis, will almost inevitably perturb data acquired closer to it. Even though the PMDFs in the reconstruction are not accurate they will identify the approximate location of the inclusion. Consider however the case where the void layer is un-even, or the object is oval and not circular (as in Figure 2.6.5). The regions probed by the photons will be different for each source-detector combination. Using PMDFs based on a diffusing-only object in this case is likely to fail to locate and localise changes accurately. More studies are required to determine the level of accuracy achievable using difference imaging in the presence of irregular non-scattering regions.

The future development of methods to overcome the effects of non-scattering regions may include:

- Further development of algorithms like void-TOAST to attempt reconstruction of absolute data by having a-priori information about the location of the voids, by using region retrieval methods to determine the location of the void regions before proper reconstruction (Kolehmainen *et al*, 2000), (Dorn, 2001), or by developing code based on the full RTE.
- Determining the limitations of differential imaging in the presence of the void, while also considering the inherent limitations of in-situ difference measurements as described in chapter 2.4.

- Accurate measurements of the refractive index of CSF and the surrounding tissues are required to determine whether the effects of refractive index mismatch needs to be accounted for in future void-models (Ripoll *et al*, 2000b).
- Development of an accurate 3D neonatal head model, with voids included, to generate forward data from which optimal methods of reconstructing known lesions (see Table 1.1.2) can be developed.
- It is suggested by (Dorn, 2000) that the early arriving light is that which has passed through the void, and that later light has probed diffusing regions. An overall mean-time or intensity signal gets overwhelmed by the large numbers of photons that propagate through the void. However it is possible that the TPSFs acquired in these experiments could be investigated to explore whether more information about the perturbation within a clear layer can be extracted by examining only the tails of TPSFs.

2.7 Clinical data acquisition

Issues relating to calibration, data processing and image reconstruction for the practical application of optical tomography have so far been explored. However a major challenge in implementing this new technique clinically is the form of the interface between the instrument and the subject. Section 1.1.3 described how optical tomography could provide valuable information about the neonatal brain and the adult breast. This section describes the development of specialist fibre holders for the breast and brain, a design challenge that requires consideration of calibration methods, mesh generation and general protocols for data acquisition on a clinical subject.

This chapter begins by detailing the original methods of holding fibres on cylindrical phantoms and subsequently on the arm. Describing development of a system for imaging the neonatal head we begin by detailing the main specifications. The progression of the design of a fibre holder, including development of the monstode, is then described followed by an account of the manufacture and ultimate use of a fibre holding helmet on two pre-term neonates. The effectiveness of the design is then evaluated and future improvements discussed. The experimental protocol is also evaluated and its future development considered.

Similarly the development of techniques for breast imaging is detailed by listing specifications and documenting two possible fibre holder designs. The protocol used to acquire data using the first of these two designs on two adult volunteers is then described. Potential improvements to the current methods are then discussed.

2.7.1 Early phantom studies and the arm

The very first set of MONSTIR imaging data was acquired on a cylindrical phantom. 2 source fibres and 8 detector bundles were held in a planar 35mm radius black Delrin ring which fitted tightly around the phantom. Each fibre was pushed into a radial hole in the ring, and a grub-screw tightened to hold it in place. The ring was twisted relative to the phantom to acquire data equivalent to 32 sources and 16 detectors (Hebden *et al*, 1999). Once the MONSTIR system was upgraded to 32 functioning source and detector channels, all 64 fibres were held in the planar ring and a full set of data could be acquired without needing to move the arrangement. Section 2.1.2.1 described the early calibration tools designed to also fit inside the 35mm radius planar ring. It was found that the two subjects could also squeeze the

35mm radius ring over the hand and up the arm. Thus the same fibre holder ring was also utilised for arm imaging measurements (Hillman *et al*, 2001b) as shown in Figure 2.7.1.

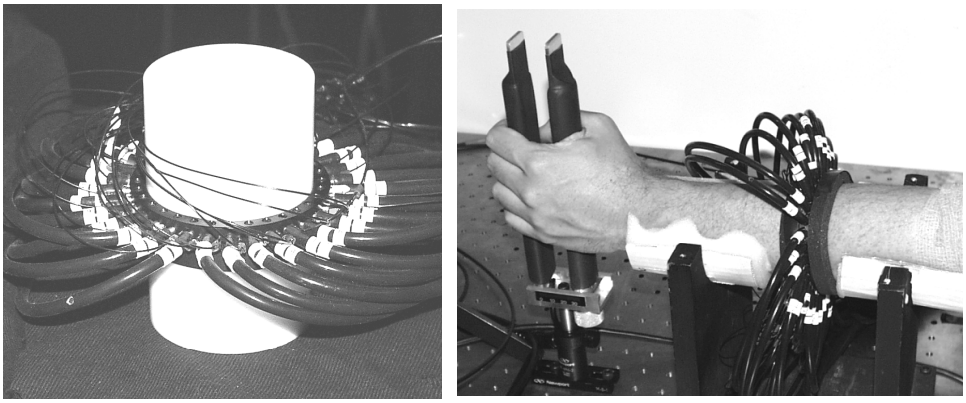

Figure 2.7.1 The first fibre holder was a black Delrin ring 35mm in radius that held 32 interspersed source and detector fibres in a single plane. It was used to acquire data on a number of cylindrical phantoms (left) and subsequently on the forearm by slipping it over the wrist (right).

Early 3D data acquisitions were performed by either translating the ring up and down a cylindrical phantom, or by using more than one ring at different heights on the cylinder (Schmidt *et al*, 2000b), (Arridge *et al*, 2000b).

As detailed in chapter 2.3, data acquired using these rings were reconstructed on simple 2D circular meshes of 70mm diameter, or later on cylindrical meshes. The fibre holder ring was assumed to have been machine accurately enough that the positions of the sources and detectors could be calculated simply by assuming they were equally spaced.

However, to develop beyond imaging objects with a 70mm diameter using a planar data acquisition, new methods of holding the fibres were required.

2.7.2 Acquiring data on a baby

From the inception of multi-channel optical tomography, practical design of holders for the multiple optical fibres was going to be a major challenge. This section introduces the issues of acquiring data on premature and term infants.

2.7.2.1 Specifications of method

The general requirements relating to positioning of fibres are:

- Sources and detectors should be distributed in a three dimensional arrangement, with each in contact with the head.
- The fibres are heavy, but must be held securely and must not move during data acquisition. The infant's head should not move during data acquisition, but it must not be restrained in a way that may cause harm.

- The position of each source and detector has to be well known to allow modelling of source-detector positions on a corresponding FEM mesh.
- Ideally the method of building the fibre holder, or measuring the positions of optodes would incorporate sufficient measurements to allow FEM mesh generation.
- The method must be flexible to allow for the fact that the shape and size of the head of each infant will be varied. Appendix A (Figure 2.9.1) shows that head diameter may vary between ~7.5 and 12cm for premature and term babies.
- It must be possible to make some form of calibration measurement within the imaging session, either a reference measurement or an absolute measurement (see section 2.1.2), preferably in situ.

Environmental factors impose certain limitations:

- § It may be necessary to exclude external illumination from room lights, photo-therapy lights (UV) or heat lamps (IR).
- § Adequate access must be available for emergency nursing care and the fibre holder should be rapidly removable if required.
- § Infants that are ventilated or on continuous positive airway pressure (CPAP) need to have their heads in a particular position (according to the type of ventilator). Also tubing for ventilation is often secured to a bonnet on the head.
- § The infant may be in an incubator or a heat-pad. The fibre holder should not interfere with temperature regulation by making the head either too hot or too cold.

Safety aspects of the fibre holder:

- § The skull is readily deformable, so pressure could damage underlying cortex. Point pressure could damage delicate skin. The ends of the fibres must therefore not protrude and apply uneven pressure to the neonate's head.
- § The holder must be such that there is no possibility of light levels beyond those safe for eye exposure could get near to the infant's eyes.
- § Skin exposure of NIR laser light must be within safe limits.
- § Neonates are kept in a very clean environment. A fibre holder that is going to be in physical contact with the infant will need to have been sterilised.

General design constraints:

- § Light leakage between active sources and detectors must be excluded.
- § The optical fibres are permanently attached to MONSTIR and so must be detachable from the fibre holder.
- Source fibres are thin and delicate, and can be broken / damaged if bent or crushed.

§ New born infants and those with sudden changes in condition would require suitable fibre holders to be available at short notice (perhaps within hours). Also the shape of the head may vary over a number of days, so manufacture must be simple and fast.

2.7.2.2 Development of the neonatal head fibre-holder

The problem of attaching optical fibres to the head for NIR topography and tomography has received relatively little attention. For single-site (global) measurements rubberised black patches that hold 2 or 3 fibres, held in place with bandages, are generally used. More rigid plates holding arrays of fibres are used in topography as shown in Figure 1.2.4 in section 1.2.1.2 (Yamashita *et al*, 1996). The only published optical tomography image of the brain of a neonate used fibres held in a simple strip which wrapped around the head, such that fibres were in a planar ring (Hintz *et al*, 1998).

A review of the background research conducted to determine suitable methods for neonatal positioning and immobilisation, surface contour mapping and considerations for laser safety is given in appendix A (2.9.9).

The optimal distribution of fibres for optical tomography of the neonatal head has yet to be devised. Studies relating to positioning optodes for transmission geometry optical breast imaging have been published (Culver *et al*, 2001). However these results do not account for the very different angular sampling achievable with true tomographic data acquisition (see Figure 1.2.5 in section 1.2.1.3). If the head were a sphere, an even arrangement of fibres would be a good way to sample the volume, however fibres cannot be placed on the face (and light would not travel in a predictable way through its air-filled spaces). A reasonable assumption is that fibres should be distributed over as much of the head as possible, allowing light to penetrate through as much of the brain, in as many different directions, as possible.

Early designs for fibre holders assumed that fibres could be positioned in a single plane. Transducers to measure the translational distance of each source or detector from a supporting ring were proposed as a reasonable way to deduce source-detector positions. However as TOAST developed into a 3D algorithm it became clear that a 3D arrangement of fibres would be required. A number of experiments involving attempts to hold fibres in multiple bands around the head demonstrated that the major problem was the weight of the fibre bundles. Without a strong anchor, fibres would not stay in their positions. The idea of creating a helmet was conceived as a method to provide a rigid frame to support the weight of the fibres.

Helmets could be constructed using thermoplastic materials more usually utilised for radiotherapy positioning as described in appendix A (Figure 2.9.18). Prototypes were moulded directly on the head of a plastic doll (see Figure 2.7.2). The rigid shell created from thermoplastic could be lined with soft foam (see appendix A (2.9.9.4)) to prevent light

leakage between sources and detectors. However the helmet was only 2-3 mm thick, making affixing fibre bundles using grub-screws (as had been done previously with the Delrin ring in Figure 2.7.1) impossible. However the short length of the brass ferule on the end of each fibre bundle meant that it could not protrude all the way through the helmet and the foam to touch the head. The solution to this problem was the monstode. Early prototype helmets were designed to fit the heads of plastic dolls and a similarly shaped epoxy resin phantom, see Figure 2.7.2.

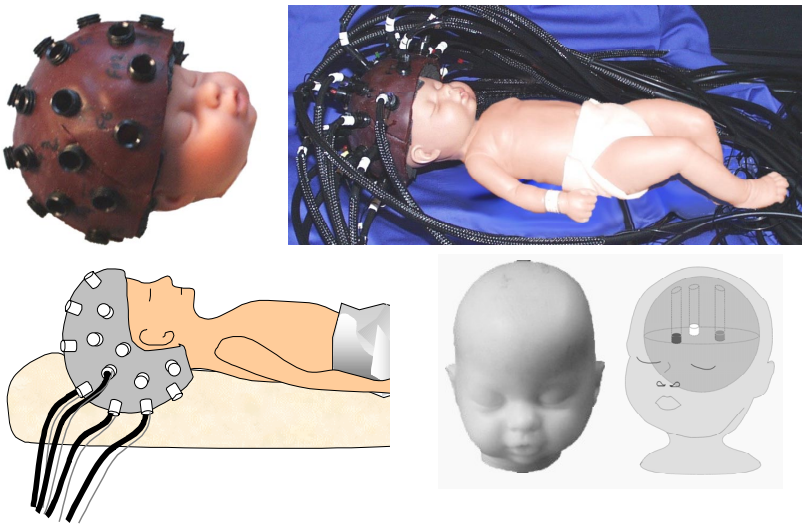

Figure 2.7.2 Early thermo-plastic helmets used to hold optical fibres on dolls and the 'baby head epoxy phantom' (bottom right)

2.7.2.2.1 Monstodes

The monstode was developed to address a number of practical problems. Initially the idea of creating a unit to hold both a source and detector was appealing since it meant that fewer positions on the head needed to be recorded, and the unit itself could be affixed to the shell of a helmet rather than the fibres themselves.

Then the idea of illuminating the tissue from a distance was conceived. This has two advantages. First, that a larger spot-size of the illumination will be less sensitive to small movements (e.g. if a hair is under the source) and secondly, by distributing the illumination over a larger area, more photons can enter the tissue without exceeding the W mm^{-2} laser safety limits (even for eye exposure). If the source fibre is held at a fixed distance from the head, the coupled detector bundle can't touch the tissue without obstructing the source illumination. While having the detector bundle also at a distance will reduce the amount of light detected slightly, it has the advantage that the measurement should also be less sensitive to small features on the surface such as hair.

The monstode was designed as a unit to hold the source and detector at a distance from the tissue, but also to allow the fibres to be rigidly held in place by a helmet made from

thermoplastic. Special three-part modules were designed as illustrated in Figure 2.7.3. Part (C) fixes into the helmet via friction, or if necessary, glue. Part (A) holds a source fibre and detector bundle pair using grub screws and is intended to remain on the fibres from measurement to measurement. As part (A) passes through part (C), Part (B) is intended as a locking device that allows the vertical position of part (A) to be varied and then fixed. This provides flexibility to ensure that the end of part (A) just touches the head and does not depend on the accuracy of the match between the shape of the helmet and the head.

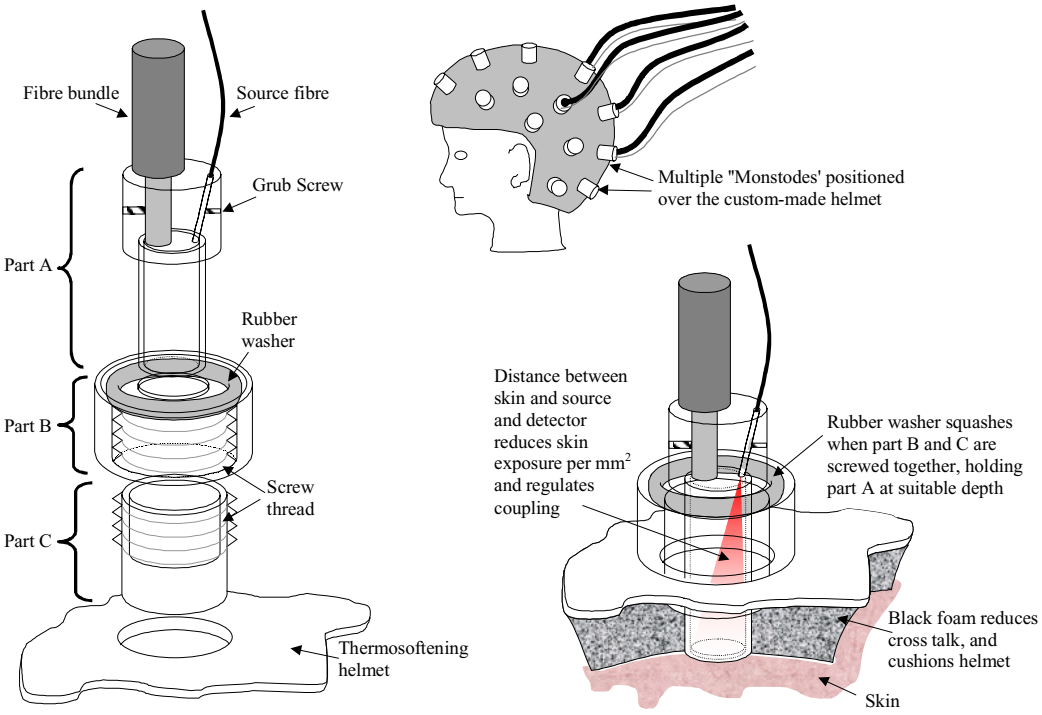

Figure 2.7.3 Monstode design

So, having manufactured a suitable helmet for an infant, and once the positions of the sources and detectors on the helmet have been decided, holes can be drilled or punched in the thermoplastic to allow insertion of part (C) (see Figure 2.7.3). The sources and detectors can then be positioned in this custom arrangement over the surface of the helmet, and readily removed if necessary.

2.7.2.2.2 In-situ monstode calibration

In addition to providing a much simpler mechanism for attaching and moving fibres, and allowing distributed illumination of tissue, the monstode also offers to opportunity to perform an in-situ calibration. As described in detail in section 2.1.2.2 it was found that source light could be attenuated such that if the source fibre were activated, the detector bundle within the same monstode could detect back-reflected and back-scattered light from the tissue. The times taken for light to travel from the end of the source to the tissue and for the light emerging from the tissue to reach the source would be the same as during data acquisition. Therefore the measurement would characterise the properties of the system IRF *and* coupling, thereby

providing a calibration measurement (Hebden *et al*, 2002c) (this technique is effective for mean-time calibration, but other datatypes may not be well accounted for, see section 2.1.2.2).

Disadvantages of the monstode include that fact that coupling sources and detectors together reduces number of sample points, which can degrade information available about regions close to the surface. Also, having the detector bundle at a distance reduces detection efficiency. Further, skin exposure per mm^2 will not be directly reduced as a function of the area illuminated since the profile of the light emerging from the current source fibres is Gaussian. This must be considered to ensure that the central peak illumination is within safety limits. Finally, in the current design, fibres emerge from the monstode at an angle perpendicular to the helmet. This prevents sources and detectors from being placed on weight-bearing parts of the helmet. Designs integrating prisms (Hintz *et al*, 1998), or a pillow with cut-outs for fibres (see Figure 2.7.2) or a device to support the helmet would be required.

2.7.2.2.3 Techniques for helmet manufacture

Following manufacture of helmets for the baby head phantom, and dolls (see Figure 2.7.2), various methods were explored to create helmets for infants attending out-patient appointments. While original helmets had been made by moulding the plastic directly onto the foam-covered doll's head, it was felt that this method would not be suitable for very tiny babies since the thermoplastic may be wet, unhygienic, hot, or uncomfortable.

Measurements made using a tile shaping template tool (appendix A (Figure 2.9.19)) over the three cross sections of the head (transverse, coronal and saggital) were found to be unreliable. Investigations into using stereoscopic photography or other methods of non-invasive 3D surface mapping (appendix A (2.9.9.3)) ensued, to allow the helmet to be manufactured as a precise fit. While a technique to map the head of a neonate for this application is being developed, it is not yet suitable for use.

As a result the design of the original 'helmet' was modified to be more flexible. Based on transverse, coronal and saggital digital photographs of the infant to be studied, the helmet was constructed in two parts; a base-plate, with an indentation based on the shape of the back of the head, and a T-shaped part which is lowered onto the forehead, joining the base-plate by the ears and the top of the cranium. Both pieces were made of thermoplastic, lined with foam, and had suitable holes to hold up to 32 monstodes (see Figure 2.7.3). While being adaptable, attempts were made to ensure that the T-shaped piece could be accurately re-positioned after data acquisition to allow the 3D co-ordinates of the monstodes to be recorded using a Microscribe 3D mapping arm (see appendix A (Figure 2.9.21). Fibre holders are currently taylor-made for each infant, manufacture and testing of the first two prototypes is illustrated below.

2.7.2.3 Neonatal studies 1 and 2

Two infants have so far been the subjects of MONSTIR imaging trials. In each case the same procedure for 'helmet' manufacture was followed (although slight modifications to the second were made in light of results from the first). The first infant was one of a pair of twins born at 23 weeks, and was the equivalent of 30 weeks gestation at the time of the imaging trial. She was fairly stable, not ventilated and quite small. The second baby had been born at 31 weeks gestation and was 1 week old when studied. She was slightly smaller than the first neonate studied but generally healthy and not ventilated. In both cases parental consent was obtained to manufacture a helmet and ultimately make measurements using MONSTIR.

Figure 2.7.4 shows the photographs of subject 2 from which saggital, coronal and transverse measurements were determined (note the tape-measure in each picture).

Figure 2.7.4 Photographs of the second infant taken to determine head shape (top left = saggital, top right = coronal and bottom = transverse). Note the tape measure held in-shot to give a reference of scale. The infant slept in the incubator throughout.

These photographs were then processed and all scaled to real-size using Corel Photopaint to yield a template as shown in Figure 2.7.5. This was then printed and the three cross sections cut out as templates (allowing 1 cm gap for the insulating foam).

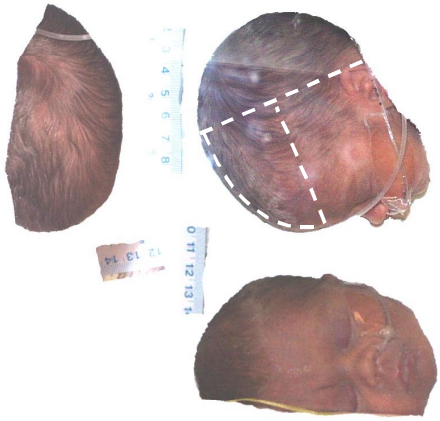

Figure 2.7.5 Processed images of subject 2, scaled and printed at true size were used as templates for 'helmet' manufacture.

An estimate of the extent to which the infant would be able to tilt her head back into a fibre holder had to be made to determine the depth of the back portion of the helmet (as illustrated by the white dotted lines in Figure 2.7.5). A flat sheet of thermoplastic was then heated using hot water, and cardboard templates based on the photographs used to make a suitable indentation for the back of the head to rest in. A T-shaped price of thermoplastic was also heated and moulded using the templates. It was found that fine adjustments to the shape of the thermoplastic could be made using heat gun (usually used for heat-shrink wire insulation) to locally heat areas and mould them by hand.

A reusable support was manufactured for the thermoplastic base-plate, to allow it to be suspended at an adjustable height. This allows fibre bundles to emerge from the underside of the base-plate perpendicular to the surface without being damaged (rather than using prisms to allow the bundles to run parallel to the base-plate) and allows access if adjustments are required (see Figure 2.7.6). The level of the base-plate can also be tilted in two directions by adjusting the relative heights of the three posts to create a comfortable position for the head.

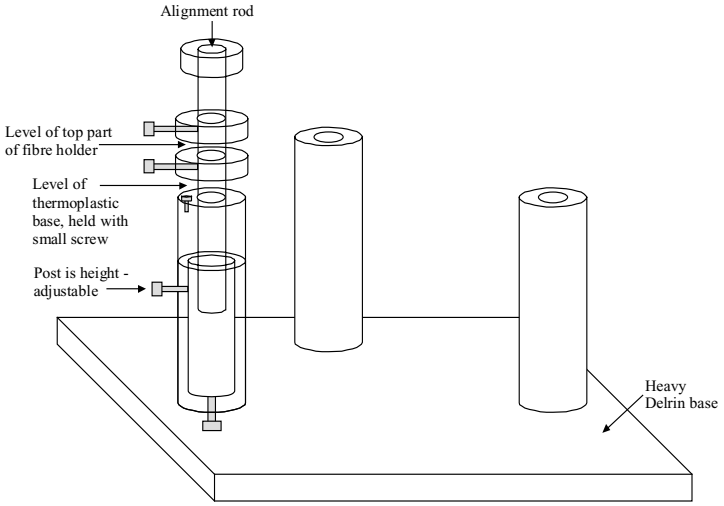

Figure 2.7.6 Reusable support for (infant-specific) thermoplastic base-plates. The three posts and their alignment rods (only one is shown for simplicity), repeatably co-register the top and bottom halves of the 'helmet'

In addition, with the thermoplastic base-plate fixed, the upper thermoplastic T-piece can be aligned with the base-plate via insertion of alignment rods. These have two adjustable washers on them to grip the upper thermoplastic piece and also to record the height of the rod (so the T-piece can be repeatably replaced). Suitable holes must be made in the thermoplastic to enable alignment. At this stage of manufacture the fibre-holder for subject 1 was taken to the NICU and during feeding was held onto the baby's head to check its fit (having been cleaned with hospital detergent, sprayed with alcohol to sterilise it and temporarily lined with sterilised foam). Figure 2.7.7 shows the helmet at this stage.

