MULTI-MODAL DIFFUSE OPTICAL TOMOGRAPHY AND BIOLUMINESCENCE TOMOGRAPHY SYSTEM FOR PRECLINICAL IMAGING

by

JAMES A. GUGGENHEIM

A thesis submitted to
University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

PSIBS Doctoral Training Centre
College of Engineering and Physical Sciences
University of Birmingham
February 2014

ABSTRACT

The development, characterisation and testing of a novel all-optical, multi-modal preclinical biomedical imaging system is presented. The system aims to provide a new way of accurately visualising the spatial distribution and activity of molecular structures and processes in small animals by combining 3D bioluminescence tomography (BLT; reconstruction-based 3D imaging of internal bioluminescent reporter distributions), diffuse optical tomography (DOT; reconstruction-based imaging of optical parameter distributions) and optical surface capture techniques. The key principle of the imaging system is to use surface capture results to enhance the accuracy of DOT image reconstruction, and to use the results of both surface capture and DOT to enhance the accuracy of BLT. If it can be done accurately, BLT addresses an important limitation of standard 2D bioluminescence imaging (BLI; in which the light flux on an animal skin surface originating from internal sources is imaged) in that BLI data is corrupted, and analyses are confounded, by the unknown and unaccounted-for light-tissue interactions that affect the signal. By using DOT-derived optical properties these are accounted for in the presented system. Presented phantom experiments show that the developed system can reconstruct luminescent source distributions and optical parameters accurately and that small animal imaging is feasible with the system. Several specifically developed parts of the methodology have highly encouraging results. The implemented surface capture method is shown to be accurate to within approximately 100μm. Multi-source BLT phantom experiments demonstrate that multiple sources can be reconstructed simultaneously with sub-millimetre positional error and single-source studies show that internal luminescent source intensity measurements are stable to within 15% with respect to varying source positions throughout a mouse-sized volume. It is shown that the pioneered DOT-BLT can provide images of higher quantitative accuracy than

BLT alone; in a presented single-view experiment the prior knowledge from DOT increased BLT reconstruction accuracy in terms of source intensity and localisation error, reducing the latter from approximately 5mm to less than 1mm. It is also shown that a developed method of reconstructing light flux on animal skin by accounting for the image formation process can reduce systematically introduced error on total flux measurements to around 3%, improving on standard BLI in which 10-fold variations are shown to be introduced in the same data set.

ACKNOWLEDGEMENTS

I thank my lead supervisor Dr Hamid Dehghani for his invaluable contribution to my studies resulting in this thesis including but not limited to excellent leadership, always taking the time to personally work through problems, challenges and ideas, and always showing a clear interest and dedication to the development of myself as a researcher as well as to the project itself. Indeed doing these things effectively, enthusiastically and memorably! Without these inputs I cannot imagine there would be a thesis here to read. Dr Iain Styles provided much the same and I thank him for his substantial input in his role as second supervisor. It really was the greatest pleasure to be supervised by both of these great academics.

I also thank my third supervisor Professor Jon Frampton for his involvement in the project, particularly with reference to getting things moving on the side of seeking collaborators in the medical school and making facilities available when needed such as the IVIS imaging system for comparative studies within the presented work. These foundations are likely to be very important to the long-term continuation of the project described in this thesis. It was always motivating to meet with Jon and gain his perspectives.

A particular thank you must go to Hector Basevi, for being a great friend and collaborator throughout this project. I feel so incredibly fortunate to have worked alongside someone so intelligent, kind and capable. His presence in the group undoubtedly improved my work no end because he could always suggest a useful idea to pursue or point out the flaw in a poor one. Hector wrote the algorithms used for the extraction of wrapped phase and unwrapped phase, and calculation of point clouds underlying the functionality of the surface capture method presented in chapter 6.

I thank Professor Ela Claridge for the discussions that we had that helped the project, par-

ticularly memorable was the encouragement and discussion when it came to devising the automated geometrric calibration methods presented in appendix A.6. Further to specifics, Ela's general enthusiasm to all things automated, imaging and analysis really were a great inspiration.

I would also like to acknowledge and thank Andy Palmer, Richard Williams, Eric Pitkeathly, Rory Steven, and Mim Dowle without whom I wouldn't know what physics is or what it is for. At least not with the level of appreciation I now have! It was great to be around such intelligent and wonderful people throughout my PhD.

I also acknowledge the remainder of the CS imaging group at Birmingham as well as the remainder of the PSIBS students who were all a great influence and created a stimulating scientific community.

Finally I'd like to thank Professor Mike Hannon and Dr Josephine Bunch (in addition to Hamid once again) for running the PSIBS course, making it happen, accepting me onto the program and helping me through it.

Financially, this work was supported by engineering and physical sciences research council (EPSRC) grant EP/F50053X/1 (funding the PSIBS Doctoral Training Centre), by national institutes of health (NIH) grant RO1CA132750, and in addition by the University of Birmingham Capital Investment Fund (CIF).

PUBLICATIONS ARISING FROM THIS WORK

Journal Papers

The design, development and initial evaluations of the novel imaging system presented in this thesis, were published in the Institute of Physics' Measurement Science and Technology journal[1], the material included the homogeneous BLT experimental phantom results presented in chapter 9 and the first heterogeneous DOT experimental phantom results presented in chapter 8, parts of the literature review in chapter 4 are also taken directly from this paper:

> James A Guggenheim, Hector R A Basevi, Jon Frampton, Iain B Styles, and Hamid Dehghani. Multi-modal molecular diffuse optical tomography system for small animal imaging. *Measurement Science and Technology*, 24(10):105405, 2013.

The novel free-space light mapping method for producing quantitatively accurate flux maps from multi-view images of diffusive surfaces (the subject of chapter 7) was described in an article published in the Journal of the Optical Society of America A:

> James A Guggenheim, Hector R A Basevi, Iain B Styles, Jon Frampton, and Hamid Dehghani. Quantitative surface radiance mapping using multiview images of light-emitting turbid media. *JOSA A*, 30(12):2572–2584, 2013.

The novel multi-view optical surface capture method presented and evaluated in chapter 6 was published in the Optical Society of America's Optics Express journal[3].

Hector R A Basevi, James A Guggenheim, Hamid Dehghani, and Iain B Styles. Simultaneous multiple view high resolution surface geometry acquisition using structured light and mirrors. *Optics Express*, 21(6):7222–7239, 2013.

The compressed sensing inspired conjugate gradient (CSCG) bioluminescence tomography reconstruction algorithm used developed in the course of this work was presented in an article in the Optical Society of America's Biomedical Optics Express journal[4].

Hector R A Basevi, Kenneth M Tichauer, Frederic Leblond, Hamid Dehghani, James A Guggenheim, Robert W Holt, and Iain B Styles. Compressive sensing based reconstruction in bioluminescence tomography improves image resolution and robustness to noise. *Biomedical optics express*, 3(9):2131, 2012.

Conference Proceedings

A conference paper was written on the results of bioluminescence tomography within experimental homogeneous phantoms in the presented imaging system[5]:

James A Guggenheim, Hector R A Basevi, Iain B Styles, Jon Frampton, and Hamid Dehghani. Bioluminescence tomography improves quantitative accuracy for pre-clinical imaging. In *European Conferences on Biomedical Optics*, pages 87990G–87990G. International Society for Optics and Photonics, 2013.

A conference paper was written on the first dual-modality experimental results obtained with a phantom, involving multi-modal hard-prior (homogeneity-assumed) DOT informing BLT[6]:

James A Guggenheim, Hector R Basevi, Iain B Styles, Jon Frampton, and Hamid Dehghani. Multi-view, multi-spectral bioluminescence tomography. In *Biomedical Optics*. Optical Society of America, 2012.

A proceedings paper was produced outlining a new idea for performing optimal wavelength

selection for BLT imaging and subsequent reconstruction most effectively by analysing the information content of wavelength combinations[7]:

> Hector R A Basevi, James A Guggenheim, Hamid Dehghani, and Iain B Styles. Information-theoretic method for wavelength selection in bioluminescence tomography. In *European Conferences on Biomedical Optics*, pages 879909–879909. International Society for Optics and Photonics, 2013.

A conference paper was written describing a new approach under investigation for the low-loss compression of sensitivity matrices as well as data used in BLT reconstruction allowing for considerably less time and resource intensive reconstruction with potentially no loss, possibly even a gain in accuracy[8]:

> Iain B Styles, Hector R A Basevi, James A Guggenheim, and Hamid Dehghani. Random matrix-based dimensionality reduction for bioluminescence tomography reconstruction. In *European Conferences on Biomedical Optics*, pages 87990J–87990J. International Society for Optics and Photonics, 2013.

A proceedings was written early on in the development of the imaging system charting some of the first results obtained with the system for BLT image reconstruction in an experimental mouse phantom[9].

> James A Guggenheim, Hamid Dehghani, Hector R A Basevi, Iain B Styles, and Jon Frampton. Development of a multi-view multi-spectral bioluminescence tomography small animal imaging system. In *European Conferences on Biomedical Optics*, pages 80881K–80881K. International Society for Optics and Photonics, 2011.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

Bioluminescence imaging (BLI) is a very popular preclinical imaging modality. It represents one method of visualising a range of biological phenomena as they occur *in vivo* in small animals. The animals used, typically mice, are modified such that some internal process or structure of interest is labelled with a bioluminescent reporter. Light is therefore emitted inside the animal when/where the target process occurs or structure is present and can be captured as it leaves the body. This typically involves imaging the light flux on the animal skin surface, in a non-contact geometry, with a sensitive camera whose field-of-view contains the whole animal. Measurements are made on a macroscopic scale but the signal origins have molecular or cellular specificity.

BLI has found a wide range of applications within various subdomains of biomedical science. Examples of applications include imaging cancer spread and tumour development over time and in response to therapies[10, 11, 12, 13], visualising immune cell responses and trafficking behaviour[14, 15], and investigating stem cell trafficking and differentiation[16]. Particularly in cases where newly proposed drugs and therapies are applied, these studies represent an essential part of the transitional stage between *in vitro* biology and human *in vivo* studies[17].

In the ideal case, the aim of the technique is to quantitatively measure and visualise effectively the spatial distribution of internal light sources. Adopting a working example of imaging cancer cells, quantitative measurements could be used for quantifying tumour burden and the spatial resolving power for identifying growth patterns and metastases. These features could be

investigated over time as a function of applied therapy, providing insight as to what therapies might work well in-man.

In pursuit of this, bioluminescence images are often used to estimate reporter concentrations and spatial distributions. However, a major limitation is encountered in that the presence of unknown and unaccounted for light-tissue interactions (strong absorption and scattering) prevents accurate knowledge of light source depth and position. This severely reduces the usefulness of the technique in general and can lead to inaccurate analyses[18].

Bioluminescence tomography (BLT) has developed as a technology whose aim is to address this limitation and reconstruct 3D images showing directly quantitative 3D distributions of the light source inside the animal. This necessitates knowledge of, and compensation for, light-tissue interactions and the precise geometry of the animal and imaging system.

In most BLT studies to date, both of these properties, the impact and extent of optical interactions through the tissue and the shape of the animal, are assumed to be known or are estimated coarsely. Particularly in the case of the optical interactions, there have been relatively few attempts to characterise and compensate for these in-situ in imaging systems or experiments.

In the presented work, a novel multi-modal, all-optical imaging system is presented which uses several modes of optical imaging to measure the shape, optical properties and bioluminescent source distributions throughout an animal in 3D. The specific optical techniques employed are a novel form of multi-view optical surface capture (which measures the animal shape), continuous-wave, multi-spectral diffuse optical tomography (DOT; which reconstructs optical properties), and multi-spectral BLT. The system uses a simple mirror-based multi-view imaging scheme that requires no moving parts.

The key underlying aim is to perform accurate BLT based on the knowledge of the animal shape and optical properties provided by the system. Additionally, the application of multi-spectral DOT, a technique more commonly applied to diagnostic breast imaging and functional brain imaging, within the small animal imaging domain provides new opportunities for exploring the scientific value of quantitative images of optical properties and key absorbers in animal studies. For example it is postulated that functional imaging of blood oxygenation could be

obtained and used to quantify angiogenesis in cancer studies. This could clearly be highly complementary to BLT and the new system can be used to explore such multi-modal optical studies.

The key achievements of the presented work revolve around the design, development, and characterisation of the new imaging system, and the devising of the novel data processing work-flow for effectively combining the obtained multi-modal data. In pursuit of this several particularly noteworthy achievements have been made.

For instance, a new surface capture method was applied, evaluated and shown to be effective at capturing surfaces of mice and phantoms with an accuracy of approximately 100μm. The method uses two mirrors, and two camera-projector pairs, to simultaneously capture surfaces from multiple perspectives. This paradigm allows a larger amount of the imaged surface to be captured than standard approaches.

Using captured surfaces, a novel approach was developed and evaluated to quantitatively recover the amount of light leaving the animal surface from acquired images. This involves considering the relative positions and perspectives of the camera and object and accounting for other system response effects. It was found that with this approach estimates of phantom surface fluxes were stabilised to within 3% with respect to arbitrary surface rotations and positions. This was shown to be a marked improvement over estimates made directly from images, as in standard BLI, where the variations were ten-fold.

The completion of tomographic imaging work-flows was achieved by using surface capture data to build geometrically accurate models in an automated fashion, and using these in image reconstruction, along with surface flux data. Image reconstructions were performed for BLT using a new algorithm developed during the course of the project. Those for DOT use established algorithms for DOT image reconstructions, applied in the new domain. Owing to these processing features, and the use of the new system data and imaging scheme, experiments and tests that are performed are novel.

Several such experiments are presented herein, many involving phantoms and one involving mice (*ex vivo*). BLT phantom experiments demonstrate that sufficient spatial sensitivity is

achieved throughout the imaged volume to allow reconstruction of compact luminescent source distributions at arbitrary positions; in the presented case of a known homogeneous phantom a value of the source strength is recovered that is robust to source position and depth for a range of positions to within 15%.

DOT phantom experiments show that anomalies in terms of optical properties and key absorber concentrations can be visualised in different heterogeneous phantoms. A multi-modal study then demonstrates the full work-flow of the system; spectral DOT is used to reconstruct chromophore concentration in a heterogeneous phantom and thereby provide a refined estimate of optical properties, these are then used to reconstruct the luminescent source distribution also present within the phantom. It was found that quantitative and qualitative improvements were made to the reconstructed luminescence images when the DOT data was incorporated.

The single presented small animal experiment stands as a proof of principle for achieving adequate signals in imaging mice in the system. Results of imaging in BLI mode in the developed system are found to be comparable to those obtained using a commercial BLI system and it is demonstrated that trans-illumination DOT data can be obtained in mice and that surface capture also works convincingly.

The presented work is highly inter-disciplinary as it addresses a specific biomedical need, that of the visualisation of bioluminescence and optical contrast *in vivo*, with solutions and approaches drawn from physical science and computer science. For example, the novel system design, construction, and characterisation would traditionally be seen as engineering and physics-based problems, whilst the data processing, such as the development, integration and utility of image reconstruction algorithms involve advance computing techniques.

The thesis is organised as follows. In the following three chapters, background information is provided in order to motivate, contextualise and inform the presented work. Chapter 2 discusses light-tissue interactions and models for simulating light transport. This builds to a discussion of optical image reconstructions fundamentals. A discussion of optical properties then leads into a clear point of motivation for the current work, namely that optical properties are hard to predict but vital to performing accurate BLT. Chapter 3 reviews preclinical small

animal imaging techniques with additional background on BLI. It is shown that many complimentary imaging techniques exist with distinct strengths and weakness and that whilst BLI is very popular and has many strengths, the biggest challenge is that of compensating for light-tissue interactions to achieve greater accuracy. The final background chapter, chapter 4 is a review of the recent developments and current state-of-the art in BLT, this justifies and informs the particular approaches taken in the current work. The next chapter (chapter 5) discusses the design and construction of the developed imaging system. Note that extensive characterisation of the system and its components is presented in appendix A.

The chapters that follow are concerned with experimental results and methods. Chapter 6 presents the implemented surface capture and automated meshing methods, and chapter 7 discusses the developed surface flux mapping technique. DOT and BLT/combined experiments are presented in chapter 8 and chapter 9 respectively, and chapter 10 presents the small animal imaging experiment. Finally the thesis is concluded in chapter 11.

CHAPTER 2

DIFFUSE OPTICS IN BIOLOGICAL TISSUES

This chapter introduces the key interactions of light with tissue which forms the basis for the interrogation of tissue using light. This is essential background to understanding the mechanisms by which the system presented in this thesis aims to perform functional imaging in terms of estimating the spatial distribution of molecular absorbing agents, as well as functional properties such as blood oxygenation fraction. Understanding these interactions is also vital to enabling the modelling of light propagation in tissue as is done in the present work to facilitate optical image reconstruction methods.

This chapter first discusses absorption and scattering phenomena, then key mechanisms that produce light inside tissue. Transport models are then discussed along with notes on their implementation and an example of modelling light propagation from an internal source to a tissue boundary is shown. The bioluminescence tomography image reconstruction problem is then introduced with some discussion of theory and this leads into a similar discussion of diffuse optical tomography. To close the chapter the optical properties of small animals are discussed including expected levels of light absorption, scattering and auto-fluorescent background in different tissues.

2.1 Interactions of Light and Tissue

2.1.1 Absorption

Absorption of light by a molecule can occur when the photon energy matches that required to raise an electron from one discrete energy level to another, providing that the matching electronic transition is permitted. When this occurs, termed *excitation*, the energy of the photon is transferred to the molecule and the photon ceases to exist. The molecule in its excited state will then return through heat-loss or photon-release (fluorescence or phosphorescence) to its ground state[19]. The fact that the electron energy step must match the energy of the light indicates that the absorption properties of a medium are dependent on both the wavelength of the light (because this is inextricably linked to the energy) and the particular molecule and its chemical properties.

The amount of optical absorption that occurs in a biological tissue is characterised by the absorption coefficient, μ_a. This coefficient is defined as the probability of a photon being absorbed per unit path length[19] which is the inverse of the absorption *mean free path*, thus a higher value of μ_a indicates a higher probability of an absorption event occurring. For light propagating through an absorbing-only medium, the exiting light is related to the incident light by the Lambert-Bouger law[20]

$$I(x) = I_0 e^{-\mu_a x}, \tag{2.1}$$

or the related Beer-Lambert law[20]

$$I(x) = I_0 e^{-\varepsilon C x}, \tag{2.2}$$

in which I_0 is the initial light intensity, and $I(x)$ is the intensity after travelling distance x in the medium. The terms ε and C indicate the extinction coefficient per unit concentration and the concentration of the absorbing agent present in the medium respectively. Figure 2.1 shows the transmittance, $I(x)/I_0$, determined using the first of the above laws for different absorption coefficients.

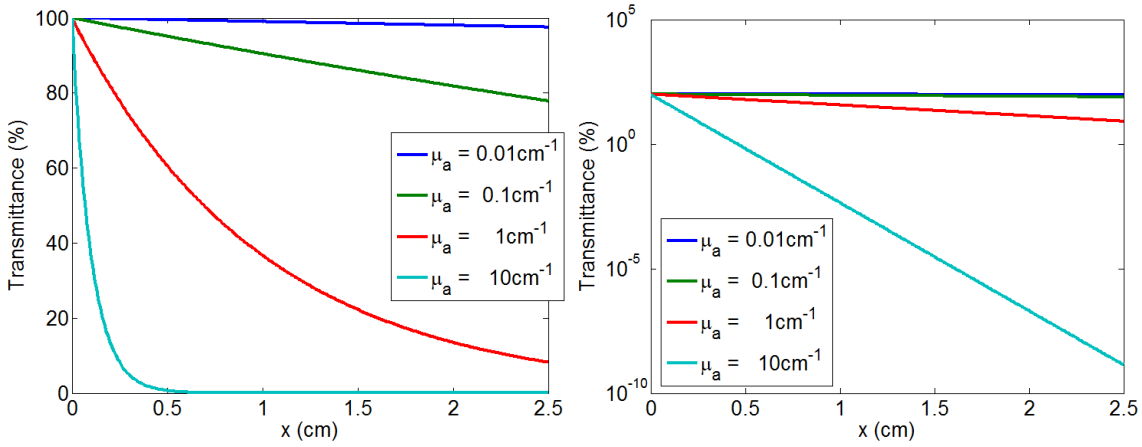

Figure 2.1: Transmittance according to the Lambert-Bouger law in an absorbing-only medium on (a) a linear scale and (b) a log-linear scale.

It can be seen that the effect of absorption is to decrease the light intensity in an exponential decay as a function of the distance travelled. In the case where μ_a is 0.1cm^{-1}, representative of tissue in the near infrared region (section 2.7.1), if absorption were the only interaction then the transmittance over a distance of 2cm (approximately the thickness of a mouse torso) would be around 80%.

The value of μ_a is strongly dependent on the wavelength of the light and the molecular composition of the tissue in terms of the amount of key absorbing agents. The key biological absorbing agents in the visible and near-infrared spectral ranges (400 to 1100nm) are water, haemoglobin (Hb), melanin and lipids. Absorption spectra for these influential chromophores are shown in figure 2.2.

All of the information in this graph is based on data compiled by the Oregon Medical Laser Center (OMLC)[21]. The fat data was first presented by van Veen et al.[23] and is based on measurements made on pig lard, whilst the water measurements were made by Hale et al.[22]. It should be noted that the values shown for melanin could vary considerably in terms of absolute absorption (approximately 10-fold) owing to the varying numbers of melanocytes that might be present[21]. The values shown for fully oxygenated haemoglobin are calculated from extinction coefficients (see below) of oxyhaemoglobin under the assumption that there is a typical 150g/L haemoglobin concentration in blood under consideration[21].

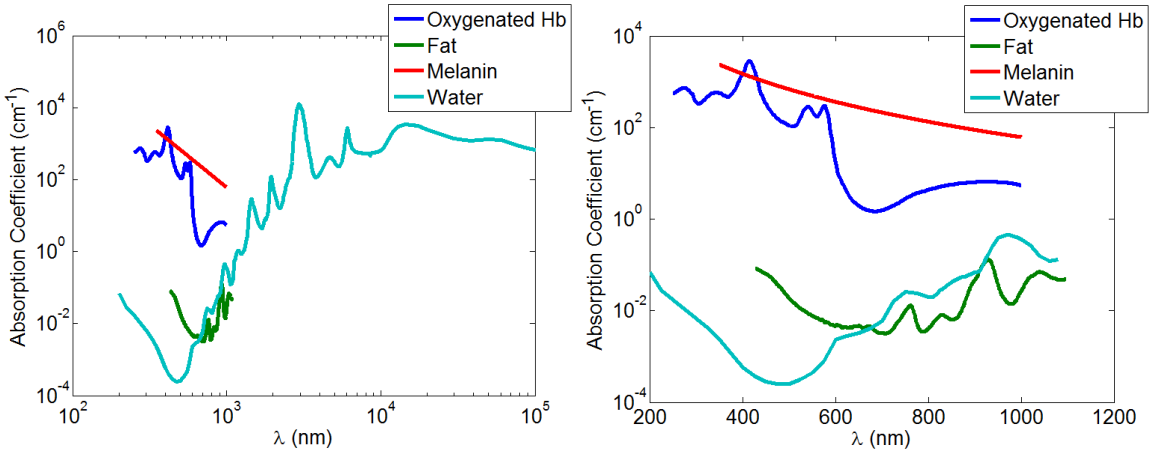

Figure 2.2: Absorption spectra for key biological absorbing agents (a) on a log-log scale for a large spectral range and (b) on a log-linear scale the VIS-NIR region of major relevance to the current work. Data is taken from a compilation by Prahl and Jacques with the Oregon Medical Laser Center (OMLC)[21]. The values for oxyhaemoglobin assume a concentration of 150g/L, typical in oxygenated blood[21]. The original source of the water data was Hale and Querry[22], and that of the fat data was van Veen et al.[23]. The melanin spectrum is an empirical fit to sparsely available data for skin melanin by Jacques[24].

The total absorption coefficient of a tissue is related to the molar concentration of present absorbing agents by

$$\mu_a(\lambda) = \sum_{i=1}^{n} \varepsilon_i(\lambda) C_i \tag{2.3}$$

where C_i is the concentration of the i^{th} chromophore (up to a total of n) in the tissue and ε is the spectral *extinction* coefficient which measures attenuation per unit molar concentration[1].

Haemoglobin is a particularly important absorber being present in most tissues, strongly absorbing in the VIS-NIR range, and a significant target for medical and biomedical observations and functional imaging. Figure 2.3 shows the spectral extinction coefficient of haemoglobin in its oxygenated and deoxygenated forms.

It can be seen that the two curves are somewhat similar in shape but clearly distinguishable across the spectrum. This is important for functional optical imaging and spectroscopy for

[1]The extinction coefficient actually indicates the probability of photon interaction (including scattering (section 2.1.2)) but for certain absorption-dominated molecules and tissue components it can be treated as representative of absorption only[19]

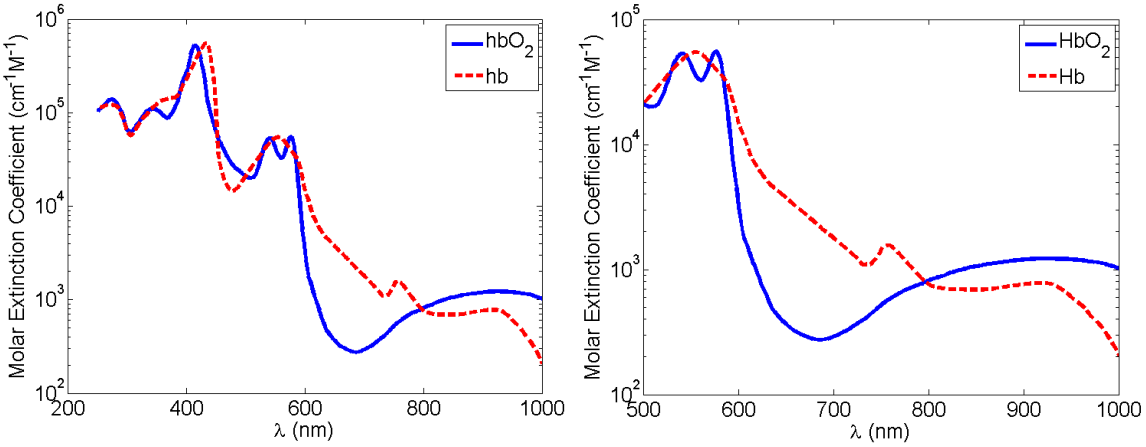

(a) (b)

Figure 2.3: Molar extinction coefficients for oxyhaemoglobin and deoxyhaemoglobin (data compiled by the Oregon Medical Laser Center (OMLC)[21]).

various applications in which the following entities, among others, are important[19].

$$C_{Hb} = C_{ox} + C_{de}, \tag{2.4}$$

and

$$SO_2 = \frac{C_{ox}}{C_{Hb}}, \tag{2.5}$$

where C_{ox} and C_{de} are the molar concentrations of oxygenated and deoxygenated haemoglobin, C_{Hb} is the total concentration of haemoglobin and SO_2 is the oxygen saturation. These are important because they can be indicative of biological processes, normal physiology or disease states[25, 19, 26].

Given the above relationship between absorption coefficient and extinction coefficients and concentrations, it is clear that if the molecular chromophores present in a tissue are known along with their respective extinction coefficients, one could solve a system of linear equations to *unmix* the concentrations of the chromophores and further estimate the total haemoglobin content and oxygen saturation. This is the key idea underlying spectral diffuse optical tomography (section 2.6).

2.1.2 Scattering

Elastic light scattering in tissue describes the event whereby a photon interacts with a biological structure and changes direction without transferring any energy. Photons are most strongly scattered in this way by particles that have a size close to the optical wavelength and a refractive index mismatch with the background medium[19]; for optical wavelengths examples of key scatterers are mitochondria and cell nuclei[27].

Like absorption, scattering can be characterised by a single coefficient. The scattering coefficient μ_s is the probability per unit distance of a scattering event occurring as light travels through the medium (the inverse of the scattering mean free path). If the absorption coefficient in eq. (2.1) is replaced by the scattering coefficient then

$$I_b(x) = I_0^{-\mu_s x}, \tag{2.6}$$

is obtained in which $I_b(x)$ represents the non-scattered or *ballistic* transmitted light. Given that a low-end value for NIR light scattering in biological tissue is 10cm^{-1} (fig. 2.4), and with reference back to fig. 2.1, it can be deduced that there is typically less than 1% ballistic light remaining after a traversal of 5mm and practically zero ballistic light after a traversal of 2cm (the approximate thickness of a mouse torso).

Scattering properties of a medium are further characterised by the scattering phase function $f(\theta)$, the normalised version of which describes the proportional scattered light intensity as a function of angle[20]. From this can be obtained the anisotropy factor, g, which is the mean cosine of the scattered light angular distribution. Values for g range from -1 (entirely back-scattered light), through 0 (isotropic scattering; equal probability for all directions), to 1 (completely forward scattered light)[28].

The reduced scattering coefficient or transport scattering coefficient, μ_s', is defined as[20]

$$\mu_s' = \mu_s(1-g). \tag{2.7}$$

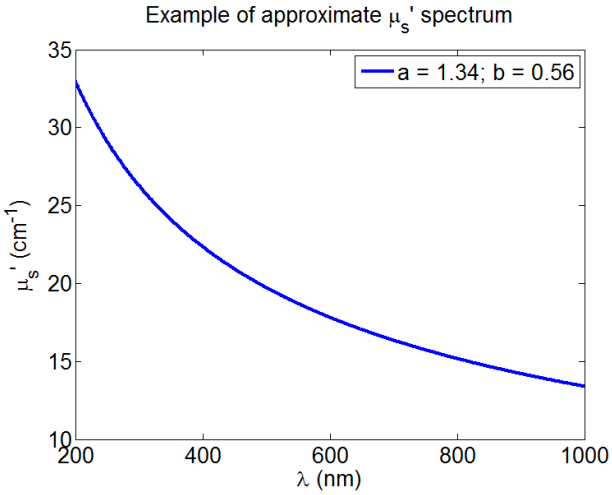

Figure 2.4: μ_s' spectrum calculated using eq. (2.8) with values for a and b that are "typical of breast adipose tissue"[25]

Far from the light source, the distribution of light in a scattering medium with anisotropy factor g (typically around 0.9 in tissue[28]) and scattering coefficient μ_s is identical to the distribution in an isotropically scattering medium with a scattering coefficient of μ_s' [20].

Under the assumptions that scattering events occur independently and only via interaction with individual spherical particles, the Mie solution to Maxwell's equations accurately models the probability of interaction[19]. Considering scattering in tissues, the reduced scattering coefficient in the spectral range of interest (the VIS-NIR) fits well to an empirical approximation to Mie scattering[25, 29, 30, 31];

$$\mu_s' = a\lambda^{-b}, \tag{2.8}$$

with the wavelength λ being expressed in μm and μ_s' in units of mm^{-1}. In this formulation a is referred to as the scatter amplitude and b as the scatter power. In general, under this approximation, larger scatterers give rise to lower a and b values[25].

Figure 2.4 shows the reduced scattering coefficient spectrum calculated using eq. (2.8) for values of scatter power and amplitude that are "typical of breast adipose tissue"[25] as an example. It can be seen that the scattering reduces with wavelength and that around 15 scattering events can be expected per centimetre in the VIS-NIR region.

2.2 Sources of Light Inside Tissue

Whilst in many circumstances light is put into tissue from external sources, there are two mechanisms namely fluorescence and bioluminescence[1] that are commonly utilised in biomedical imaging by which light is generated from inside tissue. These are now discussed.

2.2.1 Fluorescence

Fluorescence is a phenomenon whereby molecular photon absorption (causing excitation; section 2.1.1) is followed by release of energy in the form of another photon (returning the electron to the ground state). Between the absorption (which takes place in a time period on the order of femtoseconds (10^{-15}s)) and the fluorescence emission (which continues over a time period on the order of nanoseconds (10^{-9}s)), there is first a vibrational relaxation where the electron loses some energy (this lasts for a time period on the order of picoseconds (10^{-12}s))[19], as such the ejected photon is always lower energy (higher wavelength) than the exciting photon.

Highly fluorescent markers are used a great deal in many biomedical imaging applications, both in microscopy and in diffuse imaging[32, 33]. The basic working principle is that if a biological target entity can be tagged (made to be co-localised with) a fluorescent molecule then the fluorecence emissions following excitation can be captured producing images of the target. The ratio of the number of emitted photons to the number of those that are absorbed by a molecule is its *quantum yield*. For imaging applications, a higher quantum yield is desirable as it provides higher contrast with shorter exposure times, this being particularly important for watching dynamic processes and for deep tissue imaging.

As well as deliberately introduced (exogeneous) fluorescence, biological tissues exhibit endogenous autofluorescence. This can be advantageous in a small number of applications because it can be used as a sensing target. However, more typically it creates unwanted, undirected background light serving to lower the relative SNR of desired signals, autofluorescence is discussed further in section 2.7.2.

[1] there are additional mechanisms such as chemiluminescence and phosphorescence but only the most relevant mechanisms are discussed in this thesis.

2.2.2 Bioluminescence

Like fluorescence, bioluminescence is a process that involves an excited molecule emitting photons, in this case mostly in the visible range, upon relaxation to the ground state. Unlike in the case of fluorescence however, the source of energy causing the initial excitation is a biochemical reaction (rather absorption of light). As opposed to chemiluminescence in general, bioluminescence specifically involves interactions between molecules that have developed as part of living organisms. The most well-known example of bioluminescence in nature is that of the firefly, and the biomolecules underlying firefly bioluminescence (particularly those of the *Photinus pyralis*, the most common firefly in North America) are the most exploited in biomedical applications (see section 3.3).

As with the utility of fluorescent markers, bioluminescent markers can be used in imaging by co-localising bioluminescent reporters with entitites of interest. Recent developments have allowed the incorporation of bioluminescent reporter systems into normally non-luminescent hosts thus creating internal light sources. Section 3.3 discusses bioluminescence imaging in detail.

2.3 Light Propagation through Tissue

The transport of light through tissue can be described by the radiative transport equation (RTE) or Boltzmann equation[19, 34]

$$\frac{1}{c}\frac{\delta L(\mathbf{r},\hat{s},t)}{\delta t}+\nabla\cdot L(\mathbf{r},\hat{s},t)\hat{s}=-\mu_t L(\mathbf{r},\hat{s},t)+\mu_s\int_{4\pi}L(\mathbf{r},\hat{s},t)P(\hat{s}\cdot\hat{s}')d\Omega'+Q(\mathbf{r},\hat{s},t), \qquad (2.9)$$

in which $L(\mathbf{r},\hat{s},t)$ is the radiance - the energy flow per unit normal area per unit solid angle per unit time - at position $\mathbf{r}$ into direction $\hat{s}$ at time t. c is the speed of light in the medium, μ_t is the total interaction coefficient (defined as $\mu_t=\mu_a+\mu_s$), the product $P(\hat{s}\cdot\hat{s}')d\Omega'$ is the probability of light travelling in direction $\hat{s}$ being scattered into solid angle $d\Omega'$ around direction $\hat{s}'$ which, in the case where it is only dependent on the angle between the two directions considered is the

previously introduced scattering phase function, and $Q(\mathbf{r},\hat{s},t)$ is the energy of photons produced by a light source within a unit volume at $\mathbf{r}$ into a unit solid angle centred on the direction $\hat{s}$ at time t.

Owing to the complexity of the RTE, its differential-integral nature and the number of variables, solving it directly is very challenging and direct implementation is prohibitively computationally expensive[35, 36]. Monte Carlo can be used in which the paths of millions or more individual photons or photon packets are traced explicitly[37]. With sufficiently many photons this can be made representative of the RTE[19], this is still computationally expensive[38]. As such, it is conventional to use approximations which are more computationally viable[39, 34]. One method of approximation is the expansion of the source and radiance terms by spherical harmonics expansion. By truncating the expansion at N terms the result is the so-called P_N approximation, the simplest and most-used of which is the first order (P_1) approximation which gives rise to the diffusion equation:

$$\frac{\delta\Phi(\mathbf{r},t)}{c\delta t} + \mu_a\Phi(\mathbf{r},t) - \nabla\cdot[D\nabla\Phi(\mathbf{r},t)] = S(\mathbf{r},t), \tag{2.10}$$

where

$$\Phi(\mathbf{r},t) = \int_{4\pi} L(\mathbf{r},\hat{s},t)d\Omega', \tag{2.11}$$

$$S(\mathbf{r},t) = \int_{4\pi} Q(\mathbf{r},\hat{s},t)d\Omega'. \tag{2.12}$$

Here, $\Phi(\mathbf{r},t)$ is the intensity or fluence rate - "energy flow per unit area per unit time regardless of the flow direction"[19] - at position $\mathbf{r}$ and time t, as shown above this is the integral over all solid angles of the radiance, S is an isotropic source term, and D is the diffusion coefficient defined as:

$$D = \frac{1}{3(\mu_a + \mu_s')} \tag{2.13}$$

The assumptions underlying the derivation of this diffusion approximation (DA) dictate that it is strictly valid only when the light source is isotropic and where effective isotropic scattering is much greater than absorption i.e. $\mu_s' >> \mu_a$.

2.4 Simulation of Light Propagation

With models such as the DA, it is possible to perform computational modelling (simulation) of light propagation through tissue. However, accurate calculations can only be obtained locally due to the non-linearity of propagating intensity with respect to optical property changes, as such it is necessary to break an interrogated volume into small sub-volumes throughout which optical properties are considered constant. Approaches include finite element method (FEM)[25, 40], boundary element method (BEM)[41, 42] and finite difference (FD)[43] schemes. It is also possible to use analytical models when the volume has a simple shape or where the optical properties are spatially invariant[44].

In the FEM approach used in this work (chapters 8 and 9), a geometrically accurate model of a small animal or phantom is built which is broken into discrete tetrahedral elements, the vertices of which are referred to as nodes. Within this approach complex geometries can be modelled effectively and spatially variant parameters can be assigned a value per-node to represent heterogeneity.

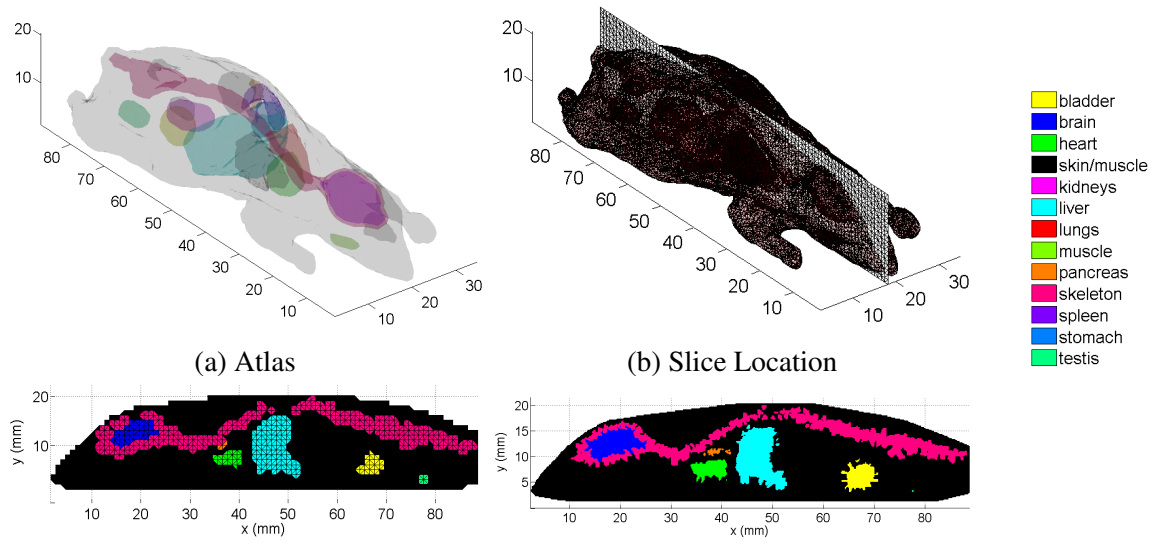

(a) Atlas (b) Slice Location

(c) Coarse FEM (d) Fine FEM

Figure 2.5: Example of an FEM mesh used for modelling light propagation. A mouse atlas[45, 46] (a) was sliced through the centre (b) to produce a 2D FEM mesh (c and d). A coarse mesh is shown (c) so that elements can be seen and a finer mesh is shown (d) to show the capability of representing detailed structures such as the labelled organs.