Figure 2.7.7 The first fibre holder, prior to the fibres being fitted.

Positions for the monstodes were based on a ~ 2cm grid for the back portion of the head, utilising as much of the area as possible. Monstodes on the T-piece were placed at separations over the forehead and the cranium. For subject 2, additional pieces were added to the T-piece to allow two more fibres to be positioned in the evident gap at the top of the head (see Figure 2.7.8). A drill was used to make holes that were small enough for part (C) of the monstode (Figure 2.7.3) to be pushed in and held by friction alone. This worked equally well with non-perforated and perforated thermoplastic (as used for subjects 1 and 2 respectively).

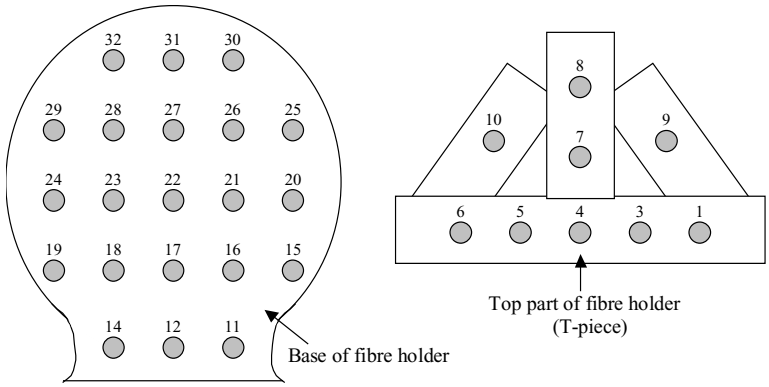

Figure 2.7.8 Fibre positions used for the second infant studied.

The thermoplastic with the monstode sockets in place was then re-sterilised using alcohol. The foam was prepared by placing strips on a sterile surface and punching holes in it using a sterilised punch. The foam was then fitted into the thermoplastic shells so that its holes aligned with the monstode holes (in a clean environment, wearing sterile gloves). The helmet was then wrapped and taken, along with MONSTIR, to the NICU.

Once in the NICU, the monstode part (A) pieces (see Figure 2.7.3) that were already attached to the sources and detectors, were dipped into a disinfectant solution and then each pushed into the appropriate monstode socket as defined in Figure 2.7.8. The monstode was advanced to a point where its tip was just below the depth of the foam so that the infant's head would rest only on the light-leakage protecting foam. Figure 2.7.9 shows photographs of the first fibre holder with all of the MONSTIR fibres in place.

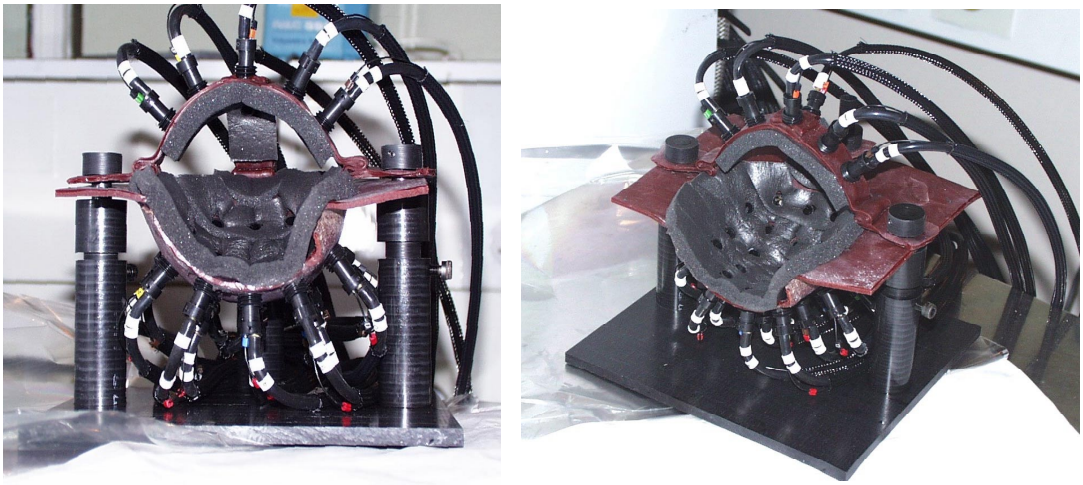

Figure 2.7.9 The first fibre holder with 30 optical fibres in place. The system was sterile at this point.

A hospital cot was then prepared. Two heated Squishon gel mattresses (see appendix A (Figure 2.9.16)) were used to build the level of the cot up to the level of the thermoplastic base-plate as shown in Figure 2.7.10

Figure 2.7.10 The second fibre holder in a neonatal cot with two warmed gel mattresses used to raise the height of the infant to be level with the helmet.

The infant was then transferred to the 'imaging cot' by a nurse as shown in Figure 2.7.11.

Figure 2.7.11 The infant was transferred into the imaging cot to acclimatise

Initially the upper part of the fibre holder was removed leaving just the comfortable foam line[illegible]euvred until its head rested comfortably in the cup. At t[illegible]gure 2.7.12).

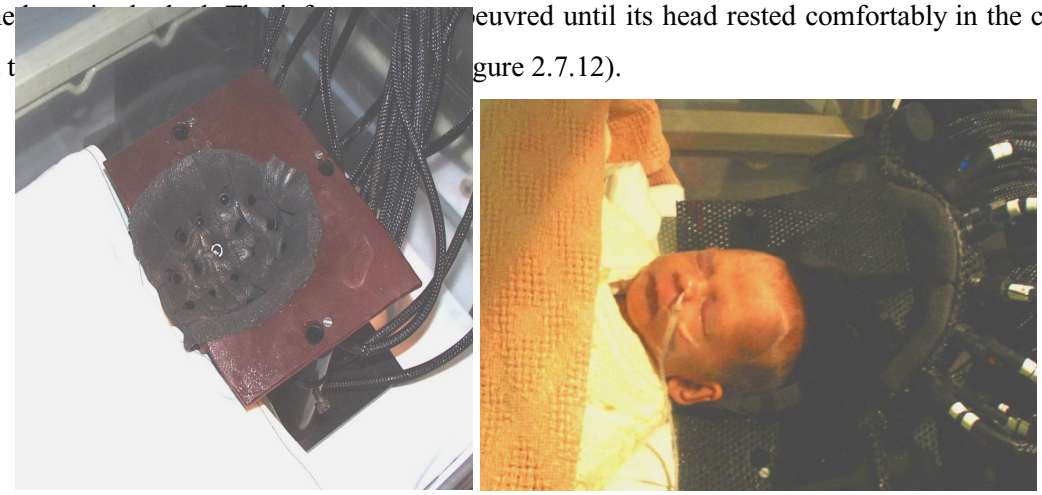

Figure 2.7.12 (left) the first fibre holder showing the comfortable foam lined cup. (right) the second infant asleep in position.

Once the infant was settled, the top part of the fibre holder was put into position and held in place using the alignment rods (Figure 2.7.6). Note that only friction actually holds these rods in place so the top piece could be removed quickly, and would be flexible if the infant moved drastically, see Figure 2.7.13.

Figure 2.7.13 The second infant with the full fibre holder in position just prior to data acquisition.

Image data were acquired as summarised in section 2.7.2.4. Following data acquisition, the infant was returned to her cot. The part (A) pieces of the monstodes were removed (since they were attached to MONSTIR's fibres). Without any further disturbance the fibre holder was taken to the UCL EIT laboratory, where their Microscribe 3D surface plotting instrument (see appendix A (Figure 2.9.21)) was used to record the 3D co-ordinates of all of the optodes. Extra measurements recorded elements of the contours of the fibre holder itself.

The second infant was in the imaging cot for a total of 2.5 hours and showed no signs of discomfort or distress. She needed to be fed every hour (through her naso-gastric tube) but this was readily performed by a nurse, without needing to move the infant from the imaging cot (see Figure 2.7.14).

Figure 2.7.14 The infant could readily receive nursing care (including feeding) during the course of the experiment.

2.7.2.4 Summary of experimental protocol

- Subject identified
- Fibre-holder manufacture commenced
- MONSTIR taken to NICU and switched on for ~ 10 hour warm up time
- ADF file written based on estimates of VOA positions given monstode geometry on fibre holder (see section 1.2.3.5). (For subject 2 this included 3 different acquisitions, one illuminating with just sources on the forehead, and two full – head data sets, 15 seconds per source acquisition time).
- Fibre holder sterilised, MONSTIR fibres attached, placed in 'imaging cot'
- Baby positioned with head in base-plate until settled, T-piece then put on.
- Image data and a monstode calibration data set were acquired.
- Infant removed from cot
- Fibres removed from helmet
- Helmet taken for measurement with Microscribe 3D co-ordinate mapping arm.
- MONSTIR returned to laser-lab.

The MONSTIR system was positioned in a side room in the NICU both for warm-up, and during the studies. Figure 2.7.15 shows the whole clinical set up during data acquisition.

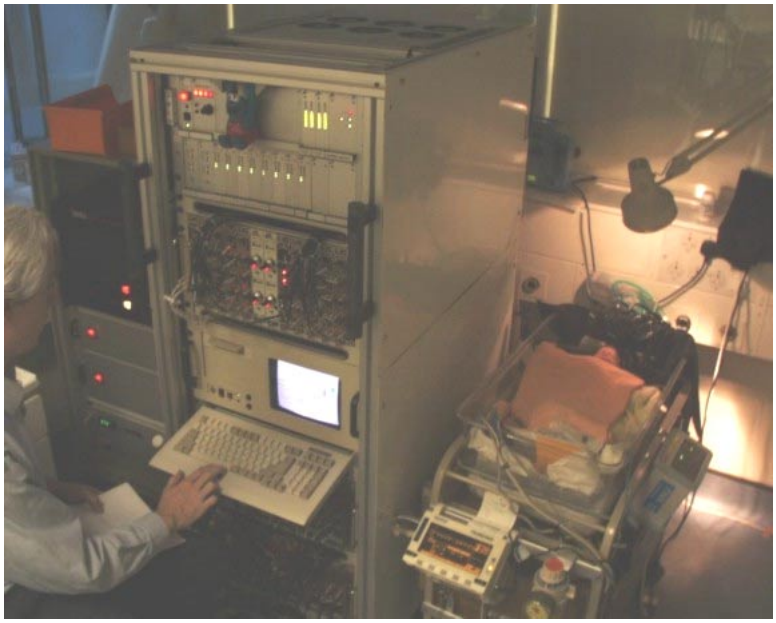

Figure 2.7.15 The MONSTIR system (including secondary rack with the laser etc.- far left) during data acquisition.

2.7.2.5 Results of neonatal studies

The measurements of the surface contours of the head in these preliminary trials were not sufficient to generate a 3D mesh, so performing a 3D reconstruction of the data is not possible at this time. However the data can be used to explore the effectiveness of calibration, the wavelength dependence of the data, and certain of its other features (see section 2.6.7). The measured 3D co-ordinates of the optode positions can be used to plot maps of the data. Figure 2.7.16 (left) shows the mean-time values measured on the back of the head (projected onto the x-y plane, crosses mark detector positions) when illumination was from the source fibres on the front. (right) shows the mean-time measured on the front of the head when serially illuminated by a row of sources on the back of the head (on subject 2, compare with the geometry of Figure 2.7.8). Note that channels 2 and 13 were not used for this data acquisition. We can see the change in the projected mean-time values corresponding to the moving source.

Figure 2.7.17 shows similar plots of the ratio of intensities detected at 815nm/780nm at the back of the head with serial illumination from the front. (Note that the dual wavelength fibre laser was used for this data acquisition, see sections 1.2.3.4 and 2.1.4.5). Since the relative magnitude of the 780nm and 815nm pulses was not specifically measured, these plots will have a constant multiplier representing the initial ratio of the input pulses. Wavelength independent systematic errors, coupling differences and effects of the VOAs should cancel.

In both cases (but particularly for mean-time) the data is smooth showing that calibration and coupling is consistent between the different detectors. Note that the detectors at the base

of the neck (11, 12 and 14) were omitted from the plots as they were obstructed by a combination of the foam lining and the cardigan of the subject.

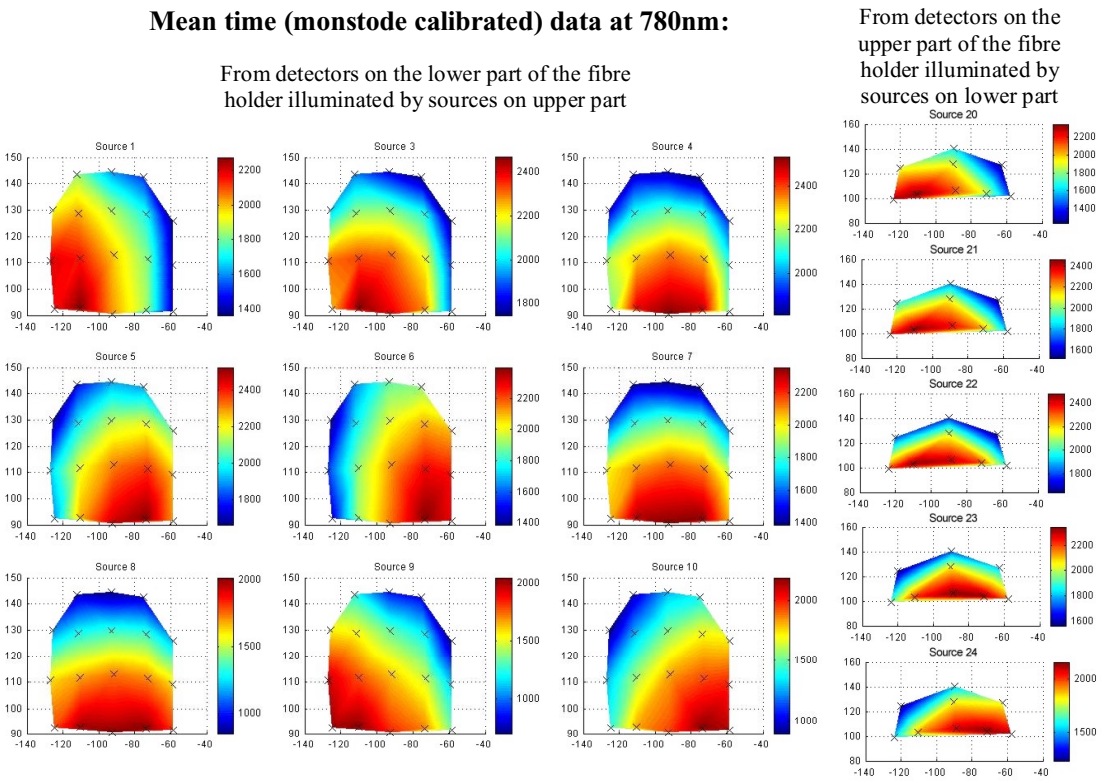

Figure 2.7.16 Absolute mean-time data projections (lef) from the front of the head to the back and (right) from the back of the head to the front. Note that channels 2 and 13 were not used. Crosses mark the positions of detectors.Compare geometry to Figure 2.7.8.

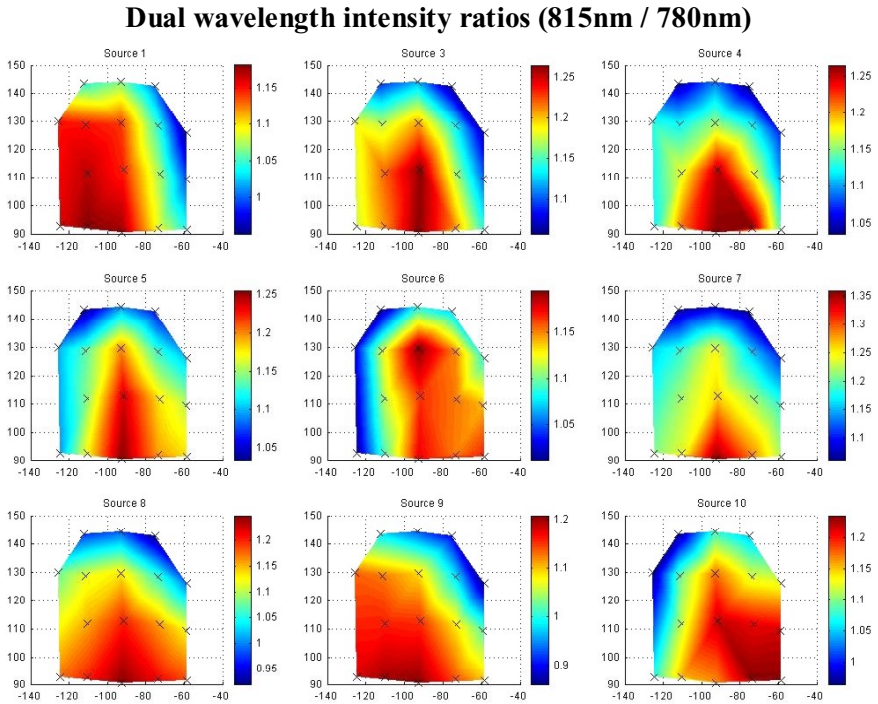

Figure 2.7.17 The ratio of intensity data acquired at 815nm / 780nm detected at the back of the head with serial illumination by sources at the front of the head, crosses mark detector positions (compare with mean-time plots in Figure 2.7.16)

The dynamic range of the measured signal over the head was large. The arbitrary ADF file was not optimised and the source illumination was kept to within eye-safe limits. Figure 2.7.18 shows raw TPSFs measured on the head between source 1 and detectors 20, 22, 24 (see Figure 2.7.8). These combinations all transilluminate the head, but detector 24 is furthest away. All were acquired with their VOA set to maximum transmission.

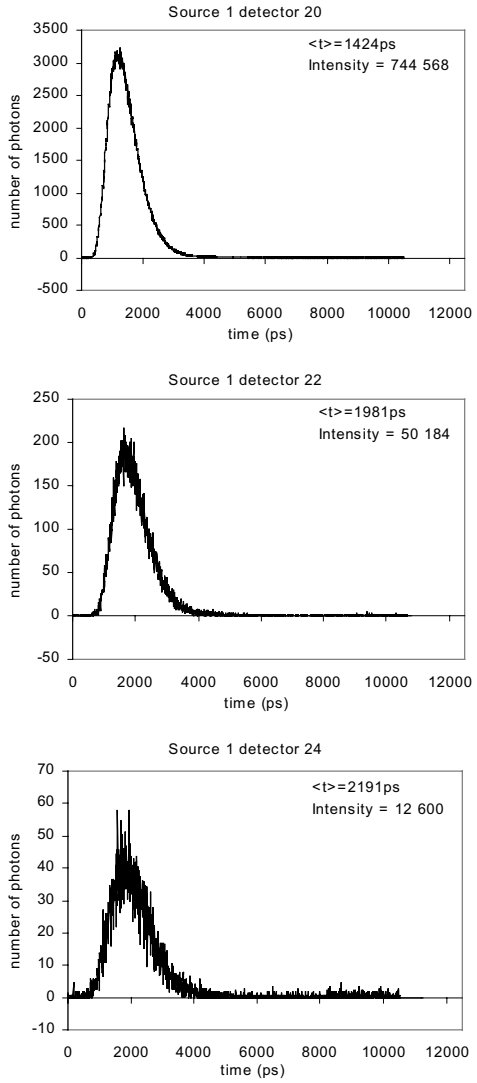

Figure 2.7.18 TPSFs acquired on subject 2 with source 1 and detectors 20, 22 and 24 where 24 is most distant from the source. All were acquired with a VOA set to maximum transmission.

The signal-to noise on the detector 24 TPSF is lower than on 22 and 20, corresponding to its distance (although the efficiency of the detectors and the coupling of the detector bundle to the head may also contribute to differing signal magnitudes).

The data were also used to perform *very* preliminary 2D reconstructions. The monstode positions on the fibre-holder were such that a ring was formed around the head by channels 6, 5, 4, 2, 1, 15, 16, 17, 18 and 19. The relative co-ordinates of the monstodes were approximated to be on the surface of an oval (see Figure 2.7.19) and a 2D mesh generated was based on this shape. It was not possible to perform a 2D:3D correction since data were

not acquired on a homogenous reference phantom and a 3D oval mesh was not available. Where wavelength-difference data are used, the effect of the 2D:3D mismatch may be reduced (see section 2.3.1.1). Absolute μ_a and μ'_s images were reconstructed from absolute mean-time (see Figure 2.7.20) starting values of μ_a = 0.027mm^{-1} and μ'_s = 1mm^{-1} were derived using the method described in appendix A (2.9.7). The absorption images were converted into saturation and volume in that same way as the arm images shown in section 2.5.1.2 (using the same haemoglobin extinction coefficients and $\mu_{a,bg}$) as shown in Figure 2.7.21. Wavelength-difference images reconstructed from mean and intensity data are also shown in Figure 2.7.22.

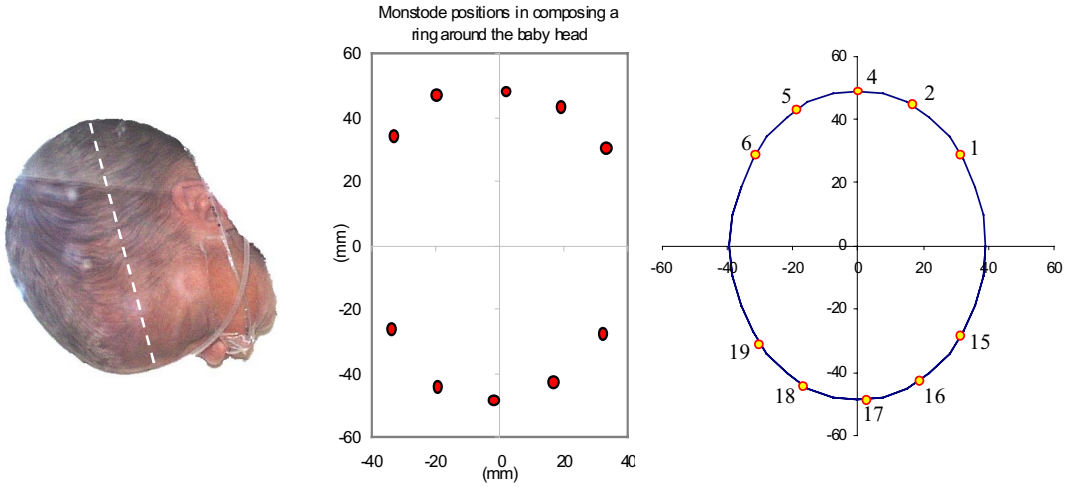

Figure 2.7.19 Sources and detectors approximately in an oval, moved to be on the surface of an oval mesh and reconstructed in 2D. The photo on the left illustrated the approximate plne of the measurement.

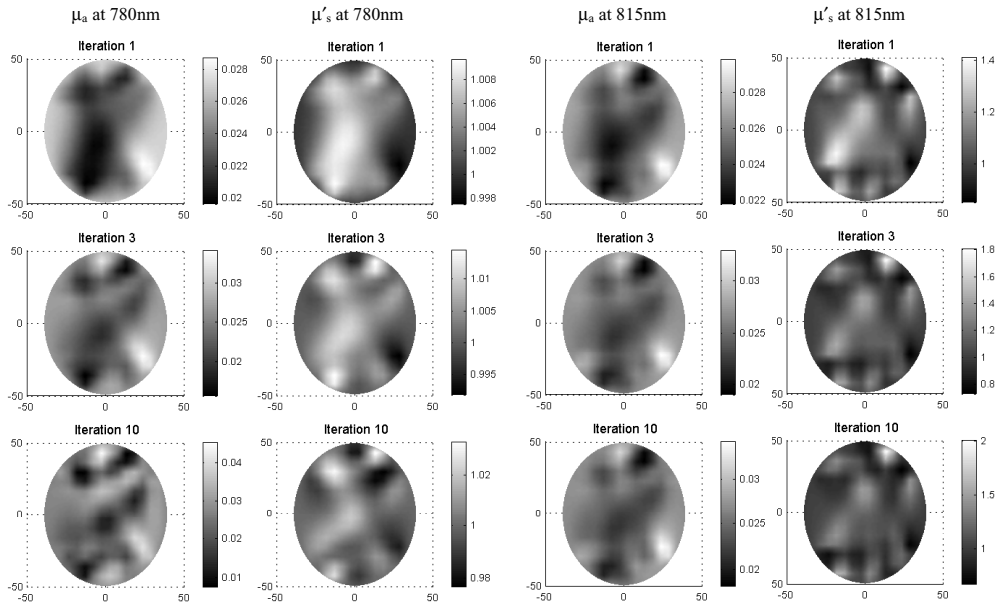

Figure 2.7.20 μ_A and μ'_s images of the neonatal head at two wavelengths, using only 10 sources and detectors, from absolute mean-time with no 2D:3D correction.

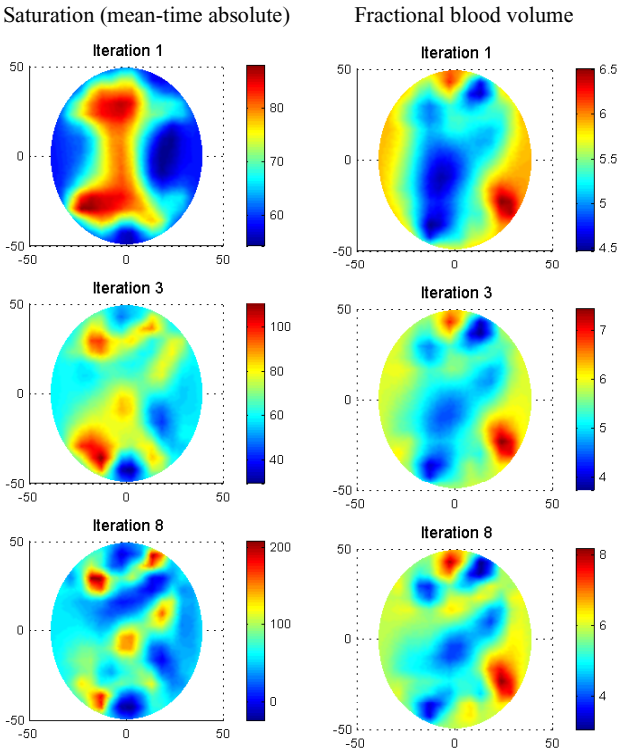

Figure 2.7.21 (left) Saturation images of the neonatal head derived from μ_a images in Figure 2.7.20 (right) blood volume images derived from μ_a images.

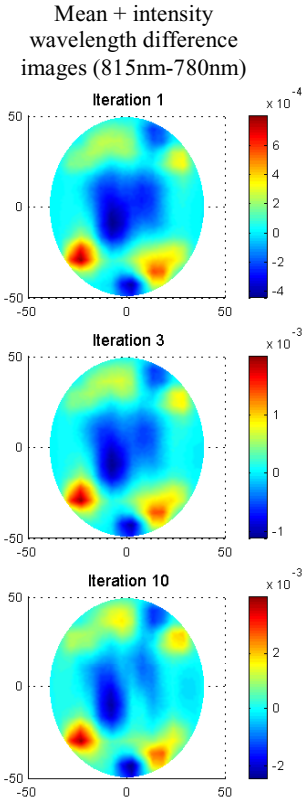

Figure 2.7.22 Wavelength-difference images of neonatal head showing change the in μ_a between 815nm-780nm from <t> and intensity difference data

2.7.2.5.1 Image summary neonatal head images (absolute)

Subject	Heathy 31 week gestation (1week old) neonate
Mesh	Oval (2D) 3196 nodes, 6205 elements
Basis (used to generate PMDFs)	Pixel 16 x 16
Starting parameters: μ_a	0.027mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Refractive index	1.4
Iterations	1, 3, 10
Sources	10
Detectors per source	5 to 7
Datatypes	Mean
Calibration	Deconv equivalent (monstode)
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s
Acquisition time (per source)	15 secs
Wavelength	780 nm + 815nm (interlaced)
μ_a background estimate	0.003mm^{-1}

Table 2.7.1 Properties of images shown in Figure 2.7.20 and Figure 2.7.21

2.7.2.5.2 Image summary neonatal head images (wavelength-difference)

Subject	Heathy 31 week gestation (1week old) neonate
Mesh	Oval (2D) 3196 nodes, 6205 elements
Basis (used to generate PMDFs)	Pixel 16 x 16
Starting parameters: μ_a	0.027mm^{-1}
Starting parameters: μ'_s	1 mm^{-1}
Refractive index	1.4
Iterations	1, 3, 10
Sources	10
Detectors per source	5 to 7
Datatypes	Mean (monstode calibrated) and raw intensity
Calibration	Difference (815 to 780nm)
2d 3d correction applied?	No
Simultaneous	μ_a and μ'_s (μ'_s not shown)
Acquisition time (per source)	15 secs
Wavelength	780 nm + 815nm (interlaced)

Table 2.7.2 Properties of images shown in Figure 2.7.22

2.7.2.6 Evaluation of fibre holder design

2.7.2.6.1 Advantages of the fibre holder design

- The holder fitted well, only 1 detector (on the forehead) showed signs of light leakage and 3 at the base of the neck were obscured.
- The infant seemed comfortable, no marks were left on the skin after data acquisition
- Manufacture took place over two days, but with all appropriate materials and tools available it would be possible to construct a fibre holder of this type within half a day.
- The monstode in-situ mean-time calibration resulted in smooth data (Figure 2.7.16).
- The rigid set-up with a heavy base, and the apparent strength of the infant's head control meant that there was very little chance of eye exposure from any of the sources. Having the base-plate raised, and the fibre bundles emerging from the underside meant that prisms weren't required (although they are being developed).

- Materials and techniques would be familiar to plaster-room / radiotherapy technicians. Most components would be available via existing hospital suppliers
- Adaptable design of the fibre holder base means that a library of the thermoplastic components could be built up to allow 'off the shelf' selection of a suitable holder for a given infant. To fine-tune a fit, the depth of each monstode can be adjusted and fine adjustments to shape could be made using a heat-gun.
- Sterilisation measures taken were simple and acceptable to clinical staff.

2.7.2.6.2 Problems and improvements

- By removing the fibres in the monstodes (piece A) from the helmet before measuring their co-ordinates with the Microscribe arm, their exact positions were not recorded. A dedicated Microscribe could be purchased and accompany MONSTIR to the NICU, which might also make it possible to acquire more surface contour information even on the head of the infant itself (see appendix A (Figure 2.9.21)).
- Using photographs was sufficient for the manufacture of the flexible holder, but insufficient information was available about the contours of the head for 3D mesh generation. Techniques of warping standard meshes into a shape that corresponds to 32 co-ordinates of the monstodes are being developed. Also techniques to create a full surface contour of the head using stereoscopic photography (or thresholded MRI data) are being investigated (although note that infants may be in a closed incubator making access difficult). 3D mapping would allow the helmet to be machined to a precise fit as well as allowing accurate mesh generation. Manufacture of a homogenous reference phantom (see section 2.1.4.3) could also be incorporated into this process.
- Data acquisition has not been attempted on ventilated infants. Oscillating ventilators have a rigid tube, which requires the infant to lie on its back with its head to one side. Conventional ventilators have flexible pipes coming from the side the infant usually lies on their back and the pipes are secured to a bonnet. CPAP (constant positive air pressure) generally has a pipe that comes down to the mouth over the top of the head, which is usually secured using a special hat (see Figure 2.7.23), the infant lies on its back.

Figure 2.7.23 Hats designed to hold the tube during CPAP are available from EME Ltd (see Figure 2.9.16)

While ventilator tubes shouldn't be connected to the fibre holder in case the infant needs to be removed quickly, a method of accommodating the tubes and holding them in place needs to be incorporated into the fibre-holder design.

- While raising the level of the cot base to the height of the base-plate was not a problem in this case, if an infant were in a specialist treatment cot, or an incubator, a less intrusive design would be needed and prisms may be required to allow for the radius of curvature of the fibres.

2.7.2.7 Evaluation of neonatal experimental protocol

2.7.2.7.1 Warm up

The required 10 hour warm-up time of the MONSTIR system is undesirable in the long term since, if an infant with a suspected condition is identified, it should be possible to acquire data as rapidly as possible. Reduction of warm-up time may be possible in future instruments, or alternatively a dedicated room in the NICU would allow MONSTIR to be running continuously.

2.7.2.7.2 Signal-to noise

The SNR of the TPSFs in Figure 2.7.18 was sufficient to calculate <t> but the number of photons detected in 15 seconds was lower than the level of 10^6 suggested to be acceptable for higher order moments in (Hebden *et al*, 1998) (see section 2.1.1.2). SNR could be improved by increasing the acquisition time for each source, or increasing the laser power (with eye-safety measures increased, and automated ADF generation so that closer detectors do not get saturated (see section 1.2.3.5)).

2.7.2.7.3 Acquisition order

If acquisitions need to be longer than 15 seconds per source, it is possible that the infant will become unsettled during full data set acquisition. The order in which the sources are illuminated could be changed so that the full volume of the head is sampled with gradually increasing density, i.e. detect from source 4 (on the forehead), then source 22 (at the back of the head), then a source to the side etc. such that the whole volume has been coarsley sampled even if the acquisition stops part way through.

2.7.2.7.4 Calibration

Calibration has been discussed at length in this thesis and it is evident that using a reference measurement will allow more information to be extracted from MONSTIR measurements. Construction of a head-shaped homogenous (well-matched) reference phantom could potentially be incorporated into manufacture of the fibre holder, or an

adaptable phantom could be made (e.g. from intralipid inside a latex balloon). Studies should be performed to evaluate the effectiveness of a homogenous reference in the clinical setting. Regardless of problems with using the head as its own reference (see chapter 2.4), unless absolute (e.g. monstode-based) calibration can be improved, using a homogenous phantom is the only way that absolute structure including scatter properties will be imaged accurately.

2.7.2.7.5 Invoked-change validation studies

Study of long-term changes such as tissue degeneration (or diagnostic tests where a change is induced) might be achieved via comparison of successive absolute images (possibly using a homogenous reference for calibration), or through reconstruction of the difference between one measurement and the next. Utilising a relative measurement means that calibration and cancellation of meshing errors etc. is far better (see section 2.1.2.3). Note however that if the brain itself is used as the reference that errors may result (see section 2.4). In order to validate the performance of MONSTIR, simple experiments are required to demonstrate its ability to measure predictable changes (in analogy to the arm experiments in (Hillman *et al*, 2001b) and section 2.5.5). Some experiments that could be included into a clinical protocol would be:

- **Hypoxic / hypercapnic swing** ventilated infants in particular can have their inspired oxygen and carbon dioxide adjusted within safe physiological limits. It may be possible to induce a slight change in oxygenation or blood flow in the brain for around 10 minutes and investigate the resulting spatial variation in reconstructed parameters.
- **Invoked functional response** while MONSTIR and TOAST are not optimised to reconstruct changes close to the surface, it may be possible to identify changes in functional activity in regions of the cortex corresponding to visual stimulus (e.g. flashing lights (Meek *et al*, 1998)) or passive motor stimulus (e.g. stroking the leg) (Boas *et al*, 2001b) (see Figure 1.1.12, section 1.1.3.1). Such studies have been widely performed on neonates using single source-detector CW instruments. Response of a particular infant to a stimulus could potentially be characterised using such an instrument (e.g. the NIRO 500, Hamamtsu KK, Japan) prior to MONSTIR data acquisition. Note however that the time-scale of the response to stimulus seen using other NIR techniques is of the order of seconds. MONSTIR must collect data for > 10 minutes assuming a steady-state. Hence a stimulus that can be sustained for longer periods must be found, or data must be acquired for each source on a successive activation cycle. E.g. invoking a response and acquiring single-source data for 20 seconds, waiting for recovery back to baseline, and then stimulating again

to acquire data for the second source and so on. However this protocol would increase total experiment time dramatically, thereby introducing errors due to drift, movement and coupling changes.

- **Contrast agents such as Indocyanine green (ICG)** a contrast agent causes regions containing blood to become more absorbing relative to an image acquired prior to administration of the contrast agent. ICG can also be used to evaluate blood *flow* in the brain (Kusaka *et al*, 2001). While ICG is not thought to cross the blood-brain barrier (Springett *et al*, 2000) it will take time to become distributed within the system. It will then be cleaned out of the blood by the liver. Image data would have to be acquired over a period of time where the concentration of the ICG in the blood is fairly constant. This might be difficult in practice. In addition getting ethical approval to use contrast agents on premature infants may be difficult.
- **Static / dynamic baby head phantom**. A validation step yet to be performed using MONSTIR is to create an image of the baby-head phantom shown in Figure 2.7.2. In addition a 'dynamic' baby head phantom has recently been developed through which blood can be pumped. Problems with meshing, calibration, and 3D mesh generation can be evaluated and optimised in a situation where the change occurring is known.

2.7.3 Acquiring data on the breast

Acquiring data on the breast presents far fewer challenges than brain imaging. However the constraints on breast imaging fibre holders are quite different and there are existing methods that are already accepted for acquiring imaging data on the breast. This section lists the specifications of a breast imaging fibre holder, explores existing approaches and details the development of two fibre holding systems for breast imaging.

2.7.3.1 Specifications of method

The general requirements relating to positioning of fibres are:

- In theory a 3D distribution of sources and detectors is best, although rings of sources and detectors may be appropriate. Compressing the breast is undesirable since it is uncomfortable, and may also reduce the amount of blood in the breast, thereby reducing potential contrast.
- The fibres are heavy, must be held securely and must not move during data acquisition.
- Since the breast is fairly deformable it should be possible to create standardised fibre holders (while allowing for the range of sizes), for easier preparation, mesh generation

and optode placement determination. Since the contours of the breast should be known for mesh generation, the breast should conform to the known shape of the holder.

- It must be possible to make some form of calibration measurement within the imaging session, either a reference measurement or an absolute measurement (see section 2.1.2), preferably in situ.

Environmental factors impose certain limitations:

§ The equipment must not appear intimidating, and should be comfortable and allow easy positioning of the breast.

§ The varied ages, shapes, sizes and mobility of women will affect their ability to position the breast into a fibre holder and keep stationary for the duration of data acquisition (e.g. back problems, tummy size).

Safety aspects of the fibre holder:

§ Skin exposure to NIR laser light must be within safe limits.

§ The holder must be such that if the subject moves, there is no way that they can be subjected to unsafe light levels for eye exposure (and that MCP-PMTs are not damaged due to direct illumination). Alternatively, eye protection must be worn.

Manufacturing constraints:

§ The breast must be in contact with all fibres simultaneously.

§ The optical fibres are permanently attached to MONSTIR and so must be detachable from the fibre holder.

- Source fibres are thin and delicate, and can be damaged if bent or crushed.

§ The fibre holder should be prepared in advance of the subject arriving, perhaps with prior knowledge of her bra size.

2.7.3.2 Development of breast fibre-holder design

2.7.3.2.1 Background

A number of methods for positioning the breast for imaging studies have been proposed. The majority of optical imaging researchers elect to use a compression geometry for breast imaging (Fantini *et al,* 1998b), (Ntziachristos *et al,* 1998), (Rinneberg *et al,* 1998), (Grosenick *et al,* 1999). The amount of light transmitted through the compressed breast provides good SNR. However compression can reduce the amount of blood in the breast, thereby reducing possible contrast, the discomfort of compression associated with x-ray mammography also

remains. Further, there is virtually no depth resolution if data are acquired in flat transmission geometry, which along with edge effects may affect quantitation and accuracy.

Another approach to breast imaging in optical tomography is to have some matching fluid between the breast and the optical fibres (Hoogenraad *et al,* 1997). In a system developed by Hoogenraad *et al*, the subject lies on her front with her breast inside a cone-shaped reservoir in the bed, filled with a scattering liquid. This has the advantage that the shape of the fibre holder does not needed to conform to the contours of the breast, and surface coupling variations don't affect intensity calibration (see section 2.1.2.4). However there are intrinsic problems with this technique if the fluid properties don't match those of the breast very well see secton 2.1.4.3). SNR will be the same for small and large breasts. Also, although an alternative to compression, the use of coupling liquid in the clinical environment is a significant challenge and may deter some women.

Further optical tomography imaging systems exist that allow the breast to hang in a natural shape and have versatile fibre mounts that can be advanced to touch the breast. (Pogue *et al,* 1997) describes an iris of fibres that can acquire data in a single circular plane with a range of diameters. The system described in (Schmitz *et al,* 2000) has a highly engineered adjustable hemispherical fibre positioning system, complemented by a number of iris rings capable of accommodating a range of pendulous shapes. However these adaptable techniques may still not suit all subjects, and would generate a vast number of different possible geometries, each of which would required corresponding 3D mesh generation for image reconstruction.

Researchers developing electrical impedance tomography (EIT) have investigated methods of holding multiple electrodes on the breast. A device described in (Holder, 1995) is shaped as a series of rings that compose a cone. Different combinations of the rings accommodate different sizes of breast. In practice this device was used with the woman lying on her back, and a vacuum pump used to draw the breast up into the holder. A review of other background research into other existing breast imaging and positioning techniques is given in appendix A (2.9.10).

2.7.3.2.2 Development of fibre-holder design

The first optical tomography breast fibre holder constructed was based on the conical design developed for EIT imaging described above ((Holder, 1995)). Composed of three rings, the cone can be adapted by removing rings to accommodate breasts of different sizes (the EIT design had two additional rings for a wider range of sizes). While when used for EIT, the subject lay on her back and a vacuum pump was used to draw the breast up into the holder, it was anticipated that for optical tomography, the subject could lean forward, letting

gravity cause the breast to conform to the shape of the cone. This should work particularly well for post-menopausal women, whose breasts will be less dense and is a more comfortable solution. Figure 2.7.24 shows the first fibre holder built, positioned on a homogenous resin phantom. This first holder accommodated individual sources and detectors, although later a new holder was constructed to accept monstodes. This allowed fibres to be more readily moved around, reduces laser-exposure and allows simple in-situ calibration measurements to be made (see section 2.7.2.2.1).

Figure 2.7.24 Conical breast fibre holder: three rings can be separated to create an annulus for various breast sizes.(right) the FEM mesh generated to allow the conical shape of the phantoms or the breast to be modelled.

Advantages of using the conical shape include the fact that a single mesh (right, Figure 2.7.24) is required for image reconstruction (and also that 2D meshes may sometimes suffice if data are acquired in rings (see section 2.3.3)). Secondly a conical calibration reference phantom (such as the one shown in Figure 2.7.25) can readily be used before or after data acquisition on a volunteer to provide calibration measurements.

The conical fibre holder was used to image a number of breast phantoms (Hebden *et al*, 2001) (see section 2.3.3). However, in order to acquire data on an adult volunteer a mount needed to be constructed to allow the subject to position her breast in the holder.

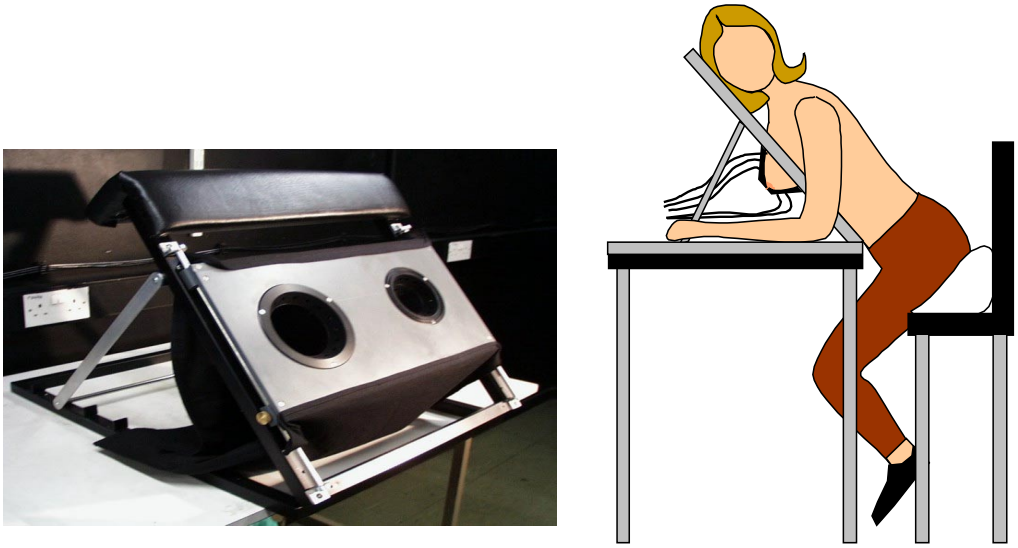

Figure 2.7.25 Photograph of the mount designed to hold the fibre holder for breast imaging and a schematic showing how the subject might be positioned.

Figure 2.7.25 shows the preliminary design of the mount of the fibre holder, along with a schematic of how the subject may be positioned. The tilt of the mount can be varied and the height of the fibre holder adjusted. Two of the rings from the conical holder can be seen in the mount.

A second design for a fibre holder was also conceived, but has yet to be tested on a clinical subject. This design is similar to the approach to the neonatal head fibre holder, using thermoplastic to create a rigid shell. Rather than moulding it directly onto the breast (as is used in radiotherapy, appendix A (Figure 2.9.26), a prototype was constructed by making a plaster cast of a bra. This cast was then used as the basis for moulding a thermoplastic shell that could then accommodate monstodes and be lined with soft light-excluding foam as illustrated in Figure 2.7.26.

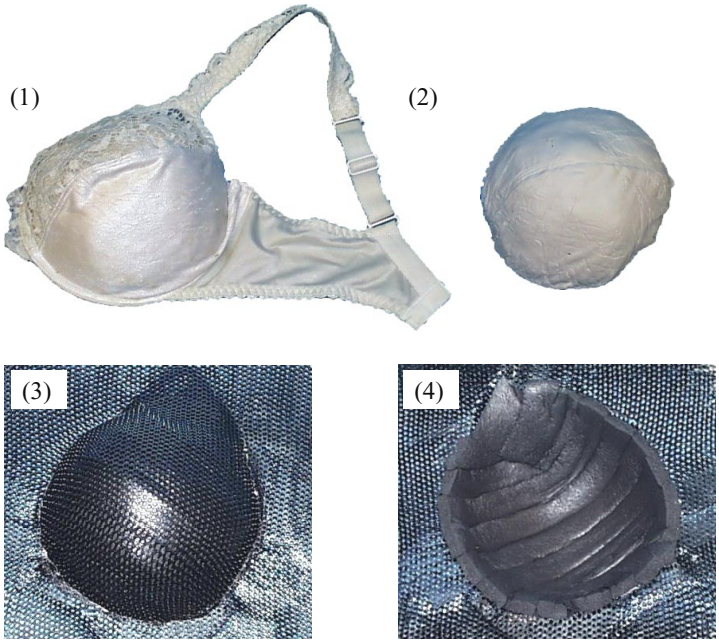

Figure 2.7.26 Scheme for manufacturing a 'thermoplastic bra' fibre holder for optical tomography.