There are several openly available FEM-based light-modelling software packages including

Time-resolved Optical Absorption and Scattering Tomography (TOAST)[47, 48, 49], Molecular Optical Simulation Environment (MOSE[50, 51]) and Near Infrared Fluorescence and Spectral Tomography (NIRFAST)[25, 52, 53], which is used in the presented work. An example of a 2D NIRFAST FEM mesh is shown in fig. 2.5. The mesh represents a 2D slice through a 3D mouse atlas[45, 46] containing several labelled organs. It can be seen that complex spatial structures are represented.

If internal structures are ignored and a spherical light source is placed near the liver as shown in fig. 2.6(a), fluence can be simulated inside this model as shown in fig. 2.6(b-d).

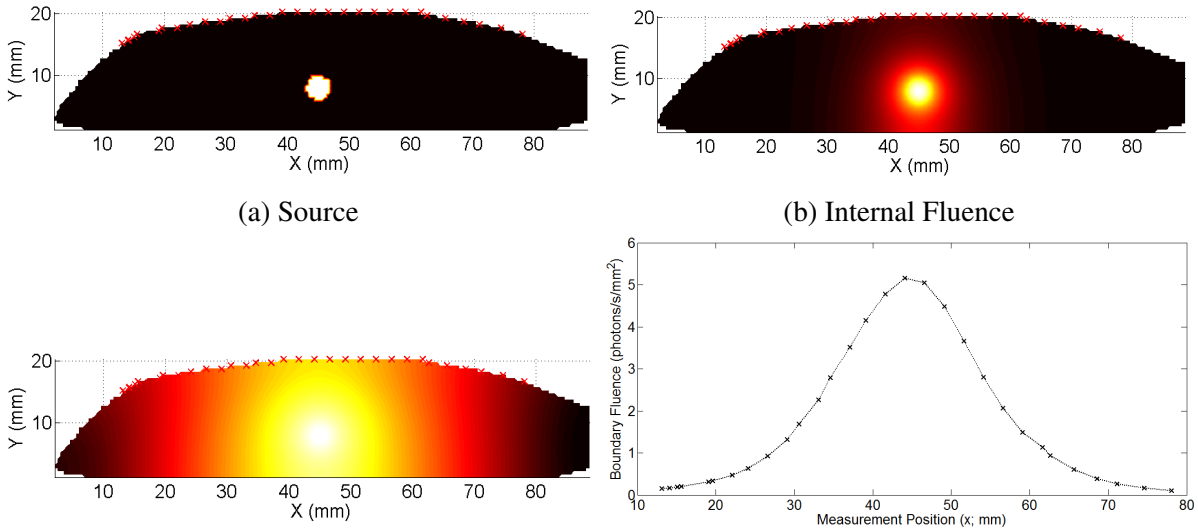

(a) Source

(b) Internal Fluence

(c) Internal log(Fluence)

(d) Top Boundary Fluence

Figure 2.6: With a simulated source added to the mesh (a), internal fluence (b,c) is calculated (per-node) using NIRFAST and boundary fluence (d) is calculated for each FEM detector position (red crosses on the mesh border). For simplicity, homogeneous optical properties of μ_a = 0.1cm $^{-1}$ and μ_s' = 10 cm $^{-1}$ are assumed.

By this approach, expected boundary measurements for a modelled system can be obtained given estimates of light source and optical property distributions. These can be compared with observations of the real system to evaluate the estimates and therefore provide an objective function for optimisation methods in image reconstruction. Two such methods are bioluminescence tomography and diffuse optical tomography, the principles of which are now discussed.

2.5 Bioluminescence Tomography

Bioluminescence tomography (BLT) is a method first proposed by Wang[54] that aims to reconstruct an internal light source distribution image such as that in fig. 2.6(a) from boundary data such as that in fig. 2.6(d). The internal sources in question are bioluminescent emitters co-localised with features such as tumours inside mice[55], as such the scale of fig. 2.6 is representative. Physical boundary measurements are obtained using a camera focused on the animal surface. Further detail including discussion of systems, applications and motivations of this method as a biomedical imaging technique are discussed in chapters 3 and 4 whilst the following serves to introduce the underlying principles of image reconstruction using FEM (section 2.4).

Given that accurate simulated boundary data can be obtained given knowledge of the source and properties, the image reconstruction problem can be formed as a least-squares type minimisation:

$$\mathbf{b} = \min_{\mathbf{b}} ||\mathbf{y} - F(\mathbf{b})||_2^2. \tag{2.14}$$

The solved for quantity $\mathbf{b}$ is the desired internal source distribution; a vector containing a source value per node in the FEM model. F is the model of light transport that accounts for optical properties and geometry, and $\mathbf{y}$ is real measurements obtained by practical experiment. In this scenario, fluence is linear with respect to source intensity. Thus the system of equations describing the relationship between internal source values and boundary fluxes can be represented by a single sensitivity (Jacobian) matrix:

$$\mathbf{y} = \mathbf{W}\mathbf{b}. \tag{2.15}$$

There are usually many more internal nodes than boundary observations and as such the Jacobian, $\mathbf{W}$, is non-square with many more columns than rows. It is therefore not generally possible to calculate $\mathbf{b}$ directly by

$$\mathbf{b} = \mathbf{W}^{-1}\mathbf{y}. \tag{2.16}$$

The problem is mathematically ill-posed, ill-conditioned and non-unique[56, 57]. Moreover, it has been shown that this is a general feature of the problem in that however many observations are made, the information content available is such that there is non-uniqueness[57]. The non-uniqueness problem is illustrated in fig. 2.7.

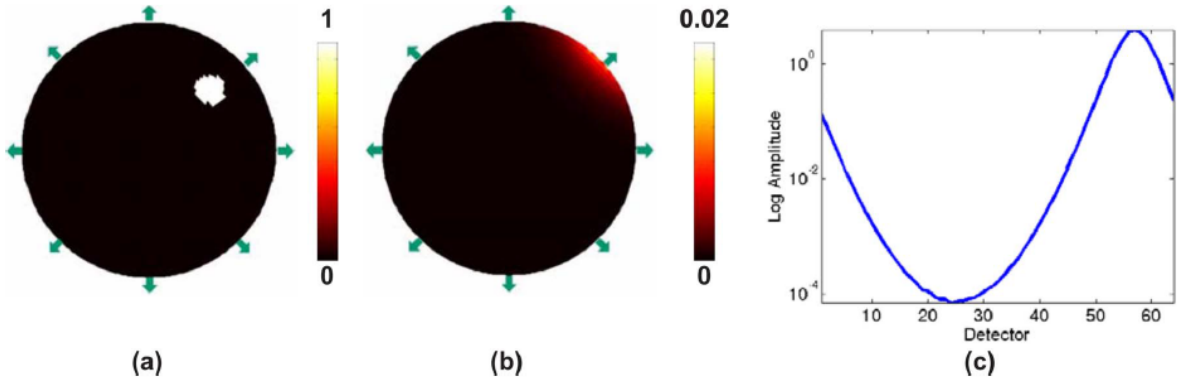

Figure 2.7: Demonstration of non-uniqueness in bioluminescence tomography, the 2D circle represents a cross-sectional slice through a mouse torso with optical boundary measurement locations marked with arrows. The light source distribution (a) gives rise to boundary measurements (c) and another distribution (b) gives rise to identical measurements. This figure is reproduced from Dehghani et al.[56]

Non-uniqueness can be reduced by measuring light of multiple wavelengths[57, 58, 56, 59, 44]; owing to the dependence of optical properties upon wavelength (sections 2.1.1 and 2.1.2) light at different wavelengths takes a different course through the volume and therefore provides different information. Multi-spectral BLT with data for k wavelengths can be represented as above by making the following assignments:

$$\mathbf{W} := \begin{bmatrix} \mathbf{W}_{\lambda_1} \\ \vdots \\ \mathbf{W}_{\lambda_k} \end{bmatrix} \tag{2.17}$$

$$\mathbf{y} := \begin{bmatrix} \mathbf{y}_{\lambda_1} \\ \vdots \\ \mathbf{y}_{\lambda_k} \end{bmatrix} \tag{2.18}$$

The inversion problem is still unlikely to have a unique solution but with appropriate choice

of regularisation and optimisation algorithm to solve eq. (2.14) reconstruction is possible[56, 59, 60]. Details of applications, specific implementations and measurement systems for BLT are discussed in chapters 3, 4 and 9 whilst a simple reconstruction is illustrated in fig. 2.8.

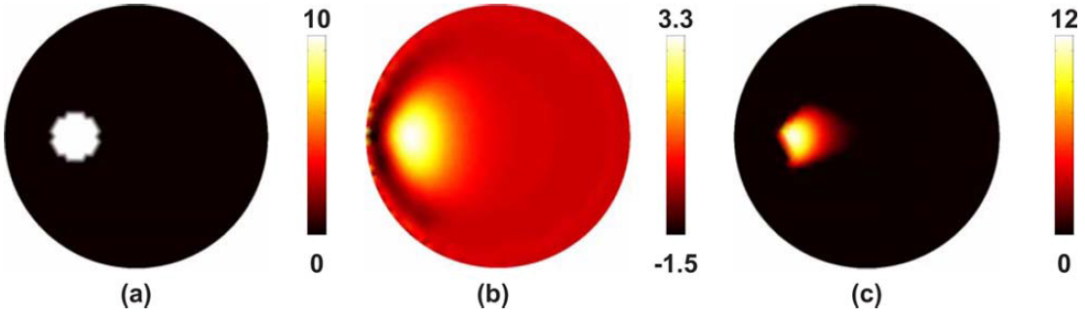

Figure 2.8: Example of multi-spectral BLT image reconstruction in a simple circular phantom using data from six wavelengths[56]. The original image (a) is reconstructed using two different algorithms; one without (b) and one with (c) a non-negativity constraint to reduce the space of possible solutions. It can be seen that the source can be recovered effectively. This figure is reproduced from Dehghani et al.[56]

2.6 Diffuse Optical Tomography

With the methods available for computing light propagation (sections 2.3 to 2.5), it has been discussed that BLT reconstructs light sources when the tissue geometry, optical properties and measurements are well-known. It is also possible to solve for the optical properties assuming that the other parameters are known. This is called diffuse optical tomography (DOT) when applied to image reconstruction[61, 62] and diffuse optical spectroscopy (DOS) when applied to single spatially-averaged or point-like reconstructions[63, 64]. In conjunction with the capability to model optical properties as a function of chromophore concentrations (section 2.1.1) and scattering models (section 2.1.2), it is possible to solve directly for chromophore concentrations and scattering coefficients making DOT and DOS into *functional* sensing techniques.

DOT has been applied widely in research. Key applications include the diagnostic imaging of breast cancer in humans[65, 66, 39, 67], and the mapping of functional parameters in the human brain[68, 69, 70, 71]. DOT has been applied more recently to small animal imaging. As

this is a key area of investigation in this thesis, a detailed overview of key developments and systems is included in chapter 4.

Unlike in BLT problems where internal light sources are present, DOT relies on interrogation light being introduced to a volume from outside. Figure fig. 2.9 shows examples of fluence simulation from external sources.

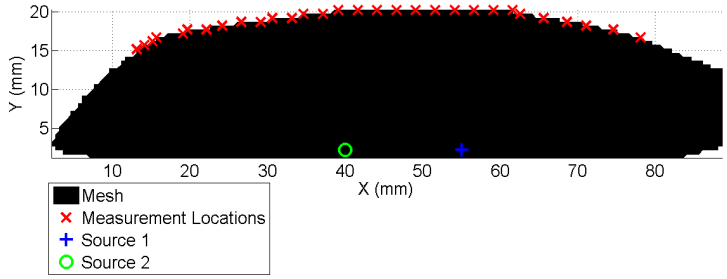

(a) Schematic showing source locations

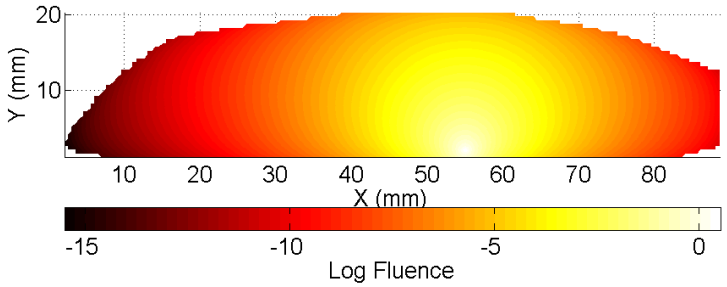

(b) Source 1 Fluence

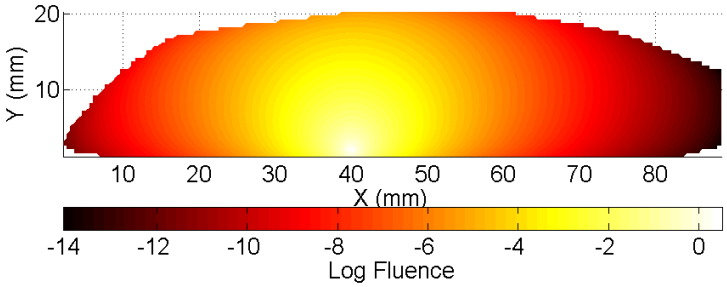

(c) Source 2 Fluence

Figure 2.9: Examples of simulated fluence from boundary sources used in DOT calculated using NIRFAST[25]. To illustrate this sources have been placed onto the same FEM mesh used in fig. 2.6, data can be extracted at the boundary as shown in fig. 2.6.

These simulations assumed that the input light was constant; this is called a continuous wave (CW) system[25, 27] and can be achieved using lasers or fiber-coupled lamps along with any detector. Another type of DOT system is time-domain (TD), in which short laser pulse (less than 1×10^{-10} seconds[27]; approximated as delta functions in time) sources are used and ultra-

high speed detectors are needed to capture time-resolved intensity. The TD data contains more information than CW data because the shape of the temporal point spread function (TPSF) in addition to the amplitude changes as a function of optical properties. The other important DOT system type is frequency-domain (FD), in which sinusoidally amplitude-modulated sources, with frequencies typically in the range of 100 to 1000 MHz[27], are used and more complex detection systems capture the phase change and de-modulation of the input light[72]. FD and TD data-types equivalently capture more information than CW systems[27] but CW data has the advantage of requiring less complex and less expensive data-capture and analysis methods. The remainder of this thesis is generally concerned with CW data as this is the datatype used within the presented imaging system.

In the standard method of DOT, in which μ_a and/or μ_s' are solved for given measurements of transmitted light through a volume, predicted and measured data are obtained for many source-detector (S-D) pairs given an initial estimate of the optical properties. The optimisation problem is then

$$\tilde{\mu} = \min_{\mu} ||\mathbf{x} - F(\mu)||_2^2, \tag{2.19}$$

assuming a least squares type optimisation, where F is the forward model of light propagation through the tissue given the optical properties μ.

As in the case of BLT, the DOT inverse problem is generally ill-conditioned and ill-posed and this is generally dealt with by regularisation. The DOT problem is also challenging because the boundary fluence values are non-linear with respect to the optical properties and as such non-linear reconstruction methods are required. As in the case of BLT, multiple-wavelengths can be used to provide additional data and this is particularly useful when combined with the approach of solving directly for chromophores and scattering properties[61, 66, 26]. In this case the space of possible solutions can be constrained such that the only allowed optical properties are those representing a linear combination of known chromophores likely to be present. This improves image reconstruction, guarantees physically possible solutions and provides functional imaging by returning chromophore abundances. This approach is used in the presented work and as such is described in detail later (chapter 8).

An example of images reconstructed using spectrally constrained DOT with an illustrative phantom are shown in fig. 2.10. It can be seen that images can be retrieved of scatter power

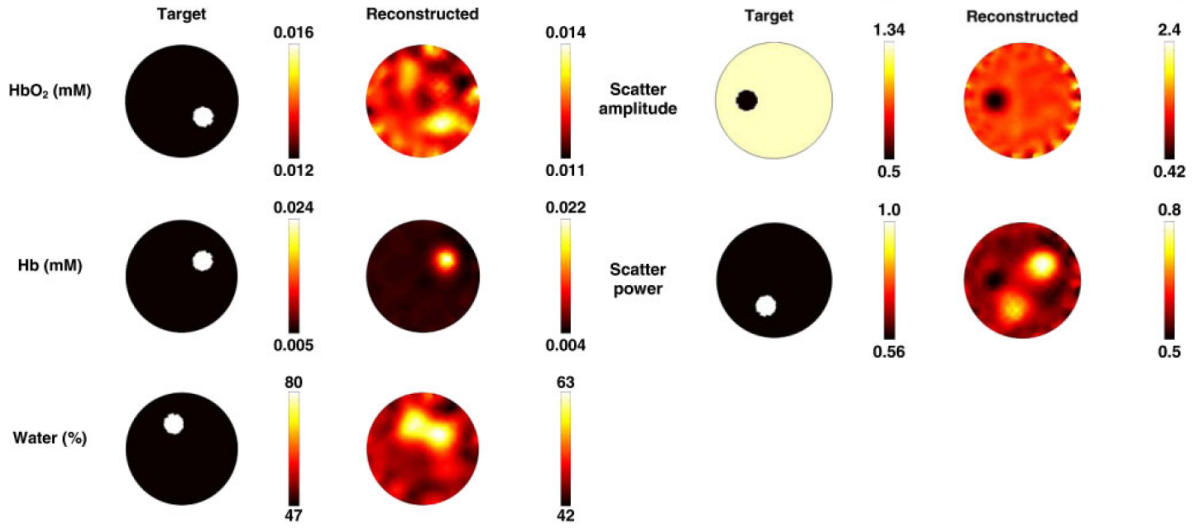

Figure 2.10: Example of reconstructed images obtained for a circular phantom composed of deoxyhaemoglobin (Hb), oxyhaemoglobin (HbO), water and scatterers. Hb and HbO image values show concentration in mM and water values are percentage by volume. Scatter amplitude is in units of mm^{-1} and scatter power has no units. This figure has been reproduced reformatted from Dehghani et al.[25].

and amplitude (eq. (2.8)) and chromophores responsible for absorption. It can also be seen that there is cross-talk between parameters i.e. artefacts in images of one parameter that resemble the characteristics of another parameter, this is a common challenge for DOT.

An opportunity arises when working with spectral DOT to re-employ the relationship between derived spectral properties and μ_a and μ_s' values at particular wavelengths (eqs. (2.3) and (2.8)) to calculate optical properties at particular wavelengths of interest from reconstructed spectral properties irregardless of whether or not these were wavelengths at which data was obtained. This is exploited in the presented work; reconstructed spectral properties from DOT based on NIR wavelengths are used to calculate optical properties at particular visible wavelengths at which BLT measurements are made.

2.7 Optical Properties of Small Animals

2.7.1 Properties of Different Tissue Types

There is no well-validated collection of mouse optical properties available. However, Alexandrakis et al.[31] compiled a library of optical properties from available literature taking values from many sources. Note that the original data was obtained using various techniques, animal species and conditions and therefore are only a rough estimate. The authors fit available μ_a values to a model that assumed absorption spectra were exclusively combinations of the known spectra of oxyhaemoglobin, haemoglobin and water (section 2.1.1) and fit scattering values to the model presented in 2.8. The resulting absorption and scattering spectra for a variety of tissue types are shown in fig. 2.11.

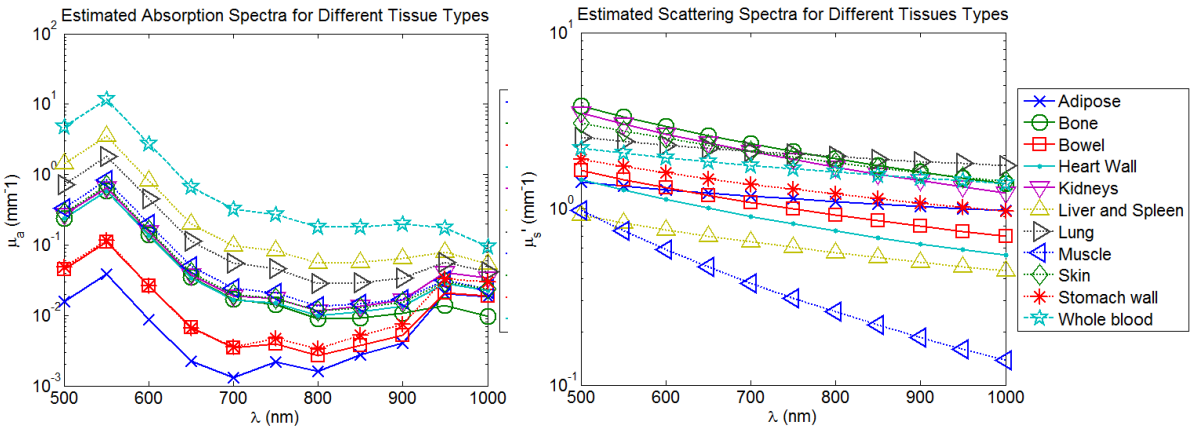

Figure 2.11: Approximate absorption and scattering spectra for different tissues calculated using the models of Alexandrakis et al.[31]

It can be seen that the range of absorption values expected in different tissues and across the wavelength range presented covers five orders of magnitude whilst the scattering spans only two. It can also be seen that absorption typically is at its lowest from around 650 to around 900nm whilst scattering is inversely proportional to wavelength across the range shown. Whole blood is the most absorbing tissue reported with a peak μ_a at around 550nm of about 10mm^{-1} followed by the liver and spleen and then the lungs. Muscle, bone, kidneys, heart and skin are mid-range with very similar absorption spectra and the stomach and bowel have lower similar absorption characteristics. Adipose tissue is the least absorbing of those considered with a

minimum μ_a of about 0.001mm^{-1} at 700nm.

The most scattering appears to occur in the bone, kidneys, lungs, skin and blood closely followed by stomach, bowel, adipose and heart. The liver and spleen have relatively low scattering and muscle is the least scattering of all. There is clearly a strong dependence on tissue type and therefore considerable heterogeneity throughout a small animal; consider as a guideline fig. 2.5 in conjunction with the presented values.

Figure 2.12 shows the same data in another format so that absorption and scattering of individual tissue-types can be compared more easily. It can be seen that many tissues, namely adipose, bone, bowel, kidneys, stomach and skin, have scattering values that are consistently an order of magnitude higher than the absorption for the range of wavelengths considered. This is important because the DA (section 2.3) is only valid when μ_s' is much larger than μ_a [19]. In the lung and heart the scattering is consistently higher but at wavelengths lower than around 600nm scattering and absorption draw close together. In the case of the liver and spleen, muscle and whole blood, the absorption is actually higher than the scattering for some wavelengths below around 600nm and as such in these cases the DA would certainly not be expected to hold.

Owing to the fact that the optical properties are based on variable parameters such as tissue oxygenation, it is clear that they will change dynamically with the environment, health and other circumstances of the animal. This makes it very difficult to reliably predict them and accurately model light propagation *in vivo*. There are undoubtedly other characteristics that have not been considered but are highly relevant to the transport of light through the body. These include the effects of void regions such as inside the inflated lungs or mouth, and the presence of unknown materials such as the variable contents of the stomach and bowel.

The poorly documented, dynamic and unpredictable nature of optical properties in general is a key motivation for performing DOT to calculate tissue properties on-the-fly in experiments that rely on modelling them such as BLT experiments. The presence of unknown materials presents another motivation for performing DOT in this way. These are clear points of motivation for the imaging method presented in this thesis which is designed to perform DOT alongside BLT.

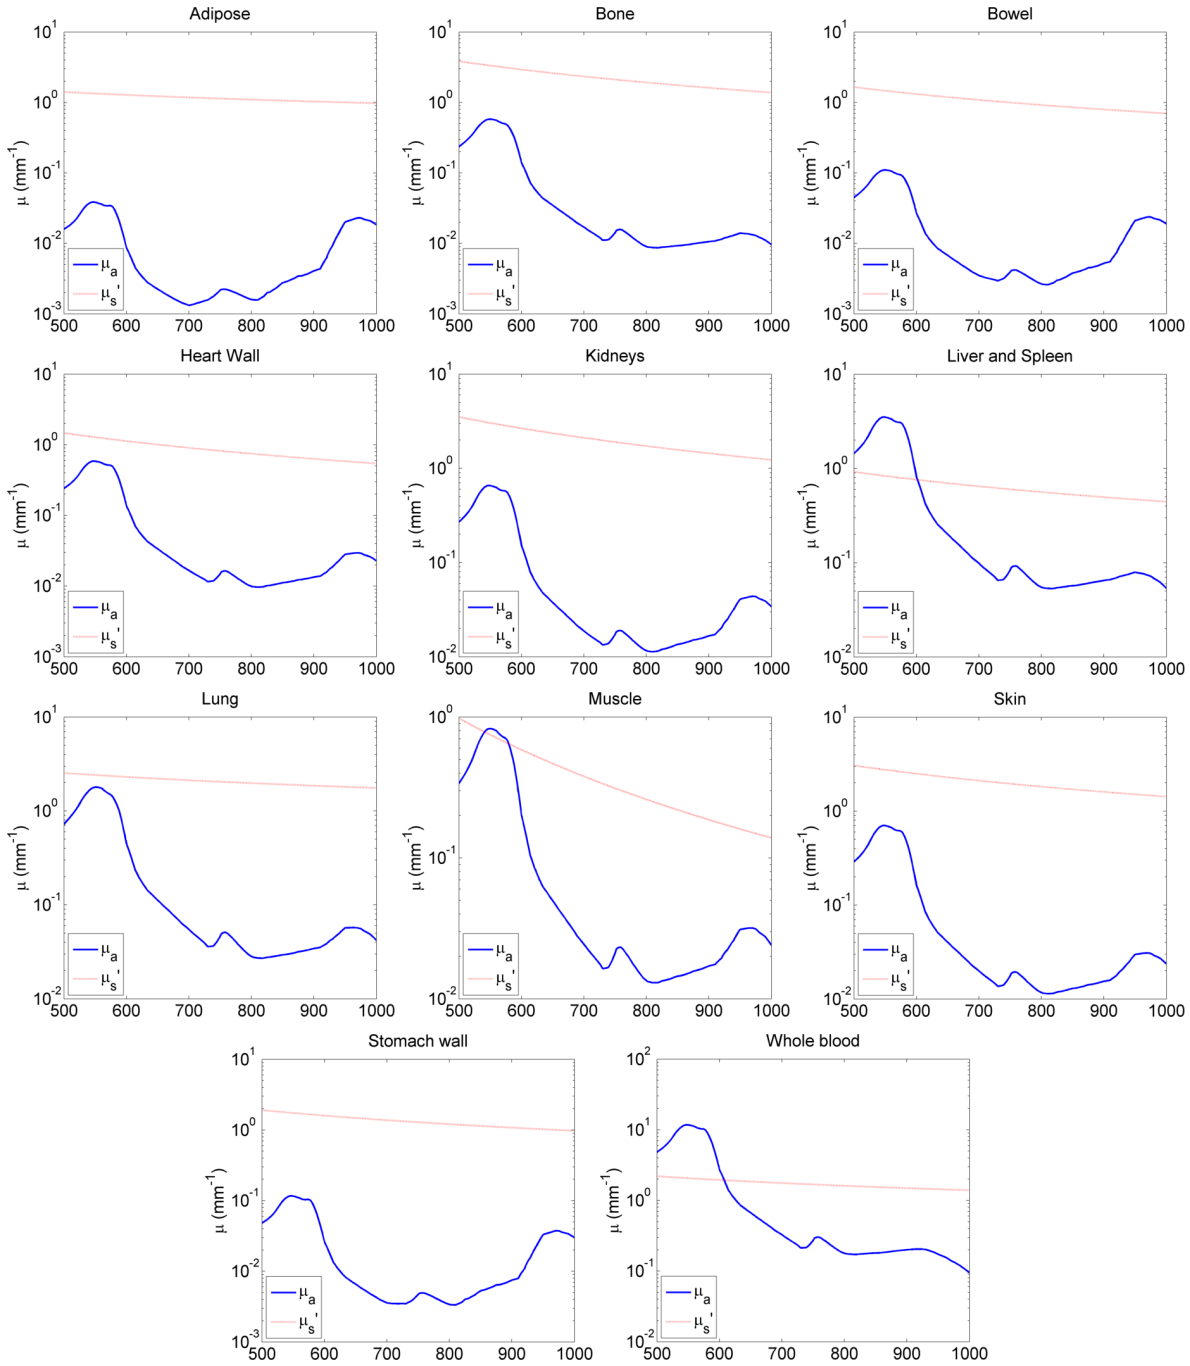

Figure 2.12: Comparison of absorption and scattering of different tissue types, the data is the same as that in fig. 2.11

It has been shown that dependent on the tissue type, the DA may or may not hold in small animals. This provides a potential complication for the present work in which the DA is used (chapters 8 and 9). However, the presented optical properties are only a guideline due to the mechanism underlying their compilation, additionally in many tissues the DA should be valid, certainly at higher wavelengths. Furthermore all of the methods developed in this thesis are transferable to involve the use of more advanced models such as higher order simplified spherical harmonics models[36]. These models are more computationally intensive (though still less-so than the RTE) and it is therefore of interest to explore how far the DA can take us in practice, it is a point of interest to investigate how well the DA might approximate conditions even if values used explicitly turn out to be incorrect, in a sense this is another potential advantage to using the same model and system to record DOT signals and estimate properties and then use them to model BLT sources; systematic errors may be self-compensating in this set-up.

2.7.2 Autofluorescence

It has been discussed that fluorophores can be used as imaging probes (section 2.2.1) and in the following chapter macroscopic fluorescence imaging is discussed in more detail (section 3.2.6). Within this context, *autofluorescence* describes unwanted fluorescent background arising from molecules other than a fluorophore of interest[73, 74, 75]. This is an issue in microscopy and macroscopic tissue imaging and it occurs because many endogeneous biomolecules exhibit fluorescence thus excitation light introduced to excite target fluorophores also excites these other molecules. The result is a background signal which effectively lowers SNRs and mitigates the sensitivity and specificity of imaging techniques.

In small animal imaging, the dominant sources of autofluorescence are collagen which is excited most strongly by wavelengths around 300 to 350nm and emits most strongly at around 400nm[76] and chlorophyll break-down products obtained from food which emit most strongly at around 670nm and excite most strongly at approximately 400nm[75]. As these are both found in the skin this is significantly autofluorescent[74]. In general autofluorescence is most strongly excited at the blue end of the visible spectrum and much less so at red to near infrared

wavelength. As such red- and near infrared-excitable probes have been pursued[77] but many of the most well-studied fluorescent probes such as green fluorescent protein (GFP) require blue-green excitation. When imaging GFP in deep tissues the autofluorescent background signal can be several times the signal arising from the GFP[78]. This necessitates the use of advanced techniques to de-couple the wanted signal from the background such as spectral un-mixing[74] or fluorescent lifetime analysis[78].

2.8 Conclusion

This chapter has introduced key light/tissue interactions, methods for modelling light propagation through tissue and the fundamental ideas of DOT and BLT image reconstruction. These imaging techniques are applied in the present work within the context of a small animal imaging system designed to perform BLT and DOT and as such this initial introduction is essential background to the remainder of the work presented. Characteristics of important mechanisms underlying the generation of light in tissue, fluorescence and bioluminescence have also been discussed as well as the issue of autofluorescence in tissue which raises important considerations for any imaging method involving the illumination of tissue. The most realistic estimates of the optical properties of small animal tissues have also been presented, these are important to work presented later in the thesis towards characterising the imaging system using phantoms; the realism of the phantoms can be compared to the real tissue optical properties to place in context the characterisations that have been done. These also allow theoretical analyses of expectations of system performance to be more accurately carried for example in simulations.

Following this introduction to key principles underlying the presented technology, the next chapter provides further background information with an emphasis on other small animal imaging technologies so as to provide a wider context for the developed system. A particular emphasis is placed upon bioluminescence imaging, the platform on which BLT builds, which discussions of the practicalities and methods of imaging experiments.

CHAPTER 3

PRECLINICAL IMAGING

3.1 Introduction

The ultimate goal of biomedical science is to understand biological systems, particularly the human body, in both normal and abnormal conditions such that the health of the system can be diagnosed, monitored and restored when abnormalities are encountered. Imaging plays an important role in obtaining this understanding by providing visually assessable, quantitative, spatially and temporally resolved measurements of parts of biological systems that allow them to be analysed and understood.

A particularly important current concern is the development of new drugs and therapies for the treatment of major diseases such as cancer. This process can be aided greatly by visualising the precise action of a prospective drug in the biological system and the precise reaction of the system to the drug.

Whilst the most relevant data is obtained when visualising such processes *in vivo* in humans, many imaging techniques are limited in that they must operate on cultured, excised, or otherwise isolated cells and tissues that are therefore no longer part of the in-tact, interacting organism of original interest. These types of imaging methods are very important because conditions can be well-controlled and variables isolated for scientific analysis but they cannot provide data that is fully representative of the dynamics of biological processes as they occur within the original intact organ structures or inside living subjects.

In addition to technical restrictions, there are sometimes procedural, ethical and safety issues that prevent many studies from being performed in humans. It is very expensive and risky to trial new drugs so there is a huge drive to acquire the most relevant data from studies that are the closest match to the ideal in-human experiment but without the high cost and risk.

Preclinical imaging describes a class of macroscopic techniques that address these issues by visualising processes in living animals, most often small animals such as mice, often *longitudinally*, i.e. imaging the same animal at many time points to understand the long-term dynamics of a process such as a cancer therapy. The living mouse is a good (though not perfect) analogue to the living human; the internal organ structure and roles, body-wide and cellular processes, immune responses and mechanisms are very similar, as such it provides an accurate depiction of the *in vivo* biology without the costs, risks and ethical issues associated with human studies. Owing to these properties, preclinical imaging is playing an increasingly prominent role in the translational stage between lab-based biomedical imaging techniques such as optical microscopy and clinical trials.

This thesis presents work towards the development of a novel optical preclinical imaging system with many potential applications throughout biomedical science but with a particular emphasis placed on imaging cancer, i.e. tumour development and metastases, in mouse models. In order to provide context for the developed system in terms of the current state-of-the-art, this chapter provides first an analysis of existing preclinical imaging techniques, describing the strengths and limitations, and then a thorough description of bioluminescence imaging; one particular emerging preclinical tool that is built upon directly by the presented system. This will illustrate context and motivation for the current work.

3.2 Preclinical Imaging Techniques

The following subsections discuss in some detail nuclear imaging techniques; Positron Emission Tomography (PET) and Single-Photon Emission Computed Tomography (SPECT), Magnetic Resonance Imaging (MRI), X-Ray Computed Tomography (CT), Ultrasound (US), Pho-

toacoustic Tomography (PAT), Fluorescence Imaging (FLI), and Bioluminescence Imaging (BLI). Discussions outline the principles of each technique and some key features and limitations.

All of the following techniques can be used to some extent for longitudinal macroscopic small animal imaging studies to assess phenomena such as tumour growth, cell trafficking and drug efficacy. An important limitation of all of the techniques is that mice generally must be anaesthetised which naturally causes physiological changes.

A key factor is the use of ionising radiation, which introduces a confounding factor to the animal study (potentially acting as a therapy in cancer studies[79], or causing new cancers or even lethal doses of radiation) and also providing an unwanted health and safety risk to the human operator.

An important differentiating factor between different techniques is the mechanism for contrast, e.g. optical absorption, acoustic impedance boundaries, or x-ray absorption, which essentially underpins the types of events, markers or processes that can be studied with each technique. Contrast agents that are specific to each contrast mechanism are available for many techniques which broadens the range of possible applications. Whilst the number of agents available or the number of different processes that can be visualised can be seen as advantages or limitations of a technique, the wide variety of mechanisms and applications reveals all of the following techniques are valuable and in this respect complimentary in terms of the data they provide.

3.2.1 Nuclear Imaging

Positron emission tomography (PET) is a molecular imaging modality that is used both clinically and preclinically to image an exogenous molecular target, or tracer. The tracer contains a radioactive isotope which emits a positron (the anti-particle to the electron, with opposite charge) when it decays. The positron travels a small distance in the tissue, ≈ 0.5mm, dependent on the isotope-specific particle energy and the tissue density, then annihilates with a nearby electron. Upon annihilation, two high energy photons (gamma rays) are produced travelling in

opposite directions at the speed of light. These are detected by scintillators positioned around the imaged subject and associated together by the very small difference in their time of arrival (on the order of nanoseconds) at the detectors. For each event identified there is therefore a line (drawn between the points of detection) upon which the origin might lie. With multiple measurements of this type, reconstruction algorithms are used to build quantitative, volumetric, 3D images of the tracer distribution throughout the body[80, 79].

Single photon emission computed tomography (SPECT) is similarly based on imaging radioactive tracers though in this case the signal-generating events are, as the name suggests, single-photon rather than 2-photon emissions and of lower energy radiation[79, 80]. Given that single photons are detected, measurements lack the directional certainty associated with a PET detection and as such systems typically use directed detection schemes which select incoming photons by their incoming angle[79]; a great deal of photons are necessarily discarded and sensitivity is reduced.

Leading preclinical PET and SPECT scanners can achieve spatial resolutions of 0.5 to 2mm[79] meaning that only relatively large-scale macroscopic spatial changes can be observed. The main advantages of the techniques are the high molecular specificity and the ability to provide quantitative measurements of tracer content giving an insight into functional information.

Applications are guided by the available tracers for imaging. The most common tracer molecule used is fluorodeoxyglucose (FDG). This takes the place of glucose and therefore is indicative of metabolism[81, 82].

An advantage of PET is that for the high energy radiation that is involved, there is orders-of-magnitude less probability of interaction with tissues and much less tissue-dependent variation in the interactions that occur compared to visible and near-infrared photons. As such, signals are more straightforwardly quantitative and sensitivity is limited only by detection efficiency[80]. This is on the order of 2-5%[80, 79]. Sensitivity is about an order of magnitude less[79] for SPECT and, owing to the lower energies involved, corrections for attenuation and scattering are necessary[79]. The field-of-view and high energy transmittance of both modalities allows them to image whole animals easily[79].

A mitigating factor for sensitivity is tracer uptake, animal short-term physiological changes, due to for example whether they are warmed, fasted or anaesthetised, are known to effect the uptake of the tracer[83].

Whilst non-invasive, nuclear imaging does involve the use of ionising radiation which, particularly over longitudinal studies can be substantial.

An advantage is total specificity to the tracer molecule, though this is not generally representative of specificity to observed biological processes because other issues such as tracer kinetics, localisation and clearing also impact upon the overall specificity. Clearly these issues also affect the effective sensitivity.

A feature of SPECT is that with particular detectors capable of sufficiently distinguishing energy levels, multiple tracers can be distinguished and imaged simultaneously[84, 85]. A limitation of PET is that multi-marker imaging is not possible in this way.

A limitation of nuclear imaging is that it typically requires a cyclotron either on-site or nearby. This is required for the manufacture of the radioactive isotope used in the tracer molecule. The cost thus introduced can be substantial and complications are also introduced in terms of timing, the tracer only being good for a finite amount of time, meaning transportation and experiment timing become important considerations. This is made further problematic when the isotope has to be combined using bespoke techniques with other reagents to produce the final tracer molecule, as is often the case in research.

Because the contrast and sensitivity in nuclear imaging techniques is entirely with respect to tracer molecules, the structural context of the measured signal is not known without use of additional imaging modalities. This can be seen as a limitation of these techniques given that it necessitates the use of multi-modality approaches to provide spatial context. SPECT or PET combined with CT is therefore often used and combined scanners are commercially available[79], though this imparts a larger ionising radiation dose and is more expensive. Additionally there is the limitation that CT does not have good soft-tissue contrast without the use of appropriate contrast agents which are an added complication. PET/MRI scanners are now commercially available and growing in popularity[86], in this case MRI provides clear anatomical

reference but without ionising radiation[79, 86], on the other hand this is still more expensive.