While the shape of this holder is irregular, the Microscribe 3D surface mapping tool (see appendix A (Figure 2.9.21)) could be used to very well characterise the holder and the corresponding optode positions. In addition, only one fibre holder would be required per bra size. Preparation for a subject would just involve fitting the appropriate thermoplastic bra cup into a mount such as the one shown in Figure 2.7.25 and inserting the optical fibres (foam could be replaced for each subject).

2.7.3.3 Summary of breast imaging protocol

The conical fibre holder has so far been used to acquire data on two (healthy) adult volunteers. Results are shown in section 2.5.4 , chapter 2.5. Once a suitable volunteer has been found:

- MONSTIR is turned on to warm up and fibres are placed into the two largest rings in the fibre holder in Figure 2.7.24 (16 in each ring). So far the rings have in fact been mounted separately rather than composing a cone.
- An ADF file for each ring is created based on the attenuation level through a homogenous phantom.
- When the subject arrives she puts on a gown and familiarises herself with the set-up.
- Sitting in an adjustable chair (or standing if more comfortable) the subject tries each ring and decides which is more suitable (in both cases each women chose to use the smaller ring for one breast and the larger one for the other).
- Data sets along with a monstode calibration data set are acquired.
- Subject leaves.
- Data are acquired on a homogenous reference phantom ($\mu_a \approx 0.007mm^{-1}$ and $\mu'_s \approx 1mm^{-1}$ @ 800nm).

2.7.3.4 Evaluation of fibre holders

2.7.3.4.1 Conical holder

- The conical holder allows fairly simple FEM mesh generation since it is composed of basic shapes. Having the fibres in rings allows data to be reconstructed as slices, if desired, using either 3D or 2D (with necessary assumptions, see 2.3.1.4).
- While in most cases good coupling was achieved between the breast and the fibres, both women found positioning quite difficult. Leaning far enough forward to allow gravity to pull the breast into the cup was difficult especially for women with a short body or a large tummy. It is planned to try incorporating a fibre holder into a bed so that the subject can lie on her front.
- The conical shape of the holder was not properly utilised in these early studies since only 1 ring was used at a time. The more anatomical shape of the thermoplastic fibre holder shown in Figure 2.7.26 should make it easier and more instinctive for the woman to position herself.
- Prior information about the bra size of the subject was not utilised in the cone studies.

2.7.3.4.2 Thermoplastic bra design

- If monstode sockets are placed all over the thermoplastic bra (Figure 2.7.26), it is possible that a fibre arrangement could be created to target particular areas of the

breast e.g. around the nipple, or closer to the chest wall, depending on the location of the suspicious lesion.

- A homogenous reference phantom would be required to fit the irregular shape of the thermoplastic fibre holder. It may be possible to manufacture such a phantom using the plaster-cast of the bra used to make the original thermoplastic shell. A more flexible phantom could also be developed e.g. using an intralipid mix (with appropriate or even subject specific optical properties) in a latex balloon (see section 2.2.1 and (McBride *et al*, 1999)).
- 3D mesh generation techniques are not yet sufficiently developed to create a mesh based on the surface co-ordinates of an irregular shape. Methods such as warping meshes to fit particular shapes are being developed. Although alternatively, the thermoplastic could be moulded to a simpler geometric shape such as a hemisphere, based on a similar volume to each bra size (or based on the radiotherapy treatment bra shown in appendix A (Figure 2.9.26)), which would simplify mesh generation.

2.7.3.5 Evaluation of breast experimental protocol

- It is not optimal that each subject chose different holders for each breast. When characterising healthy volunteers it might be better to image each breast in the same ring (or cone) so that features due to bad channels can be easily distinguished from genuine structure.
- The homogenous reference phantom measurement was found to be wavelength dependent (see section 2.5.4.1). The properties of the phantom need to be very well characterised if functional information is to be extracted.
- While it is tempting to cater for the average breast size, optical tomography may be particularly useful for imaging small breasts, and perhaps the denser breasts of younger subjects. There are a number of (usually benign) lesions that affect younger women which require investigation, and radiological methods are unsuitable due to compression and the density of the breasts (see Table 1.1.3). Ultrasound is commonly used in such cases, but optical tomography could potentially provide improved specificity to this sub-group. This suggests that clinical studies need not be entirely based on post-menopausal women, nor indeed should phantom studies be restricted to the optical properties of the ageing breast.
- 3D image reconstructions of the breast phantom shown in Figure 2.3.13 (section 2.3.3.2) took 38 hours to perform 6 iterations. 2D images of data acquired in a planar ring took only 36 minutes for 11 iterations, and a linear reconstruction would take

less than a 5 minutes (including datatype extraction and if PMDFs based on the reference phantom were pre-calculated). Since a reference phantom is used, it may be possible to perform a simple linear-reconstruction of data almost instantaneously, providing a preliminary indication of the quality of the data before the subject leaves.

- ICG is a NIR contrast agent that is already used clinically as an intravascular tracer and has been used in optical tomographic breast imaging to enhance contrast of well-perfused tumours. Use of ICG could be considered if clinical trials of subjects with lesions are unsuccessful. Note however that problems may exist with the changing background absorption of the breast in relation to the tumour ((Ntziachristos *et al*, 1999b)), and there is a very slight risk associated with ICG injection and ethical approval issues would need to be addressed.
- Validation of breast imaging results is much simpler than for brain imaging. For clinical trials, where a subject has a pre-diagnosed lesion, a wealth of information should be available from previous diagnostic scans and tests. A definite diagnosis, perhaps via subsequent tissue biopsy, is likely to be available for comparison with optical tomography results, allowing the accuracy of the new technique to be readily evaluated.

2.8 Further work

This thesis has presented a number of studies that contribute towards the development of advanced methods for optical tomographic image reconstruction. During each investigation, particular areas were identified which required further work, as summarised below:

Chapters 2.1, 2.2 and 2.3 presented early development of techniques for calibration and datatype extraction, along with an investigation into the match between TOAST and MONSTIR, and a number of preliminary phantom imaging trials. The most important aspects of these investigations can be summarised as the need to improve information extraction:

- **Information extraction.** Current methods of extracting information from TPSFs dramatically under-use the capabilities of MONSTIR and TOAST. The poor noise tolerance of TOAST and the difficulties associated with calibrating MONSTIR affect convergence and hence propagate through to affect image quantitation and μ_a / μ'_s separation. Using only (well calibrated) mean-time data does not provide sufficient *information* to separate μ_a and μ'_s accurately. A summary of possible alternatives to datatypes is given in section 2.2.7.1. Pursuing methods such as reconstructing parameters that describe fits to the time-domain profile of the TPSF may improve results, and are currently being explored. Compensation for the noise characteristics of data within TOAST may also aid convergence, especially if second or third generation TCSPC instrumentation correspondingly provides data with fewer systematic errors and less noise.

Chapter 2.4 described problems associated with reconstructing difference data acquired on an object with static heterogeneous structure. It was concluded that further work was required to quantify the significance of this error to successful breast and brain imaging.

- **Quantifying the effect of static structure on clinical difference images.** While the investigation in chapter 2.4 demonstrated that static structure *would* affect difference images, it has yet to be determined whether this error will introduce distortions that outweigh the potential advantages of utilising *in-situ* difference data. Development of more complex, anatomically-based 3D FEM models would allow specific lesions or conditions (see Table 1.1.2 and Table 1.1.3) to be simulated (possibly also containing void regions as described in chapter 2.6). Optimal methods of *reliably* and *repeatably* detecting conditions could then be developed, and tolerable levels of error determined.

- **Evaluation of correction methods.** Methods suggested for compensation of background structure could also be evaluated using a realistic simulation. Results could be compared to those achievable without correction, or the results of manipulating absolute images.

Chapter 2.5 demonstrated a number of techniques for extraction of functional parameters from multi-wavelength data. It was concluded that further work was required to validate the methods developed as described below:

- **Validating the μ_a model**. The assumptions made in deriving saturation and fractional blood volume from μ_a included the wavelength independence of $\mu_{a,bg}$, and the applicability and accuracy of the chosen haemoglobin extinction coefficients. The effects of these assumptions could be demonstrated through measurements of liquid phantoms containing haemoglobin. The oxygenation status of the solution could be changed by bubbling oxygen or nitrogen through it, and an oxygen electrode could be used to validate results (similarly to the method described in (McBride *et al*, 1999)).

- **Experimental evaluation of three-wavelength methods**. The equations derived assuming three wavelengths were available were not validated experimentally. The Tsunami Ti:Sapphire could be used to acquire data at three (closely grouped) wavelengths to determine whether using three absolute images, or two wavelength-difference images to derive functional parameters can truly improve imaging results compared to dual-wavelength acquisition.

- **3D reconstruction of clinical data**. Since the clinical examples shown utilised 2D reconstructions, it is important that the effect of 2D:3D mismatch on derived functional parameters be determined. This could be achieved by reconstructing the arm and breast data using a 3D mesh, or by simulating forward data in 3D and examining the effect on derived parameters of reconstructing using 2D mesh.

- **Diagnostic significance of blurring**. Simulations demonstrated that, even when the μ_a model is correct, images would still have intrinsic errors relating to image reconstruction. It is likely that *reliable* localisation of regions of high or low saturation or fractional blood volume would provide useful diagnostic information, regardless of quantitative accuracy. However further work must demonstrate that, in clinically realistic geometries, inherent blurring and the effects of noise will not distort images such that they imply erroneous diagnoses.

- **Development of solution constraints**. The methods of constraining μ_a solutions by assuming wavelength independent scatter, and by reconstructing directly for

saturation and blood volume showed great promise. The intrinsic wavelength dependent properties of tissues should undoubtedly be exploited, and could potentially dramatically improve functional imaging results. The techniques should be evaluated using suitable clinical data. The effectiveness of constraining solutions for non-linear image reconstructions should also be explored.

The experiment described in chapter 2.6 demonstrated the potential significance of the 'void problem'. While results showed that the presence of a spherical 2mm thick clear layer would dramatically affect light propagation across a diffusing cylinder, results do *not* prove that light will behave similarly in the neonatal head. The purpose of the experiment was to validate the 3D void-TOAST model in a simple geometry, such that more advanced simulations based on genuine anatomy could be developed and reliably interpreted in the future. The following aspects of the void problem require further investigation:

- **Determination of refractive index mismatches**. Experimental results demonstrated that the refractive index mismatch between the material composing the void and the surrounding diffusing 'tissue' would affect the severity of the void problem. Measurements of the refractive index mismatch between CSF and brain tissue, and CSF and the subarachnoid membrane are required. The magnitude of the mismatches will determine whether it is necessary to compensate for refractive index in models that are used to explore the void problem.
- **Verification that CSF affects data acquired on the neonatal head**. Preliminary attempts were made to use data acquired on a neonate to determine whether preferential propagation of light through the CSF was occurring. The results were not conclusive and demonstrate that it is very difficult to assess whether CSF will detrimentally affect imaging results. Development of anatomically accurate simulations (and perhaps even phantoms) would more readily demonstrate the effects of the void on extraction of diagnostically useful, *reliable* information (as with the effects of static structure on difference imaging described in chapter 2.4). Simulations could be based on segmented MRIs of the brains of premature infants, inclusions could mimic the common conditions summarised in Table 1.1.2. The 3D void-TOAST forward solver could potentially be used to create such simulations, with much development, and compensation for refractive index implemented if necessary.
- **Development of methods to compensate for the void problem**. The use of hybrid models to reconstruct data acquired in the presence of void layers and regions needs to be explored in more detail. It is unlikely that prior knowledge of the void boundaries will ever be sufficiently accurate to allow absolute reconstruction of data

(assuming that the effects of CSF are as severe as implied by results published so far). Region retrieval methods to locate void regions, as well as transport-based code are being developed. Difference imaging was briefly examined, although the less limited case of objects with irregular geometries and void regions *must* be explored before conclusive remarks can be made. Difference reconstructions will also be affected by errors relating to constant background structure (which will *include* the effects of voids) described in chapter 2.4.

- **Exploring late-arriving photons**. It is possible that the photons in the later parts of the measured TPSFs hold more information about the regions beyond a clear layer than those arriving at the detector early. The TPSFs measured in these experiments should be examined to determine whether images based on parameters that are preferentially sensitive to the tail of a TPSF could better locate inclusions.

The final chapter explored issues relating to the practical application of clinical optical tomography. The potential improvements to the fibre holders designed for breast and brain imaging were discussed in detail in sections 2.7.2.6 and 2.7.3.4. Overall, the major practical obstacles to successful data acquisition and image reconstruction were found to be:

- **Surface mapping and 3D mesh generation**. Surface mapping techniques are not yet suitably developed to allow automated fibre-holder manufacture, or accurate 3D mesh generation. Techniques for mapping the neonatal head are being developed, and subjectively mapping breasts is unlikely to be required, although irregularly shaped fibre holders would need to be characterised. 3D mesh generation is being developed, although creating accurate meshes, with suitable boundaries and node-densities is currently very difficult. Methods for creating 3D meshes from a number of surface coordinates need to be developed as a priority to allow proper reconstruction of clinical data.

- **Calibration of clinical data**. For brain imaging, the use of monstodes allowed fairly good calibration of mean-time data. However phantom results imply that these data are unlikely to yield good quality images. Throughout this thesis, the use of a homogenous reference phantom measurement has been found to produce much better images than attempting absolute calibration. Problems associated with using the head itself as a reference phantom were described at length. While likely to be difficult, attempts should be made to acquire homogenous reference data in the clinical setting, to determine whether to is possible to exploit the potential benefits of such a measurement.

Overall, optical tomography is likely to require future development of sophisticated models, phantoms and experimental protocols to demonstrate that it is suitable for use as a reliable clinical diagnostic tool. This thesis has demonstrated some of the important constraints and limitations of the method, and suggested a number of corrections and new techniques for overcoming such obstacles.

2.9 APPENDIX A

2.9.1 Circumference of neonatal head with age

As discussed in section 1.1.3.1:

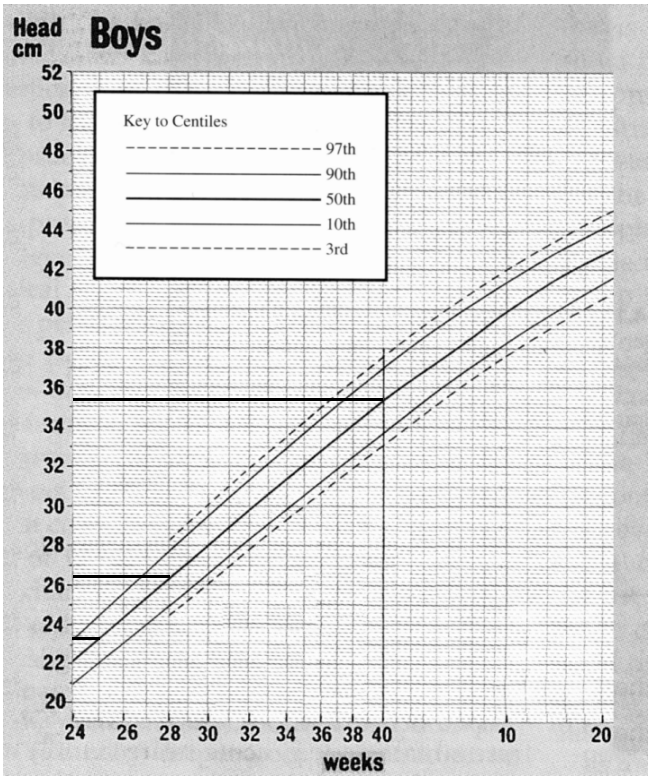

Figure 2.9.1 Male head circumference chart, gestational age until 40 weeks, then post-natal age. Shows average (assuming circular head), 40 week (term) = 11.2 cm diameter, 28 week = 8.4 cm diameter, and 25 week = 7.4 cm diameter. (From(Johnston, 1998) *p 45) Sizes for girls are fractionally smaller.*

2.9.2 Central moment datatypes in terms of moments

2.9.2.1 Variance

Central moments were introduced in section 1.3.3 as moments about the mean. Central moments can be expressed in terms of ordinary temporal moments. The general form of a central moment of order *u* is given by:

$$Central\mathcal{M}_{s,u}[\Gamma_{n,m}(t)] = \frac{1}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)(t-<t>)^{u}e^{-st}dt$$

[2.9.1]

Where s = 0 for general moments. So the 2nd central moment, variance is given by:

$$\begin{aligned}
&Variance[\Gamma_{n,m}(t)] = \frac{1}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)(t-<t>)^{2}dt \\
&= \frac{1}{E}\int_{t=0}^{t=\infty}(\Gamma_{n,m}(t)t^{2} - 2\Gamma_{n,m}(t)t<t> + \Gamma_{n,m}(t)<t>^{2})dt \\
&= \frac{1}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)t^{2}dt - \frac{2<t>}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)tdt + \frac{<t>^{2}}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)dt \\
&= \frac{1}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)t^{2}dt - <t>^{2}
\end{aligned}$$

[2.9.2]

Therefore:

$$Variance[\Gamma_{n,m}(t)] = \mathcal{M}_{0,2}[\Gamma_{n,m}(t)] - \mathcal{M}_{0,1}[\Gamma_{n,m}(t)]^{2} = 2^{nd}\,moment - 1^{st}\,moment^{2}$$

[2.9.3]

2.9.2.2 Skew

Similarly for skew:

$$\begin{aligned}
&Skew[\Gamma_{n,m}(t)] = \frac{1}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)(t-<t>)^{3}dt \\
&= \frac{1}{E}\int_{t=0}^{t=\infty}(\Gamma_{n,m}(t)t^{3} - 3\Gamma_{n,m}(t)t^{2}<t> + 3\Gamma_{n,m}(t)t<t>^{2} - \Gamma_{n,m}(t)<t>^{3})dt \\
&= \frac{1}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)t^{3}dt - \frac{3<t>}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)t^{2}dt + \frac{3<t>^{2}}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)tdt - \frac{<t>^{3}}{E}\int_{t=0}^{t=\infty}\Gamma_{n,m}(t)dt
\end{aligned}$$

[2.9.4]

Leading to:

$$\begin{aligned}
&Skew[\Gamma_{n,m}(t)] = \mathcal{M}_{0,3}[\Gamma_{n,m}(t)] - 3\mathcal{M}_{0,1}[\Gamma_{n,m}(t)]\mathcal{M}_{0,2}[\Gamma_{n,m}(t)] + 2\mathcal{M}_{0,1}[\Gamma_{n,m}(t)]^{3} \\
&= 3^{rd}\,moment - 3\times 1^{st}\,moment\times 2^{nd}\,moment + 2\times 1^{st}\,moment^{3}
\end{aligned}$$

[2.9.5]

2.9.3 Dependence of datatypes on absorption and scatter

As described in section 1.3.3.6, a 2D simulation was performed using the distribution shown in Figure 2.9.2. Noise-free forward data were calculated using TOAST. Results are shown for intensity, mean-time, variance, skew (Figure 2.9.3), Mellin-Laplace transforms for s=0.005 and n=0 to 3 (Figure 2.9.4), and for complex amplitude and phase at 50 and 100MHz (Figure 2.9.5). The percentage change in each datatypes relative to values simulated on a homogenous background (without the inclusions) is shown. Plots show data for each detector, anticlockwise, around the object.

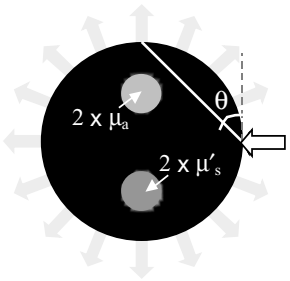

2D simulation,
radius = 35mm
$\mu_a = 0.01mm^{-1}$
$\mu'_s = 1mm^{-1}$
n=1.4
32 sources and detectors
interspersed and evenly
distributed

Figure 2.9.2 2D simulation used to explore effect of single scattering or absorbing inclusion on boundary datatypes.

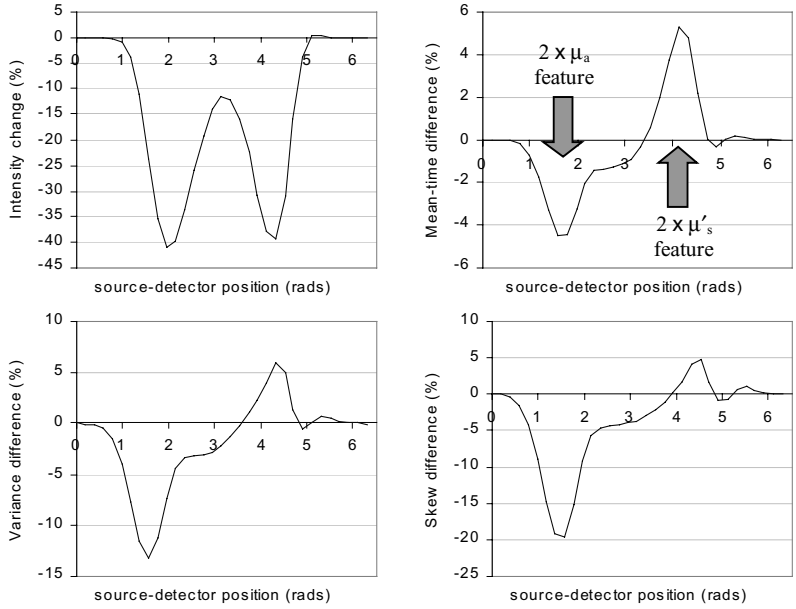

Figure 2.9.3 Intensity, mean-time, variance and skew percentage change between a homogenous distribution and one with an absorbing and a scattering inclusion present (Figure 2.9.2), for illumination by a single source with detectors positioned evenly around the perimeter of the object.

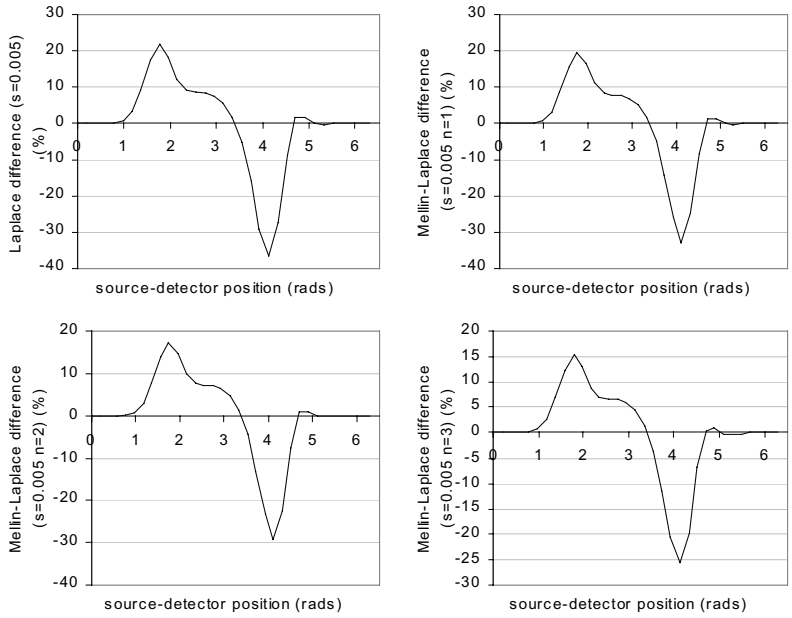

Figure 2.9.4 Mellin-Laplace (for s=0.005 an n=0 to 3) percentage change between a homogenous distribution and one with an absorbing and a scattering inclusion present (Figure 2.9.2), for illumination by a single source with detectors positioned evenly around the perimeter of the object.

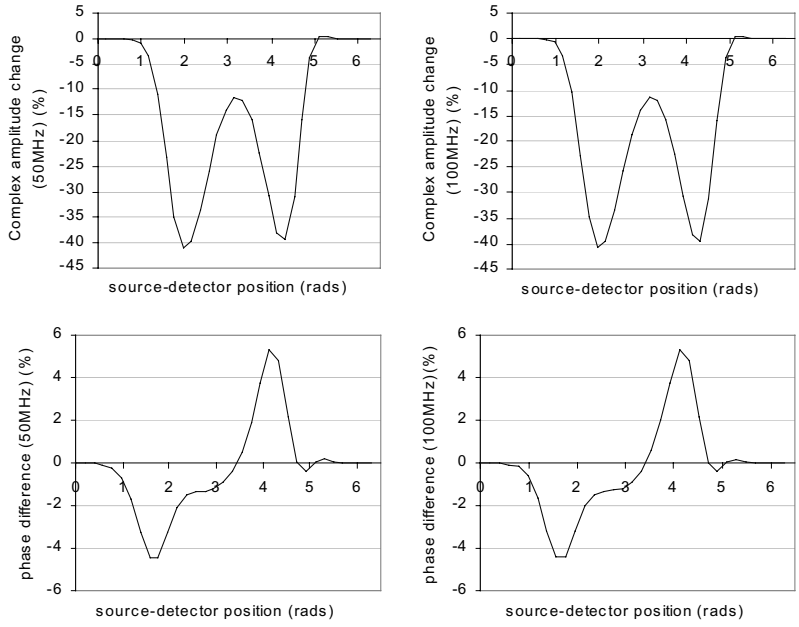

Figure 2.9.5 Complex amplitude and phase (for ω = 50MHz and 100MHz) percentage change between a homogenous distribution and one with an absorbing and a scattering inclusion present (Figure 2.9.2), for illumination by a single source with detectors positioned evenly around the perimeter of the object.

A second simulation was also performed to assess the effect on datatypes of an inclusion that was both absorbing and scattering (as shown in Figure 2.9.6). The corresponding percentage changes in each datatype are shown in Figure 2.9.7, Figure 2.9.8 and Figure 2.9.9.

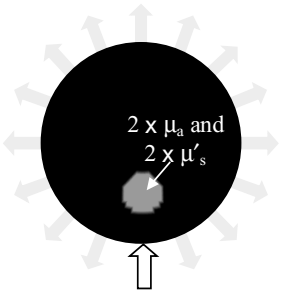

2D simulation,
radius = 35mm
$\mu_a = 0.01mm^{-1}$
$\mu'_s = 1mm^{-1}$
n=1.4
32 sources and detectors
interspersed and evenly
distributed

Figure 2.9.62D simulation used to explore effect of a scattering and absorbing inclusion on boundary datatypes.

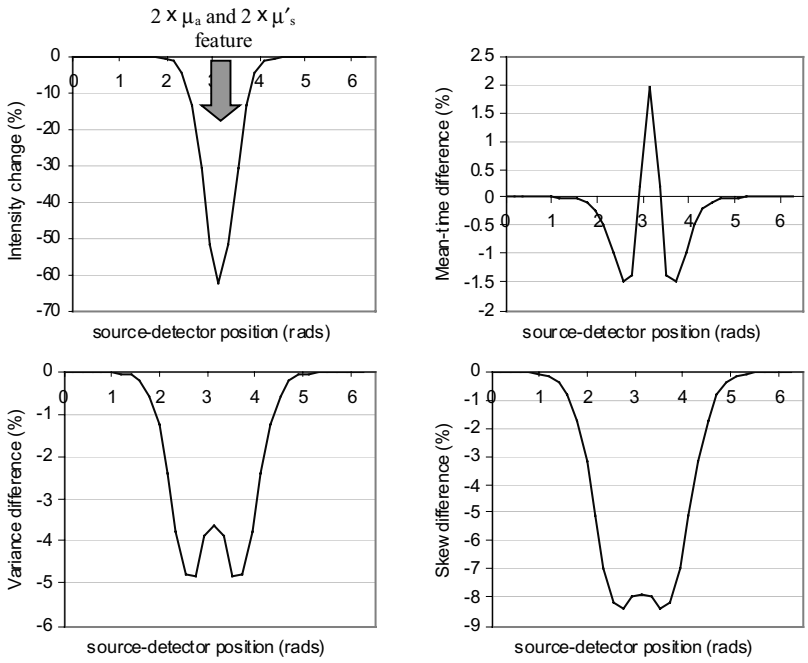

Figure 2.9.7 Intensity, mean-time, variance and skew percentage change between a homogenous distribution and one with a single absorbing and scattering inclusion present (Figure 2.9.6), for illumination by a single source with detectors positioned evenly around the perimeter of the object.

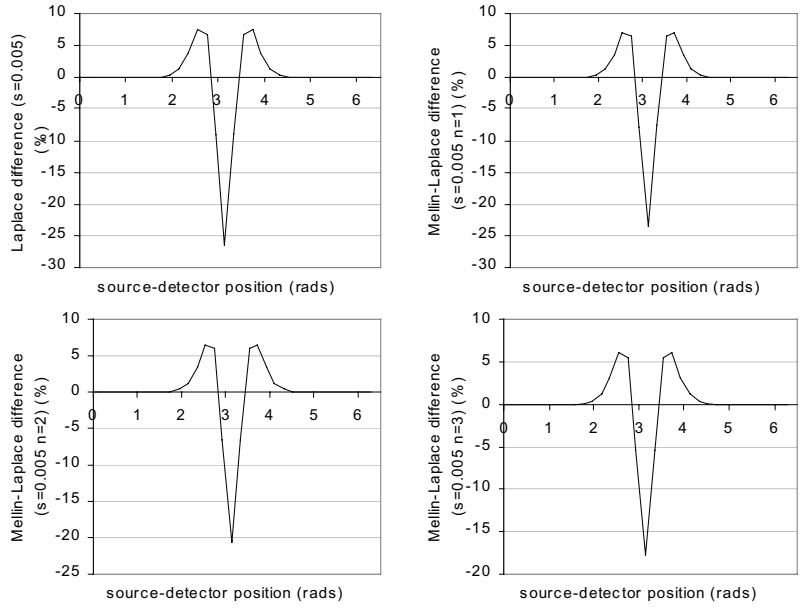

Figure 2.9.8 Mellin-Laplace (for s=0.005 an n=0 to 3) percentage change between a homogenous distribution and one with a single absorbing and scattering inclusion present (Figure 2.9.6), for illumination by a single source with detectors positioned evenly around the perimeter of the object.

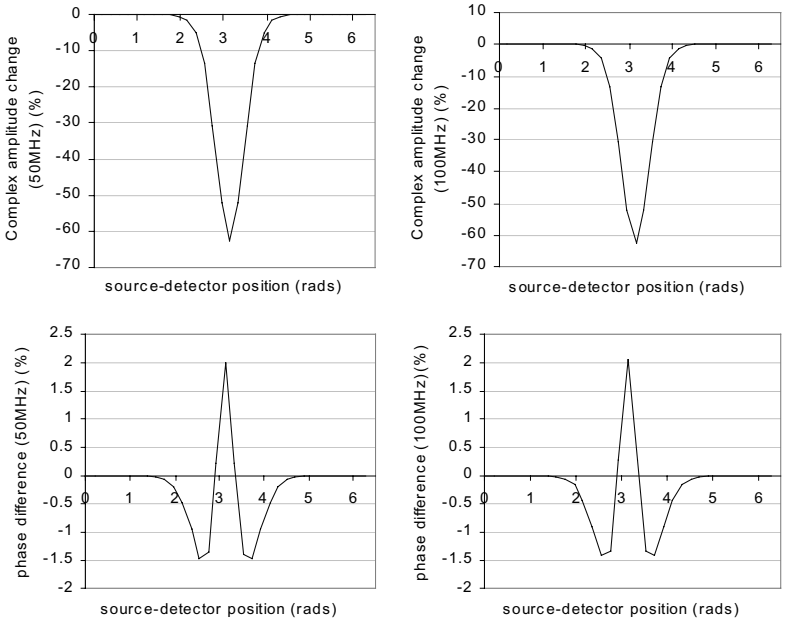

Figure 2.9.9 Complex amplitude and phase (for ω = 50MHz and 100MHz) percentage change between a homogenous distribution and one with a single absorbing and scattering inclusion present (Figure 2.9.6), for illumination by a single source with detectors positioned evenly around the perimeter of the object.

2.9.4 Equivalence of datatypes

As described in section 1.3.3.6, the significance of describing a TPSF in terms of extracted moments can be explained if we consider how the combination of moments represents a series expansion of the Fourier domain TPSF. Figure 2.9.10 shows the Fourier transform of a typical TPSF (generated from the frequency domain 3D infinite-space Green's function (see section 1.3.1.2.1)).

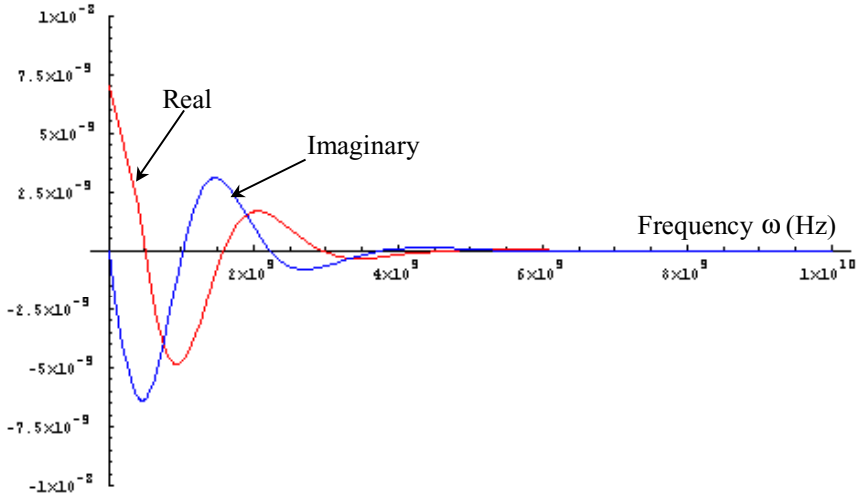

Figure 2.9.10 Real and imaginary parts of the Fourier transform of a TPSF in the range 0 to 10 000 MHz, corresponding to 100 ps temporal resolution at the maximum frequency (μ_a=0.01 mm^{-1}, μ'_s=1 mm^{-1} , n=1.56, d=70 mm)

The Taylor expansion of a function *f(a)* is:

$$f(a) = f(a_0) + (a-a_0)\frac{\partial f(a)}{\partial a}\bigg|_{a=a_0} + \frac{(a-a_0)^2}{2!}\frac{\partial^2 f(a)}{\partial a^2}\bigg|_{a=a_0} + \frac{(a-a_0)^3}{3!}\frac{\partial^3 f(a)}{\partial a^3}\bigg|_{a=a_0} \ldots.$$

[2.9.6]

2.9.4.1 Moments (s=0)

From the definition of temporal moments given in 1.3.3.2 we can express moments in terms of the frequency domain form of a TPSF. Where

$$Y_{n,m}(\omega) = \sum_{t=0}^{T} y_{n,m}(t) e^{-i\omega t} \ ,$$

[2.9.7]

we know that the integrated intensity is equal to $Y_{n,m}(0)$, however we can also see that the first differential of $Y_{n,m}(\omega)$ evaluated at w=0 is:

$$\frac{\partial Y_{n,m}(\omega)}{\partial \omega}\bigg|_{\omega=0} = -i\sum_{t=0}^{T} y_{n,m}(t)\, t\, e^{-i\omega t}\bigg|_{\omega=0} = -i\sum_{t=0}^{T} y_{n,m}(t)\, t = -i < t > .$$

[2.9.8]

Thus the mean-time of a TPSF can be expressed in terms of the first differential of the Fourier transform at ω=0. This implies that mean-time is equal to the gradient of the Fourier

transform at ω=0. We can continue to differentiate the Fourier transform, and find that higher moments can also be expressed in terms of the derivatives of the Fourier transform, namely:

$$\text{Mean-time} = \frac{i\left.\frac{\partial Y(\omega)}{\partial \omega}\right|_{\omega=0}}{Y(0)} \qquad [\ 2.9.9\ ]$$

$$2^{nd}\ moment = \frac{-\left.\frac{\partial^2 Y(\omega)}{\partial \omega^2}\right|_{\omega=0}}{Y(0)} \qquad [\ 2.9.10\ ]$$

$$3^{rd}\ moment = \frac{-i\left.\frac{\partial^3 Y(\omega)}{\partial \omega^3}\right|_{\omega=0}}{Y(0)} \qquad [\ 2.9.11\ ]$$

The general form is:

$$\mathcal{M}_{0,u}[Y_{n,m}(\omega)] = i^u \frac{\left.\frac{\partial^u (Y_{n,m}(\omega))}{\partial \omega^u}\right|_{\omega=0}}{Y_{n,m}(0)} \qquad [\ 2.9.12\ ]$$

From [2.9.6], in the case where $a_o=0$ we get:

$$f(a) = f(0) + a\left.\frac{\partial f(a)}{\partial a}\right|_{a=0} + \frac{a^2}{2!}\left.\frac{\partial^2 f(a)}{\partial a^2}\right|_{a=0} + \frac{a^3}{3!}\left.\frac{\partial^3 f(a)}{\partial a^3}\right|_{a=0} .... \qquad [\ 2.9.13\ ]$$

Thus the Fourier transform of a TPSF (Y(ω)) can be expressed as a series of the moments of the TPSF. By substituting [2.9.12] into [2.9.13] where a=ω we get:

$$\mathrm{Re}[\mathrm{Y}(\omega)] = Intensity\ (1 - \frac{\omega^2}{2!} \times 2^{nd}\ moment + \frac{\omega^4}{4} \times 4^{th}\ moment.....)$$

$$\mathrm{Im}[\mathrm{Y}(\omega)] = Intensity\ (-\omega < t > + \frac{\omega^3}{3!} \times 3^{rd}\ moment.....) \qquad [\ 2.9.14\ ]$$

This result demonstrates that combinations of datatypes will provide a description of the Fourier transform of a TPSF within the range of ω for which the Taylor series converges, although without intensity the description lacks the scaling factor describing the amplitude. It also shows that mean-time only holds information about the imaginary part of the TPSF, and that the 2nd moment describes the real part.

We can test the convergence of this series by evaluating the series in [2.9.14] using analytic expressions for the moments (see 2.9.7) and comparing the result to the 3D infinite space frequency domain Green's function. Figure 2.9.11 shows that for terms of the series up to ω^6, the fit to the Fourier domain TPSF is good up to around 800 MHz. For terms of the series up to ω^3, the fit to the Fourier domain TPSF is good up to around 400 MHz. At around

1.5GHz both series expansions diverge rapidly from the true values of Y(ω). The best fit is where ω = 0, which is just the DC intensity, since at ω = 0, the series solution is exact.

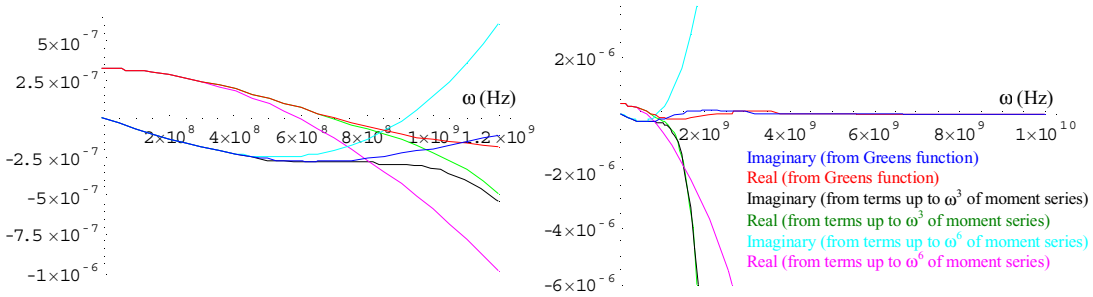

Figure 2.9.11 Plots of analytical Fourier transform of a TPSF and the Taylor series expansion of the Fourier transform in terms of moments. (Left) shows the plots up to 1.2 GHz, (Right) shows how the series diverge above 2GHz. μ_a=0.01 mm^{-1}, μ_s'=1 mm^{-1} , n=1.4, d=70 m.

However, we do not generally calculate higher order moments than mean-time. Variance and skew, the second and third central moments, are more normally used.

2.9.4.2 Central moments:

In the frequency domain, the 2nd central moment (variance) would be given by:

$$Central\mathcal{M}_{0,2}[Y_{n,m}(\omega)] = -\frac{\left.\frac{\partial^2(Y_{n,m}(\omega))}{\partial\omega^2}\right|_{\omega=0}}{Y_{n,m}(0)} + \left(\frac{\left.\frac{\partial(Y_{n,m}(\omega))}{\partial\omega}\right|_{\omega=0}}{Y_{n,m}(0)}\right)^2 \qquad [2.9.15]$$

This can be derived from expressions of central moments in terms of ordinary moments given in 2.9.2 . Similarly skew can be expressed as:

$$Central\mathcal{M}_{0,3}[Y_{m,n}(\omega)] = -i\frac{\left.\frac{\partial^3(Y_{n,m}(\omega))}{\partial\omega^3}\right|_{\omega=0}}{Y_{n,m}(0)} + 3i\left(\frac{\left.\frac{\partial^2(Y_{n,m}(\omega))}{\partial\omega^2}\right|_{\omega=0}}{Y_{n,m}(0)}\frac{\left.\frac{\partial(Y_{n,m}(\omega))}{\partial\omega}\right|_{\omega=0}}{Y_{n,m}(0)}\right) - 2i\left(\frac{\left.\frac{\partial(Y_{n,m}(\omega))}{\partial\omega}\right|_{\omega=0}}{Y_{n,m}(0)}\right)^3 \qquad [2.9.16]$$

The 4th order central moment (kurtosis) can be expressed is a similar way:

$$Central\mathcal{M}_{0,4}[Y_{m,n}(\omega)] = \frac{\left.\frac{\partial^4(Y_{n,m}(\omega))}{\partial\omega^4}\right|_{\omega=0}}{Y_{n,m}(0)} - 4\left(\frac{\left.\frac{\partial^3(Y_{n,m}(\omega))}{\partial\omega^3}\right|_{\omega=0}}{Y_{n,m}(0)}\frac{\left.\frac{\partial(Y_{n,m}(\omega))}{\partial\omega}\right|_{\omega=0}}{Y_{n,m}(0)}\right)$$
$$+ 6\left(\frac{\left.\frac{\partial^2(Y_{n,m}(\omega))}{\partial\omega^2}\right|_{\omega=0}}{Y_{n,m}(0)}\left(\frac{\left.\frac{\partial(Y_{n,m}(\omega))}{\partial\omega}\right|_{\omega=0}}{Y_{n,m}(0)}\right)^2\right) - 3\left(\frac{\left.\frac{\partial(Y_{n,m}(\omega))}{\partial\omega}\right|_{\omega=0}}{Y_{n,m}(0)}\right)^4 \qquad [2.9.17]$$

While these central moment expressions do not appear to have the form of a Taylor expansion, we find that if we consider an expansion of Log(f(a)) we get:

$$Log(f(a)) = Log(f(0)) + a\frac{\left.\frac{\partial f(a)}{\partial a}\right|_{a=0}}{f(0)} + \frac{a^2}{2!}\left(\frac{\left.\frac{\partial^2 f(a)}{\partial a^2}\right|_{a=0}}{f(0)} - \left(\frac{\left.\frac{\partial f(a)}{\partial a}\right|_{a=0}}{f(0)}\right)^2\right)$$
$$+\frac{a^3}{3!}\left(\frac{\left.\frac{\partial^3 f(a)}{\partial a^3}\right|_{a=0}}{f(0)} - 3\left(\frac{\left.\frac{\partial f(a)}{\partial a}\right|_{a=0}}{f(0)}\frac{\left.\frac{\partial^2 f(a)}{\partial a^2}\right|_{a=0}}{f(0)}\right) + 2\left(\frac{\left.\frac{\partial f(a)}{\partial a}\right|_{a=0}}{f(0)}\right)^3\right)$$
$$+\frac{a^4}{4!}\left(\frac{\left.\frac{\partial^4 f(a)}{\partial a^4}\right|_{a=0}}{f(0)} + 12\left(\left(\frac{\left.\frac{\partial f(a)}{\partial a}\right|_{a=0}}{f(0)}\right)^2\frac{\left.\frac{\partial^2 f(a)}{\partial a^2}\right|_{a=0}}{f(0)}\right) - 4\left(\frac{\left.\frac{\partial f(a)}{\partial a}\right|_{a=0}}{f(0)}\frac{\left.\frac{\partial^3 f(a)}{\partial a^3}\right|_{a=0}}{f(0)}\right) - 3\left(\frac{\left.\frac{\partial^2 f(a)}{\partial a^2}\right|_{a=0}}{f(0)}\right)^2 - 6\left(\frac{\left.\frac{\partial f(a)}{\partial a}\right|_{a=0}}{f(0)}\right)^4\right)....$$

[2.9.18]

If we substitute $f(a)$ = Y(ω), the first term of this series is Log(Intensity). The second term corresponds to mean-time, the third to variance and the fourth to skew. However the fifth term does not equal kurtosis. This means that we can only express an approximation of the Log of the Fourier transform as the series of mean-time and the two central moments. The series would be given by:

$$Log(Y(\omega)) \approx Log(Intensity) - i\omega < t > -\frac{\omega^2}{2!}Variance - \frac{i\omega^3}{3!}Skew$$

[2.9.19]

We can compare how this series expansion fits with the true Fourier transform of the TPSF in the same way as with the ordinary moment series shown in 2.9.4.1. Figure 2.9.12 shows the exponential of the series in [2.9.19] (with only the terms up to skew), compared to the true analytical Fourier transform. Despite having fewer terms than the ordinary moment series (Figure 2.9.11), the central moment series fits the Fourier transform well up to ~ 2GHz, and doesn't diverge at all over the range shown (up to 14GHz).

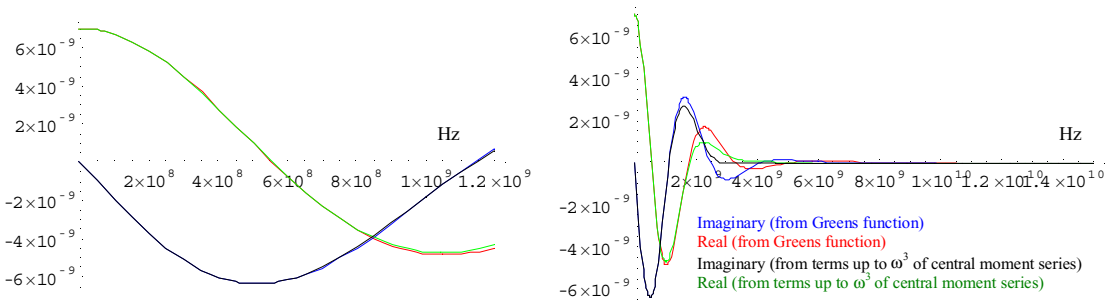

Figure 2.9.12 Plots of the analytical Fourier transform of a TPSF and the Taylor series expansion of the Fourier transform in terms of central moments. (Left) shows the plots up to 1.2 GHz, (Right) shows that the series converges fairly well up to 14GHz. μ_a=0.01 mm^{-1}, μ_s'=1 mm^{-1}, n=1.4, d=70 m.