3.2.2 Magnetic Resonance Imaging (MRI)

Magnetic resonance imaging (MRI) is well-established as a clinical imaging modality and is also used preclinically. MRI requires a strong magnetic field ($\approx$ 2 to 7 Tesla) necessitating large powerful magnets and high-end cooling systems.

The technique exploits the property that some atomic nuclei, those that do not contain an even number of both protons and neutrons such as protium (^{1}H), have a property called spin which can be thought of as the constant rotation about an axis. In general the spin axes of nuclei are randomly orientated and as such cancel each other out on a macroscopic scale[87].

In MRI, a large external magnetic field causes nuclei to precess around the field direction and for spins to become (on average) aligned together in a minimum-energy state. Radio-frequency pulses are applied to perturb the spins, and measurements are made based on their *relaxation* to the low energy state. Different molecular chemical environments produce different signals owing to chemical shielding[87].

^{1}H is the most common target for MRI because it is so abundant in tissues being a key constituent of water and fat and has a high MR response[87], though carbon, fluorine, and phosphorous can be targeted as well[88]. MR chemical imaging is possible for example using hyper-polarised carbon-14 to image PH and also cellular redox activity[89], perfusion-weighted MRI can estimate blood flow and in so doing monitor angiogenesis[89].

Contrast agents such as gadolinium (Gd)[90, 91] can be used to enhance the signal differences between tissue-types, mostly based on altering relaxation times[88]. There is also increasing work with magnetic (e.g. iron-oxide based) nano-particles providing contrast[92, 93, 94].

MRI can be seen as a gold standard for structural imaging, with spatial resolution less than 100μm possible (dependent on the field strength), and excellent anatomical contrast, but it does not usually provide functional information or molecular imaging capabilities. As such, it has been combined with many complimentary methods. Gulsen et al.[95] developed an optical and MRI multi-modal system specifically designed for cancer imaging, and there has been a great

deal of work and interest in multi-modal PET-MRI systems[86].

The major disadvantage to MRI is its very high cost and resource requirement; the initial and running costs are often prohibitive, and a great deal of space is needed to house an MR scanner.

3.2.3 X-Ray Computed Tomography (CT)

As with MRI and nuclear imaging, x-ray computed tomography (CT) is well-known and widely used clinically and also has a preclinical counterpart[96, 97, 98, 99].

In preclinical CT, x-rays (electromagnetic radiation with wavelengths approximately ranging from single nanometres down to tens of picometres) are transmitted through the animal from a source such as an x-ray tube and detected on the other side of the animal by a detector such as a scintillator. Detectors and sources are usually moved around the animal so as to make many 2D X-ray projections. These are then used with reconstruction algorithms to reconstruct 3D images.

The source of contrast is exclusively the amount of x-ray attenuation caused by the tissue which is strong in bone and weak in soft tissues. As such CT is particularly suited to imaging bones and applications such as arthritis studies[96]. Whilst X-ray scattering is relatively weak in tissue, detectors are generally designed to accept only those rays that arrive perpendicular to the detector and as such reject scattered radiation.

Soft tissues are hard to distinguish in CT scans unless contrast agents are used[98, 99]. Examples of contrast agents are iodine-based agents[100, 101] and gold nano-particles[102, 103], these can be injected into the blood stream to image blood vessels and differently vascularised organs and tumours.

CT performs structural imaging in high resolution (less than 100μm[96]); sharp, detailed macroscopic images of anatomy can be obtained, though it cannot generally be used for molecular imaging or functional imaging.

Due to these features CT is often combined, either in bespoke multi-modality systems or with post-acquisition image co-registration, with a functional imaging modality so as to allow

inspection of functional changes alongside anatomic context. Examples of multi-modal systems include a combined optical and CT system presented by Tichauer et al.[33] and an optical, nuclear and CT imaging system developed by Cao et al.[104]. An example of a multi-system, multi-modal study is that of Deroose et al.[105] in which CT, PET and optical techniques were used together to study cancer.

In small animal imaging, high resolution CT studies can quickly deliver large ionising radiation doses to the animal which is a considerable disadvantage particularly in longitudinal studies[96]. CT scanners require X-ray shielding to protect the operator, mechanical gantries to perform the rotation, and advanced x-ray sources and detectors. As such, there is substantial complexity and cost involved.

3.2.4 Ultrasound (US) Imaging

Preclinical ultrasound (US) imaging involves the generation of high frequency (on the order of tens of megahertz) mechanical (acoustic) waves that propagate through tissues. The waves are reflected at boundaries between regions of differing acoustic impedance such as tissue-bone interfaces. Reflected waves are then detected and knowledge of the speed of sound is used to calculate the total path-length of reflected waves yielding the distance to the boundary. This principle is used in imaging systems where waves are introduced and many reflections detected at many positions. 2D or 3D images are then generated with acoustic boundary contrast. Contrast agents in the form of micro-bubbles can be used, these can be introduced into the blood for example and provide a source of acoustic impedance mismatch and therefore of detectable reflections.

Doppler US is based on the same principles and also exploits the *doppler shift* that occurs when the originator of a wave is moving towards or away from the point of detection such that the wave is either stretched or compressed (causing an observable change in frequency) with time. By analysing these effects the motion of the wave sources can be assessed.

There are several advantages for US systems; it is non-ionising and low-cost compared to other techniques. The resolution can be less than a 100μm. It is also practically unique in that

it can be real-time.

The main disadvantages of ultrasound relate to contrast. Firstly, the mechanism underlying the contrast (i.e. acoustic boundaries) is limiting for many applications. The use of micro-bubble contrast agents can improve this. Whilst this is a developing technology that has so far only been applied to a limited set of problems, it has allowed US to move into the domain of molecular imaging[106, 107, 108, 109]. Secondly, the contrast of images is generally poor[106, 110, 111, 112].

3.2.5 Photoacoustic Tomography (PAT)

Photoacoustic/optoacoustic tomography (PAT)[113, 111, 112] is a relatively new imaging technique with growing interest in research and the first commercial systems recently available (such as the iThera Medical MSOT[114] and the VisualSonics Vevo LAZR[115]). The technique is a hybrid in which optical excitation generated by a pulsed laser is absorbed by tissue chromophores causing practically instantaneous heating and localised pressure increases. The energy is subsequently released in the form of broadband mechanical waves and can be detected as in US imaging. Image reconstruction techniques such as time-reversal approaches[116] can then be used to reconstruct a 3D image of the initial pressure distribution which is indicative of the spatially-varying absorption[111]. The preferential absorption by different chromophores is dependent on the wavelength of the excitation light (see chapter 2) and multi-spectral imaging is being explored for the quantitative analysis of endogeneous and exogeneous chromophore concentrations throughout a tissue volume[117].

Owing to the optical absorption-based contrast and the prominence of haemoglobin as an absorber in tissue (chapter 2), PAT has found striking early results in imaging vasculature and it has been used to show the different vasculature growth patterns within different subcutaneous tumours as well as showing the response of these to therapies[118]. Several absorbing contrast agents have been used such as indocyanin green (ICG)[119].

A major advantage to PAT is that it combines the high resolution of US with high absorption-based optical contrast. It uses no ionising radiation. Furthermore, owing to the differing absorp-

tion characteristics of different chromophores, multi-spectral approaches can be used to identify and quantify the concentrations of different chromophores spectroscopically[114, 120].

A limitation of PAT is that, because the technique is still developing, expertise is required by the user and few systems are commercially available and thoroughly tested. Whilst boasting enhanced contrast and more favourable contrast mechanisms for many applications compared to ultrasound, the need for a pulsed laser such as wavelength-tunable OPOs is a significant increase in cost, making the cost altogether more comparable to high-end optical systems, below those of nuclear imaging, CT, and MRI.

A further limitation is the penetration depth of the excitation light which, whilst able to penetrate a whole animal body at many wavelengths in the visible to near infrared, is heavily attenuated, this poses a challenge for, particularly *quantitative*, image reconstruction which is a topic of ongoing research.

3.2.6 Fluorescence Imaging

Fluorescence imaging (FLI; fig. 3.1) is an optical preclinical imaging modality in which optically-excitable fluorophores (fluorecent molecules; see section 2.2.1) such as green fluorescent protein (GFP) are imaged inside a living animal; the fluorophores are introduced into the animal system prior to imaging. These are then excited by a source of optical excitation such as a laser matched to the absorption peak of the fluorophore and the subsequent longer-wavelength emissions are captured by a camera focused on the surface of the animal.

Some amount, dependent on the fluorophore location and precise tissue properties and excitation field, of the emitted light reaches the animal surface. The light is highly diffused making image resolution poor and rendering the imaging counts only indirectly relatable to the underlying fluorophore concentration and position.

Fluorescence-mediated tomography (FMT) is an emerging extension to FLI which involves modelling the light-tissue interaction and reconstructing the internal marker distribution allowing quantitative and higher-resolution imaging, this is discussed further in chapter 4. One particularly useful approach to somewhat overcoming the difficulty of unknown underlying optical

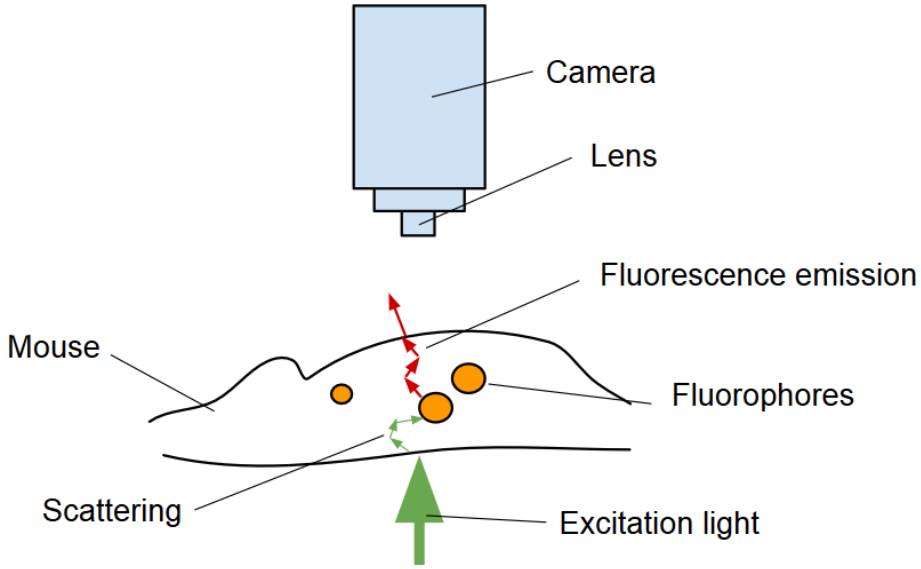

Figure 3.1: Schematic depicting fluorescence imaging (FLI). A camera focused on the animal skin surface captures light emitted from internal markers. In FLI these are excited by introduced excitation light. All the light is scattered and absorbed substantially by the tissue as discussed in chapter 2.

properties in FMT/FLI is the use of the Born ratio, in which excitation light and emission light are assumed to have the same interactions and an approximate normalisation can be performed. Of course, owing to the variability of tissue interactions between wavelengths (chapter 2) this is only an approximate technique and the success should vary depending on precise conditions.

The molecular imaging capability of these techniques allows functional imaging studies comparable to those of nuclear imaging to be undertaken (and with the latest tomographic systems with comparable resolution and quantitation) but without the high cost (running and initial investment), ionising radiation or requirement for cyclotron. There are also a great many fluorescent markers already available in part due to the wide and relatively long-running usage of fluorescence in microscopy. However, the utility of many fluorophores is limited in practice by the higher typical absorption and scattering of tissues at lower wavelengths, thus the most useful are those whose excitation and emission wavelengths fall within the red to near infrared part of the spectrum (see chapter 2).

A challenge for FLI is that there is a strong background signal generated by tissue (section 2.7.2). Whilst techniques such as spectral un-mixing[74] and fluorescence lifetime decoupling[78] can be used to separate the signal from the background, these techniques require

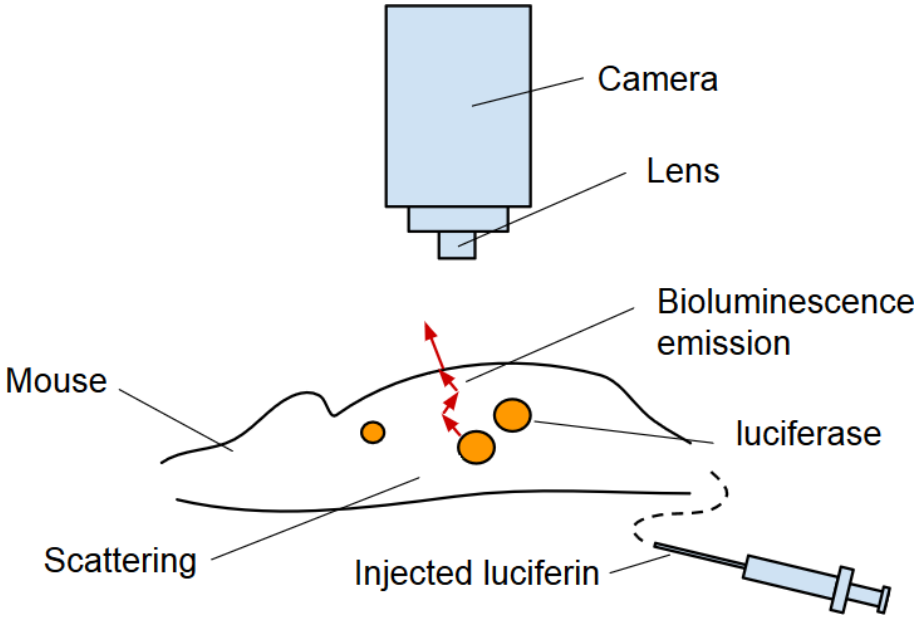

Figure 3.2: Schematic depicting bioluminescence imaging (BLI). A camera focused on the animal skin surface captures light emitted from internal markers. In BLI these are the result of a biochemical reaction involving an injected substrate. All the light is scattered and absorbed substantially by the tissue as discussed in chapter 2.

extra, more advanced data to be obtained raising cost and experimental time. Furthermore, the use of added data such as spectral or time-resolved data to remove background limits their usage in the same study in tomographic reconstruction.

Another complication in FLI is that fluorophores are subject to photobleaching by the excitation light meaning that emission signals deplete at a certain point.

3.3 Bioluminescence Imaging

Bioluminescence imaging (BLI, fig. 3.2) is another optical technique that is based on detecting photons transmitted through tissue originating from light-emitting molecules. In this case the markers are bioluminescent rather than fluorescent and as such the light is generated as a result of a biochemical reaction rather than due to excitation light (sections 2.2 and 3.3.2).

BLI shares most of the advantages and disadvantages of FLI and can also image functionally in a similar manner to nuclear methods.

A particular advantage of BLI compared to FLI is that there is no auto-fluorescent back-

ground or photo-bleaching effects to contend with as there is no excitation light. There is therefore essentially zero background making specificity and sensitivity in principle relatively high.

The current major limitation, as with FLI, is that the generally unknown optical properties of the tissue through which the light has travelled confounds quantitative analyses and scattering limits spatial resolution. Also, because there is no excitation light to use, methods such as the Born ratio cannot be used to overcome this.

BLI has a developing 3D counterpart called bioluminescence tomography (BLT; introduced in section 2.5). Developing BLT is a major part of the presented work and as such the topic is discussed in detail throughout the remainder of the thesis. A review of imaging systems, technical challenges and developments is presented in chapter 4. The remainder of this chapter provides further background on BLI with discussions of the different bioluminescent marker systems, the most common bioluminescent reaction, and typical signal levels expected for emission due to internal bioluminescence at the skin surface.

3.3.1 Bioluminescent Marker Systems

Bioluminescent reactions involve pairs of molecules comprising a *luciferase* enzyme and a small molecule substrate[1]. In the case of the firefly, the originator of the most widely-used bioluminescent markers, the biomolecules are called *Firefly luciferase* (Fluc) and D-luciferin. In BLI studies, the luciferase is incorporated into the small animal, co-located with a target process or structure and the substrate is injected prior to imaging. The light-emitting bioluminescent reaction then occurs at locations where there is luciferase.

Bioluminescent agents have been isolated from several organisms, table 3.1 shows a list of marker systems that have been used in BLI studies[121, 122]. It can be seen that several markers exist and that the substrates are usually either D-luciferin or coelenterazine. Peak emission wavelengths span a range from 490nm to 612nm, with the highest value belonging to Fluc. This is important because strong wavelength-dependant light-tissue interactions typically

[1] an exception to this is the bacterial Lux operon (see table 3.1) in which the substrate situation is different[121]

Luciferase	Substrate	Peak emission at 37°C
North American firefly (Fluc)	D-luciferin	612nm
Click beetle red (CBR)	D-luciferin	611nm
Click beetle green (CBG)	D-luciferin	544nm
Renilla reniformis (Rluc) (native Rluc)	Coelenterazine	480nm
Renilla reniformis (Rluc8.6-535)	Coelenterazine	535nm
Gaussia princeps (Gluc)	Coelenterazine	480nm
LuxAB	$FMNH_2$ + long-chain aldehyde	490nm

Table 3.1: Luciferase-luciferin reporter pairs that have been used for bioluminescence imaging, adapted from tables that appeared in Prescher and Contag[121] and Luker and Luker[122]

limit signal least in the red end of the visible spectrum (section 2.1). This is a key reason for the wide popularity of Fluc and the ongoing efforts to produce red-shifted as well as higher power reporters by researchers[123].

Figure 3.3 shows emission spectra for several luciferases as measured by Zhao et al.[124].

3.3.2 The Fluc Reaction

Figure 3.4 shows the firefly bioluminescence reaction which occurs in three stages[123]. In the first stage, the luciferase catalyses the conversion of luciferin and Mg-ATP into a luciferyl adenylate (eq. (3.1)). This reacts with local molecular oxygen to create an electronically excited intermediate (eq. (3.2)). Finally the relaxation to the ground state results in visible light emission (eq. (3.3)). The reaction has a quantum yield of around 40%[125] which is high for a chemiluminescent reaction[126, 123].

3.3.3 Typical Luminescent Signals

Table 3.2 shows maximum signal reported in images in a representative sample of ten BLI studies. These studies were deliberately chosen to cover a range of reporters and application domains, and they all used the same type of imaging system (the IVIS[44] family of systems).

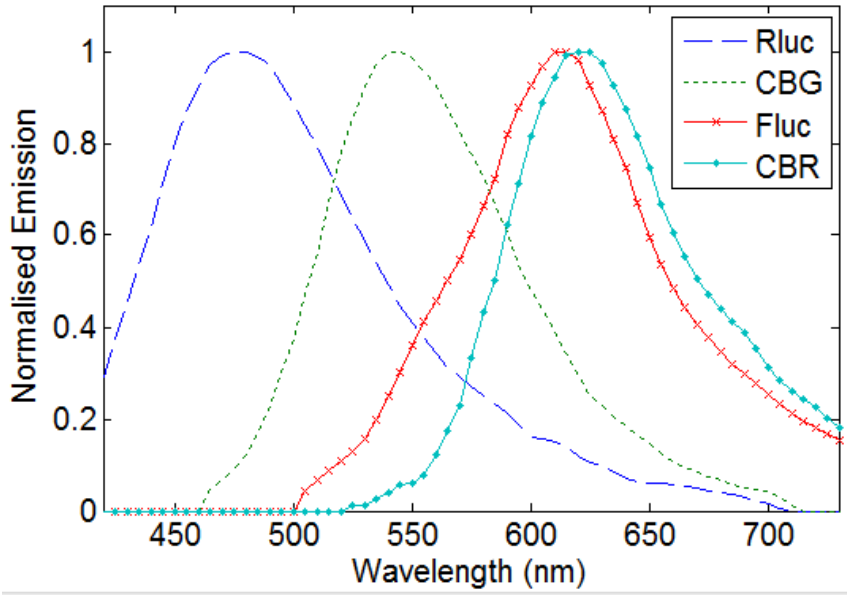

Figure 3.3: Several bioluminescence emission spectra for reactions involving different luciferases as indicated. These are adapted from Zhao et al.[124], they represent emission in mammalian cells at 35°.

$$\mathrm{E} + \mathrm{LH_2} + \mathrm{ATP} \overset{\mathrm{Mg^{2+}}}{\leftrightarrows} \mathrm{E \cdot LH_2 - AMP} + \mathrm{PP_i} \tag{3.1}$$

$$\mathrm{E \cdot LH_2 - AMP} + \mathrm{O_2} \rightarrow \mathrm{E \cdot AMP \cdot Oxyluciferin^*} + \mathrm{CO_2} \tag{3.2}$$

$$\mathrm{E \cdot AMP \cdot Oxyluciferin^*} \rightarrow \mathrm{E + Oxyluciferin + AMP} + h\nu \tag{3.3}$$

Figure 3.4: Firefly bioluminescence reaction[123]. E indicates the firefly luciferase enzyme (Fluc), $\mathrm{LH_2}$ the corresponding luciferin, ATP adenosine triphosphate, AMP adenosine monophosphate, $\mathrm{PP_i}$ pyrophosphate, and the superscripted * indicates electronic excitation. $h\nu$ indicates generated light (expressed as the energy; the product of Planck's constant h and the light frequency ν).

The values reported are nominally surface radiances, though these can only be treated as estimates as the methods used to account for imaging system perspective and animal curvature and positioning are not thorough[127] (see chapter 7).

It can be seen that a range of values spanning several orders of magnitude are reported. Typically the values are on the order of 10^5 to 10^6, though much higher values can also be found (e.g. in Barrett et al.[128]).

Given the solid angle dependence of radiance, and assuming isotropic light emission at the surface due to the many scattering events occurring in the tissue, this number can be multiplied by 2π (radians in a hemisphere) to estimate the total light flux through the surface.

Reference	Application	Reported Value (p/s/cm^2/sr)
Barrett et al.[128]	Imaging leukemia development and treatment response	10^{10}
Bhaumik and Gambhir[129]	Demonstrating the use of Rluc	10^5
Hyde et al.[130]	Studying the development of Lyme disease	10^5
Gross and Piwnica-Worms[13]	Imaging immune cell activity	10^7
Watts et al.[131]	Imaging peptide deposition associated with Alzheimer's disease	10^6
Virostko et al.[132]	Studying diabetes-related pancreatic β cell quantity	10^5
Van de Bittner et al.[133]	Imaging hydrogen peroxide production in a tumour model	10^7
Uhrbom et al.[134]	Imaging tumour (glioma) development	10^6
Tinkum et al.[135]	Imaging dynamics of a specific endogenous promoter	10^7
Stabenow et al.[136]	Imaging T cell activity in the liver	10^7

Table 3.2: Representative selection of bioluminescence imaging signal levels reported in ten peer-reviewed studies. The maximum signal reported in BLI image format is shown to the nearest power of ten. All values were reported as surface radiances and must be considered as an estimate given the methods typically used to calculate this (see chapter 7). Note that whilst several studies show results for excised tissues and isolated reporters in well-plates, the only values considered were those obtained when imaging whole mice with internal bioluminescent markers.

3.4 Conclusion

This chapter has provided an overview of prominent preclinical imaging techniques. These range from very well established techniques, such as nuclear imaging modalities and MRI, to emerging techniques such as photoacoustic tomography. An emphasis has been placed on bioluminescence imaging as this is the most relevant technique to the current work which develops its 3D counterpart. Overall this chapter has provided the broad context for the current work.

It has been discussed that the biggest challenge for BLI is that quantitative accuracy and spatial integrity are not achieved due to unaccounted-for, unknown light-tissue interactions. The presented work addresses this issue by developing a new system and work-flow for performing BLT. The developed system combines BLT with DOT to estimate optical interactions and explore new functional imaging mechanisms and surface capture to provide structural imaging of the animals exterior. This chapter has described many cases of multi-modal imaging systems

and it can be seen that this approach is justified and of great interest.

The following chapter reviews current work in the field of BLT and provides a more focused context for the developed system.

CHAPTER 4

REVIEW OF SMALL ANIMAL OPTICAL TOMOGRAPHY SYSTEMS

4.1 Introduction

This chapter reviews recent developments in non-contact small animal optical tomography systems, mostly limited to bioluminescence tomography (BLT) but with some discussion of the related fluorescence molecular tomography[1] (FMT). These are the 3D counterparts of BLI and FLI introduced in the previous chapter.

Whilst several commercial systems do some form of BLT (see Leblond[32] for a review of commercially available pre-clinical systems), there is ongoing research to improve and validate tomographic methods and accurate quantitative BLT using commercial systems is not generally reported.

4.2 The *basic* BLT system

The first, (hereafter referred to as the *basic*) BLT set-up which was utilised by several investigators[137, 54, 138, 139] involved a highly-sensitive CCD camera in a fixed position pointing at a phantom or animal held vertically on a rotating platform (fig. 4.1). Multiple, typically four, distinct angularly-resolved views of the surface can then be captured in images acquired one-at-a-time

[1] sometimes called "fluorescence-mediated tomography" or just "fluorescence tomography"

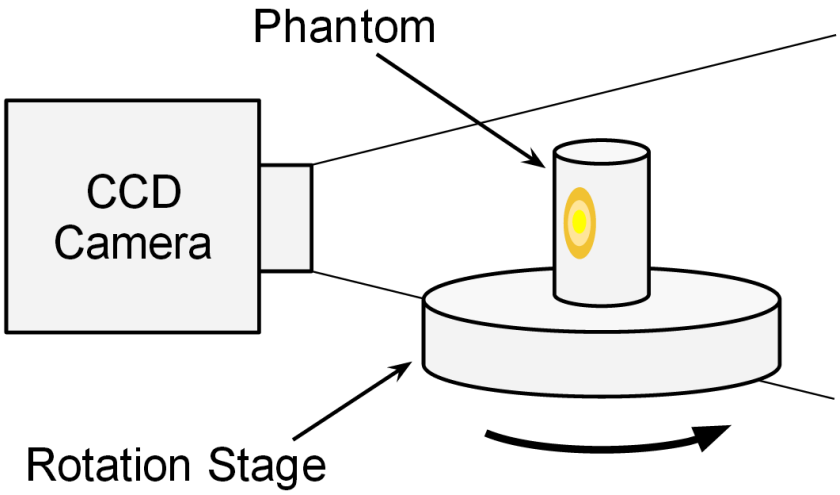

Figure 4.1: The basic BLT system design, a sensitive camera is pointed at a rotating sample stage supporting a phantom containing a light source.

following appropriate rotations of the subject providing data from all around the surface. Essentially the same set-up was used for fluorescence tomography with the addition of an excitation source[140, 141, 142, 143, 144, 145]. Whilst providing near-perfect surface coverage, such systems are limited in that multiple acquisitions are required in order to obtain a full data set. This limitation is significant because long exposure times are typically needed (on the order of minutes) in order to achieve adequate SNR when imaging deep bioluminescent sources and as such sequential imaging can result in infeasible experimental time requirements. Additionally, imaging in this particular geometry leads to difficulties with supporting the animal which necessarily must be held vertically in a therefore unnatural pose.

The major disadvantage of the basic system other than imaging time complexity (linear in number of orientations acquired from) is that there is typically not enough data to reconstruct an accurate image because the image reconstruction is so highly ill-posed (there are significantly more internal source values to find than there are boundary measurements with independent content; see chapter 2)[57, 56].

4.3 Multispectral and Multi-View Systems

It has been shown that using multi-spectral data significantly improves the accuracy of BLT image reconstruction[58, 56] by increasing measurement information content and reducing the ill-posedness of the model inversion. Whilst the basic BLT set-up can be extended to include filters and therefore collect multi-spectral as well as multi-view data sets this once again extends experimental time. Thus there is a trade-off between the amount of data and the experimental time.

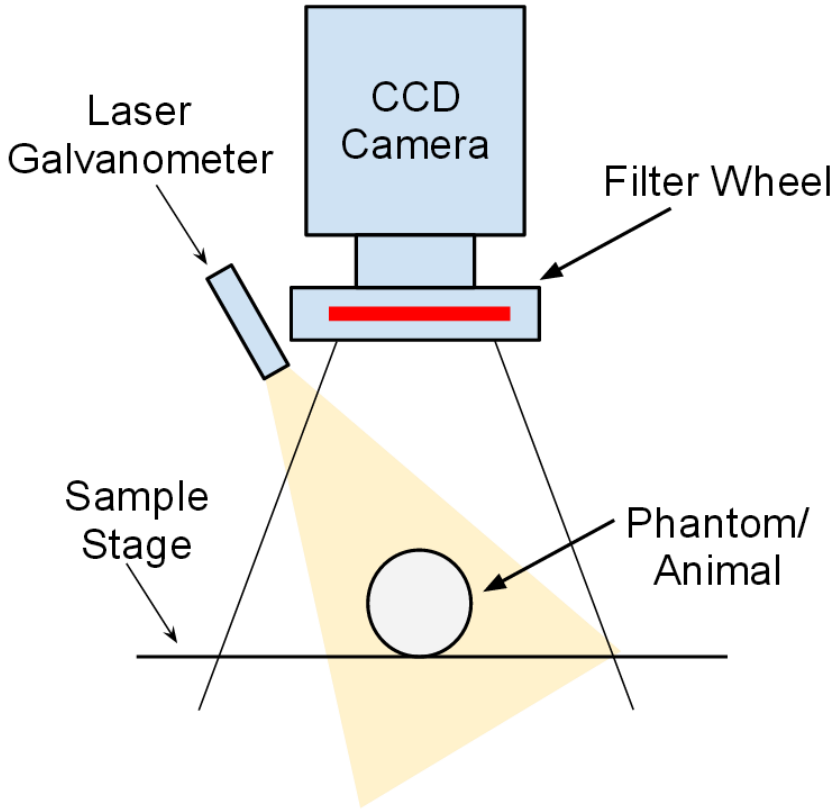

Figure 4.2: Single-view multi-spectral BLT system presented by Kuo et al.s[44]

Kuo et al.[44] presented a single-view multi-spectral imaging system utilising a filter-wheel (fig. 4.2). This system additionally included a laser galvonometer to scan the subject geometry (discussed later). An advantage to this system using the vertical light path is that a living animal can rest in a natural position and be anaesthetised lying flat on the stage.

This system was augmented by Chaudhari et al.[60] who devised a multi-spectral, multi-view system (fig. 4.3) by incorporating four 45° mirrors positioned to provide four perpendicu-

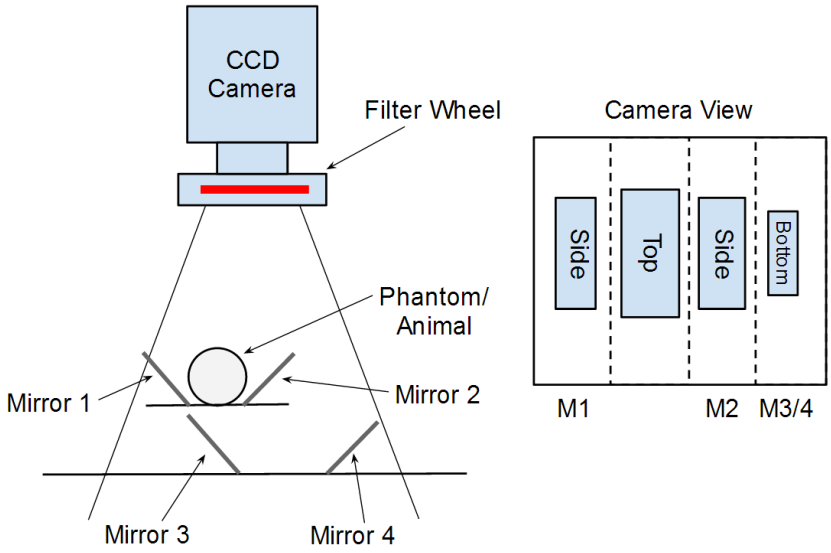

Figure 4.3: Multi-spectral, multi-view system developed by Chaudhari et al.[60] by adding a periscope-like mirror formation to the design shown in fig. 4.2. Shown on the right is the view from the camera which sees four perpendicular views around the subject, though these travel different path-lengths particularly those in the *bottom* view which are out of focus.

lar views around the subject, this is a cheap but significantly useful improvement because unlike the basic BLT system it allows multiple-view data to be acquired simultaneously, significantly reducing experimental time. A disadvantage of the approach is that it was found to require two differently focused images per wavelength owing to large optical path-length differences between views (i.e. the top and bottom views in fig. 4.3) and the images needed to be in focus due to the methods used in post-processing[1].

Li et al.[146] developed a multi-view system based on a conical mirror design with sequential spectral imaging, this allowed increased surface coverage and allowed a large amount of the CCD to be used to capture the angularly/spatially varying views making the data collection efficient compared to other designs.

Both of the above designs focus on parallelising multi-view acquisition and therefore facilitating spectral and multi-view imaging in a more feasible time-frame.

Wang et al.[147] developed a system capable of simultaneously acquiring both multi-view and multi-spectral data within a single image by the use of a *rainbow* mouse-holder and four

[1] In chapter 7 a method is introduced that can accommodate multiple path-lengths when mapping CCD image data back onto the animal/phantom surface.

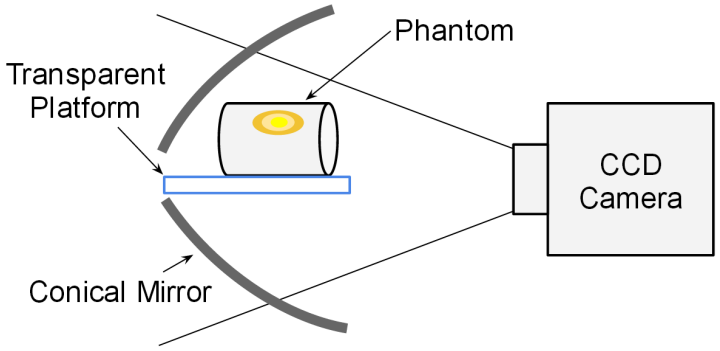

Figure 4.4: Multi-view detection system developed by Li et al.[146] to capture a great deal of the surface from a single physical view-point. Note that the actual system contained two laser scanners (one for fluorescence excitation and one for laser profilimetry) which are not shown here.

mirrors positioned around the subject (fig. 4.5). The mouse holder comprised three different coloured glass filter materials in a recurring pattern such that evenly spaced strips of the animal surface were visible at each wavelength. This approach worked well for three distinct spectral bands but if more wavelengths were required then too little of the surface might be visible at each spectral band, in addition because of the need for coloured glass filters the spectral transmittance bands were large and irregular leading to poor spectral specificity. In addition the particular placement of the mirrors and the animal meant that large parts of the CCD remained unused[147] and data was obtained in a relatively low number of image pixels. Whilst this technique is ideal in terms of acquisition time, the data is not complete spatially for each wavelength.

More recently Wang et al.[148] devised a new method for collecting multi-spectral data in single images based on a digital spectral separation device, though this has only been demonstrated for two wavelengths.

4.4 Including Surface Geometry Data

Accurate optical tomographic image reconstruction methods require knowledge of the shape of the imaged subject (sections 2.5 and 2.6). Whilst geometrically simple phantoms with known geometry are often used to validate prototype BLT systems and methods thus meeting this

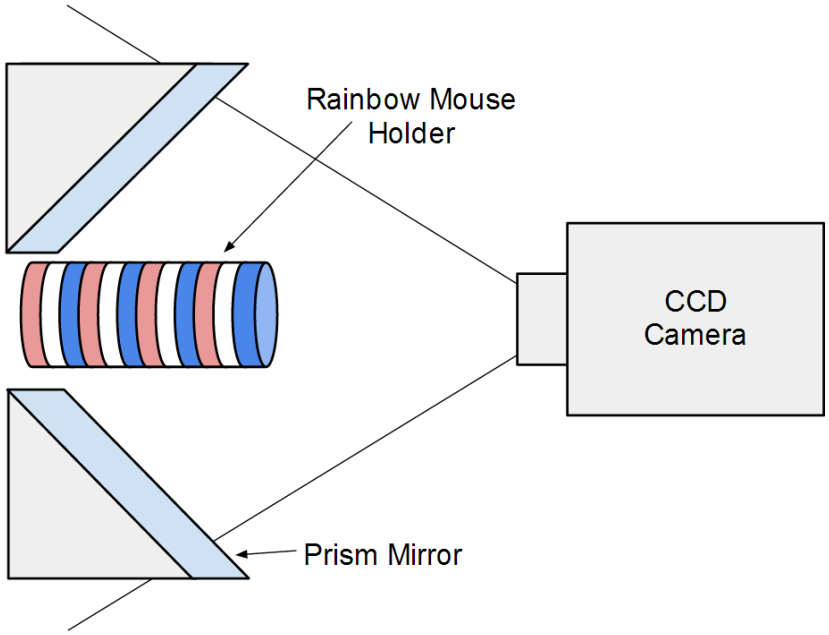

Figure 4.5: Simultaneous multi-spectral and multi-view system developed by Wang et al.[147]. Note that there are 4 mirrors in the system (of which two are shown) that are placed so as to provide two more perpendicular views behind and in front of the object in the perspective of this figure.

requirement, in general the subject shape (i.e. that of a mouse) involves complex curvature and is not known in advance. As such it is necessary to measure the shape of the animal within the scope of an imaging experiment. One set of solutions to this problem involve additional imaging using some separate structural imaging modality such as MRI (section 3.2.2) or CT (section 3.2.3)[149, 139, 150] within the work-flow shown in fig. 4.6.

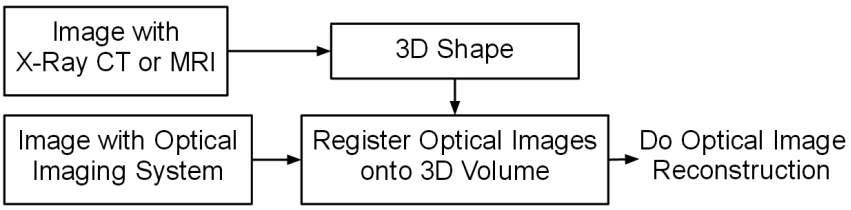

Figure 4.6: Work-flow for using separate structural imaging system for providing prior information regarding object shape to the optical image reconstruction problem.

This increases experimental cost substantially and also introduces a requirement for image registration which can be highly non-trivial. However, dual-modality visualisations can help put results into context and provide complementary data in imaging studies in addition to measuring the model geometry. Structural imaging modalities also provide an opportunity to segment

optically distinct regions to assign appropriate published optical property values so long as there is sufficient image contrast in the other modality. It has been demonstrated that this approach can improve image reconstruction[149, 139, 150], though this cannot account for optical property differences between individual imaged subjects and the published values that are assigned to recognisable, segmentable regions. The image co-registation problem can be made easier by using a mouse-holder to keep the animal in the same pose[151, 149].

Another approach is to devise multi-modal systems. Liu et al.[152] developed a dual-modality CT/BLT system in which CT data provided geometry and allowed assignment of published optical properties within a single system. Others, for example Schulz[153] and Yang[154], applied the same principles to CT/FMT. Kepshire et al.[155, 33] developed a highly sensitive time-resolved FMT and CT system in which photomultiplier tube (PMT) coupled fibres were used as photon counting detectors.

Multi-modal systems are advantageous over using separate imaging systems because the subject can stay in the same position between acquisitions making registration at least simpler and at best unnecessary. Experimental time as well as time-based variations in e.g. anatomy or functional physiology can also be minimised. Co-registration of the multi-modal data is still required though it only needs to be done once assuming that the instrument is fixed and well-understood.

A significantly simpler and lower cost alternative to using CT or MRI is using optical surface capture methods (e.g. structured light techniques[156]). Deliolanis et al.[141] developed an FMT system that utilised multiple angularly resolved optical projections to reconstruct animal geometry. This requires some added experimental time and complexity due to the need to rotate the sample and acquire many images. Li[146] utilised a laser line scanning system included in the conical mirror system (fig. 4.4) to capture geometric data. A method that is simpler and cheaper, based on sinusoidal structured light projection, was used by Kuo et al.[44] (fig. 4.2) to capture the directly visible portion of the animal surface and a similar method has recently been developed by Basevi et al.[3] that can additionally capture surfaces visible in mirrors[1]. These

[1]This is the approach used in the presented work, discussed in detail in chapter 6.

approaches are advantageous because neither optical components nor the animal have to move between images and in the latter case multiple surface-views are obtained simultaneously.