2.9.4.3 Mellin-Laplace

We can also examine the way in which Mellin-Laplace moments describe the Fourier transform of TPSF. In a similar way to the relation between Y(ω) and temporal moments derived above, we can show that Mellin-Laplace transforms can be described using:

$$\mathcal{M}_{s,u}[Y_{n,m}(\omega)] = \frac{i\left.\frac{\partial^u Y_{n,m}(\omega)}{\partial\omega^u}\right|_{\omega=-is}}{Y(0)}$$

[2.9.20]

So using [2.9.6] with a_0 = - *is* we can write the series:

$$Y(\omega) = Intensity\,(\mathcal{M}_{s,0} + i(\omega+is)\mathcal{M}_{s,1} - \frac{(\omega+is)^2}{2!}\mathcal{M}_{s,2} - i\frac{(\omega+is)^3}{3!}\mathcal{M}_{s,3} + \frac{(\omega+is)^4}{4!}\mathcal{M}_{s,4}.....)$$

[2.9.21]

Which (if *s* is real) has real and imaginary parts:

$$\mathrm{Re}[Y(\omega)] = \mathcal{M}_{0,0}\left(\mathcal{M}_{s,0} - s\mathcal{M}_{s,1} + \frac{(s^2-\omega^2)}{2!}\mathcal{M}_{s,2} - \frac{(s^3-3s\omega^2)}{3!}\mathcal{M}_{s,3} + \frac{(s^4-6s^2\omega^2+\omega^4)}{4!}\mathcal{M}_{s,4}...\right)$$

[2.9.22]

$$\mathrm{Im}[Y(\omega)] = \mathcal{M}_{0,0}\left(\omega\mathcal{M}_{s,1} - \frac{2s\omega}{2!}\mathcal{M}_{s,2} - \frac{(\omega^3-3s^2\omega)}{3!}\mathcal{M}_{s,3} + \frac{4(s\omega^3-s^3\omega)}{4!}\mathcal{M}_{s,4}...\right)$$

[2.9.23]

This is similar to the moments expansion in [2.9.14], but here, different order Mellin Laplace terms appear in both the real and imaginary series (except for Laplace $\mathcal{M}_{s,0}$, which is real). Also, where the moment series was exactly equal to the Fourier transform at $\omega = 0$, although the imaginary part of the Mellin Laplace series is equal to zero, the real part only approximates the real part of the Fourier transform, depending on the convergence for the number of terms and the value of *s*. Note that all these Mellin-Laplace results reduce to those for moments in the special case when *s* = 0.

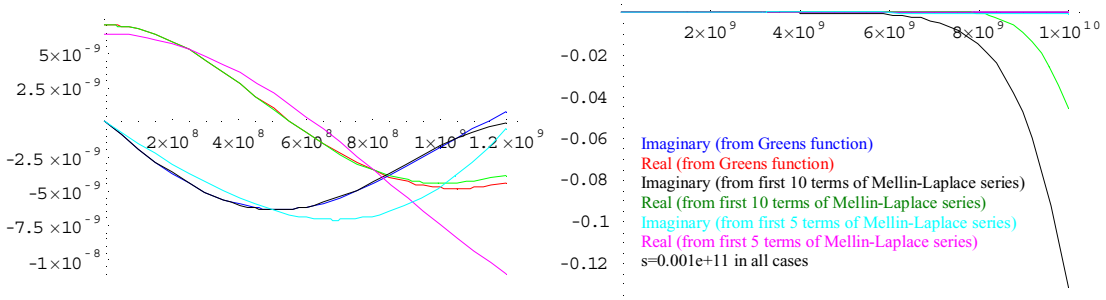

Figure 2.9.13 Plots of the analytical Fourier transform of a TPSF and the Taylor series expansion of the Fourier transform in terms of Mellin-Laplace moments. (Left) shows the plots up to 1.2 GHz, (Right) shows how the series diverge above ~ 6GHz.

2.9.5 Convolution Proofs

The relation between convolved functions when they are subject to certain integral transforms is important for calibration of data from the MONSTIR system. The integral transforms are the operators that yield the common datatypes. The convolved function represents the measured data *y(t)*, which is a convolution of the data that we want *D(t)* and the system IRF *I(v)* which we wish to remove via calibration. (see section 2.1.2).

2.9.5.1 Fourier Domain

Convolution in the Fourier domain is a multiplication:

$$Y_{n,m}(\omega) = F[y_{n,m}(t)] = F[D_{n,m}(t) * I_{n,m}(t)] = F[D_{n,m}(t)] \times F[I_{n,m}(t)] = \mathrm{D}_{n,m}(\omega) \times \mathrm{I}_{n,m}(\omega)$$

[2.9.24]

This can be proved in the following way: The Fourier transform of *y(t)* is given by:

$$F[y(t)] = Y(\omega) = \int_{-\infty}^{\infty} y(t) e^{-i\omega t} dt$$

[2.9.25]

A convolution is given by:

$$y(t) = \sum_{v} D(t-v) I(v)$$

[2.9.26]

Substituting [2.9.26] into [2.9.25] we get:

$$F[D(t) * I(t)] = Y(\omega) = \int_{t} \int_{v} D(t-v) I(v) e^{-i\omega t} dt$$

[2.9.27]

let t = x+v:

$$= \int_{x} \int_{v} D(x) I(v) e^{-i\omega(x+v)} dx dv$$

[2.9.28]

$$= \int_{x} \int_{v} D(x) e^{-i\omega x} I(v) e^{-i\omega v} dx dv = \int_{x} D(x) e^{-i\omega x} dx \times \int_{x} I(v) e^{-i\omega v} dv$$

[2.9.29]

which yields:

$$= F[D(x)) \times F(I(v)] = \mathrm{D}(\omega) \times \mathrm{I}(\omega)$$

[2.9.30]

2.9.5.2 Complex datatypes:

Writing *Y(ω)* in terms of its amplitude and phase, and using [2.9.30] leads to:

$$\begin{aligned}
&\mathrm{Y}_{n,m}(\omega) = Amp[\mathrm{Y}_{n,m}(\omega)] e^{iPhase[\mathrm{Y}_{n,m}(\omega)]} \\
&= \mathrm{D}_{n,m}(\omega) \mathrm{I}_{n,m}(\omega) \\
&= Amp[\mathrm{D}_{n,m}(\omega)] e^{iPhase[\mathrm{D}_{n,m}(\omega)]} Amp[\mathrm{I}_{n,m}(\omega)] e^{iPhase[\mathrm{I}_{n,m}(\omega)]} \\
&= Amp[\mathrm{D}_{n,m}(\omega)] Amp[\mathrm{I}_{n,m}(\omega)] e^{i(Phase[\mathrm{D}_{n,m}(\omega)] + Phase[\mathrm{I}_{n,m}(\omega)]}
\end{aligned}$$

[2.9.31]

Thus we find the relations stated in section 2.1.3.1:

$$Amp[\mathrm{Y}_{n,m}(\omega)] = Amp[\mathrm{D}_{n,m}(\omega)]Amp[\mathrm{I}_{n,m}(\omega)]$$
$$Phase[\mathrm{Y}_{n,m}(\omega)] = Phase[\mathrm{D}_{n,m}(\omega)] + Phase[\mathrm{I}_{n,m}(\omega)]$$

[2.9.32]

2.9.5.3 Intensity

Integrated intensity can be written as:

$$Intensity[y(t)] = \sum_t y(t)$$

[2.9.33]

Substituting [2.9.26] and letting t = x +v

$$\sum_t \sum_v D(t-v)I(v) = \sum_x \sum_v D(x)I(v) = \sum_x D(x) \times \sum_v I(v)$$

[2.9.34]

Leads to

$$Intensity[Y] = Intensity[D] \times Intensity[I].$$

[2.9.35]

2.9.5.4 Mean-time

The proof shown in section 2.1.3.1 for the behaviour of the mean-time of two convolved functions can also be achieved using the differential form of moment operator described in This leads to [2.9.36] which requires differentiation of the product of the frequency domain forms of *D(t)* and *I(t)*.

$$\mathcal{M}_{0,x}(D_{n,m}(\omega) \times I_{n,m}(\omega)) = i^x \frac{\left.\frac{\partial^x (D_{n,m}(\omega) \times I_{n,m}(\omega))}{\partial \omega^x}\right|_{\omega=0}}{D_{n,m}(0) \times I_{n,m}(0)}$$

[2.9.36]

So for mean-time (the first differential) we get:

$$Mean[D_{n,m}(\omega) \times I_{n,m}(\omega)] = i \frac{D_{n,m}(0) \left.\frac{\partial I_{n,m}(\omega)}{\partial \omega}\right|_{\omega=0} + I_{n,m}(0) \left.\frac{\partial D_{n,m}(\omega)}{\partial \omega}\right|_{\omega=0}}{D_{n,m}(0) \times I_{n,m}(0)}$$

[2.9.37]

$$= i \frac{\left.\frac{\partial I_{n,m}(\omega)}{\partial \omega}\right|_{\omega=0}}{I_{n,m}(0)} + i \frac{\left.\frac{\partial D_{n,m}(\omega)}{\partial \omega}\right|_{\omega=0}}{D_{n,m}(0)} = Mean[I] + Mean[D]$$

[2.9.38]

Which is stated in section 2.1.3.1.

2.9.5.5 2nd moment

The (non-central) second moment:

$$2^{nd}\,Moment[y(t)] = \frac{\sum_t y(t)t^2}{\sum_t y(t)}$$

[2.9.39]

Substituting [2.9.26] and letting t = x +v,

$$\frac{\sum_t\sum_v D(t-v)I(v)t^2}{\sum_t\sum_v D(t-v)I(v)} = \frac{\sum_x\sum_v D(x)I(v)(x+v)^2}{\sum_x\sum_v D(x)I(v)}$$

[2.9.40]

$$= \frac{\sum_x D(x)x^2}{\sum_x D(x)} + \frac{\sum_v I(v)v^2}{\sum_v I(v)} + 2\frac{\sum_x D(x)x\sum_v I(v)v}{\sum_x D(x)\sum_v I(v)}$$

[2.9.41]

Which means that for the 2nd temporal moment, the relation between convolved signals is not a simple sum but:

$$2^{nd}\,moment[y] = 2^{nd}\,moment[I] + 2^{nd}\,moment[D] + 2Mean[I]Mean[D]$$

[2.9.42]

2.9.5.6 Variance

The 2nd *central* moment (variance) is defined as:

$$Var[y(t)] = \frac{\sum_t y(t)(t-<t>)^2}{\sum_t y(t)}$$

[2.9.43]

where <t> is the mean-time of y(t). Substituting [2.9.26] and letting t = x +v yields:

$$= \frac{\sum_t\sum_v D(t-v)I(v)(t-<t>)^2}{\sum_t\sum_v D(t-v)I(v)} = \frac{\sum_x\sum_v D(x)I(v)(x-<x>+v-<v>)^2}{\sum_x\sum_v D(x)I(v)}$$

[2.9.44]

$$= \frac{\sum_x D(x)(x-<x>)^2}{\sum_x D(x)} + \frac{\sum_x I(v)(v-<v>)^2}{\sum_x I(v)} + 2\frac{\sum_x D(x)(x-<x>)\sum_v I(v)(v-<v>)}{\sum_x D(x)\sum_v I(v)}$$

[2.9.45]

If we substitute the equation for mean-time into the third term we find it equals:

$$\frac{\sum_x D(x)\left(x-\left(\frac{\sum_x D(x)x}{\sum_x D(x)}\right)\right)\sum_v I(v)\left(v-\left(\frac{\sum_v I(v)v}{\sum_v I(v)}\right)\right)}{\sum_x D(x)\sum_v I(v)}$$

[2.9.46]

$$= \frac{\left(\sum_x D(x)x - \sum_x D(x)\left(\frac{\sum_x D(x)x}{\sum_x D(x)}\right)\right)\left(\sum_v I(v)v - \sum_v I(v)\left(\frac{\sum_v I(v)v}{\sum_v I(v)}\right)\right)}{\sum_x D(x)\sum_v I(v)}$$

[2.9.47]

$$= \frac{\left(\sum_x D(x)x - \sum_x D(x)x\right)\left(\sum_v I(v)v - \sum_v I(v)v\right)}{\sum_x D(x)\sum_v I(v)} = 0$$

[2.9.48]

Thus, as stated in section 2.1.3.1:

$$Var[Y] = Var[I] + Var[D]$$

[2.9.49]

2.9.5.7 Skew

Skew is the third central moment defined as:

$$Skew[y(t)] = \frac{\sum_t y(t)(t - < t >)^3}{\sum_t y(t)}$$

[2.9.50]

where <t> is the mean-time of y(t). Substituting [2.9.26] and letting t = x +v yields:

$$= \frac{\sum_x \sum_v D(x)I(v)(x - < x > + v - < v >)^3}{\sum_x \sum_v D(x)I(v)}$$

[2.9.51]

$$= \frac{\sum_x D(x)(x - < x >)^3}{\sum_x D(x)} + \frac{\sum_x I(v)(v - < v >)^3}{\sum_x I(v)}$$
$$+ 3\frac{\sum_x D(x)(x - < x >)\sum_v I(v)(v - < v >)^2}{\sum_x D(x)\sum_v I(v)}$$
$$+ 3\frac{\sum_x D(x)(x - < x >)^2\sum_v I(v)(v - < v >)}{\sum_x D(x)\sum_v I(v)}$$

[2.9.52]

As with variance, the 3rd and 4th terms = 0 since

$$\sum_x D(x)(x - < x >) = \sum_v I(v)(v - < v >) = 0$$

[2.9.53]

So [2.9.52], as stated in section 2.1.3.1, reduces to:

$$Skew[Y] = Skew[I] + Skew[D]$$

[2.9.54]

2.9.5.8 Kurtosis

Although it is tempting to assume that the simple summing of central moments will carry through to all central moments, this relation breaks down at the 4th central moment, due to an extra non-zero term.

$$Kurtosis[y(t)] = \frac{\sum_t y(t)(t - < t >)^4}{\sum_t y(t)}$$

[2.9.55]

where <t> is the mean-time of y(t). Substituting [2.9.26] and letting t = x +v yields:

$$= \frac{\sum_x \sum_v D(x)I(v)(x - < x > + v - < v >)^4}{\sum_x \sum_v D(x)I(v)}$$

[2.9.56]

$$= \frac{\sum_x D(x)(x - < x >)^4}{\sum_x D(x)} + \frac{\sum_x I(v)(v - < v >)^4}{\sum_x I(v)}$$
$$+ 4\frac{\sum_x D(x)(x - < x >)\sum_v I(v)(v - < v >)^3}{\sum_x D(x)\sum_v I(v)}$$
$$+ 4\frac{\sum_x D(x)(x - < x >)^3\sum_v I(v)(v - < v >)}{\sum_x D(x)\sum_v I(v)}$$
$$+ 6\frac{\sum_x D(x)(x - < x >)^2\sum_v I(v)(v - < v >)^2}{\sum_x D(x)\sum_v I(v)}$$

[2.9.57]

So although the 3rd and 4th terms = 0 (as for variance and skew), the 5th term is non-zero. So that:

$$Kurtosis[Y] = Kurtosis[I] + Kurtosis[D] + 6 \times Var[I]Var[D]$$

[2.9.58]

2.9.5.9 Mellin Laplace 0

The convolution relations for Mellin Laplace can be derived via the differential form of the datatype operator or using the summation manipulation method used above for moments. Laplace is the $u = 0$ case of Mellin Laplace:

$$Laplace[y(t)] = Mellap_0[y(t)] = \frac{\sum_t y(t)e^{-st}}{\sum_t y(t)}$$

[2.9.59]

Substituting [2.9.26] and letting t = x +v yields:

$$= \frac{\sum_x \sum_v D(x)I(v)e^{-s(x+v)}}{\sum_x \sum_v D(x)I(v)} = \frac{\sum_x D(x)e^{-sx}}{\sum_x D(x)} \times \frac{\sum_v I(v)e^{-sv}}{\sum_v I(v)}$$

[2.9.60]

Leading to the result stated in section 2.1.3.1:

$$Lapace[Y] = Laplace[D] \times Laplace[I]$$

[2.9.61]

2.9.5.10 Mellin-Laplace 1

$$Mellap_1[y(t)] = \frac{\sum_t y(t)te^{-st}}{\sum_t y(t)}$$

[2.9.62]

Substituting [2.9.26] and letting t = x +v yields:

$$= \frac{\sum_x \sum_v D(x)I(v)(x+v)e^{-s(x+v)}}{\sum_x \sum_v D(x)I(v)}$$

[2.9.63]

$$= \frac{\sum_x D(x)xe^{-sx} \sum_v I(v)e^{-sv}}{\sum_x \sum_v D(x)I(v)} + \frac{\sum_x D(x)e^{-sx} \sum_v I(v)ve^{-sv}}{\sum_x \sum_v D(x)I(v)}$$

[2.9.64]

$$= \frac{\sum_x D(x)xe^{-sx}}{\sum_x D(x)} \times \frac{\sum_v I(v)e^{-sv}}{\sum_v I(v)} + \frac{\sum_v I(v)ve^{-sv}}{\sum_v I(v)} \times \frac{\sum_x D(x)e^{-sx}}{\sum_x D(x)}$$

[2.9.65]

Leading to:

$$Mellap_1[Y] = Laplace[D]Mellap_1[I] + Laplace[I]Mellap_1[D]$$

[2.9.66]

2.9.6 Values of $Tool_{Abs}$

Section 2.1.2.1 described how measurements made on a number of calibration tools can be combined to calibrate datatypes. All terms in the resulting equation ([2.1.40]) are measurable except for the effect of the absolute calibration tool on each datatype. This extra term is summarised in Table 2.9.1.

Datatype	$\mathcal{M}[Tool_{Abs}]$
Mean	23 ps (7 mm separation in air)
Variance	0 (negligible broadening in tool)
Skew	0 (negligible broadening in tool)
Complex Phase	$23e^{-12}\omega = 0.0023$rad at $\omega = 100$ MHz.
Intensity	(not suitable for absolute calibration)
Laplace	$s \times 23$ ps
Complex amplitude	(not suitable for absolute calibration)

Table 2.9.1The approximate values of the extra term in [2.1.40] which represents the impulse response of the absolute calibration tool for the relevant datatypes

2.9.7 Using averaging to derive an average optical property estimate

The averaging technique described in section 2.1.4.3 can be used to analyse data that is *not* acquired on truly homogenous or axially symmetric objects. The mean properties of an object can be estimated from plots of average mean-time or intensity v/s distance. These may be required to estimate initial conditions for reconstructions. This technique is used to determine starting conditions for image reconstruction of clinical data (see section 2.5.4 and 2.5.5). The analytic form of mean-time (derived from 3D infinite-space Green's functions and [2.9.9]), is given by:

$$< t >_{3D\_\inf} = \frac{\sqrt{3}nd(\mu_a + \mu_s')}{2c\sqrt{\mu_a(\mu_a + \mu_s')}}$$

[2.9.67]

Which means that where M is the gradient of a plot of mean-time v/s distance d:

$$\mu_s' \approx \left(\frac{4M^2c^2}{3n^2} - 1\right)\mu_a$$

[2.9.68]

where c = speed of light and n= refractive index. [2.9.68] hence allows an estimate of μ_a to be made from a set of averaged mean-time values if μ'_s is approximately known (or visa-versa).

Similarly, the analytic form of variance may be used (if well-calibrated):

$$Variance_{3D\_\inf} = \frac{\sqrt{3}n^2d(\mu_a + \mu_s')^2}{4c^2\left(\mu_a(\mu_a + \mu_s')\right)^{3/2}}$$

[2.9.69]

2.9.8 Image manipulation (volume integrals)

When we analyse images, it is tempting to assume that the peak value of the optical properties represents the quantitation. However in images that have a less than perfect resolution, the target objects will be blurred. When TOAST is seeking a solution that produces the same data as MONSTIR, areas of increased optical property outside the boundary of the target feature area will be adding to the simulated data. We could just as reasonably consider the *total* μ_a or μ'_s within the volume of the blurred feature as the retrieved optical property of the feature. Extrapolation of this for evaluation of clinical images is not direct since inclusions may not be discrete, although poor peak value quantitation is always seen, both in reconstructions from real data and simulated data.

Figure 2.9.14 shows a μ'_s image reconstructed on a 2D mesh from real data (the difference in mean-times on the conical breast phantoms see section 2.3.3.1 (Figure 2.3.11). Both features should be the same magnitude. The target values are two cylindrical objects 8 mm in diameter, 10 mm high with a μ'_s of 2 mm^{-1}. The cross sections of the target and image are plotted.

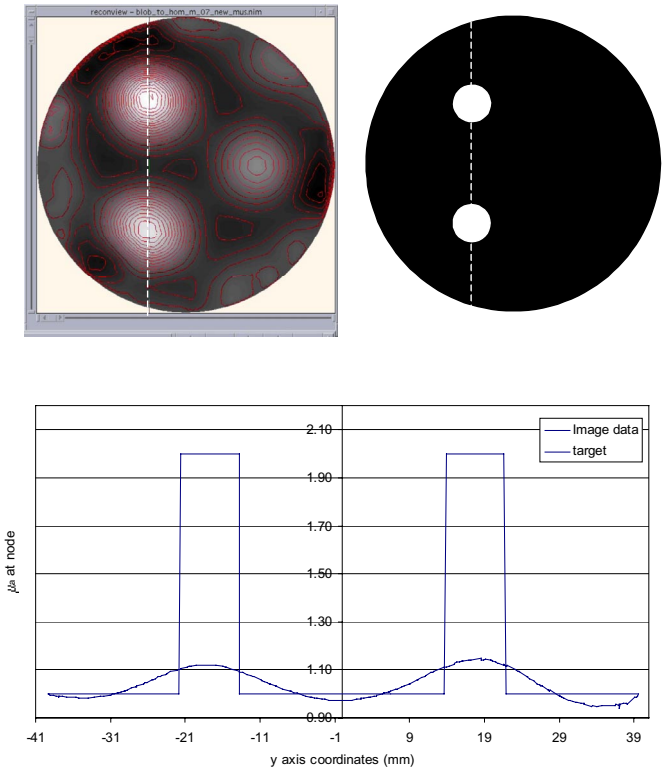

Figure 2.9.14 The μ'_s difference image of the breast phantom (left), shown in section 2.3.3.1 (Figure 2.3.11) should have two features of equal magnitude as in the target image (right). The cross section of both the target and image along the dotted line is plotted.

Clearly the amplitude of the peak value is much lower than it should be. However there is clear blurring. We want a way to quantify the 'total volume μ'_s' of this blurred area to compare it to the target. This can be done simplistically by adding together the values of μ'_s as if each point away from the centre of each feature is a sphere with those optical properties allocated.

$$total = \sum_{y=y_{peak}}^{y=y_{peak}+peak\_width/2} \mu'_s(y) \times 4\pi (y - y_{peak})^2 \qquad [\ 2.9.70\ ]$$

This assumes that both the target and the feature in the image are spherical, which is of course not true, since the target object is a cylinder. However the order of magnitude of the volume integrals calculated using [2.9.70] are more likely to agree between the target and image. For the image and target shown in Figure 2.9.14 we can test the agreement between the volume integrals when clearly there is a big discrepancy in the peak values. Allowing for the possible error in the background optical properties of the phantom, we shall assume that the background is 0.976 mm^{-1}, and thus the target should equal twice this value. The integrals for each inclusion in Figure 2.9.14 have been calculated using [2.9.70], following normalisation to the background, each feature is shown along with its value in Figure 2.9.15.