4.5 Multi-Modality Imaging

Beyond secondary systems that provide structural/anatomical priors for BLT, other multi-modality systems have been developed that provide complimentary imaging data for multi-modal studies thus providing enhanced scientific information content.

Cao et al.[104] developed a multi-modal single-photon emission computed tomography (SPECT), CT and optical system for BLT and FMT that utilises the geometric information from CT and uses SPECT to obtain prior information which informs FMT or BLT reconstruction. It was shown that reconstruction with SPECT priors was better than without. Alexandrakis et al.[31] proposed a system for combined optical and positron emission tomography (OPET) imaging which is designed so that the cylindrical (physically tomographic) detector array can detect both visible light and emitted gamma rays[157, 158].

It has been shown explicitly that BLT reconstruction performance is strongly improved by the use of accurate heterogeneous models of optical property distributions as opposed to assumptions of homogeneous or inaccurate properties[159, 152, 160]. Razansky et al.[161] showed that by utilising absorption measurement by integrated PAT, FMT image reconstruction could be improved.

It has been suggested that DOT could be used to obtain optical property measurements and shown that this is effective in simulation[162, 163, 160]. The first small animal systems using DOT have also been recently presented. Zhang et al.[163] showed that DOT using a single laser diode integrated within a basic BLT system improved reconstruction whilst Tan et al.[164] performed DOT alongside FMT using a single laser for both modalities within a basic set-up. Pekar[165] developed a CT-DOT-BLT system utilising a laser diode source. Within this system, hard and soft prior approaches to DOT were carried out using CT-segmented regions building upon methods where this data is used to assign published properties to regions. A

similar data flow concept was utilised by Yan et al.[166] who built a gantry-based fully rotating multi-modality system comprising a CT system, an optical detection system and DOT sources in the form of two lasers. Within this system CT priors were used and DOT reconstruction was performed at two wavelengths, absorber concentration was then deduced from maps of absorption based on the knowledge that only two absorbers were present. This method is an indirect approach to spectrally constrained DOT. However, this system carries the same disadvantage of the basic BLT system set-up[54] in that multi-view imaging is sequential and therefore time-consuming. In contrast to these methods using point-like excitation sources, Chen et al.[167] and Venugopal et al.[168] developed a small animal time-resolved DOT and FMT system based on a laser-coupled digital micromirror device (DMD) based wide-field illumination scheme allowing spatial patterns to be used and demonstrated that structured illumination and time-resolved detection improved upon standard methods. Multi-modality approaches have also been used to improve small animal DOT, for example Gulsen et al.[95] developed a DOT and MRI system and recent studies have shown that this fusion approach provides enhanced quantitative accuracy[169] and resolution[170].

4.6 The Presented System

In most cases where existing systems perform DOT to provide BLT with priors, they either reconstruct optical properties at particular wavelengths in which case this must be done for each wavelength for which BLT data is used, or do this and then fit the results to chromophore concentrations and scattering parameters after reconstruction[166]. The capacity for this second approach is limited in many existing systems due to the use of monochromatic sources.

In the current work, a novel combined multi-spectral DOT-BLT system is presented which additionally uses multi-view image acquisition and multi-view optical surface capture. A novel work-flow is used (figure 4.7) whereby optical surface capture is followed by spectral DOT and finally BLT providing two imaging end-points; a 3D functional image of chromophore concentration and a 3D luminescence image. The following chapter discusses the presented

system itself in detail.

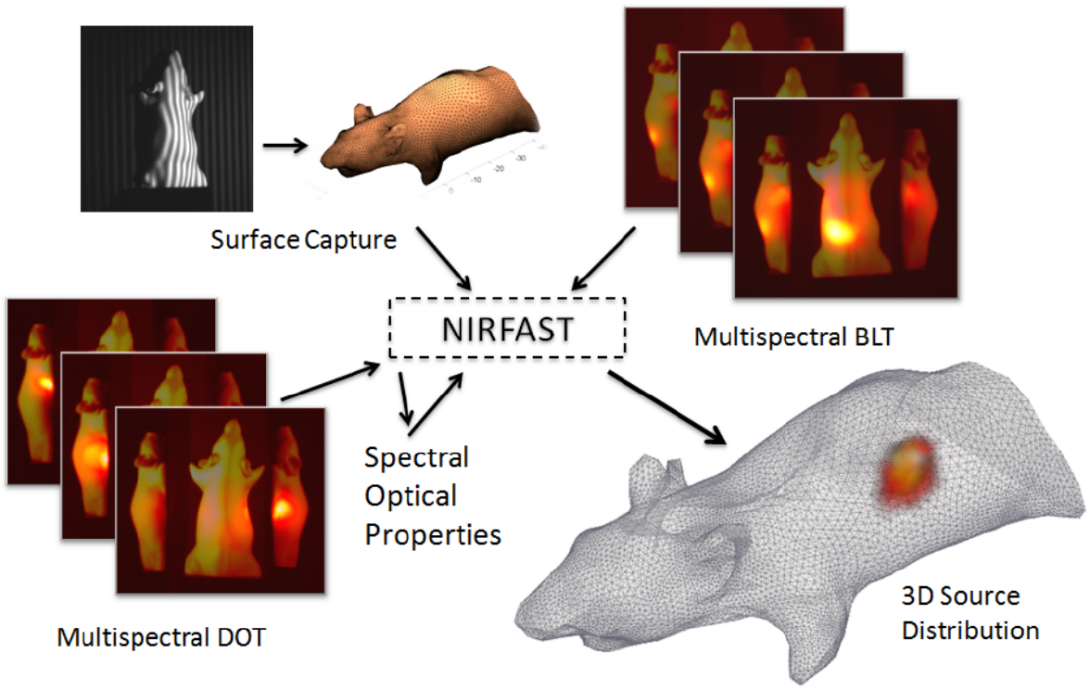

Figure 4.7: Visual representation of the system concept. A mouse is surface captured to obtain its geometry and is then imaged in spectral luminescence and in spectral near-infrared trans-illumination modes. Using NIRFAST[25], DOT is used to reconstruct chromophore, scattering and subsequently functional parameters which are additionally used to inform reconstructions of bioluminescent source distributions.

CHAPTER 5

IMAGING SYSTEM DESIGN AND CONSTRUCTION

5.1 Introduction

In this chapter, the developed novel system for performing multi-modal *in vivo* imaging of small animals is described. The system is designed to perform optical surface capture (SC; requiring an optical detector and a source of structured light; chapter 6), diffuse optical tomography (DOT; requiring an NIR detector and NIR sources; section 2.6 and chapter 8), and bioluminescence tomography (BLT; requiring a highly sensitive optical detector; section 2.5 and chapter 9). As discussed in the introduction, the purpose of this multi-modal scheme is to provide several complimentary data sets facilitating, ultimately, 3D quantitative *in vivo* imaging with overlaid intrinsic contrast (based on absorption, scattering, chromophore concentration) and extrinsic contrast from targeted bioluminescence sources.

Each part of the system is discussed and the choice of components and designs are justified such that their individual and combined performances can be fairly evaluated. The system is described in detail such that it could be reproduced. Rigorous experimental characterisation is the subject of appendix A.

5.2 System Overview

The presented system, shown in figure 5.1, follows the basic layout of several established *in vivo* optical imaging systems, such as that of Kuo et al.[44], with a vertical light path and a horizontal stage to support small animals beneath a camera. The same detection system is used for SC, DOT and BLT, minimising cost and complexity, and not requiring co-registration. The system is encased within a light tight box with components suspended within a cage-system.

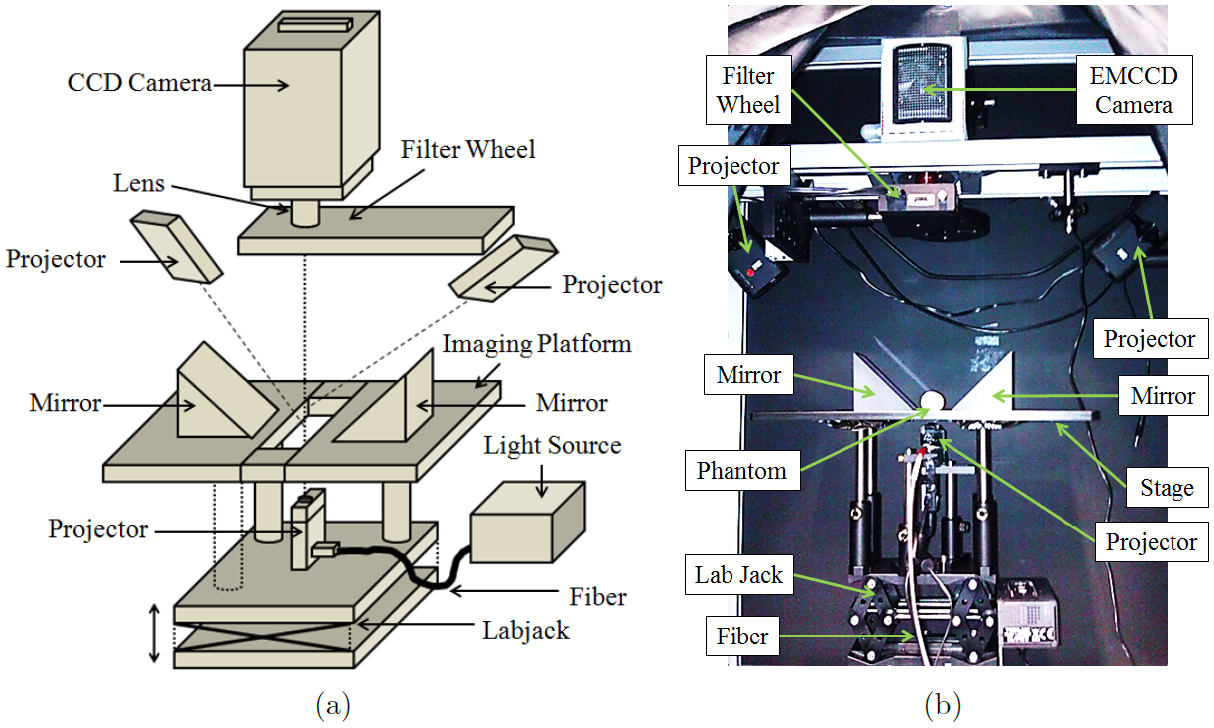

Figure 5.1: (a) Labelled schematic and (b) photograph of the developed imaging system. This figure is reproduced exactly from Guggenheim et al.[1].

5.2.1 Multi-View Detection Scheme

In a novel addition compared to most standard systems, mirrors are included to facilitate the imaging of three perpendicular views of the animal surface in all three imaging modalities.

The importance of maximising surface coverage and, therefore, number of measurements, is that this reduces the ill-posedness and non-uniqueness of BLT and DOT image reconstruction problems (chapter 2). Imaging from several sides also reduces the minimum possible source

depth (and therefore the minimum attenuation) across all perspectives, improving sensitivity. Imaging more of the surface also allows a more complete surface capture and therefore a more accurate geometric representation of the imaged subject[3] (chapter 6).

In the simplest case more data could be obtained by including more detectors, but this would increase the complexity of the system and significantly raise the cost. It is also possible to rotate the animal or detector as in the system of Yan et al.[166] though this adds mechanical and operational complexity to the system and means that imaging must be sequential. Owing to the long exposure times required this could lead to very long total imaging times. It was chosen instead to use mirrors, an approach also used effectively by other investigators[147, 171, 60], because simultaneous, multi-view imaging is achieved with very little added cost ($< 1\%$ of the cost of the camera).

In preliminary experiments and simulations, it was found that this multi-view detection scheme improved the results of BLT image reconstruction[9].

5.2.2 Multi-Spectral Detection Scheme

In addition to multi-view data, the presented system is designed to obtain multi-spectral data for both BLT and DOT for which there are several reasons. Firstly, multi-spectral data provides a highly significant improvement in BLT image reconstruction (chapter 4 and section 2.5). Secondly, spectral data can improve DOT reconstruction and enables functional imaging (section 2.6). Finally, spectral DOT reconstructs spectral properties which can be used to calculate μ_a and μ_s' at wavelengths where data was not originally obtained (section 2.6). This feature is used within the presented system in that NIR wavelengths are used for DOT, and reconstructed properties are used to calculate μ_a and μ_s' at visible wavelengths at which BLT measurements are obtained to provide prior information in BLT reconstructions. This approach was chosen because generally there is significantly less attenuation at NIR wavelengths (section 2.7.1) meaning that sensitivity is higher, exposure times can be shorter and the diffusion approximation is more likely to work accurately (section 2.3). Multi-spectral imaging is enabled by the addition of a filter wheel to the system design.

5.2.3 Wide-Field DOT Illumination Scheme

Mounted beneath the sample stage is a small optical projector coupled to a near-infrared (NIR) light source for the projection of spatially resolved NIR light, i.e. DOT sources (section 2.6), onto the underside of the animal.

The projector acts as a wide-field source[167] that allows the projection of point sources as well as spatially modulated distributed sources onto the underside of the animal within the region of illumination. The design was chosen for the flexibility in terms of possible wide-field illumination schemes that can be studied in the future, though in the scope of the presented work only point-like sources are used.

5.2.4 Two-Sided Surface Capture System

The system includes two further optical projectors suspended above the stage which are used for optical surface capture (chapter 6). The two-projector set-up was chosen to maximise the surface coverage by the illumination and to make the best use of both system mirrors in capturing the sides of the surface.

5.3 System Components

5.3.1 Enclosure and Cage System

The cage system and housing serve two key purposes, (1) to provide a structural basis for the system and rugged physical support for the components, and (2) to render the system dark and thereby remove unwanted sources of background light that could reduce SNR in imaging experiments. Note that the latter is particularly important considering the low-light levels expected when imaging bioluminescence in general[121], and especially through highly attenuating tissue (section 2.7.1).

The suspended components of the system are fastened directly to a cage system constructed

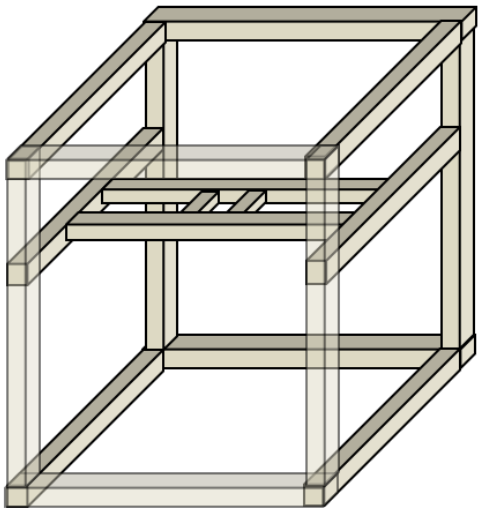

Figure 5.2: Diagram of the aluminium cage system to which components of the imaging system are fastened. The cross-beam structure in the upper-middle part is designed to support the camera very securely due to its high value and relatively large weight; the camera rests on this structure as well as being mounted to the cage by a screw. Other components are mounted directly or via steel posts to the cage.

from aluminium posts (RS Components, Corby, UK) which is itself attached securely to an optical bench (Thorlabs, Cambridgeshire, UK), an outline is shown in fig. 5.2. These posts are convenient because they allow components to be securely attached to them at any point. The design of the cage was chosen to securely support the camera in particular with two cross-beams, this is important because of the high value of the camera and because it ensures there is no movement of the system components relative to one another ergo geometric calibrations remain valid and the field-of-view is constant.

The light-tight enclosure was constructed by attaching aluminium panels to the outer cage struts so as to fully cover the sides, top, and back of the structure. Panels were painted matte black to minimise light reflection inside the box. The front of the cage was covered by two large double-folded sheets of blackout cloth material (Thorlabs, Cambridgeshire, UK) attached so as to allow that the front be either open (allowing interaction with the system) or closed (rendering the box light-tight). The blackout cloth was also placed over each of the joins between panels so as to block any stray light.

5.3.2 Optical Detection System

The optical detection system is designed to capture multi-spectral, multi-view images of light radiating through the surface of a mouse.

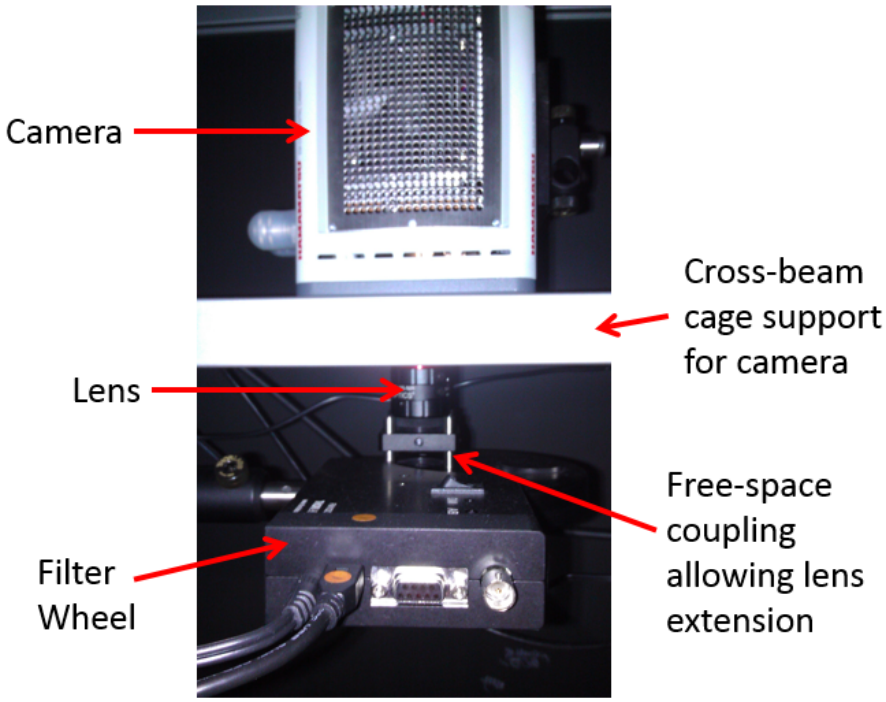

Figure 5.3: Close-up photograph of the coupling between the camera, lens and filter wheel in which the lens is partially obscured by the cross-beam supporting the weight of the camera. The free-space coupling system connecting the lens to the filter wheel can be seen.

The optical detection system comprises the camera, lens, filter wheel and prism mirrors. The back of the lens is connected directly to the camera whilst the front is connected by a small cage-plate and rod system to the filter wheel as shown in fig. 5.3. The rods surround a short free-space coupling which allows the physical extension of the lens housing (required when the lens focus is adjusted) without movement of the filter-wheel. By virtue of the lens, the camera is focused through the filter wheel roughly at the height of the sample stage. The camera captures images containing a direct view of the imaged surface and two side-views through the mirrors placed on the stage as illustrated in fig. 5.4.

The long front focal distance of the lens ($\approx 300mm$) means that the difference in path length from the lensing to the surface when imaged via the mirrors as compared to directly is a small fraction of the total distance and as such the focusing directly and via mirrors can be made

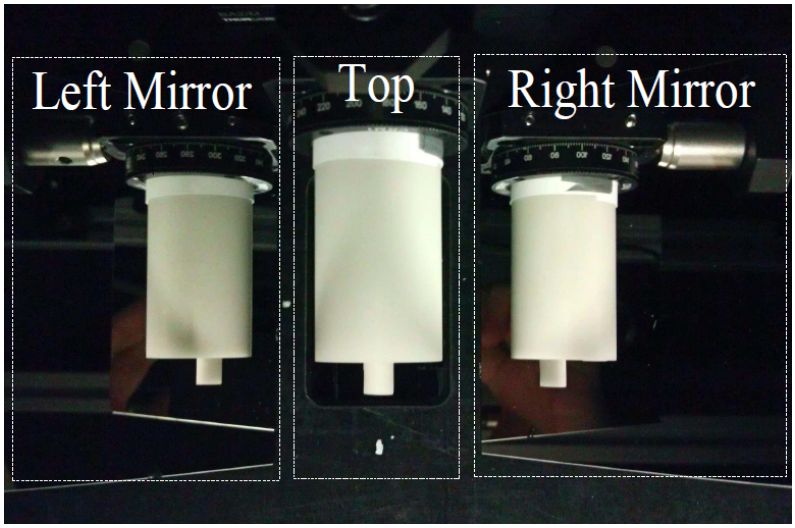

Figure 5.4: Approximate detection system perspective when mirrors are placed around a cylindrical phantom (appendix A.1.1) on the sample stage. Three views are visible; one direct and two mirror views. This image was acquired by placing a colour camera directly in front of the optical detection system, the figure is replicated from Guggenheim et al.[5].

similarly effective.

The two-mirror set-up was chosen as opposed to other approaches such as the conical mirror of Li et al.[146] or the periscope design of Chaudhuri et al.[60] for simplicity- benefiting understanding and modelling- and because it leaves the underside of the animal free from obstruction allowing DOT illumination. Furthermore, three perpendicular views can provide coverage of practically all of the animal surface that is not in contact with the stage[3] and, given the lack of independence of nearly co-localised detectors particularly when light is diffuse, there is an argument to be made that there is little new information acquired by imaging from more projections than three.

The mirrors are freely placed on the sample stage; they are not fixed down. Instead their locations are measured on-the-fly in imaging sessions (section 6.2.6). This approach was chosen to allow flexibility when imaging subjects of different size and position, and to allow optimal placement near a region of interest.

Several recent results[6, 1] (chapter 9) suggest that the addition of the multi-view data is indeed beneficial to image reconstruction in practice within the presented system.

5.3.3 Camera

The camera (figs. 5.1 and 5.3) is a Hamamatsu ImagEM-1K (Hamamatsu Photonics, Hamamatsu City, Japan), which is an air-cooled, 1MP, back-thinned, electron-multiplying (EM) CCD camera, another version of which was used in the FMT system of Li et al.[146] (chapter 4).

The physical CCD pixels are $13\mu\text{m} \times 13\mu\text{m}$ in size and can be binned in hardware $1\times$, $2\times$ or $4\times$ creating effective imaging pixel areas of up to $2704\mu\text{m}^2$ within a total detection area of approximately $13.3\text{mm} \times 13.3\text{mm}$. The camera quantum efficiency is $> 40\%$ across the spectral range of interest ($500\text{nm} - 900\text{nm}$) and $> 85\%$ in the luminescence region ($500\text{nm} - 700\text{nm}$) where low-light conditions are expected.

The camera was chosen for its high sensitivity over the VIS-NIR range and the opportunity to investigate the use of EM gain.

In EM-mode, EM gain allows the amplification of signals before readout, nominally by up to $1200\times$ (though this was not achievable in practice; appendix A.2.4), effectively reducing read noise[172]. However, though in some imaging scenarios this can bestow an overall SNR increase, it also introduces noise and is not always beneficial. Following the analysis of the system noise characteristics (appendix A.5) in which it is concluded that the SNR will not improve in typical use of the presented system, it was decided that EM mode would not be used.

The camera has several parameters in addition to the *binning*, *CCD mode* (EM or normal) and *EM gain* level. The *read mode* selects the speed at which image values are read out from the CCD; faster readout comes at the expense of greater noise (appendix A.5.5). *Analogue gain* applies a linear amplification post-readout. *Exposure time* can be varied between a few milliseconds to an hour.

Camera characterisations such as the validation of readout, dark, and EM-related noise characteristics (appendix A.5), establishing of the non-linearity of the EM gain (appendix A.2.4) and compilation of a library of mode-specific digitizer offsets for background subtraction (appendix A.2.3) are detailed in appendix A.

5.3.4 Lens

The lens is a 25mm fixed focal length VIS-NIR lens (Techspec, Edmund Optics, York, UK) which has a variable aperture ranging from $f/1.4$ to $f/17$ (always fixed within the system to $f/1.4$ - the largest possible - so as to collect the maximum signal possible). It was chosen for its high transmittance in the visible and near-infrared (NIR) spectral regions. Its minimum working distance is 100mm and the field-of-view is 19.8°, which allows the full region of interest on the sample stage to fit into an image at a working distance of approximately 300mm. It was considered that a telecentric lens might be used as this could make modelling simpler, however it is hard to find a lens that could manage the necessary field of view. It would be interesting to compare performance with a physically larger lens that could collect more light as a sensitivity increase could be gained.

5.3.5 Automated Filter Wheel

The automated filter wheel (FW102c; Thorlabs, Ely, UK) has six positions and accommodates $1''$ diameter circular filters. The housing allows the wheel to be removed and replaced quickly allowing for fast swapping of whole filter-sets, when required. This filter wheel was chosen for simplicity and flexibility; the ability to incorporate any $1''$ filter means that different combinations of wavelengths can be used, investigated, and tested. Additionally the ability to quickly change whole filter sets provided a suitable method for switching the system between spectral luminescence and spectral NIR trans-illumination modes (chapter 9, chapter 8).

5.3.6 Interference Bandpass Filters

In the presented work, 10nm full-width-half-maximum (FWHM) interference-based bandpass filters (Thorlabs, Cambridgeshire, UK; fig. 5.5) with central wavelengths in the range 500nm – 850nm are used in the filter wheel for spectral imaging. The choice of 10nm band-pass was made to obtain a good balance of spectral specificity and total transmittance and intereference-based filters were chosen because of the wide variety of available pass-bands; it being possible

to obtain a filter for essentially any central wavelength in the range of interest. However, the use of interference filters does introduce a challenge in that the pass-band characteristics have an angular dependence (appendix A.2.5). The filters used in this work have a peak transmittance of around 50 to 60% which can be improved upon with added cost, this provides an opportunity for the future for improving system sensitivity.

Figure 5.5: Photograph of a subset of the interference-based bandpass filters used in this work. From left to right the central wavelengths are 540, 700, 570, and 630nm.

5.3.7 Right-Angle Prism Mirrors

The mirrors (fig. 5.6) are 75mm right-angle prism mirrors with enhanced aluminium coating (N-BK7; Edmund Optics, York, UK). The coating was chosen for its high reflectance which is $R_{mean} > 95\%$ in the luminescence region ($500 - 700$nm) where low-light conditions are expected. The mirror size was chosen to match approximately the length of a typical mouse body.

5.3.8 NIR Light Source

The light source (fig. 5.7) consists of a pocket projector (PK-102; Optoma, London, UK) fixed beneath the stage and coupled via a 1000μm diameter, 2m long optical fibre (QP1000-2-VIS-BX; Ocean Optics, Oxford, UK) to a tungsten-halogen lamp (HL-2000-FHSA, Ocean Optics, Oxford, UK; fig. 5.7). The use of the modified projector as a source of wide-field illumination for small animal imaging follows a published design[167, 168].

The projector was modified in that the LEDs and dichroic mirrors that usually provide its coloured illumination were removed, and its housing was drilled and fitted with a fibre-adapter

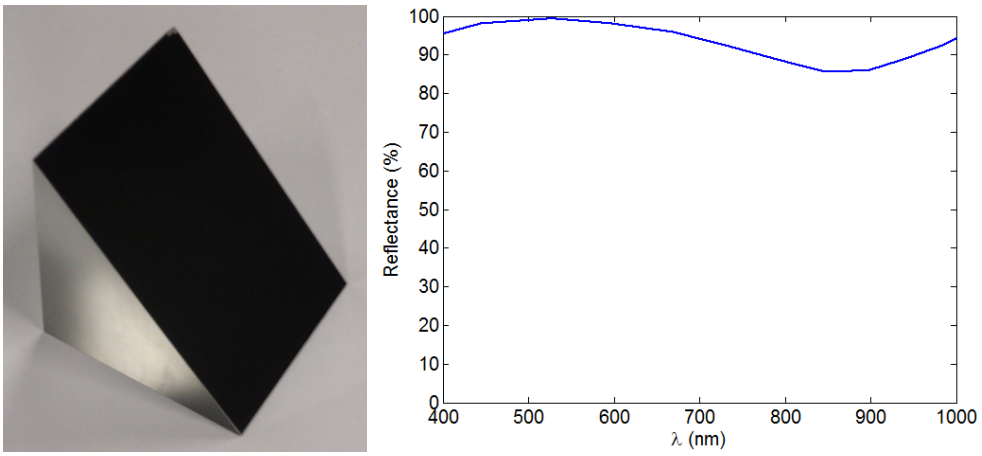

(a) (b)

Figure 5.6: Photograph of one prism mirror and nominal reflectance spectrum (manufacturer data).

to allow the reception of the optical fibre. The fibre output directly replaces the original sources; it is incident first upon a diffuser then a micro-mirror array within the unit. Any desired pattern of NIR excitation can be selected using a graphical input which is then projected under the sample by the unit.

5.3.9 Adjustable Sample Stage

The sample stage (fig. 5.7b) holds the imaged subject between the mirrors above the DOT source. It is a $400 \times 300 \times 10$mm black acetal sheet with a 30×50mm hole machined in the centre to allow transmission of light from the projector directly underneath. The material was chosen because it is adequately rigid to support objects, has low reflectance to minimise stray light, and can be easily cleaned.

The stage is mounted on top of a motorised lab jack (L490MZ; Thorlabs, Ely, UK) by four 160mm vertical posts. The lab jack has a 52mm range of travel with a repeatability of 5μm. It was included to allow the adjustment of the effective detection system focusing automatically and precisely. It was also found that the precise height adjustment could be used in geometric calibrations (appendix A.6).

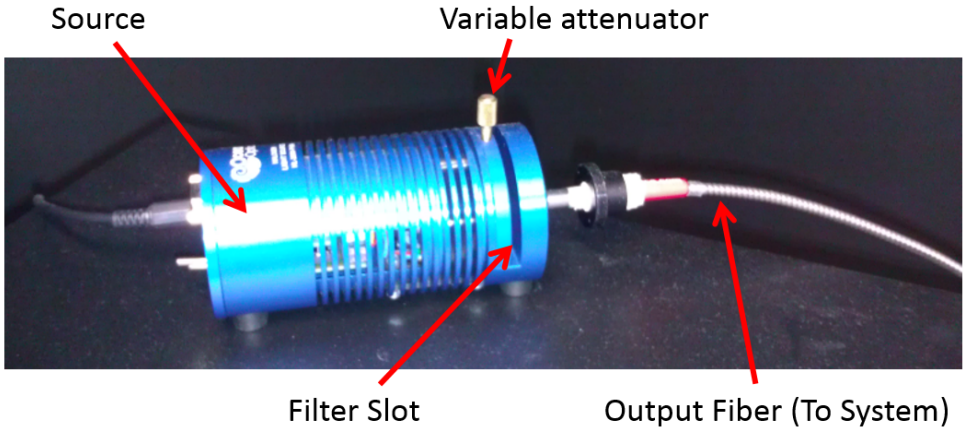

(a)

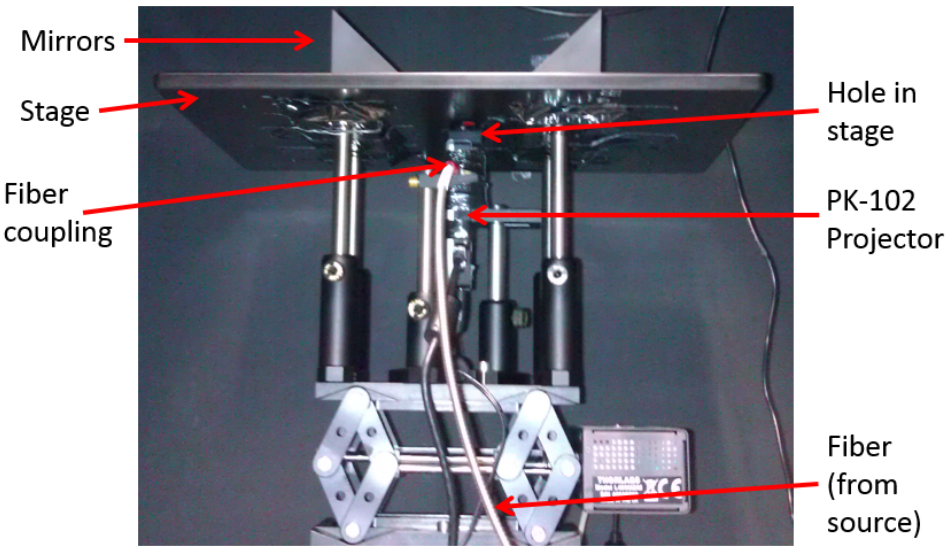

(b)

Figure 5.7: Photographs of (a) Tungsten-Halogen source and (b) PK-102 projector in-situ.

5.3.10 Surface Capture System

The surface capture (chapter 6) system comprises the optical detection system and two projectors (MPro120; 3M, Bracknell, UK) shown in fig. 5.8. The projectors are fixed to the cage (refer back to fig. 5.1) and powered independently with their batteries removed, they are arranged so as to point roughly at the centre of the sample stage and angled so as to illuminate opposite sides of an imaged subject to allow maximum surface acquisition in conjunction with the use of the mirrors. The projectors were chosen for their low cost and ready availability. They have an adjustable focus which is fixed such that they are roughly focused on the centre of the stage.

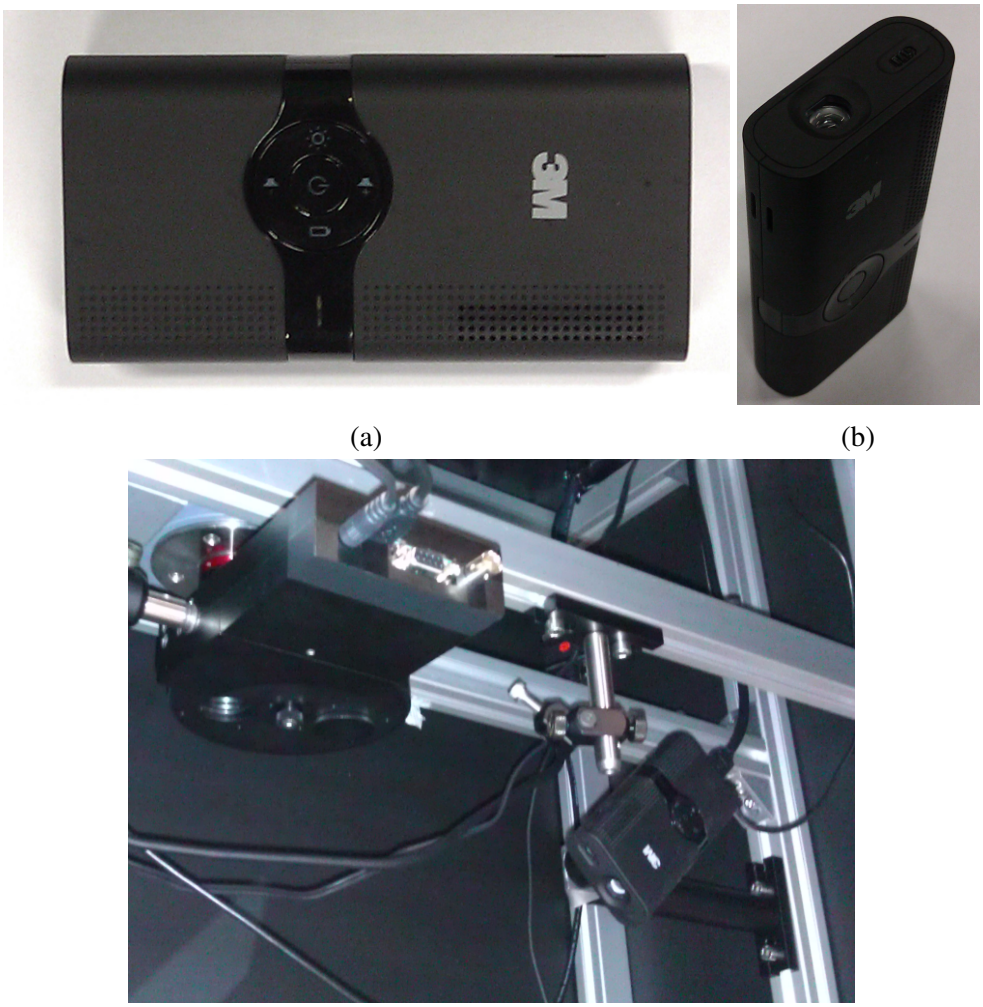

(a) (b)

(c)

Figure 5.8: Photographs of one of the surface capture projectors from two views and mounted in the system.

5.3.11 Computer

The camera, filter wheel, projectors and lab-jack are connected to a computer (Viglen Genie with Intel DQ67SW Motherboard, Intel Core i7 Processor i7-2600 (3.40GHz), Quad Core with 8MB Cache, 16GB of RAM and 2TB hard-disk drive) running 64-bit Windows 7 Enterprise (Microsoft). The computer has an NVIDIA GeForce GT520 graphics card installed so that in total (including the on-board graphics) it has 4 graphics outputs which are used to connect a monitor and the three system projectors. The filter wheel and lab-jack are connected via USB whilst the ImagEM camera is connected through a dedicated video capture card (PHOENIX-D24CL-PE1; Active Silicon Ltd, Iver, UK).

5.4 System Image Processing

Images acquired by the imaging system are generally processed in the following way in order to extract quantitative and reliable measurements.

5.4.1 Conversion from Image Grey-Levels to Real-World Units

The number of photons that irradiated the CCD are related to the values of an acquired image by[172]:

$$P = \frac{\zeta(\rho)}{G_a G_s Q}(I - I_o(\rho)) \tag{5.1}$$

where ρ is a parameter term for the current CCD mode, read mode, and binning mode. $\zeta(\rho)$ (e^-/count) is a scalar conversion factor dependent on CCD mode which is 0.6 in normal mode and 6.3 in EMCCD mode for the system camera (appendix A.2.2), G_a (no units) is the analogue gain applied which is always set to 1 in experiments, Q (e^-/photons) is the quantum efficiency of the camera, G_s (no units) is the sensitivity (or EM-) gain applied which is also always set to 1 in experiments (appendix A.5), I (counts) is the image, $I_o(\rho)$ (counts) is the counts that are present not due to incident light comprising the digitiser offset and any characteristic thermal background signal. The digitiser offset is a property of the camera that is dependent on read-mode, CCD mode and binning mode. By taking many repeated images with the lens cap on, it was measured in each mode and committed to a library for recall and subtraction (appendix A.2.3). The quantum efficiency is not known, it is set to 1 in this work and measurements that result from the above conversion are in units of e^- read out from the CCD until further calibrations or calculations are applied dependant on the imaging mode. With no analogue gain applied and working with electrons, $S = QP$, so as not to consider quantum efficiency, the relationship is:

$$S = \frac{\zeta(\rho)}{G_s}(I - I_o(\rho)), \tag{5.2}$$

or equivalently,

$$I = \frac{SG_s}{\zeta(\rho)} + I_o(\rho). \tag{5.3}$$

5.4.2 Removal of Image Artefacts

The ImagEM camera technical note[172] indicates that unwanted bright spots will appear in images on some occasions due to stray cosmic rays. In order to remove these from images, images are compared to a median filtered version and points that are found to differ in value strongly from the median filtered version are thrown away in the image processing phase. It is conjectured that this is a processing step that will not result in throwing away genuine signal because diffuse optical images tend to result by definition in smooth value distributions on the CCD. So far this technique has appeared useful for removal of erroneous signal that could substantially mitigate image reconstruction performance, though in some situations it is possible that genuine signal could be removed and more thorough testing throughout future studies will be required.

5.5 Automated Acquisition

The system is controlled by a custom-made Labview program (National Instruments, Newbury, UK), which manages all aspects of data acquisition and on-line processing and is designed to be flexible and easy to use when imaging. Imaging runs are specified using run-files (simple .csv files adhering to a common pre-defined format) which specify system parameters for arbitrarily many images that will be acquired in the order specified. The adjustable fields include: CCD mode, readout mode, analogue gain, sensitivity gain, binning level, exposure time and projector image (NIR excitation pattern). The sequence of an example imaging run is shown in figure 5.9.

It is necessary that certain operations are performed in sequence in the order shown as several camera parameters (indicated by an asterisk ('$*$') in figure 5.9) affect the range of available values and the default value for other parameters. For example setting the CCD mode changes the applicability of the sensitivity gain feature, the range of possible exposure times, the range of possible readout modes and the current exposure time and readout mode. Imaging sessions consist of a simple loop in which parameters are set and images are acquired and subsequently

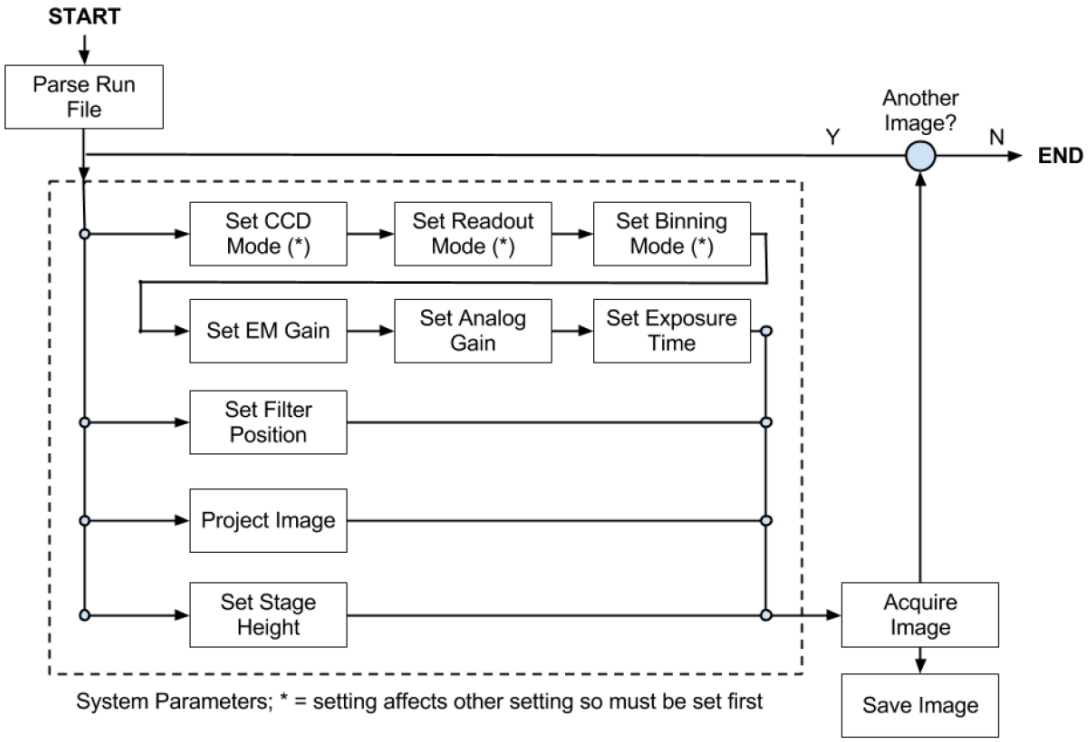

Figure 5.9: General imaging run work-flow. Note that whilst "Project Image" is only shown once, it represents a total of three parallel operations in which a projection is done with any or each of the three projectors in the system. This figure is reproduced from Guggenheim et al.[1]

saved. Images are saved in sequence and cleared from virtual memory so that long imaging sessions can be performed without exceeding system memory. The image management does introduce some temporal overhead and as such there is approximately 500 ms delay between successive image acquisitions (which is a small fraction of the time taken in most imaging runs, which is typically several tens of seconds per image).

Image data is saved in Matlab (.mat) format along with all corresponding imaging parameters which is useful for de-bugging, clarity and data processing and analysis. The .mat format was also found to be the most efficient lossless compression scheme as compared to .PNG and .TIFF formats for typical images acquired with the system.

5.6 Limitations of the Current Design

The optical detection system in its current form could be made more sensitive. Current high-end hardware includes filters that are almost twice as transmissive at the peak, cameras that have four times the physical detection space, and higher power, larger aperture lenses. An example of a system that is substantially more sensitive due to better hardware in these terms is that discussed in Behrooz et al.[173]. Nevertheless, as an instrument with which to investigate new ideas and the potential of new processing methods the presented system is perfectly adequate, it is also still very sensitive but it should be remembered when discussing the characteristics such as necessary experimental acquisition times that these can be improved. The reason that the current system is not upgraded is because of the high monetary cost involved.

A limitation of the current form of the system is that it does not presently include a facility for constant anasthesia supply or thermal regulation for maintaining the appropriate body temperature in small animals. Such systems are included as a standard in several commercially available systems such as the IVIS[44]. These should be added at such a time as it is used for *in vivo* imaging. Before this is the case there is little benefit to the addition, however, the system has been found to heat up (as the camera cools via a Peltier system but is wholly contained when the system is closed) therefore in the mid-term some form of ventilation would be a beneficial addition (see section 10.5).

An unwanted technical issue is that the projector used in the DOT light source is somewhat unreliable in that it has been observed to flicker occasionally (i.e. abandon the current projected image and emit a short flash). This is believed to be due to some instability in the hardware which is not designed for scientific use, although it could feasibly be an issue with the graphics card too. The stability has been studied for current typical acquisition times (appendix A.5.8) and was found to be reasonably stable over these time periods, however in general it would be better to replace the light source with one that is more stable and therefore does not restrict experimental settings or raise unwanted doubts.

Another known issue is that there is some weak light leakage inside the filter wheel housing, this was detected as a background in BLT experiments where long exposures were used.

Known potential sources of the light were removed or ruled out meaning that the source is unknown. As a result, between acquisitions in important experiments the filter wheel is manually powered off to ensure image integrity. This undermines the automated system acquisition and therefore should be addressed by further work. An option is to install an extra switch that can be controlled by the computer to power it down automatically. Another possibility is to replace the filter wheel with another model without the issue.

5.7 Conclusion

The presented system is designed to provide a new combination of optical data sets; those of optical surface capture, bioluminescence tomography and diffuse optical tomography. This chapter has introduced the key design considerations and constructions that are brought together to achieve this.

In order to be used routinely there are adaptations that should be made to the system to overcome the limitations discussed. These are mostly hardware alterations that do not compromise the intellectual developments made with the system as-is.

This chapter has described the system and explained the choice of components and their general behaviour and use. As has been indicated, in-depth system component characterisations and tests are presented in appendix A, which therefore provides further insight into the strengths and weakness of the system and components. The remainder of the thesis is concerned with presented imaging methods, key innovations and system-wide technique evaluations. This begins in the following chapter which discusses surface capture and meshing.

CHAPTER 6

SURFACE CAPTURE

Surface capture refers to the measurement of the surface profile of an object; the acquisition of many points that lie on the surface as 3D co-ordinates. This is used in the presented work to measure animal and phantom shapes in 3D. This allows the use of accurate geometry when performing FEM-based image reconstructions (sections 2.5 and 2.6 and chapters 8 and 9) and when mapping light measured by the camera back onto the imaged surface (chapter 7). The first half of this chapter discusses the surface capture method used within the current system explaining the underlying methods and reasoning, and presenting experimental results and corresponding analysis. The second half of the chapter presents the automated meshing method used to create FEM meshes from surface capture data.

The initial design, implementation and some evaluations of the surface capture method were presented by Basevi et al.[3]. The author of this thesis contributed ideas to the algorithm development and carried out the practical experiments that tested it as it was developed and evaluated. The author also designed the physical set-up and built it, wrote control software to run the hardware to implement the data acquisition, shared the design of the particular pattern set used for imaging, and designed, implemented and tested the camera and projector geometric calibration routines (appendix A.6) that were needed to use the method.

The method has allowed important achievements to be made particularly the acquisition of accurate animal and phantom surface measurements with freely placed projectors which are normally constrained in where they can be placed when using comparable methods[156]. In the

current method the projector locations can be tuned to suit the set-up, e.g. maximising surface coverage and illumination. In a related result, it has been shown that surfaces can be captured via mirrors which has not previously been achieved with similar methods[3].

6.1 Introduction

The geometry of the animal being imaged is important for two main reasons. Firstly, it is needed to build an accurate full-system model within which to compute light propagation facilitating image reconstruction. Secondly, it allows the visualisation of results within the correct physical context because 3D images can be rendered containing the surface as a reference allowing for clear visual interpretation and for quantitative analysis of recovered source and property distribution locations.

It has been discussed that there are systems that utilise additional imaging modalities such as CT or MRI to obtain geometry data (section 3.3 and chapter 4), however these approaches require significant additional engineering, expertise and cost, also increasing imaging time and complexity and introducing ionising radiation in the case of CT. Optical surface capture methods provide a safe, relatively simple, and cheap alternative. A variety of systems utilising optical surface capture approaches have been discussed already (chapter 4) such as those of Deliolanis et al.[141], Li et al.[146], and Kuo et al.[44].

The presented method is similar to that implemented by Kuo[44] in that it is based on fringe pattern profilimitry, however it is improved in that it achieves greater surface coverage and measurement accuracy. The presented method places no constraints on the system geometry which is advantageous for two main reasons. Firstly, because it allows unrestricted placement of the projectors allowing optimisation of sample coverage by the illumination. Secondly, because it allows surface capture using mirror views, which is achieved by utilising two *virtual* cameras (effective camera locations given reflection in each mirror) with each projector as well as the direct view. This latter point is especially important in the present system because of the use of mirrors to collect other data; the free-space light mapping function (chapter 7) integrity relies on

having an accurate surface knowledge at the points where emitted light is visible to the system including points visible in the mirror.

6.2 Method

6.2.1 Sinusoidal Fringe Pattern Profilimetry

Fringe pattern profilimetry requires a source of structured light and an imaging device. In the presented system it is done using two pico-projectors and an EMCCD camera (chapter 5). The projectors illuminate the subject and the camera acquires images of the subject under patterned illumination. The presence of the object surface within the field of projection deforms the pattern by an amount dependant upon the surface profile; by extracting the amount of deformation, the height of the object can be found.

In order to establish the amount of deformation that occurs, projected images are encoded so that pattern positions can be recognised in captured images. This allows the use of geometric rules to establish the position of a point under observation and illumination assuming the system geometry is very well known. Figure 6.1 illustrates these ideas.

This is applicable to methods using any encoding of projections along one direction; in the simplest case this could be projection of a linear function of x (e.g. encoding the column number explicitly by the projected intensity), however there are advantages to the use of more complex patterns.

Sinudoisal fringe pattern profilimetry involves the projection of images that are sinusoidal functions of position in one direction, such as in the *x* direction (fig. 6.2), in this case the *phase* (the argument of the projected co/sine function) is directly related to the distance along the pattern direction and the analysis of this value along with the geometry of the system can be used to produce a map of heights for the illuminated surface.

One advantage of sinusoids is that they are robust to blurring due to the pattern being out-of-focus because they exhibit local linearity. For sharper patterns blurring can give unwanted

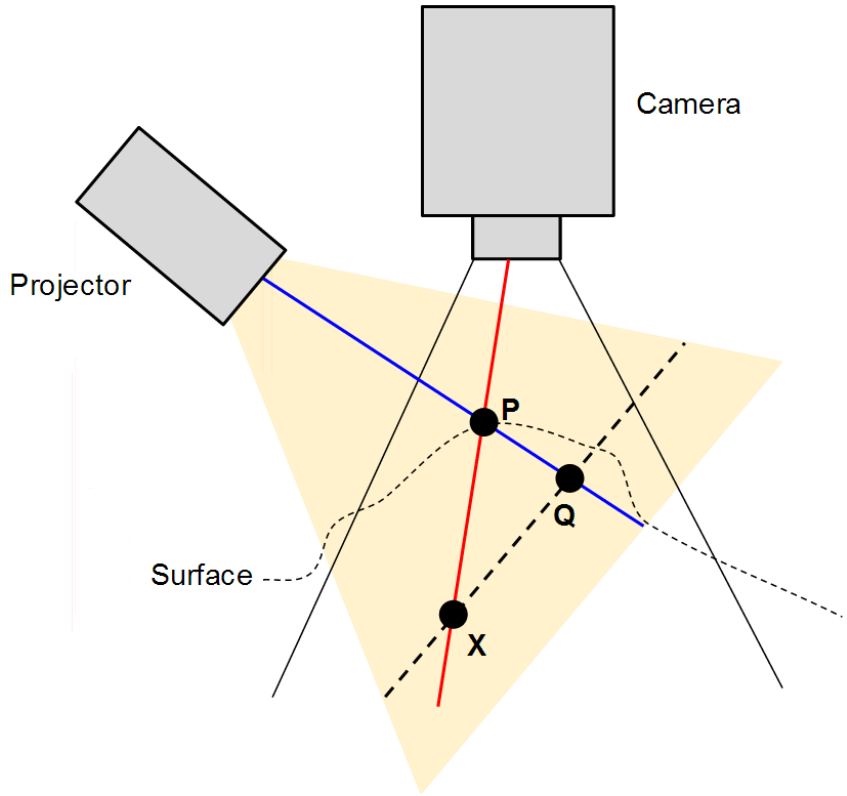

Figure 6.1: Depiction of projector-camera system used in surface capture. For a pixel viewing the surface at **P**, the imaged pattern value is that that would have been projected to point **Q** were the projection uninterrupted. The *expected* pattern value for the pixel is that that would have been projected to point **X**. There is therefore a *deformation* in the pattern (from the camera point of view) that is a function of the distance **QX** which can be recovered if the pattern and set-up are known and then can be used to deduce **PQ**; the *height*.

results that are ambiguous and/or erroneous. Another advantage is that there is low impact of noise upon recovered heights. This is because noise acts upon the measured sinusoidal function value which is quickly varying compared to the encoded position and as such the error on the position is relatively low.

The major challenge that is introduced by using sinusoids is that the phase must be extracted from the measurement which is non-trivial because of non-uniqueness arising from the one-to-many mapping of the inverse-cosine (or sine) function.

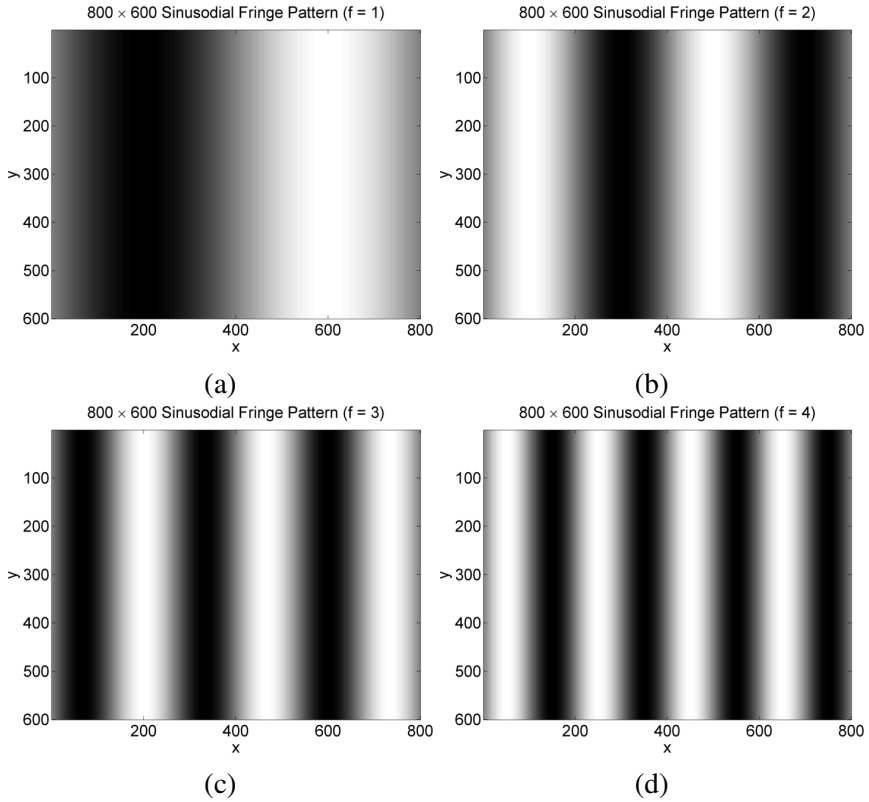

Figure 6.2: 800×600 images of sinusoidal fringe patterns in the x-direction with different frequencies; $f \in \{1, 2, 3, 4\}$ (waves-per-image).

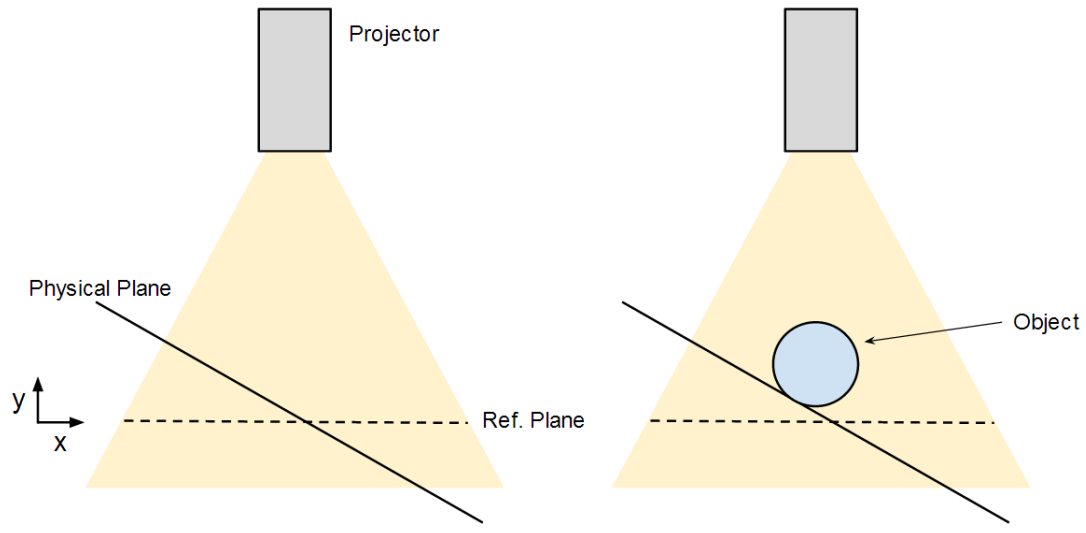

Figure 6.3: Projection onto an arbitrary plane, with co-ordinate system defined w.r.t. to the projection direction and projection onto an object on the plane. The projected pattern is deformed in a predictable fashion by the presence of the plane and the object.

6.2.2 Phase Measurement

In two dimensions, a projected sinusoidal pattern varying in the x-direction forms an image on an imaginary reference plane perpendicular to the optical axis of projection. The intensity of the projected pattern on an arbitrary physical plane (fig. 6.3(a)) is

$$I(x_p) = r(x_p)\frac{1}{2}\left(1+\cos(2\pi f x_p + \Psi_b(x_p))\right), \tag{6.1}$$

where $I(x_p)$ is the intensity as a function of x_p; the distance along the projection image, $r(x_p)$ is the reflectance of the plane, f is the spatial frequency of the pattern, and $\Psi_b(x_p)$ is a phase-shift caused by the presence of the physical plane. With an object resting on the plane (fig. 6.3(b)), the pattern intensity becomes

$$I(x_p) = r(x_p)\frac{1}{2}\left(1+\cos(2\pi f x_p + \Psi_b(x_p) + \Psi_o(x_p))\right), \tag{6.2}$$

where $\Psi_o(x_p)$ is an additional phase shift introduced by the object, and $r(x_p)$ now represents the reflectance of the object or plane dependant upon the particular configuration.

It is possible to extract the height of a point on the object from knowledge of $\Psi_o(x_p)$, but extraction of this term is challenging because of the multiple unknown terms of the equation and the presence of the periodic cosine function encapsulating several parameters. As such, the quantity that can typically be accessed is the wrapped phase,

$$(2\pi f x_p + \Psi_b(x_p) + \Psi_o(x_p)) \mod 2\pi, \tag{6.3}$$

and non-trivial phase-unwrapping is needed to recover the desired quantity:

$$2\pi f x_p + \Psi_b(x_p) + \Psi_o(x_p). \tag{6.4}$$

Many methods then make use of a background scan (of a reference plane) in order to separate $\Psi_b(x_p)$ from $\Psi_o(x_p)$.

Several techniques exist for wrapped phase extraction based on analysing the spatial relationships of values in images including Fourier filtering techniques[174, 175] and phasor-based methods[176], however these have limitations particularly where there are spatial discontinuities or different fundamental frequencies in visible patterns[3]. It is instead possible to obtain phase by using phase-shifting, the technique used in the presented system, in which images are acquired for many patterns with fixed spatial frequency but different phase offsets. Every point under illumination then samples a range of projected values as the pattern is shifted through the scene and the phase is obtained on a per-pixel basis as follows.

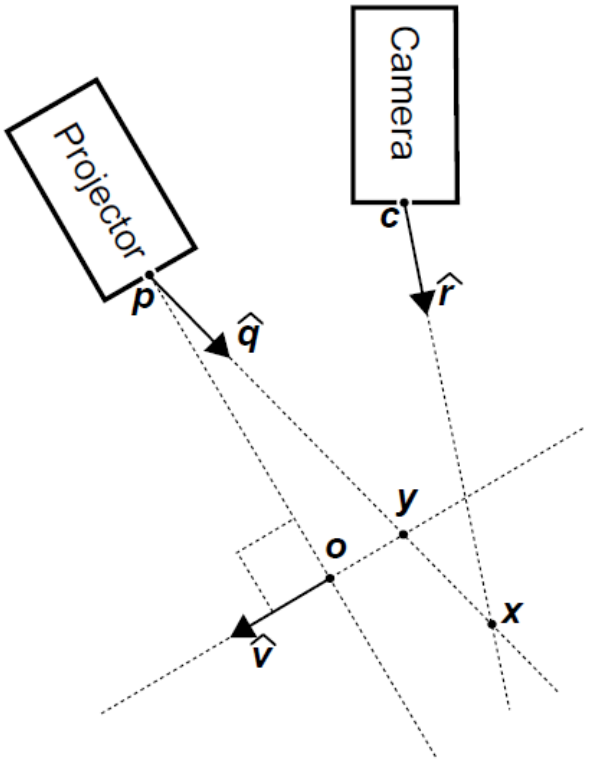

Figure 6.4: 2D projection of generalised interaction of camera-projector system assuming pinhole models, reproduced from Basevi et al.[3]. The projector projects from point **p** into direction $\hat{\mathbf{q}}$, the camera at point **c** collects the reflected ray with direction $-\hat{\mathbf{r}}$. The origin **o** is placed on the optical projection axis in a plane in which the pattern direction $\hat{\mathbf{v}}$ and spatial frequency f are known. The unknown reflecting surface point is **x** at which the projected pattern value is equal to that expected at **y**.

Consider the generalised camera-projector system shown in fig. 6.4. Projected values in the

n^{th} pattern (out of a total N) are

$$p_n(\mathbf{y}, f) = \frac{1 + \cos(2\pi f(\mathbf{y} - \mathbf{o}) \cdot \hat{\mathbf{v}} + \phi_n)}{2}, \tag{6.5}$$

where ϕ_n is the n^{th} spatial phase shift applied. The image on the camera is

$$g_n(\mathbf{x}, f) = b(\mathbf{x}) + r(\mathbf{x}) \frac{1 + \cos(\Psi(\mathbf{x}, f) + \phi_n)}{2}, \tag{6.6}$$

where $b(\mathbf{x})$ is stable background illumination and the un-shifted phase function is

$$\Psi(\mathbf{x}, f) = (2\pi f(\mathbf{y} - \mathbf{o}) \cdot \hat{\mathbf{v}}. \tag{6.7}$$

The wrapped phase can be extracted by the following formula[3]:

$$\Psi(\mathbf{x}, f) \mod 2\pi = -\arctan\left(\frac{\sum_{n=1}^{N} g_n(\mathbf{x}, f) \sin(\frac{2\pi n}{N})}{\sum_{n=1}^{N} g_n(\mathbf{x}, f) \cos(\frac{2\pi n}{N})}\right). \tag{6.8}$$

As has been briefly discussed, there is some error on the value found due to noise in the system (the combined noise of the projection and detection systems, in this case dominated by the projection system), the use of multiple phases means that there is little dependence of the error on the initial phase because each point samples a similar range of values. The error in the phase translates into an error in the resulting recovered surface point position but importantly the error on the phase is independent of the frequency whereas the range that the phase spans in physical space is not, i.e. higher frequency patterns have less error in position assuming that the phase is *correctly unwrapped*[1].

6.2.3 Phase Unwrapping

There are existing methods for unwrapping phase, however most are based on the identification of wrapping events directly from single images which makes for fast data collection but is a

[1] except for in the extreme case where the frequency is so high that whole waves are blurred out by the optics

heuristic process involving somehow separating wrapping events from genuinely large height changes which is difficult and often inaccurate[3]. It is not generally possible to tell how many times the phase has been wrapped even when a wrapping event is identified correctly which also leads to errors. Furthermore these approaches require selection of a reference point from which wrapping events are traced in an iterative, outwardly propagating manner which means that errors are propagated cumulatively to data points analysed following the occurrence. This significantly impairs results because wrapping errors appear as large step-like errors in final recovered height.

It is possible to use additional images to provide more data and improve the accuracy of the phase unwrapping process. For example, a series of binary patterns can be projected, illuminating in turn areas corresponding to each full cycle of 2π in the initial projection in order to uniquely identify them and therefore highlight positions of wrapping borders, this approach can be made more efficient by use of Gray-coding methods[177, 178], however these methods suffer from boundary-uncertainty effects because of the sharp edges.

Another possibility is to use fringe patterns of multiple frequencies to provide additional information, an approach called temporal phase unwrapping[179]. High and low frequencies provide complimentary information because lower-frequency patterns lead to less wrapping events (with no wrapping events when the projected frequency falls below one pattern-per-image) but greater sensitivity to noise.

The disadvantage of these methods is that more images are required meaning that imaging time and data processing time is higher. In the presented system accuracy is a higher priority than speed because relative to the total experimental time, which is dominated by luminescence and trans-illumination image acquisition, even a relatively long surface capture acquisition is a negligible extra time cost (section 6.2.5).

Phase unwrapping is achieved here using spatial patterns of multiple frequencies and multiple phases[3]. A frequency is first chosen that is less than one pattern-per-image after which successively higher frequencies are used. In the lowest frequency image there are no wrapping events, so although the noise is high, a noisy measurement of the absolute phase is obtained.

This noisy estimate is then used to constrain the unwrapping performed on the next (higher frequency) phase measurement allowing for identification of the correct level of wrapping and then obtaining a less noisy measurement of the absolute phase[3]. This is repeated through several frequency increments. Sample images of phase are shown in fig. 6.5 along with the unwrapped phase map obtained, it can be seen that the number of wrapping events increases with frequency and also that the initial low frequency wrapped phase is similar to the final unwrapped phase[1].

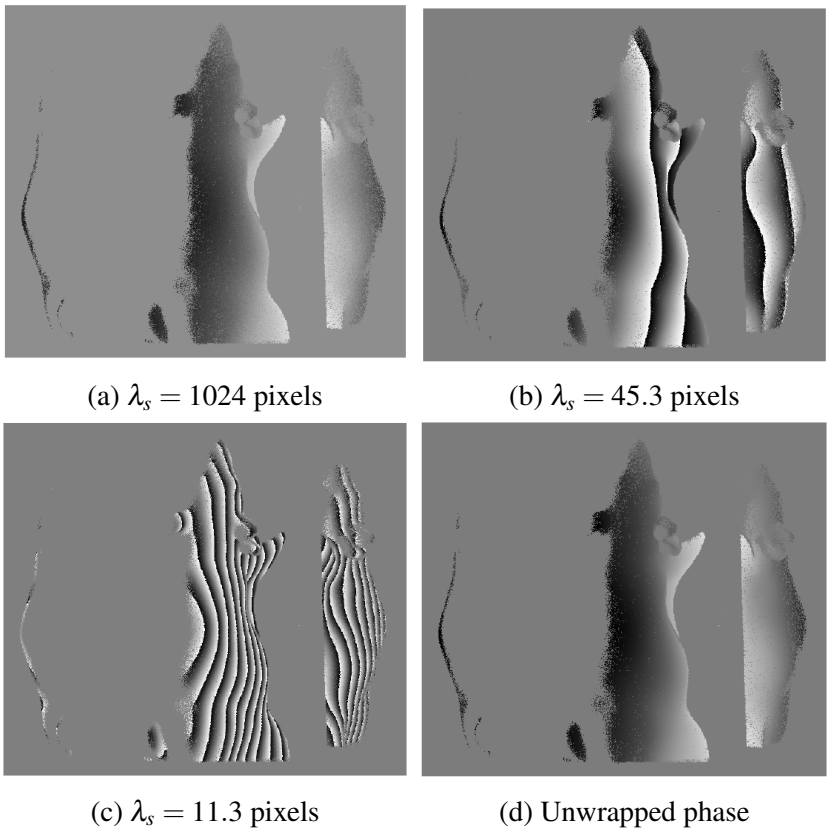

(a) $\lambda_s = 1024$ pixels (b) $\lambda_s = 45.3$ pixels

(c) $\lambda_s = 11.3$ pixels (d) Unwrapped phase

Figure 6.5: Example of (a-c) *wrapped* phase (measured on the XPM-2 phantom; appendix A.1.1) at increasing spatial frequency ($1/\lambda_s$) and (d) *unwrapped* phase, for one of the projectors. The low frequency image (a) has no wrapping events, the higher frequency image (b) has more and the highest frequency image (c) has even more. The unwrapped phase looks qualitatively similar to the low frequency phase but is more accurate. Note that background pixels have been masked out using thresholding applied to an image obtained with a full-field illumination pattern. These images are taken from Basevi et al.[3]

[1] though it is not identical as it contains noise.

6.2.4 Calculation of Absolute Spatial Co-ordinates from Phase

Height maps are constructed from unwrapped phase using knowledge of the imaging system geometry. Methods commonly use a simplifying approximation that places the projector at infinity, assumes a *cross-axis* geometry and uses $\Psi_o(x)$ explicitly to extract height values. However, this constrains the imaging system set-up such that mirrors cannot be used and requires that reference plane phase shifts $\Psi_b(x)$ have already been isolated, furthermore these approaches only provide heights relative to the reference plane.

In the presented method, extraction of co-ordinates is calculated directly from phase measurements using thorough knowledge of the system geometry in the following way.

Consider again the set-up shown in fig. 6.4. Assume a sinusoidal pattern with underlying phase values centred on zero, i.e. phase equals zero at all points on the line $\overrightarrow{\mathbf{po}}$. Assume also that the directions $\hat{\mathbf{r}}$ and $\hat{\mathbf{q}}$ are known for all camera and projector pixels. The point at which the projected ray hits the imaginary reference plane is

$$\mathbf{y} = \mathbf{p} - \frac{|\mathbf{p}-\mathbf{o}|^2}{\hat{\mathbf{q}}\cdot(\mathbf{p}-\mathbf{o})}\hat{\mathbf{q}}, \tag{6.9}$$

the projected phase at this point in the pattern is

$$\Psi = 2\pi f(\mathbf{y}-\mathbf{o})\cdot\hat{\mathbf{v}}. \tag{6.10}$$

From this and the knowledge that the point $\mathbf{x}$ lies on the line defined by the camera ray, it follows that

$$\mathbf{x} = \mathbf{c} + \frac{\alpha\cdot(\mathbf{p}-\mathbf{c})}{\alpha\cdot\hat{\mathbf{r}}}\hat{\mathbf{r}}, \tag{6.11}$$

where

$$\alpha = \frac{\Psi}{2\pi f}\frac{\mathbf{p}-\mathbf{o}}{|(\mathbf{p}-\mathbf{o})|} + |(\mathbf{p}-\mathbf{o})|\hat{\mathbf{v}}, \tag{6.12}$$

thus given comprehensive knowledge of the geometry of the system in the form of the location of effective pin-holes for the camera and projector and the directions of all rays leaving the projector and entering the camera under a straight-line approximation, unwrapped phase can be

converted into absolute spatial co-ordinates directly. This is very convenient as the operation is not dependent on values of any other pixels or on any heuristic process. The cost is that the requirement for geometric knowledge necessitates accurate geometric system calibration (appendix A.6).

6.2.5 Implementation

The components of the surface capture subsystem have been described in chapter 5. Figure 6.6 shows approximate component positions and view directions showing the scale of the system and the general set-up.

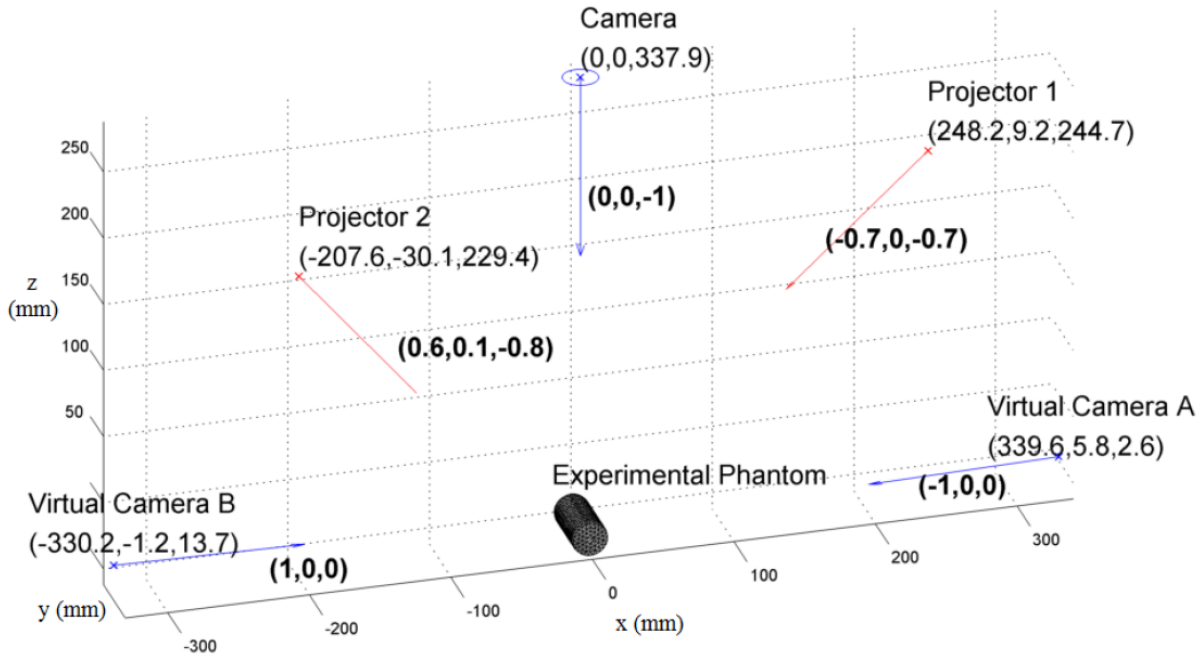

Figure 6.6: The geometry of the imaging system illustrated in terms of the positions and view-directions of the two projectors used for surface capture and the camera used for detection along with accompanying virtual cameras which are the reflection of the camera in each of the mirrors. Note that $z = 0$ is the height of the stage when the labjack is fully retracted.

The two projectors are placed so as to illuminate opposite sides of an imaged subject and between them cover most of the available part of a mouse-shaped surface. When the mirrors are added and used to collect surface capture data this coverage is increased still further. The coverage is illustrated in fig. 6.7 which shows the full illumination of the cylinder phantom (appendix A.1.1) in the system (fig. 6.7(a)) and the same phantom under pattern illumination

from each projector. It can be seen that the coverage is complimentary and generally high although does not illuminate the *whole* of the lower part of the phantom owing to the high curvature in the lower part of the cylinder (which is not present in small animal cases where the subject lies flat on the stage).

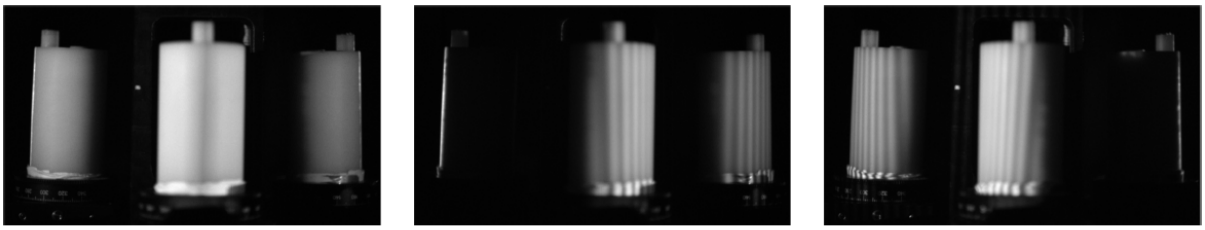

(a) (b) (c)

Figure 6.7: Sample surface capture images comprising (a) the maximum of bright images for both projectors (full-field white projection from each projector), (b) a high frequency pattern projected with projector 1 and (c) a high frequency pattern projected with projector 2.

A series of images are projected using each of the two projectors in series. 14 different spatial frequencies are used starting at 0.78125 fringes/image and increasing by a factor of $\sqrt{2}$ to a maximum of 70.7107 fringes/image which corresponds to a range of approximately 0.003 to 0.3 fringes/mm projected onto the stage. For each spatial frequency 6 evenly spaced phase shifts are used throughout the range 0 to 2π. In addition, "bright" and "dark" projections are collected meaning that in total 86 (6 phases $\times$ 14 frequencies + 2 extras) images are collected for each projector.

For surface capture imaging, the camera parameters are: EMCCD mode; read mode 3; exposure time 0.12 seconds; no binning. These modes provide the fastest imaging possible on the camera without binning which is important due to the relatively large number of images that need to be taken. With these modes, surface capture takes approximately 40 seconds per projector including overhead (from saving images and driving devices). The lack of binning means that the point-cloud is more spatially dense and therefore accurate for acquiring the curvature of the imaged sample.

Following data acquisition, the surface capture algorithm introduced above is applied to the image set and a point cloud is obtained. An example of the result when imaging the cylinder phantom is shown in figure 6.8(a) which shows three partial point clouds obtained (one for each

detector modelled; both real and virtual via mirrors). It can be seen that there is some overlap between the views which is due to the placement of the mirrors (and is beneficial; section 6.2.6).

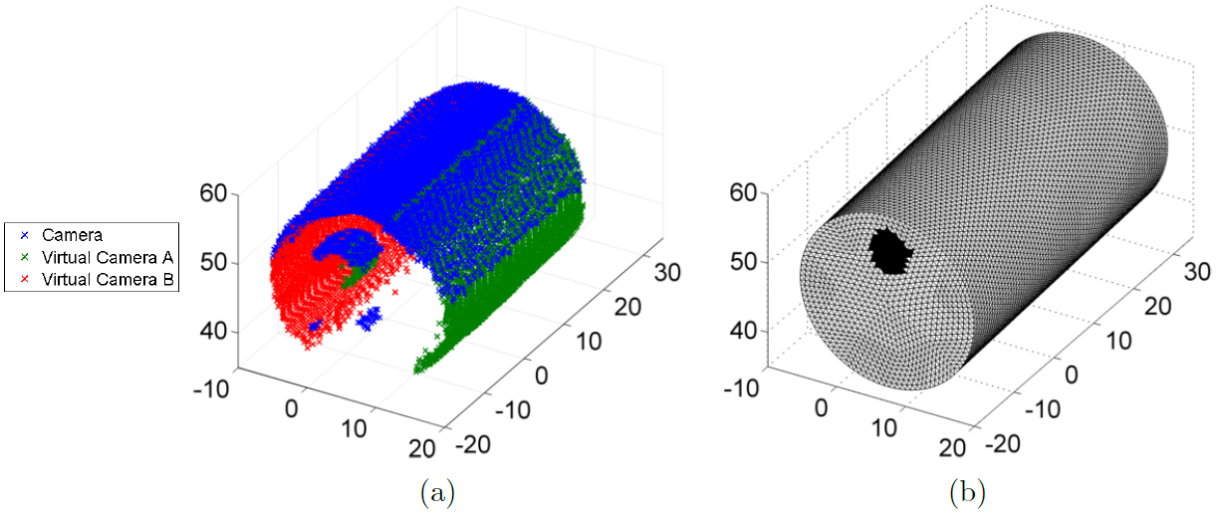

Figure 6.8: (a) Example surface capture point cloud in which points acquired at different views are indicated by different colours; and (b) pre-made FEM mesh following registration to the surface capture point cloud; black elements indicate the location of an inclusion rod that is left protruding slightly from the main cylinder and used as a reference point for finding the correct rotation (note the corresponding points visible in the point cloud for the rod). The point cloud appears truncated compared to the mesh because the back portion of the cylinder phantom was mounted into a rotation mount as part of the study that was being carried out (chapter 9).