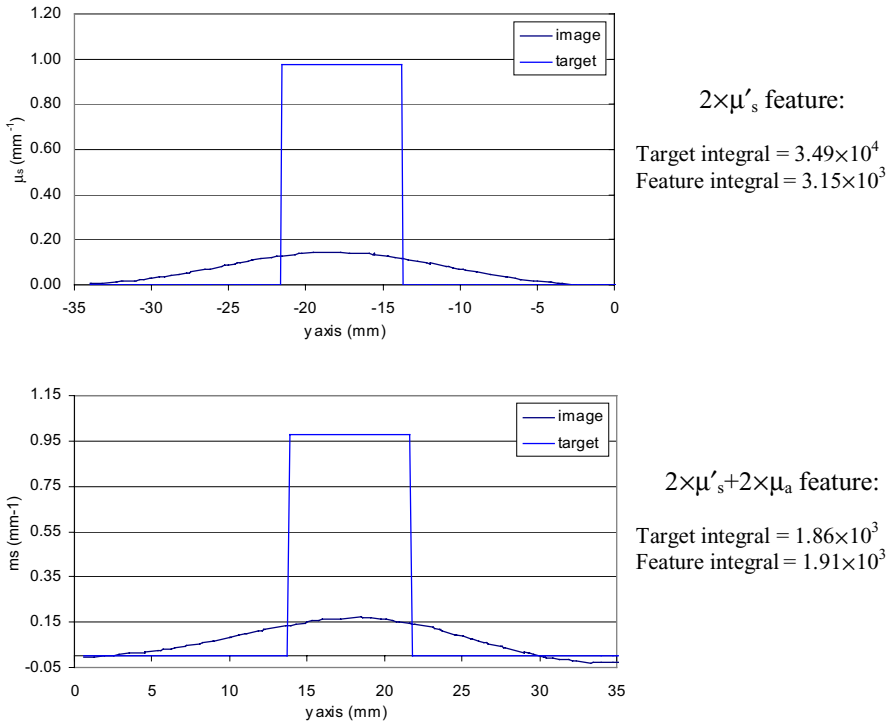

Figure 2.9.15 The two features in Figure 2.9.14, with their 'integrals' assuming they are both spherically symmetric.

The values found for the μ'_s only inclusion (which is slightly smaller than the 'both' feature in the image) has a volume integral that is < 10% that of the target. The 'both μ_a and

μ'_s' feature (which should look the same as the $2\times\mu'_s$ feature), has a volume integral very close to that of the target.

This ad-hoc method at least demonstrates that we shouldn't expect features in images that have poor resolution to have peak values that match those of the target. However further consideration is required to establish the *true* meaning of the values derived, particularly if we are analysing a 2D image. Using this method to calculate true optical property values in a clinical image is inappropriate since the volume over which you integrate must envelope the region which contains a discrete feature. Propagation of errors due to blurring, and their effects on extraction of functional parameters is investigated in chapter 2.5.

2.9.9 Research into techniques for neonatal head fibre-holder design

Chapter 2.7 described the development of methods to hold 64 optical fibres on the head of a neonate or on the adult breast. The background research relating to this development is detailed below.

2.9.9.1 Holding a distribution of *electrodes* on the head

Elecroencephalograpghy (EEG) involves placing large numbers of *electrodes* over the surface of the brain to measure tiny electrical signals from active neurones in the brain. Generally for EEG, a 10-20 arrangement is used. The electrodes are sometimes held by a special net, which stretches to maintain a constant distance between electrodes. Alternatively, glue is sometimes used to affix electrodes to the scalp for extended periods of monitoring. However electrodes are only attached to very fine light wire. Optical fibres are heavy and cannot be held in place reliably using systems designed for EEG electrodes.

2.9.9.2 Neonate positioning

Commercial devices are available for positioning neonates in a variety of positions for comfort and treatment. Figure 2.9.16 shows a “Tucker sling”, “Wedgies” and a “Squishon gel mattress” supplied by EME Ltd. The sling allows an infant to be held securely in a tilted position. Their range of other shaped cushions and gel-filled pads and mattresses (which can be heated) allow the infant to be safely positioned and, to a certain extent, restrained. Many neonatal intensive care units (NICUs) already have this equipment. Also, since the equipment is usually used in the clinical environment it will be designed to withstand the sterilisation measures necessary for use with neonates.

Figure 2.9.16 Commercial products for neonatal positioning (left) Tucker sling, (middle) Wedgie gel filled pads (right) Squishon gel filled mattress). EME (Electro-medical equipment) Ltd, 60 Gladstone Place, Brighton, Sussex, BN2 3QD, ENGLAND Telephone +44 (0) 1273 645100 Fax +44 (0) 1273 645101, http://www.eme-med.co.uk/dev.asp

For both high resolution MRI imaging and radiotherapy restraining the patient is vital. For radiotherapy some form of cast is required to allow the subject to be repeatedly positioned in exactly the same way over a course of weeks to ensure that doses are pinpointed accurately. One device, called a Vak-Lok cushion, is a plastic or rubber bag containing polystyrene

beads. The subject is in position on the device a vacuum pump removes the air from the bag causing it to become rigid (shown in Figure 2.9.17). It will hold its shape until air is let back in. These devices have been used to position neonates for MRI scanning.

Figure 2.9.17 Vak-lok patient positioning system from MED-TEC, Inc - 1401 8th St - Orange City, IA 51041 - 800.842.8688 http://www.medtec.com/products/vacuumcushions/vac_lok.asp.

The availability of such immobilisation devices implies that an infant could be restrained in this way to avoid movement during data acquisition. The whole body could be held steady, or the devices could be used simply to position the head. Similarly another material commonly used in radiotherapy known as thermoplastic could be used for both restraint and for taking a cast of a shape. Figure 2.9.18 shows how this material is used to position a patient's head.

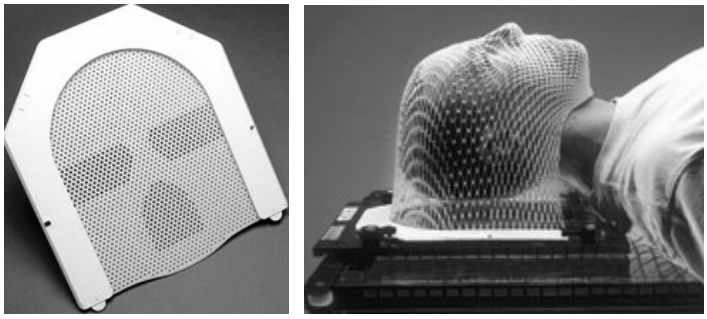

Figure 2.9.18 Thermoplastic for moulding the shape of the head and providing repeatable positioning. Available from MED-TEC (see Figure 2.9.17) http://www.medtec.com/products/headandneck/thermoplastic/comfort_perf.asp and also from WFR/Aquaplast Corporation - 30 Lawlins Park - Wyckoff, NJ 07481, Telephone: (201) 891-1042, Fax: (201) 891-2329, http://www.wfr-aquaplast.com/

Thermoplastic is available in different thicknesses and colours, as well as being either perforated or not. It needs to be heated to the manufacturer's specified required temperature in a water bath. It becomes flexible and remains flexible for a short while once cool. The perforated type stretches over 3D contours better than the plain type.

Measurements of the head are required to allow mesh generation and potentially to allow custom-made fibre-holders to be manufactured. In some cases a mould can be made using thermoplastic, although it may not be suitable for use on neonates.

2.9.9.3 Surface mapping options

Methods to measure the shape of the head for mesh generation, determination of optode placement and fibre-holder design are required. The infant's head can be physically measured using a 'tile shaping template' as shown in Figure 2.9.19.

Figure 2.9.19 A tile shaping template used for early head measurements, conforms to shape when pushed onto an object.

Prima-facie is a laser-scanning wand that uses laser triangulation to determine a 3D contour of an object for applications including radiation treatment planning (Figure 2.9.20).

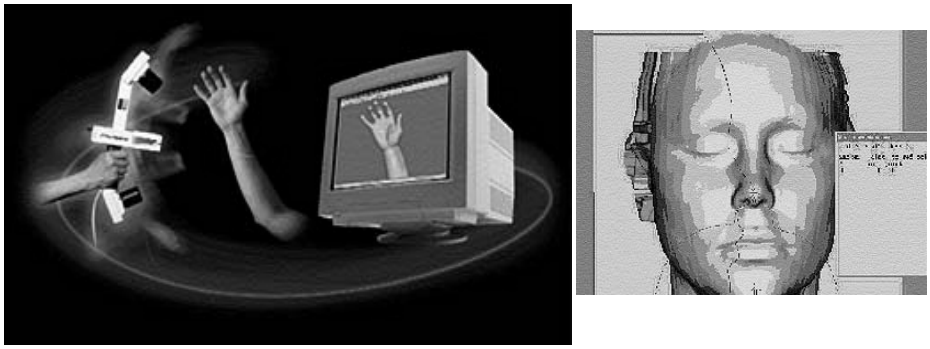

Figure 2.9.20 Prima-facie based on Polhemus' FastSCAN 3D Handheld Laser Scanner from Duncan Hynd Associates Limited, First floor offices, 53 High Street, Chinnor, Oxon, OX9 4DJ, UK http://www.dha.co.uk/3d_imaging.html.

The Prima-facie system can also be attached to a machining system that will create a scale 3D model of the scanned object. Similar results can be obtained via processing CT or MRI images of the subject. If MRI or CT images of a neonate were available, thresholding of the 3D image set can allow surface contours to be extracted.

Stereoscopic photography could also be used in a similar way to laser triangulation to derive 3D measurements of the neonate's head non-invasively. However currently problems with this technique involve the need for texture on the head that can be co-registered in images taken from different directions. Hair on the head of infants does not provide sufficient texture. A stereoscopic photography based system for mapping neonatal heads is currently being investigated.

Finally, a slightly different method of acquiring 3D co-ordinates of an object is to use a 'Microscribe' 3D pointing arm. This device has multiple joints containing transducers such that when the fine point of the arm is positioned on the surface of an object and the foot pedal

pressed, the x, y, z co-ordinates of the point are recorded. This device could be used to record the 3D positions of each optode, and potentially be used to map the contours of the whole head (see Figure 2.9.21).

Figure 2.9.21 ***Microscribe 3D co-ordinate recording arm used to map the surface of a head. From:*** *Immersion Medical, 55 West Watkins Mill Rd, Gaithersburg, Maryland 20878 USA, Telephone: 301.984.3706, Fax: 301.984.2104, http://www.immersion.com/products/3d/overview.shtml. Graphics from: http://www.rsi.gmbh.de/ms-max_e.htm*

2.9.9.4 Laser safety

Laser safety limits depend on the grade of laser, the pulse rate and wavelength and are evaluated in Watts mm-2 (see (Schmidt, 1999)). For eye-exposure in neonates we must consider that they have a longer focal length than adults. Fibre holder design and illumination power must be sufficiently low that eye damage can not occur, or some form of eye protection is required. Commercially available eye protection for use when an infant is undergoing phototherapy (for jaundice) is a possible option (see Figure 2.9.22).

Figure 2.9.22 'Wee-specs' available for premature to term infants for phototherapy UV eye protection. (EME Ltd (seeFigure 2.9.16))

Light exclusion between sources and detectors can be quite readily achieved using adhesive black polyuretheane foam “sound insulation” strips (available from http://rswww.com). Found to be very absorbing in the NIR and non-irritant, rolls of this material can be sterilised and used to line fibre holders. The foam is available in different grades of compressibility, but in general can conform to a shape while supporting weight.

2.9.10 Research into techniques for breast imaging fibre-holder design

2.9.10.1 Imaging positioning

Figure 2.9.23 shows a conventional x-ray mammogram. The woman stands and the breast is compressed between two moveable plates. Images can be acquired with either horizontal or vertical compression.

Figure 2.9.23 For an x-ray mammogram, the breast must be compressed between two flat plates.

When breast imaging is performed using an MRI scanner, in general the woman lies on her front with the breasts positioned inside holes in the bed as shown in Figure 2.9.24. RF coils are contained in the holes to allow better resolution.

Figure 2.9.24 Acquiring and MRI of the breast requires the subject to lie prone on the bed with the breasts hanging into holes which house RF coils.

For ultrasound examination of the breast the patient generally lies on her back and the ultrasound transducer is moved over the breast by hand.

Figure 2.9.25 Ultrasound examination of the breast.

As with the neonatal head, radiation therapy technology offers solutions to positioning problems. The thermoplastic described above can be used to conform to the breasts to allow repeatable positioning for radiotherapy (see Figure 2.9.26). In addition a ‘treatment bra’

(Figure 2.9.26, right) is available that claims to be able to hold the breast in a highly repeatable way and is available in a library of 40 cup sizes.

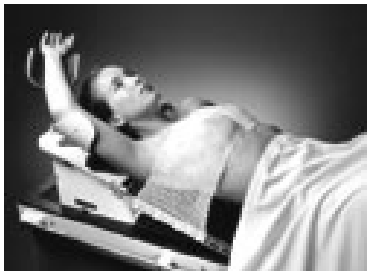

Figure 2.9.26 (right) Breast positioning for radiotherapy using thermoplastic (right) radiotherapy treatment bra. Available from MED-TEC, see Figure 2.9.18. http://www.medtec.com/products/breastandthorax/supports/breast-support-systems.asp#SIDE

The treatment bra shown above highlights the important fact that bra sizes are a fairly well regulated measurement system for breasts. If a set of re-useable fibre holders were to be developed for optical tomography, a subset of the most common bra sizes would probably be sufficient for most women. Bra size could be confirmed prior to attending an imaging session to ensure that the correct size fibre holder were available.

2.9.10.1.1 Contrast agents

Contrast agents are sometimes used in other imaging modalities including MRI of the breast. PET is sometimes used in breast cancer detection, which requires injection of a radioactive isotope. It is possible that using ICG (a NIR contrast agent) could reduce the need for calibration measurements during breast imaging (see section 2.1.2.3).

2.10 APPENDIX B

Papers detailing previous arm imaging results:

Pages 332-334:

Hillman E.M.C, Hebden J.C, Schmidt F.E.W, Arridge S.R , Fry M.E, Schweiger M, Delpy D.T, *"Initial Clinical testing of the UCL 32 channel time-resolved instrument for optical tomography":* in Biomedical Tops, OSA Tech Digest. 100-102 (2000).

Pages 335-347:

Hillman E. M. C, Hebden J. C, Schweiger M, Dehghani H, Schmidt F. E. W, Delpy D. T, Arridge S. R. *"Time Resolved Optical Tomography of the Human Forearm":* Phys. Med. Biol. 46, 4, 1117-1130 (2001).

2.11 Bibliography

Arridge, S R (1995): "Photon-measurement density functions. Part 1: Analytical forms," *Appl. Opt.* **34** (31): p 7395-7409.

Arridge, S R (1999): "Optical tomography in medical imaging," *Inverse Problems* **15** p 41-93.

Arridge, S R, Cope, M and Delpy, D T (1992): "The theoretical basis for the determination of optical pathlengths in tissue: temporal and frequency analysis," *Phys. Med. Biol.* **37** p 1531-1559.

Arridge, S R, Dehghani, H, Schweiger, M and Okada, E (2000a): "The finite element model for the propagation of light in scattering media: A direct method for domains with nonscattering regions," *Med. Phys.* **27** (1): p 252-264.

Arridge, S R, Hebden, J C, Schweiger, M, Schmidt, F E W, Fry, M E, Hillman, E M C, Dehghani, H and Delpy, D T (2000b): "A method for 3D time-resolved optical tomography," *Int. J. Imag. Sys. Tech* **11** p 2-11.

Arridge, S R and Schweiger, M (1993): "The use of multiple data types in time-resolved optical absorption and scattering tomography (TOAST)," *Proc. SPIE.* **2035** p 218-229.

Arridge, S R and Schweiger, M (1995a): "Photon-measurement density functions. Part 2: Finite-element-method calculations," *Appl. Opt.* **34** (34): p 8026-8037.

Arridge, S R and Schweiger, M (1995b): "Direct calculation of the moments of the distribution of photon time-of-flight in tissue with a finite-element method," *Appl. Opt.* **34** (15): p 2683-2687.

Arridge, S R and Schweiger, M (1998a): "A gradient based optimisation scheme for optical tomography," *Opt. Express* **2** (6): p 213-226.

Arridge, S R and Schweiger, M (1998b): "A general framework for iterative reconstruction algorithms in optical tomography, using a finite element method.", in *The IMA volumes in mathematics and its applications: Computational radiology and imaging, therapy and diagnostics*. Edited by F.N. C. Börgers, Springer, p 45-70.

Arridge, S R and W. R. B. Lionheart (1998c): "Nonuniqueness in diffusion-based optical tomography," *Opt. Lett.* **23** p 882 -884.

Barbour, R L, Graber, H L, Pei, Y, Zhong, S and Schmitz, C H (2001): "Optical tomographic imaging of dynamic features of dense-scattering media.," *J. Opt. Soc. Am. A* **18** (12): p 3018-3036.

Barret, H H and Swindell, W (1981):"Radiological Imaging - Detection and Processing,". Vol. 1. New York: Academic Press.

Bigio, I J, Mourant, J R, Boyer, J, Johnson, T and Lacey, J (1996): "Elastic scattering spectroscopy as a diagnostic tool for differentiating pathologies in the gastrointestinal tract: preliminary testing," *J. Biomedical Optics* **1** p 192-199.

Bluestone, A, Abdoulaev, G, Schmitz, C, Barbour, R and Hielsher, A (2001): "Three dimensional optical tomography of hemodynmamics in the human head," *Opt. Express* **9** (6): p 272-286.

Boas, D A (1996): "Diffuse photon probes of structural and dynamical properties of turbid media: theory and biomedical applications", *PhD thesis,* University of Pennsylvania, Department of Physics & Astronomy, Pennsylvania.

Boas, D A, Gaudette, T and Arridge, S R (2001a): "Simultaneous imaging and optode calibration with diffuse optical tomography," *Opt. Express* **8** (5): p 263-270.

Boas, D A, Gaudette, T, Strangman, G, Cheng, X, Marota, J J A and Mandeville, J B (2001b): "The accuracy of near infrared spectroscopy and imaging during focal changes in cerebral hemodynamics," *Neuroimage* **13** p 76-90.

Boas, D A, O'Leary, M A, Chance, B and Yodh, A G (1993a): "Scattering of diffuse photon density waves by spherical inhomogeneities within turbid media: Analytic solution and applications," *Pro. Natl. Acad. Sci. USA.* **91** p 4887-4891.

Boas, D A, O'Leary, M A, Chance, B and Yodh, A G (1993b): "Scattering and wavelength transduction of diffuse photon density waves," *Phys. Rev. E.* **47** (5): p R2999-R3002.

Boas, M L (1983):"Mathematical methods in the physical sciences," 2nd ed: John Wiley and Sons.

Brizel, D M, Scully, S P, Harrelson, J M, Layfield, L J, Bean, J M, Proznitz, L R and Dewhirst, M W (1996): "Tumour oxygenation predicts the likelihood of distant metastases in human soft tissue sarcoma," *Cancer. Res.* **56** p 941-943.

Cerussi, A E, J., A, Bevilacqua, F, Shah, N, Jakubowski, D, Butler, J, Holcombe, R F and Tromberg, B J (2001): "Sources of Absorption and Scattering Contrast for Near-Infrared Optical Mammography," *Acad Radiol* **8** (3): p 211-218.

Chance, B (1966): "Spectrophotometric observations of absorbance changes in the infrared region in suspensions of mitochondria and in submitochondrial particles," *In: The Biochemistry of Copper, Eds: Peisach et al., Academic Press, N.Y (U.S.A.).* .

Cheng, X and Boas, D A (1999): "Systematic diffuse optical image errors resulting from uncertainty in the background optical properties," *Opt. Express* **4** (299-307): .

Cheong, W F, Prahl, S A and Welch, A J (1990): "A Review of the Optical Properties of Biological Tissues," *IEEE J. Quantum Electron.* **26** p 2166-2185.

Choukeife, J E and L'Huillier, J P (1999): "Measurements of scattering effects within tissue-like media at two wavelengths of 632.8nm and 680nm," *Lasers Med. Sci* **14** p 286-296.

Cope, M (1991): "The application of near infrared spectroscopy to non invasive monitoring of cerebral oxygenation in the newborn infant", *PhD thesis,* University College London, Department of Medical Physics and Bioengineering, London.

Cope, M and Delpy, D T (1988): "System for long-term measurement of cerebral blood and tissue oxygenation on newborn infants by near infrared transillumination," *Med. Biol. Eng. Comput.* **26** p 289-294.

Crossman, A R and Neary, D (1995):"Neuroanatomy an illustrated colour text," 1st ed. Edinburgh: Churchill Livingstone.

Culver, J P, Ntziachristos, V, Holboke, M J and Yodh, A G (2001): "Optimization of optode arrangements for diffuse optical tomography: A singular-value analysis," *Optics Letters* **26** p 701-703.

Dehghani, H, Arridge, S R, Schweiger, M and Delpy, D T (2000b): "Optical tomography in the presence of void regions," *J. Opt. Soc. Am. A* **17** (9): p 1659-1670.

Dehghani, H, Arridge, S R, Schweiger, M and Delpy, D T (2000c): "Image reconstruction in the presence of a void region," *OSA Technical Digest, Biomedical Topical Meetings (OSA, Washington DC)* p 351-353.

Dehghani, H, Barber, D C and Basarab-Horwath, I (1999a): "Incorporating a priori anatomical information into image reconstruction in electrical impedance tomography," *Physiol. Meas* **20** p 87-102.

Dehghani, H and Delpy, D T (2000a): "Near infrared spectroscopy of the adult head: effect of scattering and absorbing obstructions in the CSF layer on light distribution in the tissue," *Appl. Opt.* **39** (25): p 4721-4729.

Dehghani, H, Delpy, D T and Arridge, S R (1999b): "Photon migration in non-scattering tissue and the effects on image reconstruction," *Phys. Med. Biol.* **44** p 2897-2906.

Delpy, D T and Cope, M (1997): "Quantification in NIR spectroscopy," *Phil. Trans. R. Soc. Lond. B.* **352** p 649-659.

Dorn, O (2000): "Scattering and absorption transport sensitivity functions for optical tomography," *Optics Express* **7** (13): p 492-506.

Dorn, O (2001): "Shape reconstruction in scattering media with voids using a transport model and level sets," *Canadian Applied Math Quarterly (CAMQ)* .

Duderstadt, J J and Martin, W R (1979):"Transport Theory,". New York, Chichester: Wiley.

Duncan, A, Whitlock, T L, Cope, M and Delpy, D T (1993): "A multiwavelength, wideband, intensity modulated optical spectrometer for near infrared spectroscopy and imaging," *Proc. SPIE.* **1888** p 248-257.

Dunn, A and Boas, D (2000): "Transport-based image reconstruction in turbid media with small source–detector separations," *Optics Letters* **25** (24): p 1777-1779.

Eda, H, Oda, I, Ito, Y, Wada, Y, Tsunazawa, Y and Takada, M (1999): "Multi-channel time-resolved optical tomographic imaging system," *Rev. Sci. Instrum.* **70** (9): p 3595-3602.

Elwell, C E, Springett, R, Hillman, E M C and Delpy, D T (1998): "Oscillations in Cerebral Haemodynamics - Implications for Functional Activation Studies," *Adv. Exp med & Biol* **471** p 57-65.

Fantini, S, Walker, S A, Franceschini, M A, Kaschke, M, Schlag, P M and Moesta, K T (1998a): "Assessment of the size, position and optical properties of breast tumours in vivo by non-invasive optical methods," *Appl. Opt* **37** p 1982-1988.