The dense point cloud recovered has absolute 3D co-ordinates in the camera co-ordinate system (appendix A.6) and can be used to create a surface and volume mesh for image reconstruction. The cylindrical phantom is a challenging case because of its significant curvature on the underside that is blocked from the projectors; thus the point cloud is not sufficient to generate a mesh, but it can be used to pre-register an existing mesh such as is shown in fig. 6.8(b). It has been found that this can be an effective method to pursue for luminescence and diffuse image reconstructions[1].

6.2.6 Automated Fitting of Mirror Location

It is an advantage of the method that components (projectors, detectors, and therefore virtual detectors) can be placed arbitrarily. This is advantageous when imaging objects of varying size, shape, and region of interest because by moving the mirrors the area under observation can

be tuned. For this reason and because as a research instrument the imaging system undergoes changes regularly, it was chosen not to fix the mirrors to the platform, rather to place them freely and develop a method for identifying their position on-the-fly from surface capture data.

This was achieved by performing the whole surface capture imaging protocol twice using two different patterns-sets with pattern-direction perpendicular to the original set (i.e. an x and y set). Unwrapped phase is obtained for both data sets providing a unique pixel-wise encoding for each projector meaning that projected phases visible in the mirror can be identified also in the main view (where there is view overlap). An example of a subset of corresponding point pairs found using the method is shown in fig. 6.9.

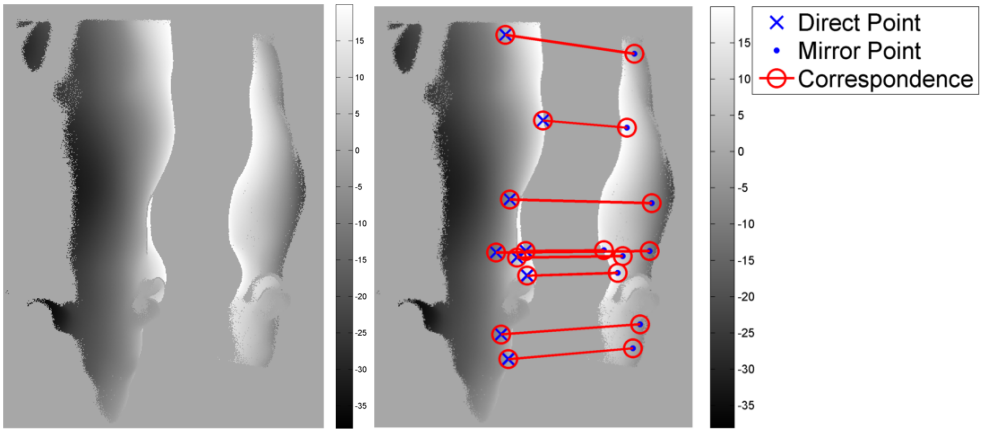

(a) Phase Map (b) Phase Map and Matched Phases

Figure 6.9: Illustration of point correspondences identified by finding similarity phase values.

Using the geometric model and the derived pairs of points, the location of the mirror surface is determined by minimising

$$\mathbf{x} = \min_{\mathbf{x}} \sum_{i=0}^{n} ||\mathbf{d}_i - \rho(\mathbf{x}, \mathbf{d}_i, \mathbf{m}_i)||_2^2, \tag{6.13}$$

where $\mathbf{x}$ is a vector containing the solved-for mirror position and normal, d_i is the co-ordinate of the i^{th} direct-view point for which a corresponding mirror point m_i has been found. $\rho(\mathbf{x}, \mathbf{d}_i, \mathbf{m}_i)$ is a function that calculates the nearest point on the camera ray corresponding to the point m_i following it having been reflected about the proposed mirror plane. Figure 6.10 shows an example result of the method.

Note that this approach to locating the mirror can be used per-subject or as a one off if the mirrors are fixed. In the latter case a phantom can be surface capture imaged having placed the mirrors such that the location can be fitted for from the data. The phantom can then be removed and other subjects can be introduced without having to re-fit the mirror properties.

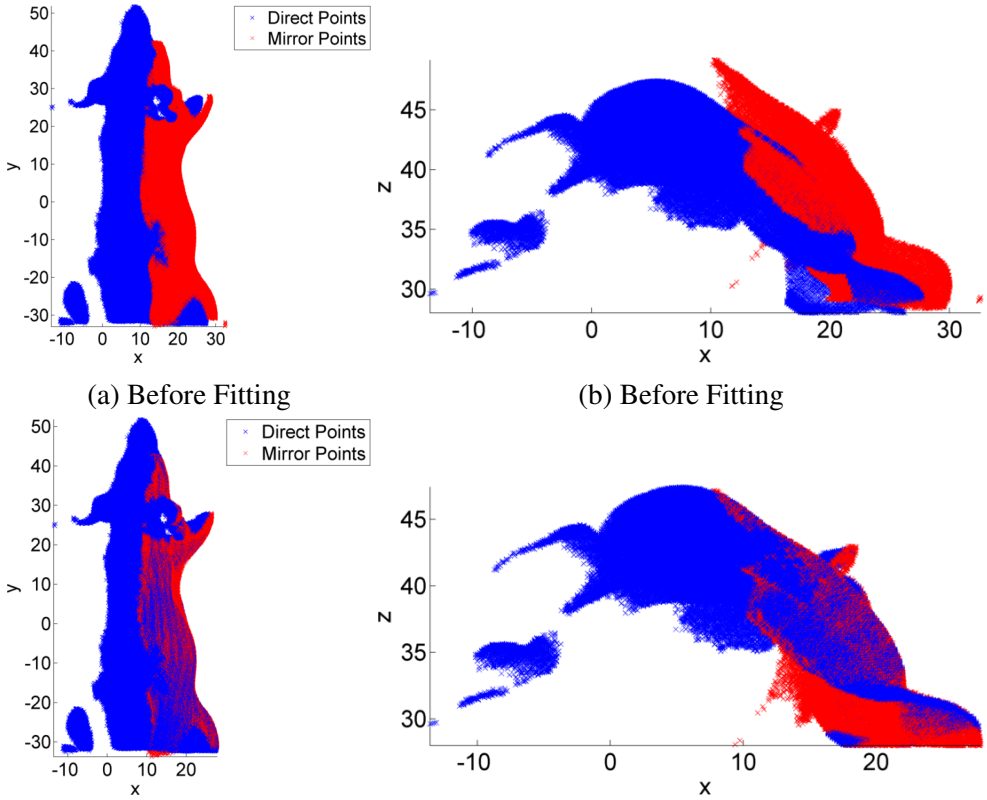

(a) Before Fitting (b) Before Fitting

(c) After Fitting (d) After Fitting

Figure 6.10: Example of mirror-view and direct-view surface capture points before and after fitting for the mirror position using the presented method. It can be seen that after the fitting the the point clouds are overlaid correctly as compared to the initial erroneous state.

6.2.7 Automated Masking

Multi-view surface capture data has been found to contain error for points on the object that are under illumination indirectly from the projector via the mirror. This is due to the incident intensity being the combination of the known direct projection and the unknown mirror-reflected projection.

In order to avoid this, a method was developed to mask the projection (i.e. set certain

projector pixels to *off* for the entire duration of an imaging session) such as to avoid unwanted projections onto the mirror based on using principles from dual photography[180] in which perspectives can be shifted in camera and projector systems such that it appears that the projector acquired the image with illumination provided by the camera.

A double-set of surface capture data (with x and y patterns as in section 6.2.6) is first acquired and processed providing a unique pixel-wise phase encoding. Values from the measured 'bright' image are then taken and inserted into a projection-sized (800×600) image at the pixels corresponding the measured phase at the originating camera pixel. As such an image of the scene as 'seen' from the projector is constructed. This is then thresholded and dilated to produce a mask that only contains 'on' pixels where they will hit the object.

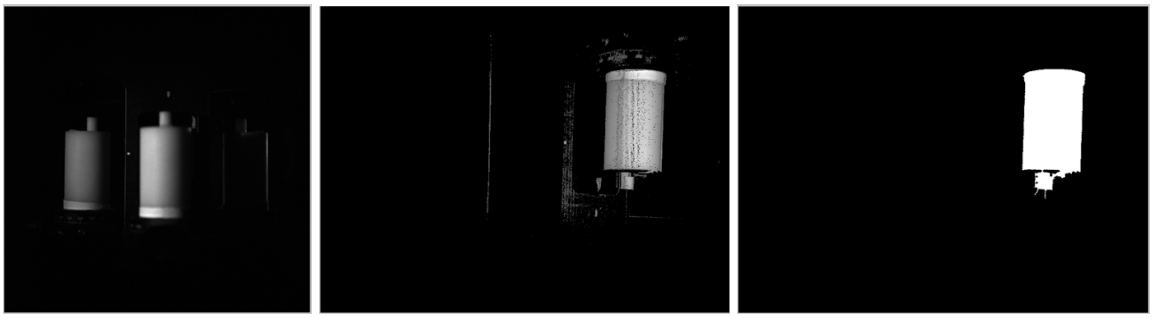

(a) Camera Image (b) Projector Image (c) Projector Mask

Figure 6.11: Highlights from the masking process with a cylinder phantom data set used as an example, showing data for a single projector. First, (a) the 'bright' image from the surface capture image set for which the projected phase measured at each camera pixel is known in both x and y directions. Second, (b) projector-view dual photography image constructed by inserting values from (a) at projector pixels indexed by the measured phase; the cylinder is seen from the perspective of the projector. Finally, (c) shows the mask generated by thresholding (b) and applying a dilation morphological image processing operator to fill holes.

6.3 Evaluation

6.3.1 Quantitative Evaluation w.r.t. XPM-2 Phantom Data

Co-ordinate Accuracy

A single manually-processed data set acquired for the XPM-2 phantom (appendix A.1.1) was used to evaluate the system in quantitative terms. The evaluation was performed with reference to a set of 878 points that lie on the true phantom surface that were provided by the manufacturer[3]; these are treated as a ground truth.

(a) (b) (c)

Figure 6.12: XPM-2 evaluation tools summary (a) photo of the XPM-2 phantom (appendix A.1.1), (b) high quality surface mesh made manually in Meshlab[181] using all the surface capture points obtained (no binning or merging), and (c) reference points for comparison provided by the manufacturer that should lie on the real geometry. These figures are taken from Basevi et al.[3]

.

Figure 6.12 shows the phantom again for reference, a high quality surface mesh generated manually in Meshlab[181] and the manufacturer-provided reference points. A histogram of the distances of reference points from the created mesh was calculated and is shown in fig. 6.13

The mean distance was 0.14mm, with a standard deviation of 0.14mm. 50% of the reference points were within 0.1mm of the surface, 82% within 0.2mm and 96% within 0.4mm. It is suggested that the rare points of greater distance are most likely due to meshing artefacts introduced in areas such as the ears[3].

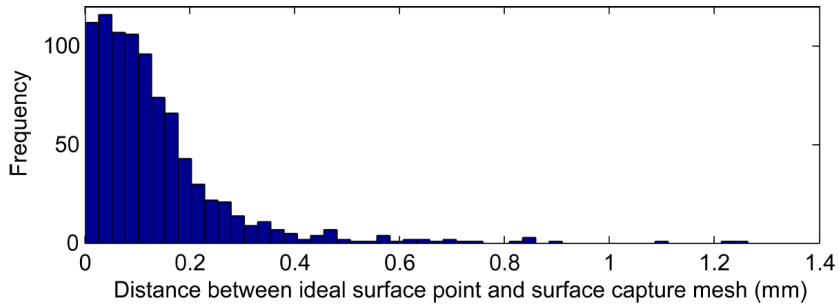

Figure 6.13: Histogram of distances of captured surface from manufacturer-provided reference points. This figure is taken from Basevi et al.[3]

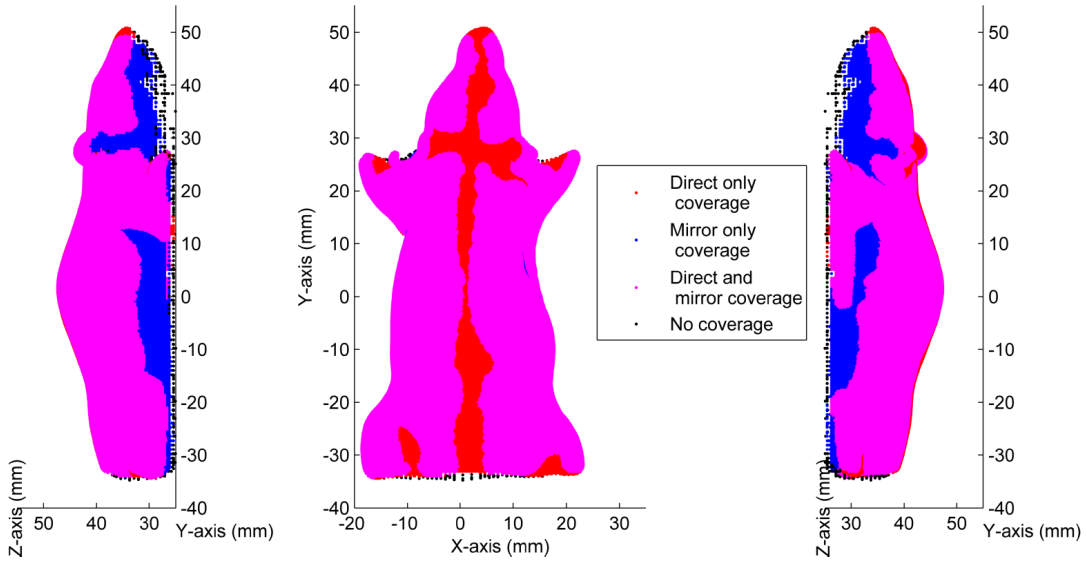

Figure 6.14: Illustration of the coverage (i.e. the visible area) of the XPM-2 surface with direct and mirror views. The mirror-views increase proportional surface coverage by 15.2% in this case. This figure is taken from Basevi et al.[3]

Coverage

Figure 6.14 shows an anlysis of the coverage of the surface in terms of what part is recoverable by the surface capture when using only direct, only mirror, and both direct and mirror views. Considering the entire surface (including the obscured underside), the coverage is increase by 15.2% by the inclusion of the mirrors. Of course the benefit provided by the mirrors will vary in general dependant upon the particular mirror and object configuration and geometry.

6.3.2 Quantitative Analysis w.r.t the Cylinder Phantom

To further test the surface capture accuracy, the cylinder phantom (appendix A.1.1 was placed and re-placed in the imaging system 9 times and surface captured using both projectors but without mirrors. The masked points were then registered rigidly to a reference perfect cylinder surface of the same size as the phantom and the distance from the surface was then calculated for all points. Results are summarised in fig. 6.15.

Figure 6.15(b) shows that the errors in the points appear to be normally distributed around 0 and that the error is in a similar range to that observed in the XPM-2 test case. Figure 6.15(c) shows that the mean and standard deviation in the absolute distances are similar across the 9 repeats; the mean mean value across the data sets is 0.0992mm and the mean standard deviation value is 0.0751mm. Figure 6.15(d) shows that the error is low as observed in the XPM-2 test; in this case specifically there are 60% of points within 0.1mm of the surface, 87% within 0.2mm and 99% within 0.4mm. This generally suggests more accuracy than in the XPM test, though this is possibly due to the fact that there is no added meshing error here as the error calculations were made with reference to an ideal cylinder (with the mesh in fig. 6.15(a) being only for visualisation reference).

6.3.3 Qualitative Analysis of Small Feature Visibility

To test the capability of the method to image fine features, two sides of a coin were imaged and rendered with intensity based solely on the measured height (fig. 6.16). It can be seen that fine

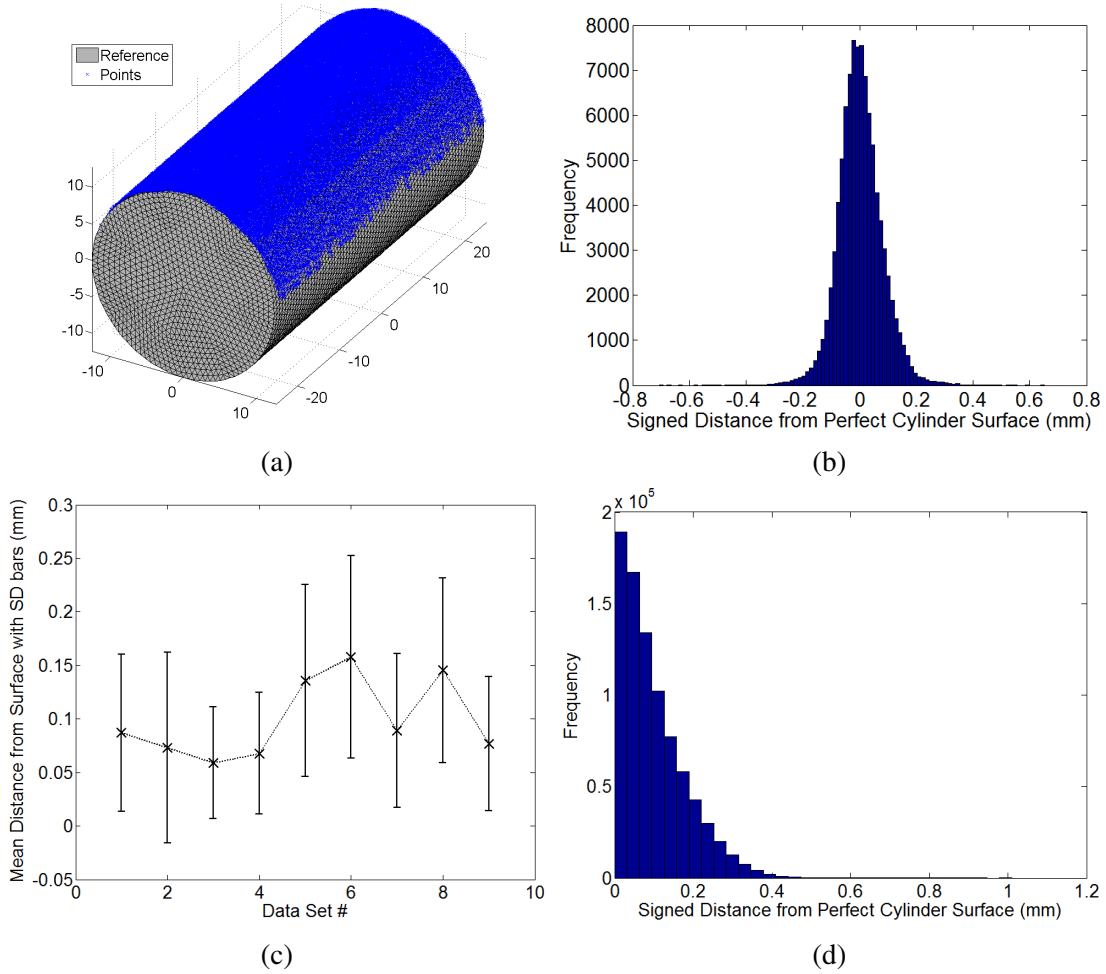

Figure 6.15: Quantitative analysis of 9 repeats of cylinder pantom surface capture: (a) representative illustration of the reference mesh and surface capture points from one set, (b) representative histogram showing the signed distance from the surface for all points in the same data set, the sign indicates whether the point is outside (positive) or inside (negative) the reference cylinder, (c) the mean un-signed (absolute) distances with standard deviation bars plotted for all sets, and (d) histogram of un-signed (absolute) distance for points from all sets considered together.

features are recognisable by only height-based contrast.

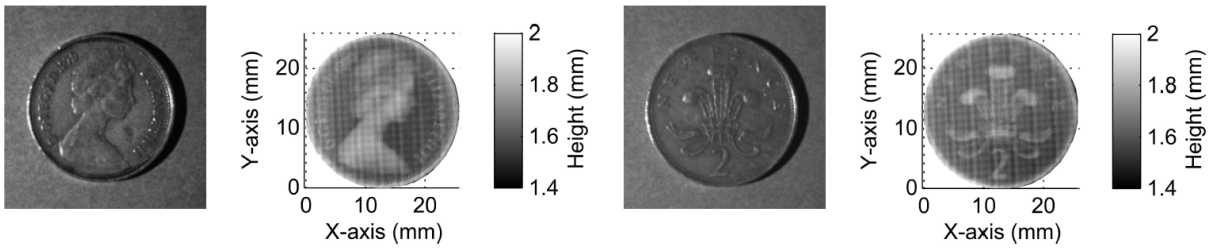

Figure 6.16: Surface capture of coins with no added contrast other than from recovered height values. Fine detail can be seen.

6.3.4 Small Animal Imaging Demonstration

Figure 6.17 shows surface capture results for two dead mice that were imaged in the system (see chapter 10). It can be seen that the reconstructed surfaces are qualitatively accurate though no ground truth was available for quantitative validation in these cases. Further analyses are presented for the first mouse shown here in terms of the coverage of the animal with the mirrors in Basevi et al.[3].

Figure 6.18 shows a surface capture of a rat that was imaged, the rat is larger than anything else that has been imaged in the system and it can be seen that the surface is still reconstructed qualitatively well. It is easy to see complex key features extracted effectively including small clumps of hair and the shape of the head. The final image shown has the texture from the original bright image mapped onto the surface, this illustrates a method that could be useful for visualising results in an interpretable way.

6.4 Conclusion

The presented surface capture method has been shown to accurately measure dense point clouds from an imaged surface with errors on the order to 100μm. Given that the typical target resolution for optical tomography instruments is on the order of mm it is proposed that this level of accuracy is more than sufficient for the system. It is also achieved very cheaply, and with little

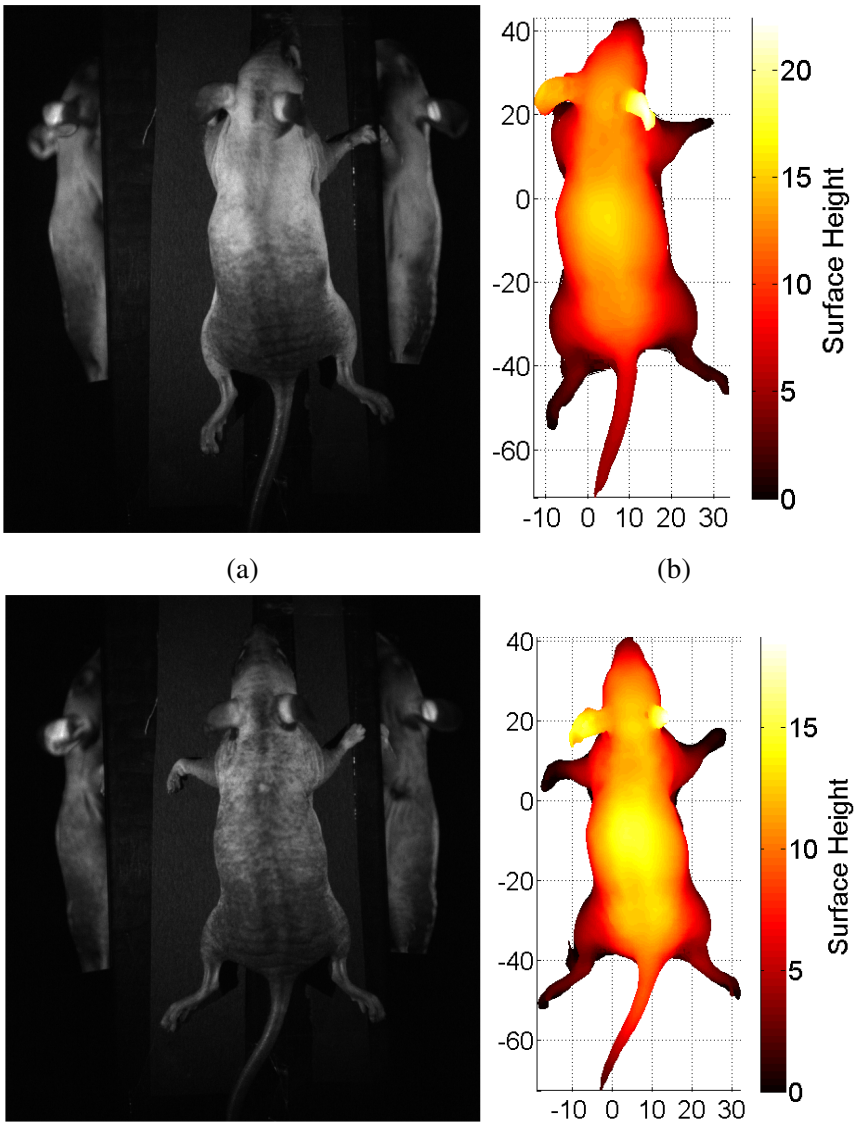

Figure 6.17: Surface capture of two mice *ex vivo* with (a,c) maximum of 2-projector bright images showing top and mirror views of the animal in the system and (b,d) surface meshes made for the mice from surface capture points coloured by height in millimetres.

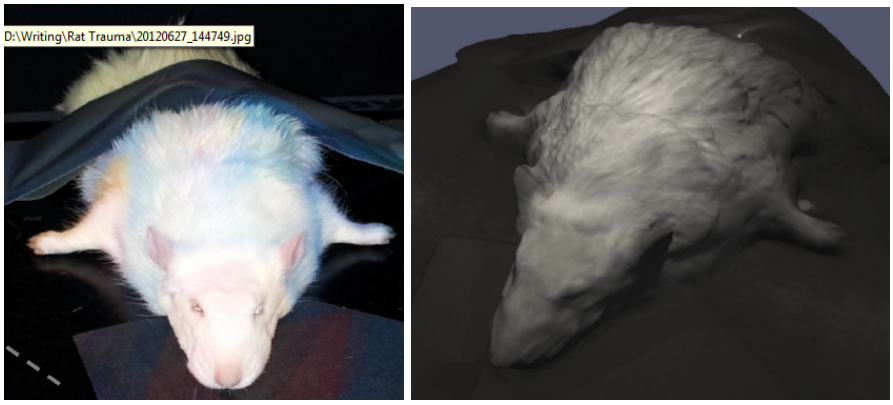

(a) (b)

Figure 6.18: Surface capture of a large rat *ex vivo*, shown as (a) photo of the rat in the imaging system taped down to support it on the stage and (c) reconstructed surface of the rat from the extracted point cloud.

additional experimental time and expertise required for its use. These are significant advantages over the use of MRI and CT for surface capture.

Additionally, the current method is registration-free because the same detection system is used for surface capture as is used for all other optical imaging. This is advantageous because it means that in conjunction with the free-space light propagation model presented in chapter 7 boundary measurements for FEM can be produced in an automated fashion direct from imaging system data. In other systems that use CT, MRI or optical capture methods using other detection systems registration is an issue that requires additional work and represents an unwanted potential source for errors.

6.4.1 Directions for Future Work

Pattern Optimisation

Preliminary data suggests that the current number of images (pattens) used for surface capture is significantly more than is necessary if optimal patterns are chosen. Exhaustive experimentation or possibly simulation could establish an optimal subset of the patterns that are used and increase the speed of the imaging and processing. It is estimated that the number of patterns could be reduced to around 10% of those used here without loss of accuracy.

Mirror Optimisation

Another idea for future work is the optimal positioning of mirrors based on the surface coverage. Additionally the use of mirrors with more complex shapes, such as the conical mirror used by Li et al.[146]. The method (including the mirror position fitting with some small modifications) can deal with arbitrary mirror configurations and as such these can be explored further to gain added signal. In an elaborate but potentially useful example, multiple mirrors could be used to redirect light paths from the sample such that multiple images appears on the CCD at different positions thereby utilising more of the CCD in imaging and getting more signal, this could be taken further by placing filters in some of the mirror paths to image multi-spectrally simultaneously similar to the idea of Wang et al.[147] but without confining the filters to a cylindrical mouse holder.

Application to the under-side

Surface capture could be performed in the present imaging system for the bottom of the animal too, or at least the portion of it visible through the stage-hole assuming that a camera (such as a small CMOS sensor or web-cam) could be positioned appropriately with appropriate field-of-view this would allow imaging of DOT source projections directly and therefore potentially the more accurate modelling of diffuse light input patterns.

6.5 Automated Mesh Generation

6.5.1 Introduction

A finite-element model (FEM) volume mesh is needed to perform image reconstruction (section 2.5 and section 2.6). Additionally, a surface mesh is needed to perform free-space calculations that relate boundary flux to CCD measurements made in the imaging system (chapter 7). The above usages of meshing involved time-intensive manual methods that needed experts to execute, and this is not sustainable in the long-term use of the system. In order to make the pro-

cess higher-throughput and therefore more readily usable and testable, automated approaches are needed.

A method was first developed whereby a pre-existing mesh of an imaged subject (easily obtained for well-known geometries and rigid phantoms that can be imaged in other scanners) was registered to surface capture points (rigidly) using a simple minimisation of distances between nearest points. This method was used to place the mesh shown in fig. 6.8 and has been used in imaging studies[1, 6, 5]. However, the approach is not applicable when the geometry is not known in advance and it does not fully exploit the surface capture system which should ideally provide a mesh of arbitrary objects imaged on a per-session basis. Then object shape is found directly in the camera co-ordinate system without prior knowledge of the specific shape and can be immediately used for modelling.

Performing meshing manually on a per-experiment basis is possible with many existing software packages such as NIRVIEW[182, 52], Meshlab[181, 183], and NETGEN[184, 185] and this is generally effective but time-consuming and labour-intensive. This section describes a work-flow developed for automatically producing meshes directly from surface capture data (i.e. from partial point clouds) that has been integrated into a single MATLAB program which makes calls to Meshlab and NETGEN in batch mode.

6.5.2 Meshing Method

There are different priorities for the surface mesh needed for free-space calculations and for the volume mesh needed for FEM. Specifically, for the FEM mesh a reasonably high node density, greater than say 1mm^{-3}, is needed so that light modelling is accurate. Conversely, the key requirement for the surface mesh is that the surface element areas are roughly equal and of an appropriate size which is somewhat larger than that used in the FEM mesh. The reason for this is that it has been found in preliminary studies that the free-space mapping works better when there are fewer faces to handle which may be because the free-space inverse problem is then less ill-posed.

As such the following method comprises three parts. The first step is the generation of a fine,

roughly uniform surface mesh which is then used as the basis for creation of other meshes. The second step is the re-meshing of this surface to provide a coarser, roughly uniform surface mesh with a specific average face area for us as the free-space mesh. The third step is the volume meshing of the fine mesh for use in FEM.

Generation of Approximately Uniform, Fine Surface Mesh

Starting with point clouds supplied by two projectors, the first step is to merge the points; there is overlap between the projector views, this can be done by taking the mean co-ordinate for each CCD pixel for which a co-ordinate has been found. Following this, the number of points is reduced by binning the co-ordinate images 4×4. This means that there is a single pixel remaining corresponding to each CCD pixel in binned luminescence or diffuse images and serves to smooth the point cloud whilst also reducing the size of the processing problem that follows. Some small amount of precision is lost here but this has not been quantified. Figure 6.19 shows an example of an original point cloud and the binned, merged point cloud acquired during imaging of the XPM-2 phantom (appendix A.1.1).

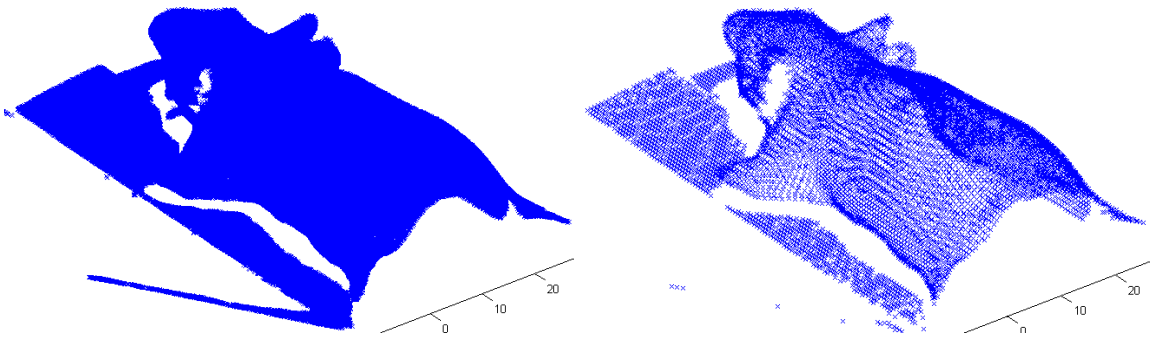

(a) (b)

Figure 6.19: Merging and binning of multi-view surface capture points illustrated by (a) original points from all views and (b) merged points following additional 4×4 co-ordinate binning.

Typically, unprocessed surface capture point clouds contain some points on the platform as well as directly on the object, as can be seen in fig. 6.19. Thus in a next step the object is segmented from the stage (and any remaining noise that is far from the object). This is achieved by a region-growing technique in which a single point is first assumed to lie on the

object (this *seed* for the algorithm is chosen by taking the point nearest to the centre-of-mass of the measured point cloud in x and y, working from the assumption that most of the points lie on the object and it is central to the cloud). A region is then grown by starting with the seed point and adding to the region any points within a Euclidean distance x_t. This is performed iteratively. Figure 6.20 demonstrates the region-growing. It can be seen that at the end the object is segmented effectively. The reason that a simple distance threshold works here is that the mirrors have a small gap where the reflective face meets the stage that is not coated meaning that there is always a small line of blind-spot between the object and the stage which halts the region growth.

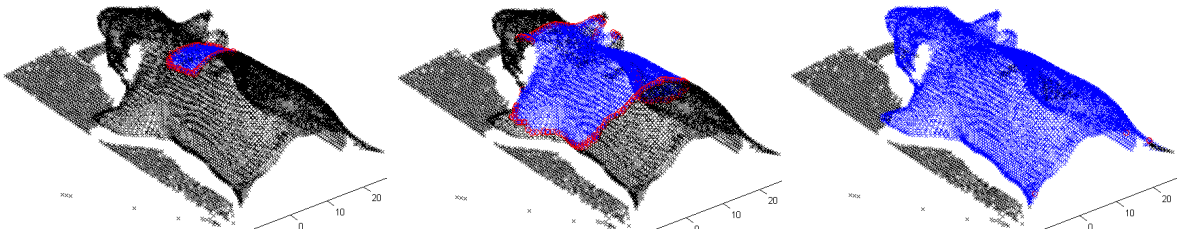

(a) Iteration 10 (b) Iteration 30 (c) End

Figure 6.20: Illustration of region growing for the separation of the object from the stage.

Following segmentation of the object, a new set of points are added to the cloud that have equivalent co-ordinates to the existing points in x and y but are given a z co-ordinate corresponding to the known height of the imaging stage. Thus the effect is to project a shadow of each point directly below each of the existing ones (fig. 6.21). These points are needed to close the surface from the perspective of meshing. The assumption made is that the object is practically flat upon the stage - if this is the case then these points should lie on the object. This is expected to be a good approximation in small animal imaging scenarios in which the malleable nature of the animal and the effects of gravity should make this happen. In the case of phantoms this is less often true because they are usually rigid and do not sag onto the stage in the same way.

Surface meshing is then performed using the MyRobustCrust (MRC) function which is distributed with NIRFAST[53] which produces a surface mesh based solely on provided points with none added or removed. This is a useful first step in meshing because it does not introduce any error by smoothing out the points so the structure is well-preserved. The result is a

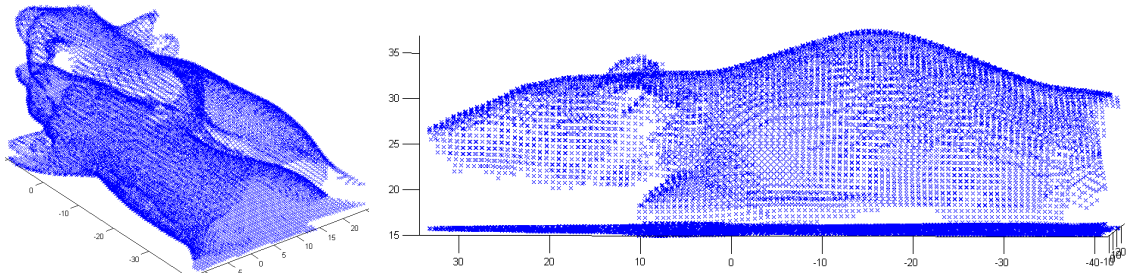

(a) (b) Side-View

Figure 6.21: Addition of 'shadow' points below the point cloud representing the bottom of the object when it is assumed to be in contact with the stage.

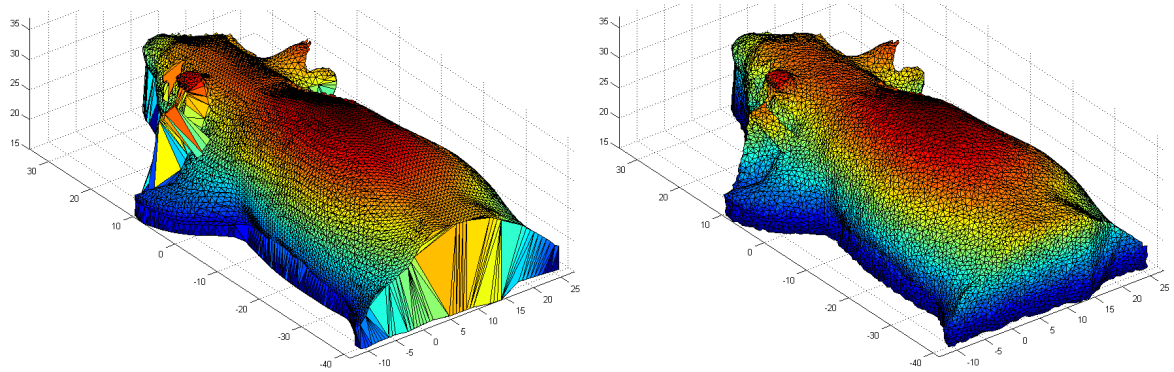

(a) First Point-Driven Mesh (b) Final Fine Uniform Mesh

Figure 6.22: Initial surface meshes created using the MyRobustCrust function. The first (a) is created directly using the shadow-added surface capture points (fig. 6.21). The second (b) is a re-meshed version created using a high-resolution uniform re-sampling of 10,000 points upon the first mesh to produce a fine and uniform mesh.

highly non-uniform mesh. To produce the final uniform, hi-resolution mesh, the non-uniform mesh is re-sampled using a poisson disk sampling method[186] (implementation provided by Meshlab[181]) with 10,000 points being generated approximately uniformly over the surface. These points are then used as input to a second round of MRC meshing which now produces a finer, more uniform mesh with very little change in actual structure or curvature. Figure 6.22 shows meshes generated for the example data set and fig. 6.23 shows the final mesh with the original binned surface capture points fig. 6.19 with edges not printed so that the points are clearly visible. It can be seen by visual inspection that the integrity of the mesh with respect to maintaining the measured shape is good.

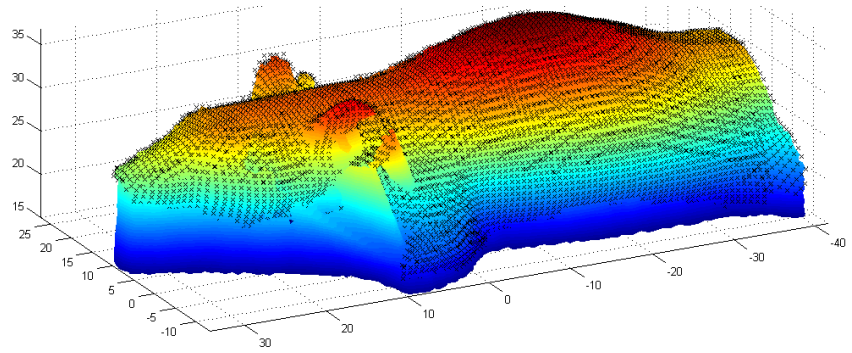

Figure 6.23: Surface capture points overlaid on fine-mesh outline.

Generation of Free-Space Mesh

The free-space mapping (chapter 7) appears from preliminary studies to perform well when there are relatively large faces; areas on the order of 2 to 4mm^2. It is possible (though requires further analysis to be established) that smaller areas would increase the complexity of the free-space inversion such that it would fail, they would certainly increase the processing time. Similarly, larger areas may become too non-specific spatially even if the average irradiance over the representative area were correctly recovered.

For this reason, it is desirable to specify the free-space mesh in terms of face area. This is achieved by using the same poisson-disk sampling/ MRC re-meshing method used to obtain the fine mesh above. In this instance the re-meshing input is the high resolution mesh. This approach relies on a *number of samples* as the only control over the resulting faces though it has

been observed that the resulting distribution of face sizes is reasonably uniform (in the presented example the mean faces area is 1.9mm^2 with a standard deviation of 0.5mm^2; fig. 6.24). The resultant face area is non-trivially dependent on the number of samples made so a simple next descent search algorithm is used to find a number of samples that produces a mesh with a desired mean face area to within some pre-defined tolerance. A resultant free-space mesh for the running example is shown in fig. 6.24 for which the target face area was 2mm^2.

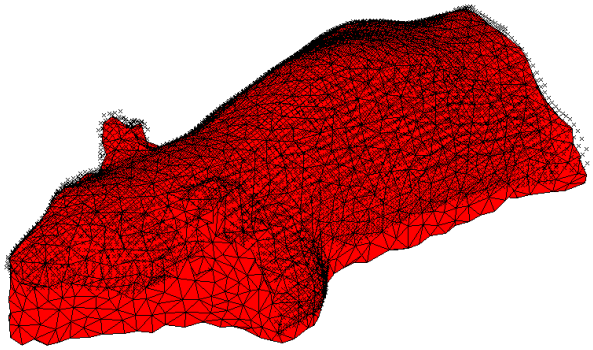

Figure 6.24: Free-space surface mesh with mean faces area is 1.9mm^2 with a standard deviation of 0.5mm^2 for which the target was area 2mm^2. Black crosses show original (merged/binned) surface capture points.

Generation of FEM Mesh

Volume meshes are created by calling NETGEN in batch mode with a meshing fineness parameter setting of "very fine" with an input surface mesh (generated equivalently to the free-space mesh from the first but with a different number of sample points as input). The actual fineness of the resultant volume mesh in this case, as in the surface meshing case, has a non-trivial dependence on the number of points drawn in the initial sampling. As such a next decent method is again used to search for an input sampling that produces a final volume mesh with a required number of nodes to within a specified tolerance. In the running example, a mesh density of 40,000 nodes was targeted, the actual resultant mesh has 40,661 (fig. 6.25a).

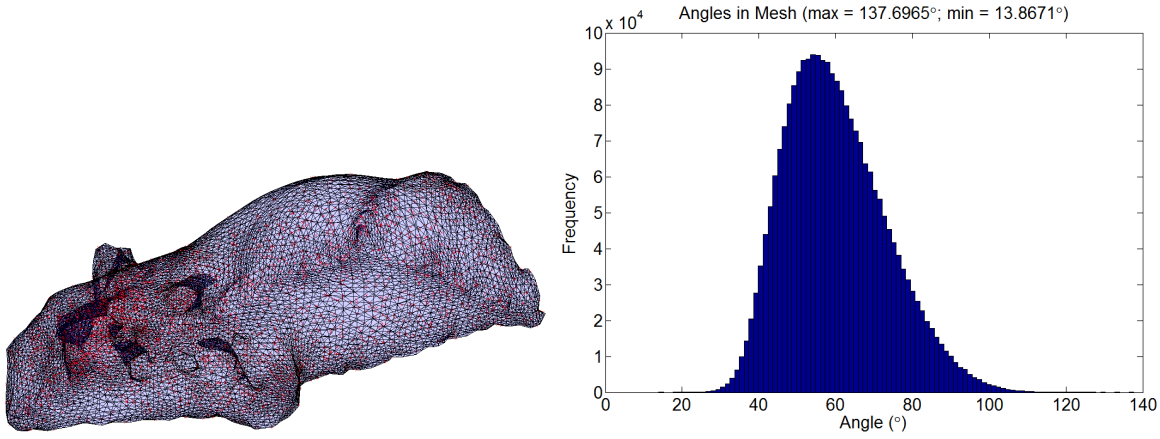

(a) (b)

Figure 6.25: (a) FEM mesh with surface shown semi-transparent blue and every tenth internal node plotted as a red cross. It can be seen that the surface is finer (somewhat representative of the internal node structure also) than in the free-space mesh reflecting the different requirement for the mesh (b) Histogram of angles in triangles throughout the FEM mesh. The majority are around 60 degrees indicating regularity of the mesh.

6.5.3 Mesh Quality

For FEM modelling, it is important to have approximately regular triangles. The current mesh was analysed in terms of the triangle angles and the results are shown in fig. 6.25b. It can be seen that most of the angles are around 60 degrees with very few being very large or small.

Node density is also important, it has been observed that the FEM analysis is only accurate if the density is sufficiently high in all regions. Histograms of the node density in terms of number of neighbouring nodes within a 3mm radius for all others, for boundary nodes and for non-boundary nodes are shown in fig. 6.26 along with a color-coded map of the density distribution. It can be seen that the average node density appears to be higher in the head and towards the edges where there are sharp surface changes than in the main body of the mesh (in the centre); also in the presented example the minimum number of neighbours within a 3mm radius was 4 which is a very small number (corresponding to approximately 0.04 per mm^3). This is a concern because the most important part of the mesh in the FEM analysis is the deep parts of the volume. A partial solution to this problem is presented below.

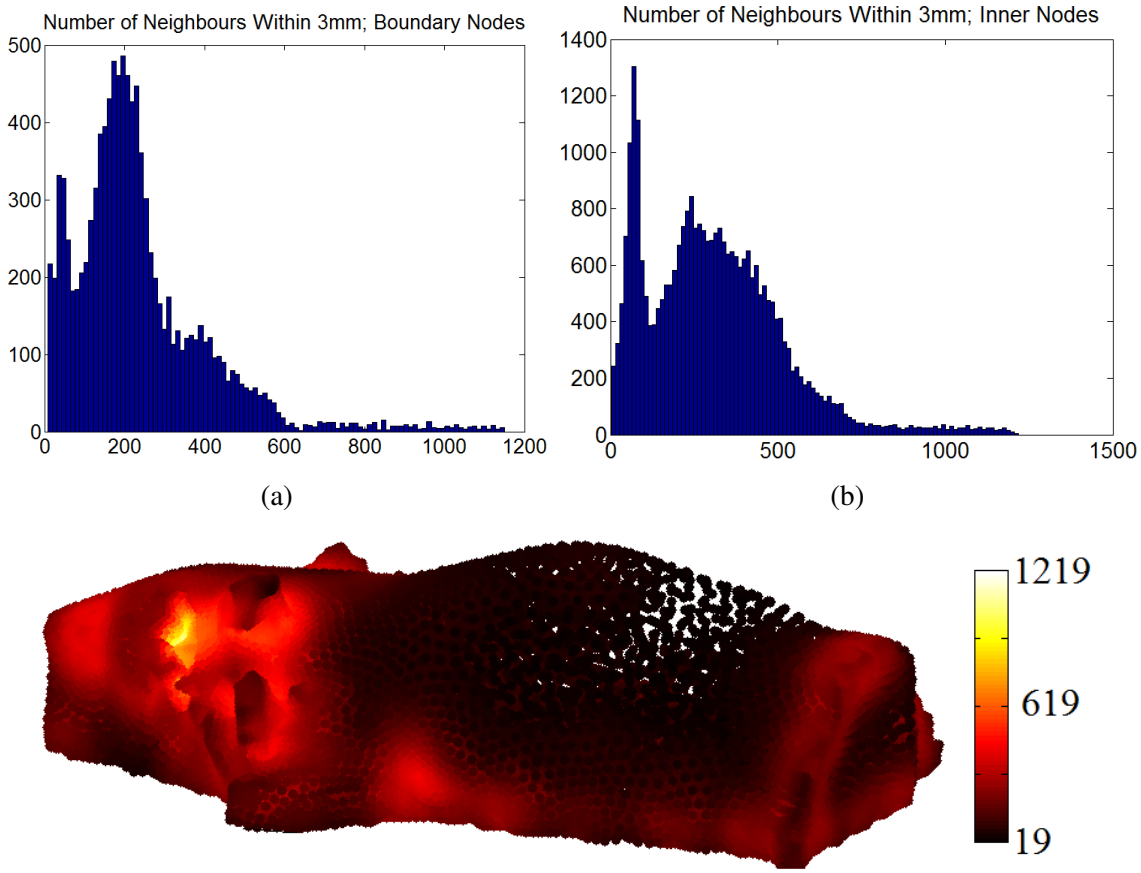

Figure 6.26: Analysis of node-density in the FEM mesh: (a) histogram of number of neighbours of each mesh boundary node within a 3mm radius, (b) histogram of number of neighbours of each mesh boundary node within a 3mm radius, (c) the spatial distribution of the same mesh density metric shown as a colour-coded node map.

6.5.4 Truncated Geometry Re-Meshing

It was observed that the node density in meshes generated with the above method is typically disproportionately high near sharp boundaries or where there is high curvature, such as in the ears in the presented example. It is assumed that this is because of the difficulty of matching the geometry well with larger triangles (fewer nodes) in these cases. This means that a great many nodes are used to represent these complex structures as compared to the main body (central part of the mesh) which means that more nodes need to be added to get accurate FEM data. This increases the complexity of the FEM analysis.

A solution to the problem is to remove parts of the mesh that are of little interest (in the presented case the head and ears) so as to use the same number of nodes to better represent the

part of the body under observation without wasting resources on those that are not of interest. In order to do this, the methods above are re-used with the mesh nodes first being truncated by a bounding box defined as the region of interest in the co-ordinate space. The methods are identical thereafter. Note that this is the one point in the routine at which manual intervention is required (to specify the bounding box having looked at the points/mesh) though it is not strictly *necessary* it is beneficial in terms of reducing computational complexity by requiring fewer nodes because the alternative is to add a lot of unnecessary nodes. Figure 6.27a shows the resultant re-meshing of the above data, a bounding box having been set that clips either end of the phantom particularly removing the head. It can be seen that the mesh has retained the detail and successfully made a flat surface where truncation was desired.

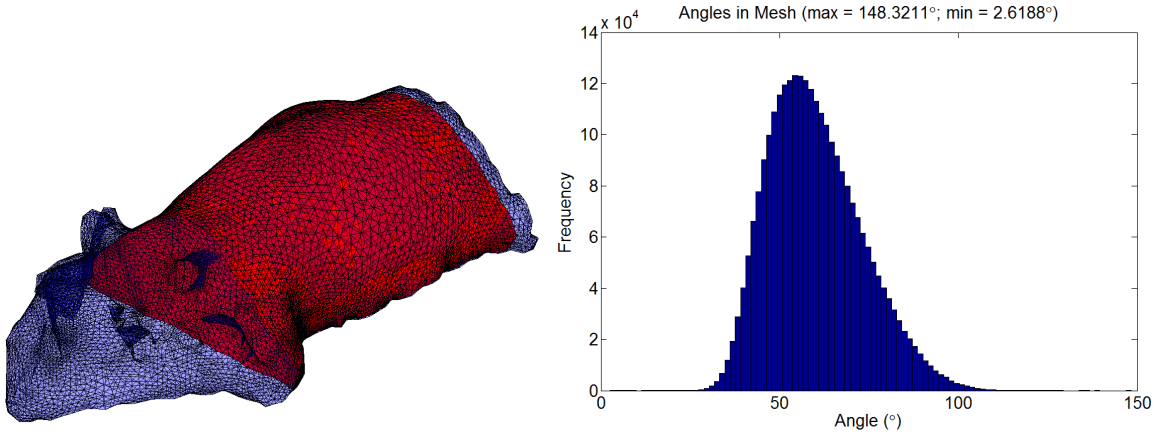

(a) (b)

Figure 6.27: (a) Re-meshing a truncated version of the generated XPM-2 volume mesh. The original mesh is shown in blue (semi-transparent) whilst the new mesh is in solid red, and (b) histogram of angles in triangles throughout the truncated FEM mesh. The majority are around 60 degrees indicating regularity of the mesh once again.

The new mesh was subjected to the same analyses as the original and the these are summarised in figs. 6.27b and 6.28. It can be seen that the triangle angles are still appropriate for FEM and that the distribution of nodes in the mesh is once again biased towards the higher-curvature areas in the sharp outer corners. However, with the new mesh (which has approximately the same number of nodes as the previous one), the minimum number of nodes within 3mm of any node is 39 which is a substantial improvement.

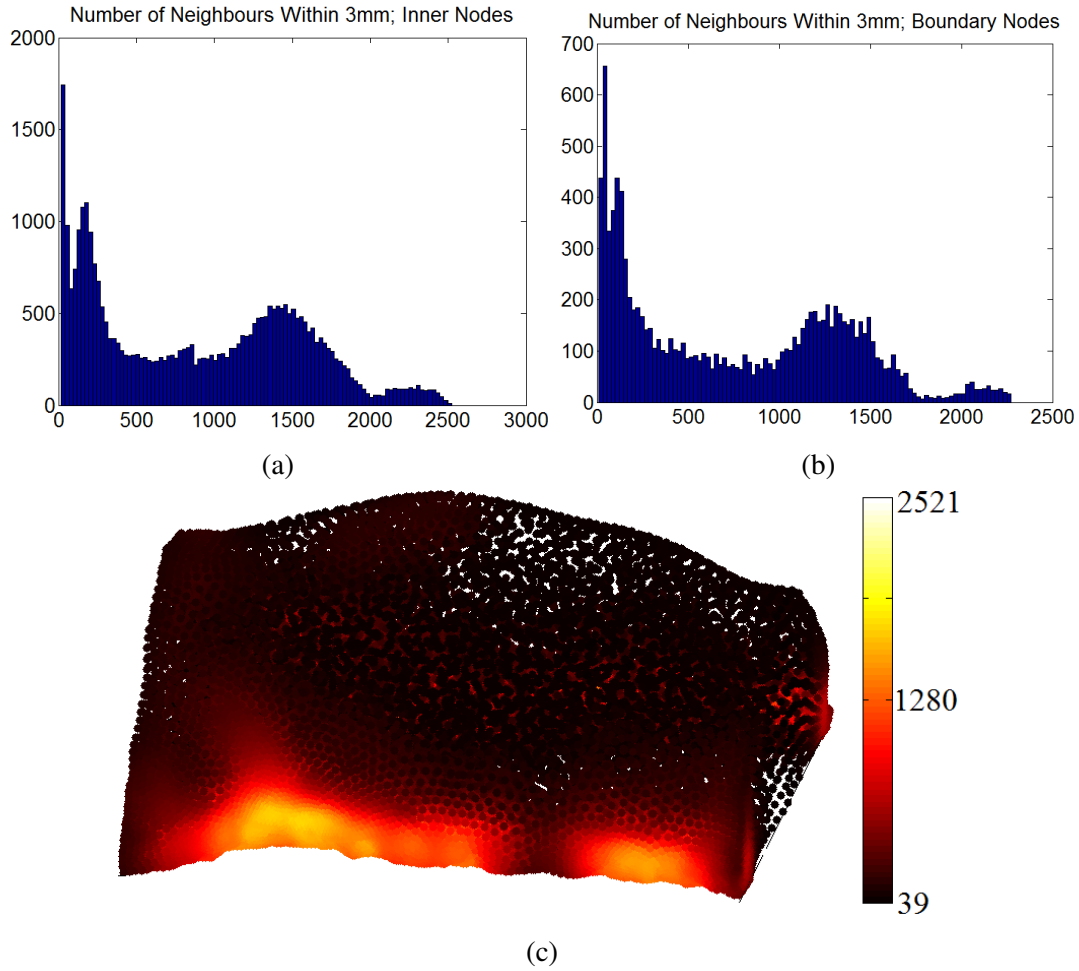

Figure 6.28: Analysis of node-density in the truncated FEM mesh: (a) histogram of number of neighbours of each mesh boundary node within a 3mm radius, (b) histogram of number of neighbours of each mesh boundary node within a 3mm radius, (c) the spatial distribution of the same mesh density metric shown as a colour-coded node map.

6.5.5 Placement of FEM Detectors

In diffuse optical tomography, the FEM 'detectors' are placed on the boundary of the medium representing points at which flux measurements are made. In the present system CCD measurements are mapped onto the boundary of the free-space mesh such that an average flux value is recovered per-face (chapter 7). In order to use these mapped measurements in FEM image reconstruction, the centroids of the faces are taken from the free-space mesh and these co-ordinates are used as the locations of the FEM detectors in the FEM mesh; in practice there are very slight variations in the curvature owing to the method above so the detectors are placed at the nearest point available on the surface of the FEM mesh. The resultant configuration for the running example is shown in fig. 6.29.

A particular advantage of constructing the meshes and FEM detectors in this approach is that the free-space transfer matrix T (chapter 7) now links FEM measurement locations directly to CCD measurements whilst the optical tomography (BLT and DOT) Jacobians (chapters 8 and 9) J and W relate internal properties to boundary measurements at the same locations. This means that the whole system can be modelled by combining the matrices, i.e.

$$S_{BLT} = TW, \tag{6.14}$$

and

$$S_{DOT} = TJ, \tag{6.15}$$

where these resulting *system matrices* describe the whole system.

6.5.6 Meshing Method Conclusions

The presented automated meshing method provides a free-space mesh with which to calculate the free-space transfer function for the system (chapter 7) relating CCD measurements to the surface faces. It also produces an FEM mesh with triangles that have angles appropriate for FEM analysis, however it has been found that the mesh node density using the present method

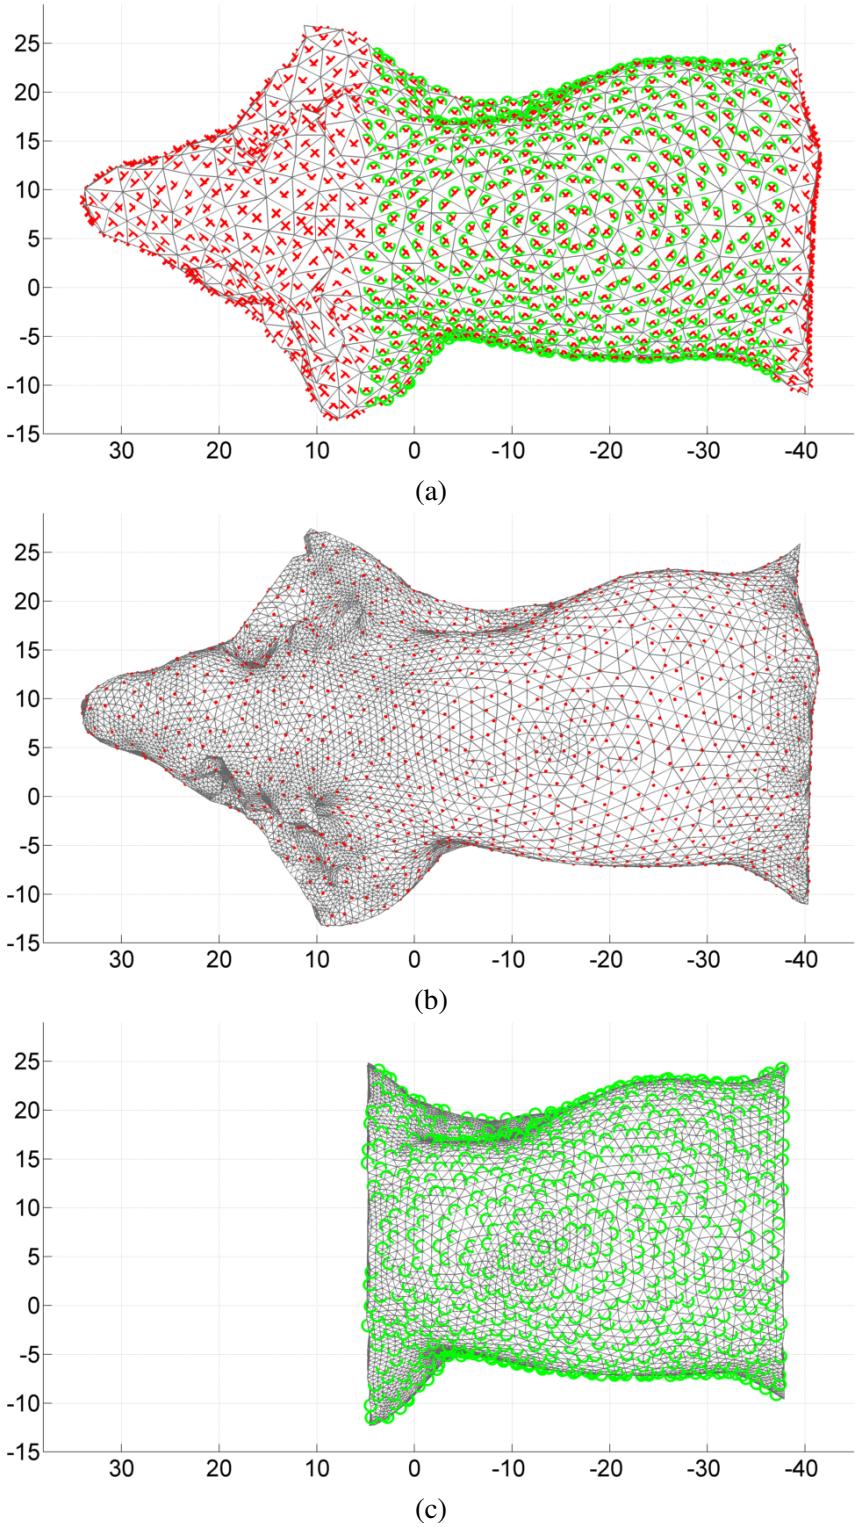

Figure 6.29: Illustration of FEM detector placement based on free-space mesh faces for which the free-space mapping function is established. Sub-figures show (a) the free-space mesh with face centroids indicated in red crosses and repeated in green crosses where they fall within the mesh truncation bounding box, (b) the original FEM mesh with face centroids placed as detectors, (c) the truncated FEM mesh showing the same.

is biased towards the edges where there is high curvature, this issue has been addressed to some extent by mesh truncation but is still an issue that bears further investigation.

It has been demonstrated that surface and volume meshes can be made using the presented method that have a specified mean face size and a specified number of nodes respectively. This allows control of the mesh generation simply in terms of the quantities that are most important to their particular usage.

Finally it has been shown how the centroids of the free-space mesh can be used as FEM detector locations which allows fully automatic generation of a full system matrix that relates CCD measurements directly to the target quantities in image reconstruction. This allows the investigation of methods of image reconstruction direct from images (which has been done by others previously in fluorescence tomography and has proved effective[187]) and also makes two-step image reconstruction (free-space mapping then image reconstruction) simpler as well.

Automated meshing along with automated data processing from the imaging system represent collectively an important step towards high throughput imaging studies.

CHAPTER 7

OBTAINING BOUNDARY LIGHT FLUX MEASUREMENTS FROM IMAGES

Within the paradigms introduced in sections 2.5 and 2.6, the image reconstruction methods in both diffuse optical tomography (DOT) and bioluminescence tomography (BLT) require light measurements on the boundary of the imaged volume. However, the raw data obtained by the presented system is in the form of images whose values are complex functions of the system and surface properties, as well as the desired surface flux. In this chapter, a new method[2] is presented for the reconstruction of surface fluxes from images acquired by the system, using a geometric system model (appendix A.6) and surface data obtained by surface capture (chapter 6).

The novel mapping method[2] is achieved by applying model-based reconstruction using an improved version of a previously published forward model of light propagation through free space from the surface to the camera[188], which accounts for the effects of perspective, focus and imaging geometry imposed by the system and achieves quantitative accuracy.

This flux mapping step is very important to the work-flow in the presented imaging system in which it is used in DOT (chapter 8) and BLT (chapter 9) to gather boundary data from images.

7.1 Introduction

It has been discussed that *in vivo* optical imaging studies of small animals typically involve the acquisition of 2D images (chapter 3). These are hoped to be somewhat representative of the light flux on the animal surface, which in turn is intended to be representative of internal luminescent activity. This may then be related back to biology.

Chapter 2 explained that the behaviour of light in tissue renders the relationship between light inside the volume and boundary measurements complex, and that modelling and inversion of this behaviour is the role of BLT and DOT image reconstruction as described in sections 2.5 and 2.6 and chapters 8 and 9.

The propagation of light from the surface to the camera, through free-space and through the optical system, is also complex. Identical surfaces seen from different perspectives or having different positions or orientations within identical imaging systems will appear different and provide different measurements. Similarly if the properties of the system such as the lensing are changed the measurements will also change. As an example, consider figure 7.1 in which a change in measurements due to a change in sample orientation is shown with a phantom. This is

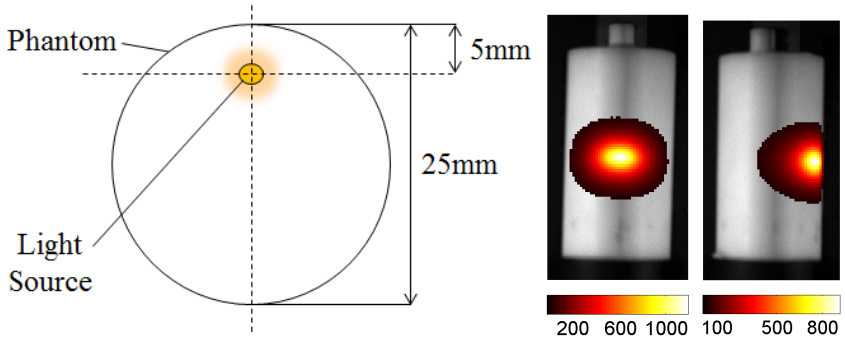

(a) Cross-section (b) (c)

Figure 7.1: (a) Cross-sectional schematic and (b-c) surface images (counts/s) of a diffuse light-emitting phantom taken from two viewing angles separated by 60°. This figure is reproduced from Guggenheim et al.[2].

representative of having a mouse lying differently within an imaging system when performing an *in vivo* imaging experiment. This effect must be understood and accounted for in order to obtain accurate boundary flux data from images.

Simple calibrations can be performed to convert image counts directly into surface measurements for well-known imaging domains (see for example Troy et al.[127]), but these cannot be generalised without more advanced techniques.

There also exists published work that models the process explicitly. Ripoll et al.[189, 187] analysed the non-contact imaging scenario and devised models of light propagation from external and internal sources to diffuse boundary fluxes and the propagation of those fluxes to non-contact detector measurements. Optical fibre detection schemes and in and out of focus camera-based systems were considered and an implementation was tested in which a camera focal plane was treated as an array of virtual detectors with an empirically determined numerical aperture term in-place to account for the angular sensitivity of the detection system[189]. By combining the models of internal and external free-space light transport into a single transformation matrix, Ripoll et al.[187] proceeded to perform model inversion to reconstruct internal fluorescence inclusions from experimental image data by use of an algebraic reconstruction technique (ART). The authors were therefore able to show that internal fluorescence distributions could be reconstructed directly from non-contact measurements.

The above work was extended by Chen et al.[190] who added an explicit model of a lens (treated as a thin lens) into the free-space forward model. This involved the introduction of a binary visibility coefficient for many discrete positions spanning the surface and detector. The visibility calculation and, to a lesser extent, the discretisation operations are computationally expensive but the model provides a thorough, general description of a lensing system. The same group further improved the model by accounting for additional apertures within the system[188]. The resulting linear forward model relates surface flux to images by a matrix multiplication.

Invoking the principle of optical path reversibility, the same group applied a simple model inversion method to solve the inverse problem; mapping fluxes from images. They used transfer matrix back-projection (multiplication by the transpose of the model matrix) and obtained qualitatively accurate results[191, 192, 193]. However, quantitative accuracy was not shown and is not generally obtainable with this approach. It was nevertheless shown that the mapped

flux could be used as the basis for BLT reconstruction in phantoms and animals[192, 193].

In this chapter methods are reported and evaluated that first improve the Chen[188] forward model allowing more generalised imaging geometry. Specifically a limitation is overcome involving the positioning of the system focal plane which was restricted to lie above the animal, but can now be positioned arbitrarily.

A different inversion method is then applied to the flux reconstruction problem in an attempt to achieve quantitative accuracy. The method is an iterative, regularised non-negative least-squares algorithm which has been applied to the BLT image reconstruction problem[58]. Results obtained by the new approach are compared with those obtained using back-projection and it is found that the new approach is indeed quantitatively far superior.

7.2 Theory

7.2.1 The Forward Problem

Consider a system containing a thin lens imaging a surface emitting diffuse light as shown diagrammatically in fig. 7.2. This thin lens set-up can act as an accurate model of more complex lensing systems providing that it is constructed with the correct front focal distance and focal length being maintained (the effective detection plane position must then be calculated according to the thin lens equation)[190, 188, 2].

Chen et al.[188] devised a forward model of the imaging process within this paradigm to map surface fluxes onto the lensed detector as follows. Assuming that the exiting light creates an effective Lambertian source at all surface points, the image intensity is:

$$P(\mathbf{r}_d) = \frac{1}{\pi} \int_S J(\mathbf{r}) T(\mathbf{r}, \mathbf{r}_d) dS, \tag{7.1}$$

where $P(\mathbf{r}_d)$ is the power on an infinitesimal detection area dA_d at detector point $\mathbf{r}_d$, and $J(\mathbf{r})$ is the surface flux on dS, a differential area element at surface point $\mathbf{r}$. The transfer function

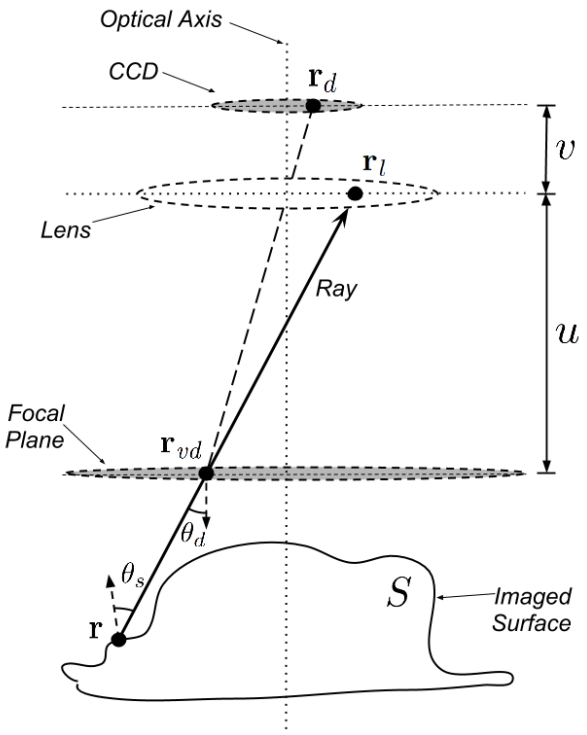

Figure 7.2: Schematic of the system model where the lens, detector and focal plane follow the thin lens equation. Surface points are either visible or invisible to detector (CCD) points dependant on whether the ray cast through the corresponding virtual detector (focal plane) point intersects the lens or aperture and is not blocked by part of the surface. In the illustration shown, the point $\mathbf{r}$ is visible to point $\mathbf{r_d}$. This figure is reproduced from Guggenheim et al.[2]

$T(\mathbf{r},\mathbf{r}_d)$ represents sensitivity between the detector and surface points

$$T(\mathbf{r},\mathbf{r}_d) = \alpha(\mathbf{r},\mathbf{r}_{vd})\,\beta(\mathbf{r},\mathbf{r}_{vd};\Omega_D)\frac{\cos\theta_s\cos\theta_d}{|\mathbf{r}_{vd}-\mathbf{r}|^2}dA_{vd}, \tag{7.2}$$

where $\mathbf{r}_{vd}$ is the focal plane point brought to focus at $\mathbf{r}_d$, θ_s is the angle between the line from $\mathbf{r}$ to $\mathbf{r}_{vd}$ and the normal to the surface at $\mathbf{r}$, and θ_d is the angle between the line from $\mathbf{r}$ to $\mathbf{r}_{vd}$ and the normal to the detector at $\mathbf{r}_d$ (figure 7.2). $dA_{vd} = dA_d/t^2$ is the focal plane element area with t being the lens magnification $t = v/u$. The object distance u and image distance v adhere to the thin lens equation $1/f = 1/u + 1/v$. The $\alpha(\mathbf{r},\mathbf{r}_{vd})$ and $\beta(\mathbf{r},\mathbf{r}_{vd};\Omega_D)$ terms are visibility

coefficients describing the visibility of surface points to detection points

$$\alpha(\mathbf{r},\mathbf{r}_{vd}) = \begin{cases} 1 & \text{If } (\mathbf{r}_{vd} \in \Omega'_E) \text{ AND } (\mathbf{s}_{\mathbf{r}\to\mathbf{r}_{vd}} \cap \mathbf{S} = \{\mathbf{r}\}) \\ 0 & \text{Otherwise} \end{cases} \tag{7.3}$$

and

$$\beta(\mathbf{r},\mathbf{r}_{vd};\Omega_D) = \begin{cases} 1 & \text{If } (\mathbf{s}_{\mathbf{r}\to\mathbf{r}_{vd}} \cap \Omega_D \neq \varnothing) \\ 0 & \text{If } (\mathbf{s}_{\mathbf{r}\to\mathbf{r}_{vd}} \cap \Omega_D = \varnothing) \end{cases}. \tag{7.4}$$

Here, $\mathbf{s}_{\mathbf{r}\to\mathbf{r}_{vd}}$ is the line segment joining $\mathbf{r}$ and $\mathbf{r}_{vd}$, Ω'_E is the region representing the lens and Ω_D is the region representing the aperture (Chen et al.[188] illustrates these parameters graphically). The α term has a value of 1 only if two conditions are met. Firstly, the virtual detector point must fall within the part of the focal plane that is visible through the lens; the line passing the surface and virtual detector points must intersect the thin lens plane within the bounds of the lens. Secondly, the line segment joining the surface and virtual detector points must not intersect the surface at any other point, i.e. there must be a clear line-of-sight, free from obstacles, between the surface and virtual detector points. The β term tests the same condition as the first α term but in this case the line must pass through an aperture. In the presented work the β term is used to model mechanical vignetting by considering the aperture at the first entrance to the optical system some way ahead of the lens front-face. Under this analysis, it is possible to establish visibility or non-visibility using simple ray-tracing within a geometric system model.

This method for establishing visibility is sufficient for scenarios in which the focal plane lies somewhere between the lens and the surface. It fails, however, when the focal plane is located behind or through the surface[2]. A modification is therefore applied in this work to correct the first visibility test. This becomes

$$\hat{\alpha}(\mathbf{r},\mathbf{r}_{vd};\Omega'_E) = \begin{cases} 1 & \text{If } (\mathbf{r}_{vd} \in \Omega'_E) \text{ AND } (\mathbf{s}_{\mathbf{r}\to\mathbf{r}_l} \cap \mathbf{S} = \{\mathbf{r}\}) \\ 0 & \text{Otherwise} \end{cases}, \tag{7.5}$$

in which $\mathbf{r}_l$ has been introduced, representing the point of intersection between the line through $\mathbf{r}$ and $\mathbf{r}_{vd}$ and the thin lens plane. The term now tests for clear line-of-sight between the surface and the lens plane as opposed to between the surface and the focal plane. This corrects the visibility test in general as illustrated in fig. 7.3. The modified forward model is

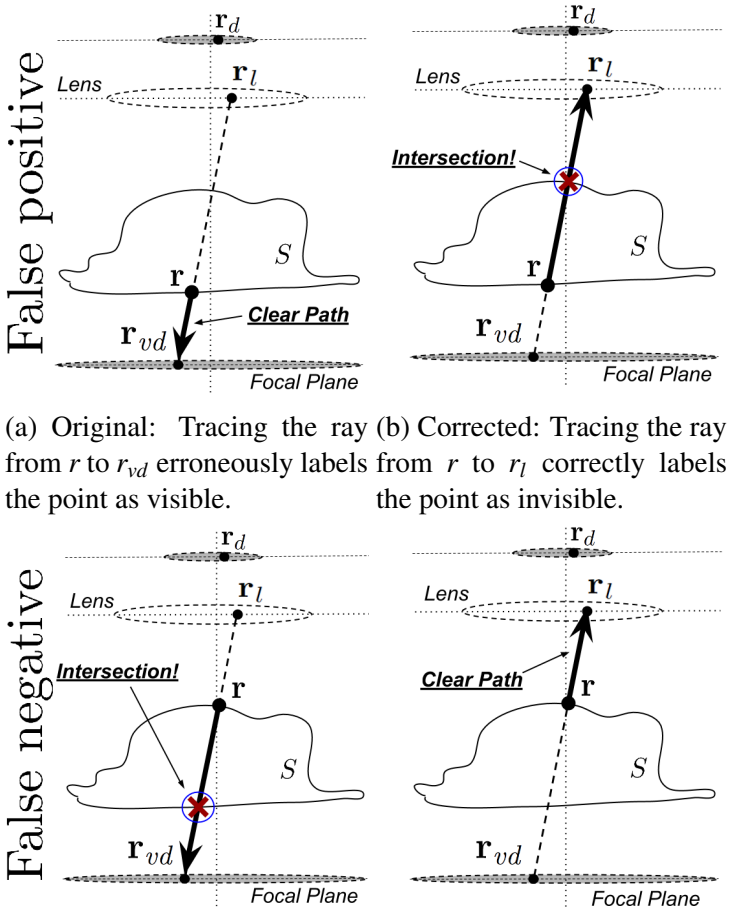

(a) Original: Tracing the ray from r to r_{vd} erroneously labels the point as visible.

(b) Corrected: Tracing the ray from r to r_l correctly labels the point as invisible.

(c) Original: Tracing the ray from r to r_{vd} erroneously labels the point as invisible.

(d) Corrected: Tracing the ray from r to r_l correctly labels the point as visible.

Figure 7.3: Illustration of the improved visibility factor $\hat{\alpha}$. This figure is reproduced from Guggenheim et al.[2].

$$P(\mathbf{r}_d) = \frac{1}{\pi}\int_S J(\mathbf{r})\Gamma(\mathbf{r},\mathbf{r}_{vd};\Omega_D)dS, \tag{7.6}$$

with

$$\Gamma(\mathbf{r},\mathbf{r}_{vd};\Omega_D) = \hat{\alpha}(\mathbf{r},\mathbf{r}_{vd};\Omega'_E)\beta(\mathbf{r},\mathbf{r}_{vd};\Omega_D)\frac{\cos\theta_s\cos\theta_d}{|\mathbf{r}_{vd}-\mathbf{r}|^2}dA_{vd}. \tag{7.7}$$

With a finely discretised surface, detector and focal plane, this can be written

$$\mathbf{a} = \mathbf{T}\mathbf{b}, \tag{7.8}$$

in which elements of the matrix $\mathbf{T}$ and the vectors $\mathbf{a}$ and $\mathbf{b}$ are

$$\begin{cases} \mathbf{a}_i = P(\mathbf{r}_d^i) \\ \mathbf{T}_{ij} = \frac{1}{\pi}\Gamma(\mathbf{r}^j, \mathbf{r}_{vd}^i; \Omega_D) \\ \mathbf{b}_j = J(\mathbf{r}^j)dS^j \end{cases} \tag{7.9}$$

For a known surface and imaging system that are unchanging, the relationship between boundary flux and image intensity is linear.

7.2.2 The Inverse Problem

The discussed forward problem is the calculation of image values from the boundary flux on an imaged surface. Correspondingly, the inverse problem is the computation of the surface flux from the image values.

Inspired by the reciprocity theorem and the principle of light path reversibility in an imaging system, Chen et al.[193] proposed to solve the inverse problem by simply using the forward model[188] with the surface flux and detected power terms switched around, i.e.

$$J(\mathbf{r}) = \frac{1}{\pi}\int_A P(\mathbf{r}_d)\,\alpha(\mathbf{r}_{vd}, \mathbf{r})\,\beta(\mathbf{r}_d, \mathbf{r}; \Omega_D)\frac{\cos\theta_s \cos\theta_d}{|\mathbf{r} - \mathbf{r}_{vd}|}dA, \tag{7.10}$$

where the visibility terms are identical to those used in the forward model[188]. The idea was that the detector could be treated as a diffusive light source and the surface as a detector. By updating the visibility conditions to include the correction introduced above, this can be written as

$$\tilde{\mathbf{b}} = \mathbf{T}^T\mathbf{a}. \tag{7.11}$$

This approach is therefore a simple backprojection via the forward transfer matrix. As such, this will be referred to as the backprojection method from here on. Whilst this was found to be effective in certain experimental situations, particularly when considering only normalised results[191, 193, 192], it does not achieve quantitative accuracy (see section 7.4).