Fantini, S, Walker, S A, Franceschini, M A, Moesta, K T, Schlag, P M, Kaschke, M and Gratton, E (1998b): "Optical characterization of breast tumors by frequency-domain optical mammography", in *OSA Trends in Optics and Photonics*. Edited by J.G. Fujimoto and M.S. Patterson, OSA, Washington DC, p 143-147.

Feng, S C, Zeng, F and Chance, B (1995): "Photon migration in the presence of a single defect: a perturbation analysis," *Appl. Opt.* **34** (19): p 3826-3837.

Ferrari, M, Giannini, I, Sideri, G and Zanette, E (1985): "Continuous non invasive monitoring of human brain by near infrared spectroscopy," *Adv. Exp. Med. Biol.* **191** p 873-882.

Ferrari, M, Marchis, C D, Giannini, I, Nicola, A, Agostino, R, S. Nodari, S and Bucci, G (1986): "Cerebral blood volume and haemoglobin oxygen saturation monitoring in neonatal brain by near IR spectroscopy," *Adv. Exp. Med. Biol.* **200** p 203-212.

Firbank, M, Elwell, C E, Cooper, C E and Delpy, D T (1998): "Experimental and theoretical comparison of NIR spectroscopy measurements of cerebral hemoglobin changes," *J Appl Physiol* **85** (5): p 1915-1921.

Firbank, M, Hiraoka, M, Essenpreis, M and Delpy, D T (1993): "Measurement of the optical properties of the skull in the wavelength range 650-950 nm," *Phys. Med. Biol.* **38** p 503.

Firbank, M, Oda, M and Delpy, D T (1995a): "An improved design for a stable and reproducible phantom material for use in near-infrared spectroscopy and imaging," *Phys. Med. Biol.* **40** p 955-960.

Firbank, M, Schweiger, M and Delpy, D T (1995b): "Investigation of "light piping" through clear regions of scattering objects," *Proc. SPIE.* **2389** p 167-173.

FitzGerald, M J T (1996):"Neuroanatomy, basic and clinical," 3rd ed. London: W.B. Saunders.

Fox, S I (1999):"Human Physiology," 6th ed. Boston, London: WCB/McGraw-Hill.

Franceschini, M A, Toronov, V, Filiaci, M E, Gratton, E and Fantini, S (2000): "On-line optical imaging of the human brain with 160-ms temporal resolution," *Optics Express.* **6** (3): p 49-57.

Gao, F, Zhao, H and Yamada, Y (2002): "Improvement of Image Quality in Diffuse Optical Tomography by use of Full Time-Resolved Data," *Applied Optics* **41** (4): p 778-791.

Graber, H L, Chang, J, Lubowsky, J, Aronson, R and Barbour, R L (1993): "Near infrared absorption imaging of dense scattering media by steady-state diffusion tomography," *Proc. SPIE.* **1888** p 372-386.

Graber, H L, Zhong, S, Pei, Y, Arif, I, Hira, J and Barbour, R L (2000): "Dynamic imaging of muscle activity by optical tomography," *Biomedical Topical Meetings, OSA Technical Digest (OSA, Washington DC, 2000)* p 407-408.

Grosenick, D, Wabnitz, H, Rinneberg, H H, Moesta, K T and Schlag, P M (1999): "Development of a time-domain optical mammograph and first in vivo applications," *Appl. Opt.* **38** p 2927-2943.

Hale, G M and Querry, M R (1973): "Optical constants of water in the 200nm to 200mm wavelength region," *Appl. Opt* **12** (3): p 555-563.

Hall-Craggs, E C B (1995):"Anatomy as a basis for clinical medicine," 3rd ed: Williams and Wilkins Waverly Europe.

Hebden, J C, Arridge, S R and Schweiger, M (1998): "Investigation of alternative data types for time resolved optical tomography," *OSA Technical Digest, Biomedical Topical Meetings (OSA, Washington DC)* **21** p 162-167.

Hebden, J C, Bland, T, Hillman, E M C, Gibson, A, Everdell, N, Delpy, D T, Arridge, S R and Douek, M (2002a): "Optical tomography pf the breast using a 32-channel time-resolved imager," *OSA Technical Digest, Biomedical Topical Meetings (OSA, Washington DC)* p in press.

Hebden, J C, Gonzalez, F M, Hillman, E M C, Yusoz, R M and Delpy, D T (2002c): "Assessment of an in situ temporal calibration method for time-resolved optical tomography," *J Biomed Opt* p in preparation.

Hebden, J C, Hall, D J, Firbank, M and Delpy, D T (1995): "Time-resolved optical imaging of a solid tissue-equivalent phantom," *Appl. Opt.* **34** (34): p 8038-804.

Hebden, J C, Hillman, E M C, Gibson, A, Everdell, N, Yusoz, R M, Delpy, D T, Arridge, S R, Austin, T and Meek, J H (2002b): "Time-resolved optical tomography of the newborn infant brain: initial clinical results," *OSA Technical Digest, Biomedical Topical Meetings (OSA, Washington DC)* p in press.

Hebden, J C, Schmidt, F E W, Fry, M E, Schweiger, M, Hillman, E M C and Delpy, D T (1999): "Simultaneous reconstruction of absorption and scattering images by multichannel measurement of purely temporal data," *Opt. Lett.* **24** p 534-536.

Hebden, J C, Veenstra, H, Dehghani, H, Hillman, E M C, Schweiger, M, Arridge, S R and Delpy, D T (2001): "Three dimensional time-resolved optical tomography of a conical breast phantom," *Appl. Opt.* **40** (19): p 3278-3287.

Hielscher, A H, Klose, A and Beuthan, J (2000): "Evolution strategies for optical tomographic characterization of homogeneous media," *Opt. Express* **9** p 507-518.

Hielscher, A H, Klose, A D and Hanson, K M (1999): "Gradient-based iterative image reconstruction scheme for time-resolved optical tomography," *IEEE Trans. Med. Imag.* **18** p 262-271.

Hielsher, A H, Alcouffe, R E and Barbour, R L (1998): "Comparison of finite-difference transport and diffusion calculations for photon migration in homogenous and heterogeneous tissues," *Phys. Med. Biol* **43** p 1285-1302.

Hillman, E M C (2001a): "Development of optical tomography techniques for functional imaging of the neonatal brain", *MPhil to PhD transfer thesis,* University College London, Department of Medical Physics and Bioengineering,

Hillman, E M C, Arridge, S R, Hebden, J C and Delpy, D T (2002): "Processing multi-wavelength time-resolved optical tomography data to derive functional parameters," *OSA Technical Digest, Biomedical Topical Meetings (OSA, Washington DC)* p in press.

Hillman, E M C, Dehghani, H, Hebden, J C, Arridge, S R, Schweiger, M and Delpy, D T (2001c): "Differential imaging in heterogeneous media: limitations of linearization assumptions in optical tomography," *Proc SPIE.* **4250** p (in press).

Hillman, E M C, Hebden, J C, Schmidt, F E W, Arridge, S R, Fry, M E, Schweiger, M and Delpy, D T (2000b): "Initial Clinical testing of the UCL 32 channel time-resolved instrument for optical tomography," *OSA Technical Digest, Biomedical Topical Meetings (OSA, Washington DC)* p 100-102.

Hillman, E M C, Hebden, J C, Schmidt, F E W, Arridge, S R, Schweiger, M, Dehghani, H and Delpy, D T (2000a): "Calibration techniques and datatype extraction for time-resolved optical tomography," *Rev. Sci. Instrum.* **71** (9): p 3415-3427.

Hillman, E M C, Hebden, J C, Schweiger, M, Dehghani, H, Schmidt, F E W, Delpy, D T and Arridge., S R (2001b): "Time Resolved Optical Tomography of the Human Forearm," *Phys. Med. Biol.* **46** (4): p In press.

Hintz, S R, Benaron, D A, Houten, J P v, Duckworth, J L, Liu, F W H, Spilman, S D, Stevenson, D K and Cheong, W-F (1998): "Stationary headband for clinical time-of-flight optical imaging at the bedside," *Photochem. Photobiol.* **68** p 361-369.

Hintz, S R, Cheong, W F, Houten, J P v, Stevenson, D K and Benaron, D A (1999): "Bedside imaging of intracranial hemorrhage in the neonate using light: Comparison with ultrasound, computed tomography, and magnetic resonance imaging," *Pediatr. Res.* **45** p 54-59.

Holder, D S (1995): "Design and electrical characteristics of an electrode array for electrical impedance tomography of the female breast," *Innov. Tech. Biol. Med.* **16** p 144-150.

Hollis, V S, Binzoni, T and Delpy, D T (2001): "Noninvasive monitoring of brain tissue temperature by near-infrared spectroscopy," *Proc. SPIE* **4250** p 470-481.

Hoogenraad, J H, van-der-Mark, M B, Colak, S B, 't-Hooft, G W and E. S. van-der-Linden (1997): "First results from the Philips optical mammoscope," *SPIE Proc.* **3194** p 184-190.

Horecker, B L (1943): "The absorption spectra of hemoglobin and its derivatives in the visible and near infra-red regions," *J. Biol. Chem.* **148** p 173-183.

Hueber, D M, Franceschini, M A, Ma, H Y, Zhang, Q, Ballesteros, J R, Fantini, S, Wallace, D, Ntziachristos, V and Chance, B (2001): "Non-invasive and quantitative near-infrared haemoglobin spectrometry in the piglet brain during hypoxic stress, using a frequency-domain multidistance instrument," *Phys. Med. Biol.* **46** p 41-62.

Jacques, S L and Hielscher, A H (1995): "Discriminating the size and density of objects within tissue: frequency-domain versus steady- state measurements," *SPIE Proc.* **2389** p 228-239.

Jobsis, F F (1977): "Non invasive, infrared monitoring of cerebral and myocardial oxygen sufficiency and circulatory parameters," *Science.* **198** p 1264-1267.

Johnston, P G B (1998):"The newborn child," 8th ed. New York: Churchill Livingstone.

Kelnar, C J H, Harvey, D and Simpson, C (1995):"The sick newborn baby," 3rd ed. London, Philadelphia: Baillière Tindall.

Klose, A, Hielscher, A, Hanson, K and Beuthan, J (1998): "Two- and three-dimensional optical tomography of finger joints for diagnostics of rheumatoid arthritis," *Proc. SPIE* **3566** p 151-159.

Kohl, M, Nolte, C, Heekeren, H R, S. Horst, U S, Obrig, H and Villringer, A (1998): "Determination of the wavelength dependence of the differential pathlength factor from near-infrared pulse signals," *Phys. Med. Biol.* **43** p 1771-1782.

Kolehmainen, V, Arridge, S R, Vauhkonen, M and Kaipio, J P (2000): "Simultaneous reconstruction of internal tissue region boundaries and coefficients in optical diffusion tomography," *Phys. Med. Biol.* **45** (11): p 3267-3284.

Kusaka, T, Isobe, K, Nagano, K, Okubo, K, Yasuda, S, Kondo, M, Itoh, S and Onishi, S (2001): "Estimation of regional cerebral blood flow distribution in infants by near-infrared topography using indocyanine green," *Neuroimage* **13** p 944-952.

Madsen, S J, Anderson, E R, Haskell, R C and Tromberg, B J (1995): "High-bandwidth frequency-domain photon migration instrument for clinical use," *SPIE Proc.* **2389** p 257-263.

Maris, M, Gratton, E, Maier, J, Mantulin, W and Chance, B (1994): "Functional near-infrared imaging of deoxygenated hemoglobin during exercise of the finger extensor muscles using the frequency-domain technique," *Bioimaging* **2** p 174-183.

McBride, T M, Pogue, B W, Gerety, E D, Poplack, S T, Österberg, U L and Paulsen, K D (1999): "Spectroscopic diffuse optical tomography for the quantitative assessment of hemoglobin concentration and oxygen saturation in breast tissue," *Appl. Opt.* **38** (25): p 5480-5490.

McBride, T O, Pogue, B W, Jiang, S, Osterberg, U L and Paulsen, K D (in press 2001): "Frequency domain parallel detection system for near-infrared breast imaging," *SPIE Proc.* **4250** .

Meek, J H, Elwell, C E, Khan, M J, Romaya, J, Wyatt, J S, Delpy, D T and Zeki, S (1995): "Regional Changes in Cerebral Haemodynamics as a Result of a Visual Stimulus Measured by Near Infrared Spectroscopy," *Proc. Roy. Soc. London B* **261** p 351-356.

Meek, J H, Firbank, M, Elwell, C E, Atkinson, J, Braddick, O and Wyatt, J S (1998): "Regional haemodynamic responses to visual stimulation in awake infants," *Pediatrics Research* **42** p 840-843.

Merenstein, G B and Gardner, S L (1998):"Handbook of Neonatal Intensive Care," 4th ed. St. Louis, London: Mosby.

Mourant, J R, Bigio, I J, Boyer, J, Johnson, T and Conn, R (1995): "Spectroscopic diagnosis of bladder cancer with elastic light scattering," *Lasers in Surgery and Medicine* **17** p 350-357.

Mourant, J R, Hielscher, A H, Eick, A A, Johnson, T M and Freyer, J P (1998): "Evidence of intrinsic differences in the light scattering properties of tumorigenic and nontumorigenic cells," *Cancer Cytopath* **84** p 366-374.

Ntziachristos, V, Chance, B and Yodh, A G (1999b): "Differential diffuse optical tomography," *Opt. Express* **5** p 230-242.

Ntziachristos, V, Ma, X and Chance, B (1998): "Time-correlated single photon counting imager for simultaneous magnetic resonance and near-infrared mammography," *Rev. Sci. Instrum.* **69** p 4221-4233.

Ntziachristos, V, Ma, X, Yodh, A G and Chance, B (1999a): "Multichannel photon counting instrument for spatially resolved near infrared spectroscopy," *Rev. Sci. Instrum.* **70** (1): p 193-201.

Ntziachristos, V, Ripoll, J, Culver, J P, Yodh, A G, Chance, B and Nieto-Vesperinas, M (2001): "Effect of multiple layers on diffuse optical," *Proc. SPIE* **4250** p 125-129.

Okada, E, Firbank, M, Schweiger, M, Arridge, S R, Cope, M and Delpy, D T (1997): "Theoretical and Experimental investigation of near-infrared light propogation in a model of the adult head," *Appl. Opt* **36** (1): p 21-31.

Okada, E, Firbank, M, Schweiger, M, Arridge, S R, Hebden, J C, Hiraoka, M and Delpy, D T (1995): "Experimental measurements on phantoms and Monte Carlo simulations to evaluate the effect of inhomogeneity on optical pathlength," *Proc. SPIE* **2389** p 174-181.

Okada, E and Delpy, D T (2000): "Effect of scattering of arachnoid trabeculae on light propagation in the adult brain," *OSA Technical Digest, Biomedical Topical Meetings (OSA, Washington DC)* p 256-258.

Okunieff, P, Hoeckel, M, Dunphy, E P, Schlenger, K, Knoop, C and Vaupel, P (1993): "Oxygen tension distributions are sufficient to explain the local response of human breast tumors treated with radiation alone," *Int. J. Radiat. Oncol. Biol. Phys.* **2** p 631-636.

Patterson, M S, Chance, B and Wilson, B C (1989): "Time resolved reflectance and transmittance for the non-invasive measurement of tissue optical properties.," *Appl. Opt.* **28** (12): p 2331-2336.

Patterson, M S, Wilson, B C and Wyman, D R (1991): "The propagation of optical radiation in tissue I. Models of radiation transport and their application," *Lasers. Med. Sci.* **6** p 155-168.

Pei, Y, Graber, H L and Barbour, R L (2001a): "Influence of systematic errors in reference states on image quality and on stability of derived information for DC optical imaging," *Appl. Opt* **40** p in press.

Pei, Y, Graber, H L and Barbour, R L (2001b): "Normalized-constraint algorithm for minimizing inter-parameter crosstalk in DC optical tomography," *Opt. Express* **11** p 97-109.

Pogue, B, McBride, T O, Österberg, U and Paulsen, K (1999): "Comparison of imaging geometries for diffuse optical tomography of tissue," *Opt. Express* **4** (8): p 270 -286.

Pogue, B, Österberg, U, Paulsen, K and McBride, T (1997): "Instrumentation and design of a frequency-domain diffuse optical tomography imager for breast cancer detection," *Opt. Express* **1** (13): p 391-403.

Pogue, B, Patterson, M S, Jiang, H and Paulsen, K (1995): "Initial assessment of a simple system for frequency domain diffuse optical tomography," *Phys. Med. Biol.* **40** p 1709-1729.

Riley, J, Dehghani, H, Schweiger, M, Arridge, S, Ripoll, J and Nieto-Vesperinas, M (2000): "3D optical tomography in the presence of void regions," *Opt. Express* **7** (13): p 462-467.

Riley, J D, Arridge, S R, Chrysanthou, Y, Dehgahani, H, Hillman, E M C and Schweiger, M (2001): "The Radiosity Diffusion Model in 3D," *SPIE Proc* **4431** p 153-164.

Rinneberg, H, Grosenick, D, Wabnitz, H, Danlewski, H, Moesta, K and Schlag, P (1998): "Time-domain optical mammography: results on phantoms, healthy volunteers, and patients," *OSA Technical Digest, Biomedical Topical Meetings (OSA, Washington DC)* **21** p 278-280.

Ripoll, J (2000a): "Light Diffusion in Turbid Media with Biomedical Application", *thesis,* Universidad Autonoma of Madrid (UAM), Department of Condensed Matter Physics, Madrid.

Ripoll, J, N.-Vesperinas, M, Arridge, S R and Dehghani, H (2000b): "Boundary conditions for light propagation in diffusive media with nonscattering regions," *J. Opt. Soc. Am. A* **17** (9): p 1671-1681.

Ripoll, J, Ntziachristos, V, Culver, J P, Pattanayak, D N, Yodh, A G and N.-Vesperinas, M (2001): "Recovery of optical parameters in multiple-layered diffusive media: theory and experiments," *J. Opt. Soc. Am. A* **18** (4): p 821-830.

Schmidt, F E W (1999): "Development of a Time-Resolved Optical Tomography System for Neonatal Brain Imaging", *PhD thesis,* University College London, Department of Medical Physics and Bioengineering, London.

Schmidt, F E W, Fry, M E, Hillman, E M C, Hebden, J C and Delpy, D T (2000a): "A 32-channel time-resolved instrument for medical optical tomography," *Rev. Sci. Instrum.* **71** (1): p 256-265.

Schmidt, F E W, Hebden, J C, Hillman, E M C, Fry, M E, Schweiger, M and Delpy, D T (2000b): "Multiple slice imaging of a tissue-equivalent phantom using time-resolved optical tomography," *Appl. Opt.* **39** p 3380-3387.

Schmitz, C H, *et al.* (2000): "Instrumentation and calibration protocol for imaging dynamic features in dense-scattering media by optical tomography.," *Appl. opt.* **39** p 6466-6486.

Schöberl, J (1997): "NETGEN - An advancing front 2D/3D-mesh generator based on abstract rules.," *Comput. Visual. Sci* **1** p 41-52.

Schweiger, M (1994): "Application of the Finite Element Method in Infrared Image Reconstruction of Scattering Media", *PhD thesis,* University of London, Department of Medical Physics and Bioengineering, University College London, London.

Schweiger, M and Arridge, S R (1995): "Direct calculation of the moments of the distribution of photon time of flight in tissue with a finite-element method," *Appl. Opt.* **34** (15): p 2683-2687.

Schweiger, M and Arridge, S R (1996): "A System for Solving the Forward and Inverse Problems in Optical Spectroscopy and Imaging," *OSA TOPS, Advances in Optical Imaging and Photon Migration* **2** p 263-268.

Schweiger, M and Arridge, S R (1997a): "Direct calculation with a finite-element method of the Laplace transform of the distribution of photon time of flight in tissue," *Appl. Opt.* **36** (34): p 9042-9049.

Schweiger, M and Arridge, S R (1997b): "Optimal data types in optical tomography", in *Lecture Notes in Computer Science 1230 : Information Processing in Medical Imaging*. Edited by J. Duncan, and Gindi, G, Springer-Verlag: Berlin, p 71-84.

Schweiger, M and Arridge, S R (1998): "Comparison of two- and three-dimensional reconstruction methods in optical tomography," *Appl. Opt.* **37** (31): p 7419-7428.

Schweiger, M and Arridge, S R (1999a): "Optical tomographic reconstruction in a complex head model using a priori region boundary information.," *Phys. Med. Biol.* **44** p 2703-2721.

Schweiger, M and Arridge, S R (1999b): "Application of temporal filters to time-resolved data in optical tomography," *Phys. Med. Biol.* **44** p 1699-1717.

Schweiger, M and Arridge, S R (in preparation 2001): "Image reconstruction in optical tomography using local basis functions," .

Sevick, E M, Burch, C L, Frisoli, J K and Lakowicz, J R (1994): "Localization of absorbers in scattering media by use of frequency-domain measurements of time-dependent photon migration," *Appl. Opt.* **33** p 3562-3571.

Simpson, C R, Kohl, M, Essenpreis, M and Cope, M (1998): "Near infrared optical properties of ex-vivo human skin and sub-cutaneous tissues measured using the Monte-Carlo inversion technique," *Phys. Med. Biol.* **43** p 2465-2478.

Springett, R, Wylezinska, M, Cady, E B, Cope, M and Delpy, D T (2000): "Oxygen dependency of cerebral oxidative phosphorylation in new born piglets," *J. Cereb. Blood Flow & Metab.* **20** p 280-289.

Takahasi, S and Yamada, Y (1998): "Simulation of 3D light propagation in a layered head model including a clear CSF layer," *OSA TOPS* **21** p 2-6.

Tromberg, B J, Shah, N, Lanning, R, Cerussi, A, Espinoza, J, Pham, T, Svaasand, L and Butler, J (2000): "non-invasive in-vivo characterisation of breast tumours using photon migration spectroscopy," *Neoplasia* **2** (1/2): p 26-40.

Vaithianathan, T, Tullis, I D C, Meek, J, Austin, T and Delpy, D T (2002): "A portable near infrared spectroscopy system for mapping functional activation in infants," *Human Brain Mapping (HBM2002) Sendai, Japan* (*in press*).

van-der-Zee, P (1992): "Measurement and modelling of the optical properties of human tissue in the near infrared", *PhD thesis,* University College London, Department of Medical Physics and Bioengineering, London.

van-Staveren, H J, Moes, C J M, van-Marle, J, Prahl, S A and van-Gemert, M J C (1991): "Light scattering in intralipid-10% in the wavelength range of 400-1100nm," *Appl. Opt.* **30** (31): p 4507-4514.

Webster, J G (1990):"Electrical Impedence tomography,". Bristol: Adam Hilger.

Weinstein, C (May 6 1994): "Angiogenesis in breast cancer," *NIH guide* **23** (17): .

Woodard, H Q and White, D R (1986): "The composition of body tissues," *Br. J. Radiol.* **59** p 1209-1219.

Xu, Y, Iftimia, N, Jiang, H, Key, L and Bolster, M (2001): "Imaging of in vitro and in vivo bones and joints with continuous-wave diffuse optical tomography," *Optics Express* **8** (7): p 447-451.

Yamashita, Y, Maki, A and Koizumi, H (1996): "Near-infra-red topographic measurement system: Imaging of absorbers localized in a scattering medium," *Rev. Sci. Instrum.* **67** (3): p 730-732.

Zint, C V, Gao, F, Torregrossa, M and Poulet, P (2001): "Near-infrared optical tomography of scattering cylindrical phantoms using time-resolved detection," *Proc. SPIE* **4250** p 109-119.