The alternative approach presented here is based instead upon recovering **a** by solving the inversion

$$\tilde{\mathbf{b}} = \mathbf{T}^{-1}\mathbf{a}. \tag{7.12}$$

Generally this cannot be directly computed because **T** is non-square and thus without a precise inverse. In an attempt to overcome this, a reconstruction method that has been applied to BLT image reconstruction[58] is investigated. The problem is first stated as a least-squares minimisation

$$\tilde{\mathbf{b}} = \min_{\mathbf{b}} ||\mathbf{T}\mathbf{b} - \mathbf{a}||_2^2. \tag{7.13}$$

It is then solved using a single-step Tikhonov-regularised inversion

$$\tilde{\mathbf{b}} = (\mathbf{T}^T\mathbf{T} + \tilde{\lambda}\mathbf{I})^{-1}\mathbf{T}^T\mathbf{a}, \tag{7.14}$$

in which $\tilde{\lambda}$ is a regularisation parameter defined as the product of a chosen value and the maximum of the diagonal of the Hessian matrix: $\tilde{\lambda} = \lambda \times max(diag(\mathbf{T}^T\mathbf{T}))$. The single-step inversion is performed, then any negative values in the solution vector $\tilde{\mathbf{b}}$ are set to zero. The remaining error is then calculated and solved for using the above equation. This step is repeated, solutions being obtained and negative values set to zero iteratively, until the stopping criterion is satisfied. In this work, a stopping criterion was used whereby the relative improvement in the solution fit between iterations must be less than 1%. In addition, iterations are limited to a chosen maximum value.

7.3 Materials and Methods

To implement the light-modelling described above, **T** matrices relating finely discretised detection and surface elements must be calculated. As a precursor, the system and surface geometries must be known allowing the focal plane and object surface to be broken into myriad tiny elements and allowing ray-tracing to be performed within the model domain. This section describes methods used to perform several simulation and practical experiments to test the introduced new mapping method. In all cases, computations were performed in MATLAB (MathWorks, Cambridgeshire, UK) on a desktop PC[1].

Prior to practical experiments, the imaging system geometry was established using the geometric calibration method described in appendix A.6.

7.3.1 Simulations

Two sets of simulation experiments were performed, the first in a 2D plane with a minimalist version of the presented imaging system being considered, and the second modelling the system in 3D.

The 3D simulation experiments were structured as follows. Firstly, an FEM mesh of the subject was constructed and NIRFAST[25] was used to calculate boundary data at the centre-points of all surface triangles. Secondly, a simulated image was obtained by multiplication of the boundary data with a calculated **T** matrix. Finally, flux reconstruction was performed for the simulated image data using another calculated **T** matrix. In each case, an inverse crime was avoided by the use of two versions of the **T** matrix, based on different surface/detector discretisations. The 2D experiments were similarly conducted but fluxes were chosen explicitly to be simple functions.

For simplicity, spectral lens transmittance efficiency, camera quantum efficiency and image digitisation were not considered in simulation.

[1] Specified in detail in Guggenheim et al.[2]

7.3.2 Practical Phantom Imaging

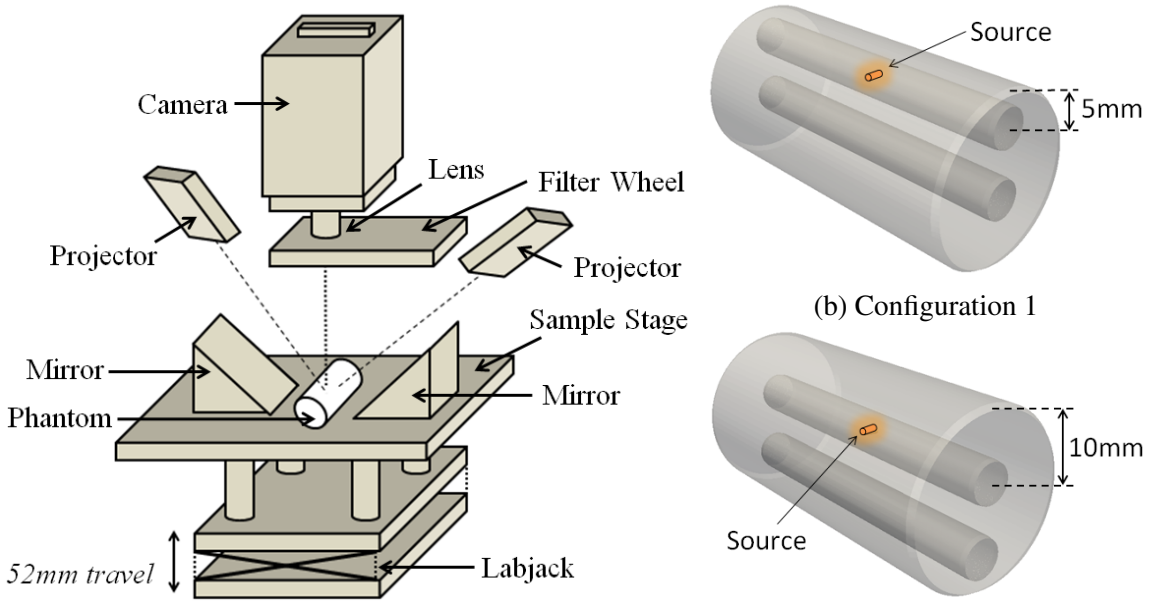

(a) Imaging system diagram.

(b) Configuration 1

(c) Configuration 2

Figure 7.4: Illustration of phantom positioned in system and schematic of two experimental configurations with different source locations. These figures have been recombined from Guggenheim et al.[2].

Physical imaging experiments were carried out in the presented imaging system (chapter 5) using the cylindrical phantom introduced in appendix A.1.1. Bioluminescence was modelled by an orange trigalight (appendix A.1.2) placed roughly half-way along one of the cylinder phantom tunnels so as to render the phantom homogeneous. Background matching rods were placed around the trigalight and in the other tunnel. Resultant experimental configurations are shown in fig. 7.4.

In practical experiments, the phantom was positioned in the geometrically calibrated system, surface captured, and then imaged at a single wavelength. Transfer matrices were calculated and fluxes were mapped from acquired images.

7.3.3 Cylinder Mesh Registration

Whilst in general it is intended that the automated combined surface capture and meshing method described in chapter 6 be used to create meshes on-the-fly by the system, in the case of

the cylinder phantom (appendix A.1.1) used in this chapter, as well as in parts of chapter 8 and chapter 9, the method is inappropriate because it assumes that the imaged object is flush with the stage. For the cylinder this is not the case and, furthermore, the underside curvature cannot be accounted for. As such, pre-made meshes are used to represent the cylinder phantom and simply registered to surface capture data.

To facilitate this, one rod is left protruding slightly from the phantom on one side such that it is visible to the camera and therefore points are found on it in surface capture. These points can then be used to establish the rod location, uniquely determining the cylinder rotation.

The specific details of the registration method can be found in Guggenheim et al.[1]. Briefly, it is a three-step process in which the first step is a global fit based on minimising the distance of surface capture points from an infinitely long perfect cylinder of appropriate size, the second step establishes the rotation of the cylinder from rod points, and the third step establishes the axial limits of the cylinder based on the assumption that at least one end is visible in the field-of-view.

7.4 Results

7.4.1 Two-Dimensional Simulation

In the first 2D-like simulation experiment, a central slice through the imaging system was considered (figure 7.5). The imaged surface was passed over a range of heights to investigate the impact of global position change on forward modelled data and reconstructed flux. Various signal strengths were considered along with added shot noise to investigate the effects of SNR on the reconstruction. Backprojection and the new mapping method are compared quantitatively.

Simulated components were 3D objects but the surface and detector were made thin (1-element-thick) to create a problem domain that can be effectively visualised in 2D, but for which the 3D analysis is still valid.

The simulated detector comprised 1024 $13 \times 13\mu$m detection elements. A thin lens ($f =$

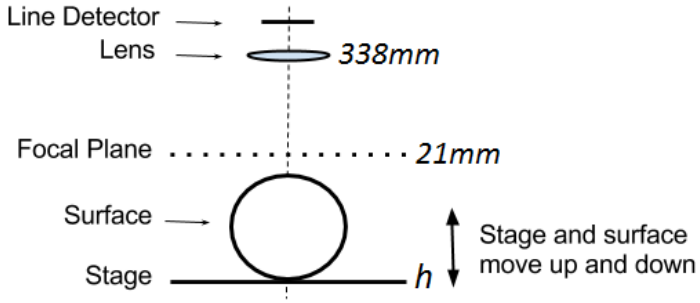

Figure 7.5: Diagram of simulated scene. This figure is reproduced from Guggenheim et al.[2].

23.9mm) was positioned 25.8mm ahead of the detector with a front focal plane 317mm ahead of the lens (at $z = 21$mm). The simulated subject was a 1mm thick cylinder with a 25mm radius, consisting of 200 equal-sized elements with areas of 0.79mm^2. Detection elements were faced perpendicular to the optical axis of the system as was the top-most cylinder element. The detector and surface elements were then discretised $6\times$ and $18\times$ respectively in the x direction only. The cylinder was assumed to rest on a horizontal stage, as such the maximum surface height was determined by the stage-height which was given one of several values ($z = -120, -80, -40, 0, 40$mm).

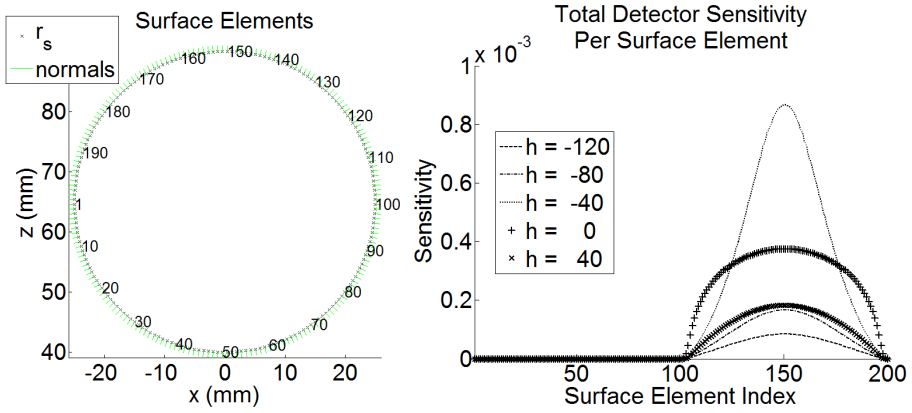

(a) (b)

Figure 7.6: (a) Diagram showing all surface points used in the simulation along with the normals. (b) Plot of total system sensitivity to the surface for different stage heights, h. This figure is reproduced from Guggenheim et al.[2].

Figure 7.6 shows the summed sensitivity across the entire detector to each surface element for the range of stage heights. Approximately half of the surface elements have non-zero sensitivity which is expected because the underside of the curved surface is invisible to the detector.

Strongest total sensitivity is seen at the top of the circular surface (also the centre of the field-of-view) and depletes symmetrically towards the sides. The depletion is caused primarily by the increasing surface curvature in this case. In general it is governed by equation 7.6 and depends therefore on the angles between the surface and detector normals and the rays that join surface and detector points, the distance of the surface from the focal plane and the visibility function.

The results demonstrate that the maximum sensitivity value is clearly influenced by the surface offset from the focal plane and it can be seen that the response function broadens as the object becomes out-of-focus. This is in-keeping with the blurring that is observed in the system.

In order to represent a high but realistic flux seen in diffuse imaging experiments, a simple flux was considered on the boundary of the cylinder; a Gaussian function of the circumferential position the maximum value of which was 50 million photons (corresponding to a radiance of approximately 10^9p/s/cm^2/sr, see section 3.3.3). The flux was multiplied by $\mathbf{T}$ matrices, simulating the measurement process in each height-resolved scenario. Figure 7.7 shows resultant simulated detector measurements alongside the initial flux.

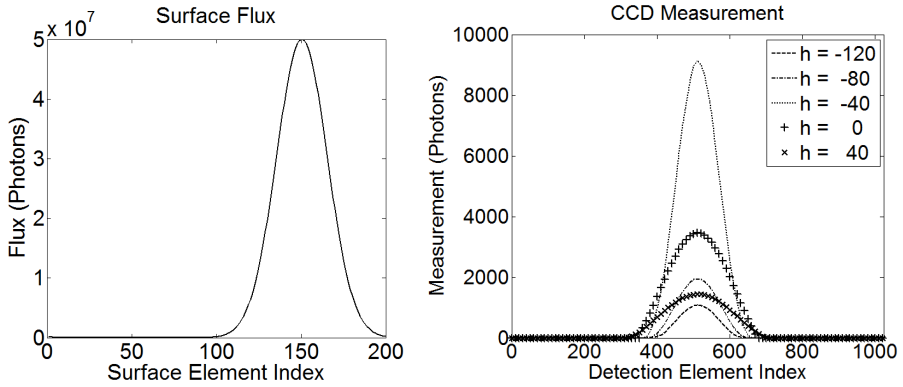

(a) (b)

Figure 7.7: (a) Simulated surface flux (ground truth), (b) measurement on line-detector simulated in each height-resolved set-up. This figure is reproduced from Guggenheim et al.[2].

The figure shows that the flux is qualitatively reproduced, i.e. imaged, on the detector. The image is best-focused when the stage height is at $h = -40$, as expected because this is the scenario in which the upper-half of the surface is closest to the focal plane of the system. As expected, the signal becomes more blurred as the distance from the surface to the focal plane is increased.

Normally distributed noise with $\sigma(r_d) = \sqrt{P(r_d)}$, emulating shot noise, was added to the simulated data so as to more accurately model practical measurement. **T** matrix calculation was repeated using a coarser discretisation level (9× for the surface and 3× for the detector) and flux was mapped using backprojection and the non-negative iterative least squares algorithm. In the case of the latter, the regularisation parameter was fixed at $\lambda = 10$ and the maximum allowed number of iterations was 10.

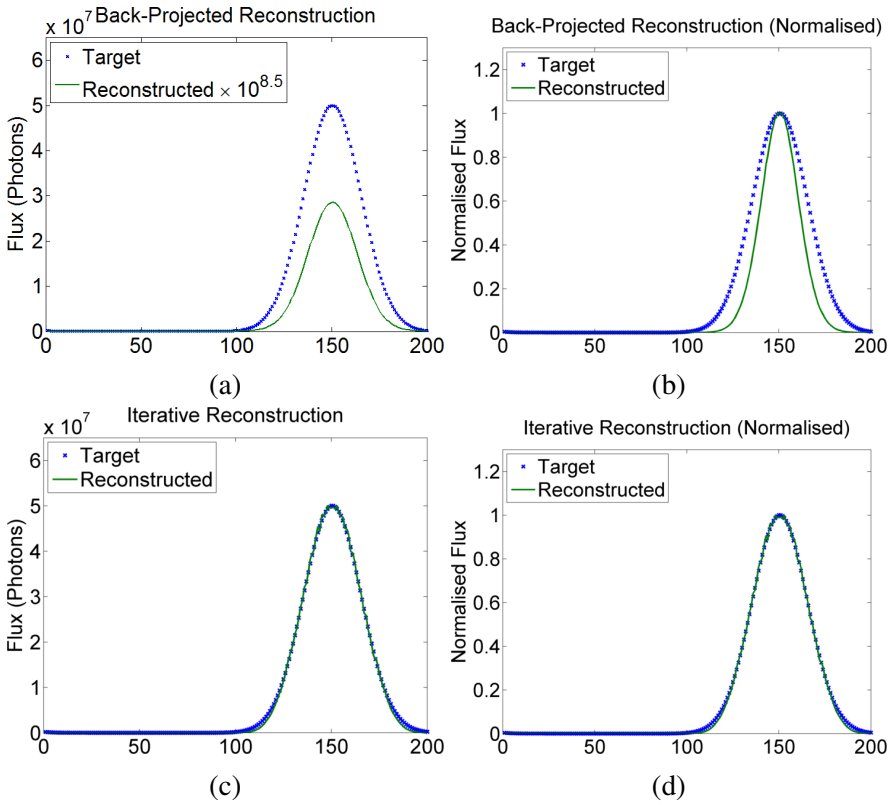

Figure 7.8: Sample reconstructed fluxes overlaid on ground truth values for $h = -40$mm and 50 million photons: (a-b) backprojection approach with absolute or normalised comparison, (c-d) proposed approach with absolute and normalised comparisons. Note particularly the 8.5× factor indicated in (a). This figure is reproduced from Guggenheim et al.[2].

Figure 7.8 shows the results for one test case with a stage height of −40mm with a peak surface flux value of 50 million photons. Whilst the normalised back-projection result is qualitatively somewhat accurate (the flux is centred on the correct location, the distribution is smooth) consistent with the literature[193], the raw backprojection result is not accurate. In fact, the backprojected peak flux value is around 8.5 orders of magnitude lower than the original signal. The reason for this is that the **T** matrix multiplication reduces signal by around 4 orders of

magnitude and this is effectively applied twice here (once for the measurement, once for the backprojection). This is clearly a fundamental flaw in the backprojection approach in this form. The new reconstruction method on the other hand provides good qualitative and quantitative accuracy in this case without normalisation.

Several more reconstructions were performed for a range of signal strengths ($\in [5\times10^7, 5\times10^6, 5\times10^5, 5\times10^4]$) and for each stage height. The percentage total error was calculated for each of the reconstructions to evaluate the results:

$$error(\mathbf{b},\mathbf{b}') = \frac{\sum_{j=1}^{n} |b'_j - b_j|}{\sum_{j=1}^{n} b_j} \tag{7.15}$$

In which $\mathbf{b}$ is the ground truth values and $\mathbf{b}'$ is the reconstructed surface flux. The reconstructions were performed 50 times with different instances of shot noise added. Figure 7.9 shows the results.

The figure shows that the backprojection approach generally lacks quantitative accuracy as it has an effective error of 100% in all cases. By contrast, the new inversion method produces very low errors in cases of high SNR (maximum SNR is shown in table 7.1 for added clarity). In general, the error for this method increases as the SNR gets lower. The mean error is less than 10% where the maximum SNR is greater than 10:1, less than 5% where it is greater than 30:1, and less than 3% in the highest signal case (with SNR $\approx$ 95:1). In this case the standard deviation on the error was less than 5% of the mean value. In some cases mean error was less than 30% despite maximum SNR being less than 10:1. To put this in context, a maximum SNR of around 200:1 could be achieved in practice (section 7.5). As such, these test cases are difficult.

7.4.2 Three-Dimensional Simulation of Multi-View Data

A 3D simulation experiment was performed in which the imaging system and physical cylinder phantom (appendix A.1.1) were modelled. A bioluminescent source was simulated at a depth of 5mm within the cylinder (as per figure 7.4b) and simulated flux was computed for a single

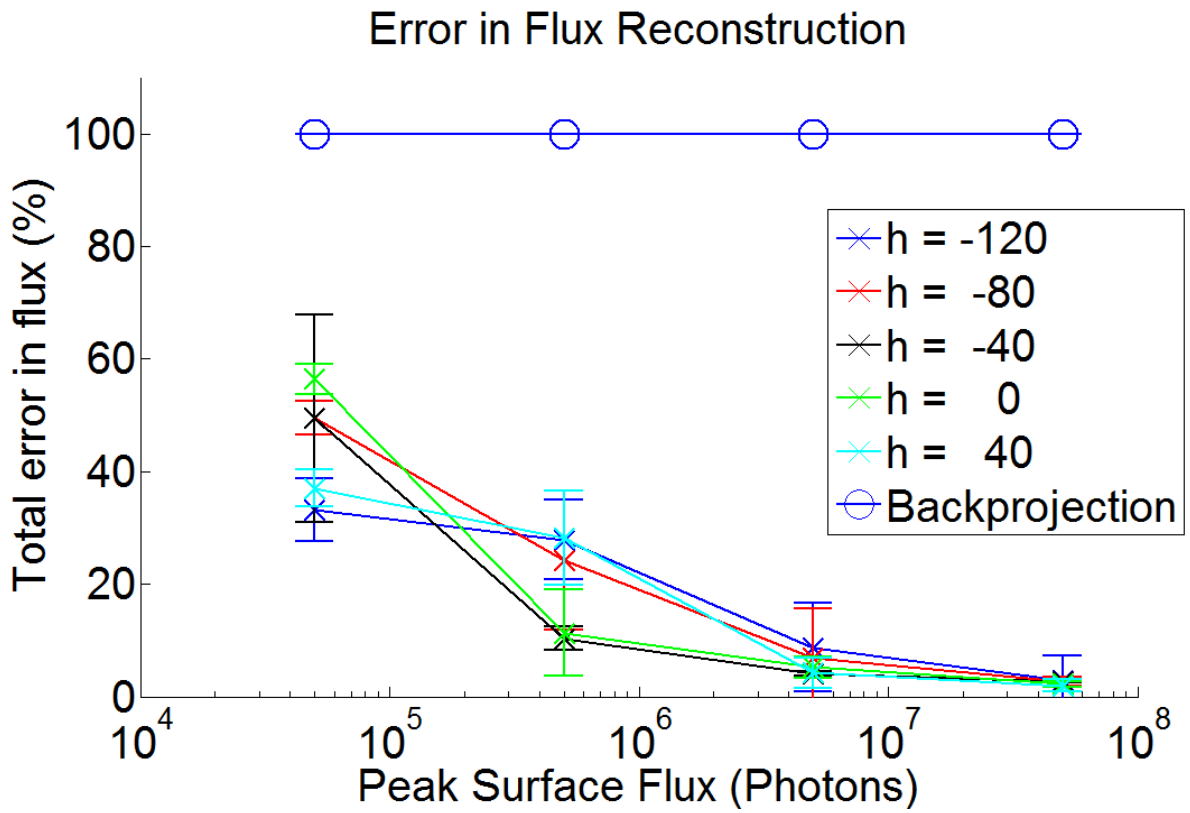

Figure 7.9: Percentage total error (eq. (7.15)) in reconstructed flux for different stage heights (h) and signal strengths. The values are the mean of 50 repeats with standard deviation bars. A single data set is shown for the backprojection method because the error was practically 100% in all cases. This figure is reproduced from Guggenheim et al.[2].

wavelength, 580nm.

The system mirror-views were modelled in addition to the main view. Three separate free-space transfer systems were considered together to model the multi-view imaging. To this end, the camera and lens geometries were reflected about each of the mirrors (treated here as infinite reflective planes) to create virtual detection systems. Figure 7.10 shows the real and virtual focal plane positions in the system. The side-view focal planes are nearer to the object indicating that the side-views were somewhat better focused on the surface.

There are three times more detector elements in the multi-view system. However, once the necessary geometric knowledge is established (the positions and orientations of real and reflected components), the modelling process continues as before. Multi-view images are then constructed by the addition of values recorded for the direct or reflected versions of detection elements.

Peak Flux (Photons)	SNR as a function of stage height h (mm) h = −120	h = −80	h = −40	h = 0	h = 40
5×10^7	32.93:1	44.23:1	95.71:1	58.83:1	38.01:1
5×10^6	10.41:1	13.99:1	30.26:1	18.60:1	12.02:1
5×10^5	3.29:1	4.42:1	9.57:1	5.88:1	3.80:1
5×10^4	1.04:1	1.40:1	3.03:1	1.86:1	1.20:1

Table 7.1: Maximum SNR in simulated noisy measurements used for reconstructions. Noise was added as shot noise so SNR is a function of signal; the listed values are therefore $\sqrt{m}$ where m is the maximum-valued simulated measurement. This table is reproduced from Guggenheim et al.[2].

The system CCD was represented by an array of 256×256 detection elements; deliberately matching the number of physical pixels in the system with the 1MP detector binned 4×4. The cylinder surface was broken into a 50×32 grid (circumferentially by axially in a local cylindrical co-ordinate system) to create 1600 discrete elements. In computing forward data, further discretisations of 26×26, 41×41, and 21×21 were applied to the surface respectively for the direct, left and right mirror views. Discretisations of 2×2, 4×4, and 2×2 were applied to the CCD. The surface was subjected to higher discretisation levels than the detector because the initial size of the surface elements was around $10\times$ the size of the virtual detector elements in the focal planes. Different levels were applied for different views because there were different distances to the surface from each of the focal planes and it has been found empirically that more discretisation is required to achieve high accuracy when the focal plane is near to the surface.

Simulated data was calculated for seven cylinder positions, representing seven cumulative rotations of the cylinder about its axis by approximately 30 degrees. Slight movements in global position were also introduced because positions were matched to the practical validation experiment presented in section 7.4.3.

The imaging system parameters were otherwise identical for each simulation. 100% mirror reflectance was assumed and mechanical vignetting effects were not considered. Simulated images are shown in figure 7.11. Surface flux was reconstructed using these simulated images. In this case no noise was added, λ was set at 0.1 and the maximum number of iterations was five. Discretisation factors used in reconstruction were equal to and one less than those used in

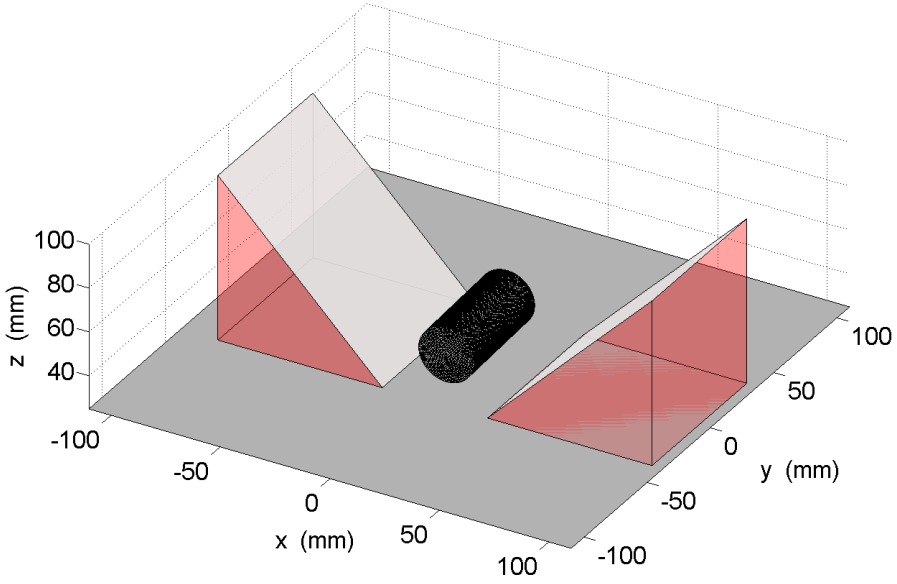

(a) Model of surface and mirrors

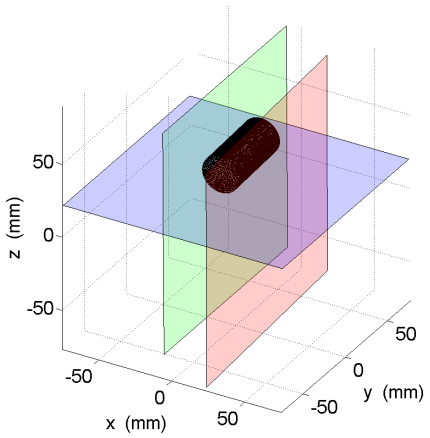

(b) Diagonal view of focal planes

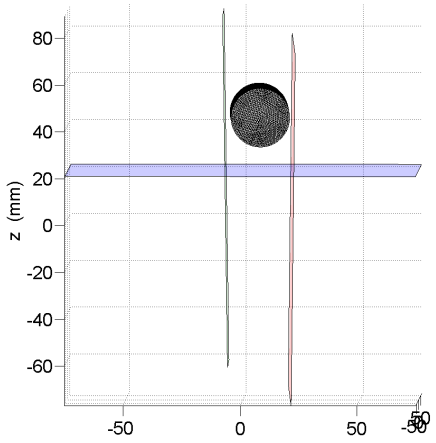

(c) Frontal view of focal planes

Figure 7.10: Geometric model of cylinder phantom, mirrors and focal planes (the direct focal plane of the lens-camera system and the two focal planes of the virtual systems through the mirrors). These figures are reproduced from Guggenheim et al.[2].

the forward model for detector and surface elements respectively.

An example of the resultant mapped flux (from the third image in figure 7.11; $\theta = 60°$) is shown along with the target flux in figure 7.12. The figure also shows a slice through the reconstructed flux data. The reconstructed flux is practically indistinguishable from the target values in this case. The equivalent slice plots corresponding to the other six reconstructions are shown in figure 7.13. The quantitative accuracy is visibly excellent across all data sets, except for the data for which the relevant part of the surface is physically invisible to the detector (see

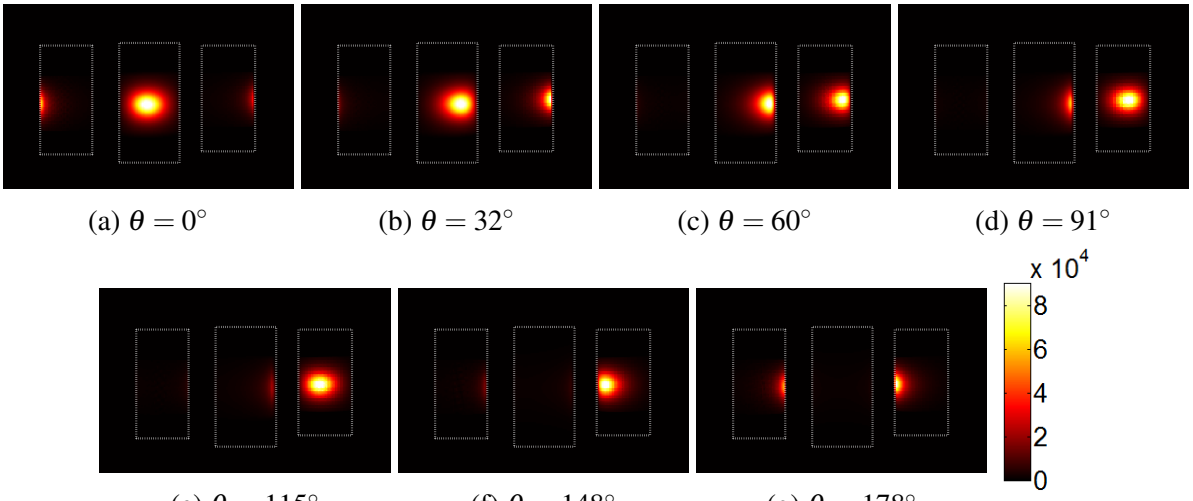

(a) $\theta = 0°$ (b) $\theta = 32°$ (c) $\theta = 60°$ (d) $\theta = 91°$

(e) $\theta = 115°$ (f) $\theta = 148°$ (g) $\theta = 178°$

Figure 7.11: Simulated multi-view images of cumulatively rotated phantom. Approximate phantom outlines for each view are shown as white dotted lines for reference. These figures are reproduced from Guggenheim et al.[2].

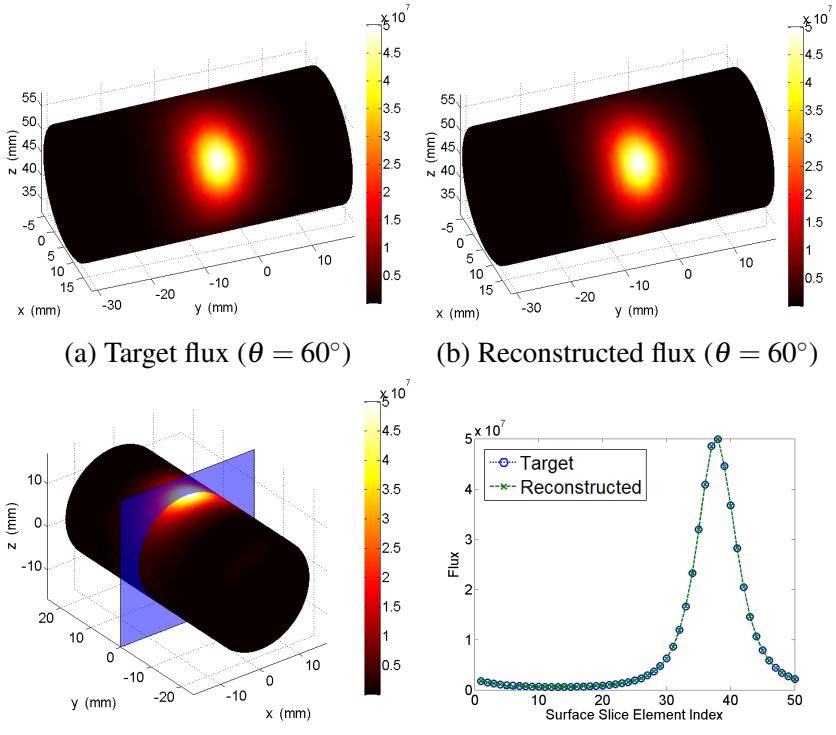

(a) Target flux ($\theta = 60°$) (b) Reconstructed flux ($\theta = 60°$)

(c) Flux slice location (d) Reconstruction slice ($\theta = 60°$)

Figure 7.12: (a) Ground truth and (b) mapped fluxes for the third simulation image along with (c) an indication of the plane at which circumferential data was extracted and plotted (d). These figures are reproduced from Guggenheim et al.[2].

figures 7.11f and 7.11g). This is unavoidable.

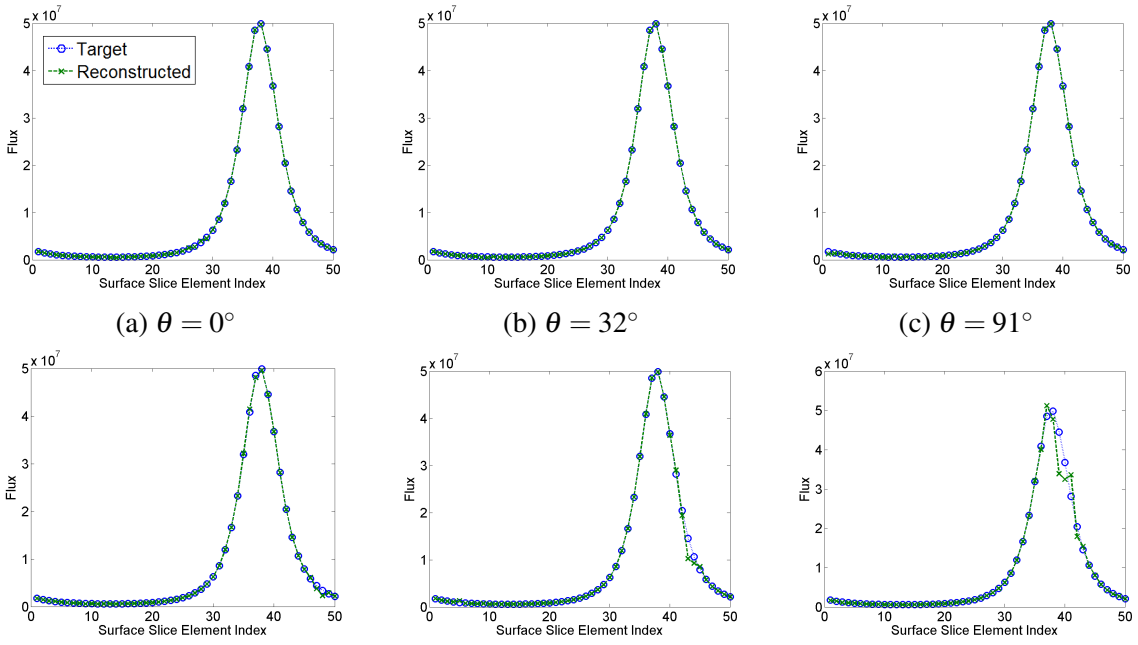

(a) $\theta = 0°$ (b) $\theta = 32°$ (c) $\theta = 91°$

(d) $\theta = 115°$ (e) $\theta = 148°$ (f) $\theta = 178°$

Figure 7.13: Circumferential slice data for reconstruction and target. The flux data is extracted following registration of the surface to a common pre-defined orientation as such the target and reconstructed fluxes would be expected to appear in a common location. These figures are reproduced from Guggenheim et al.[2].

7.4.3 Practical Validation with Bioluminescence Phantom

In the first practical experiment, the cylindrical bioluminescence phantom was prepared as shown in fig. 7.4b with the source 5mm deep. It was positioned within a rotation mount fixed between the mirrors such that it could be rotated without changing the global position. The phantom was imaged at 580nm at seven orientations resulting in the images shown in fig. 7.14.

As in simulation, a single representation of the cylindrical surface was used in computations and this was registered in each imaging case to the surface capture points (rather than building a surface from the surface capture). Consequently, in each case the surface elements were identical but located at different positions relative to the system.

Although the intention was to turn the phantom cumulatively by 30 degrees between images, there was some imprecision and as such the orientation of the cylinder was established after

imaging based on the result surface capture (section 7.3.3).

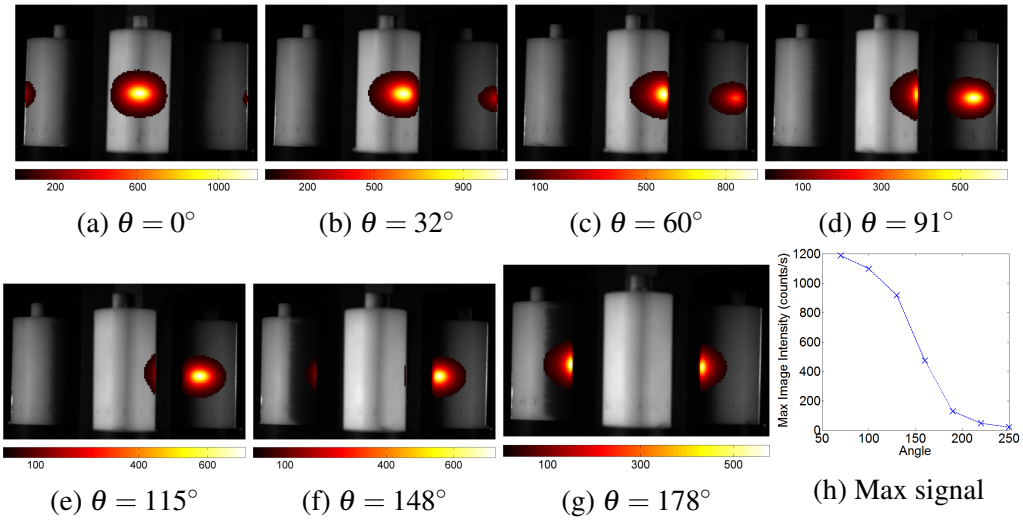

(a) $\theta = 0°$ (b) $\theta = 32°$ (c) $\theta = 60°$ (d) $\theta = 91°$

(e) $\theta = 115°$ (f) $\theta = 148°$ (g) $\theta = 178°$ (h) Max signal

Figure 7.14: Imaging data from the first set of practical experiments with (a-g) recorded images (colour scale shows electrons per second from the CCD) overlaid on back-lit images for spatial context, and (h) a plot of maximum signal recorded in the direct-views against the turned angle. These figures are reproduced from Guggenheim et al.[2].

Figure 7.14h demonstrates the angular dependence of the maximum signal in the images (excluding mirror data for clarity); the only signal typically available in BLI or FLI studies. It can be seen that there is a dramatic change owing to the changing perspective effects and visibility.

The flux mapping method was applied to the practical data with $\lambda = 1$ and a maximum number of iterations of five. Discretisation levels used were the same as in the 3D simulation reconstructions. Mechanical vignetting was accounted for by modelling the 25mm diameter aperture at the opening to the filter wheel with the β term in eq. (7.6). The vignetting causes some dimming of the mirror-view data as can be seen in the images (fig. 7.14), particularly when compared to the simulated images (fig. 7.11). A single empirically determined scaling factor was also applied to all mirror-view data to compensate for the differing transmittance properties of the interference filter (appendix A.2.5) which also dims the mirror-view data.

Figure 7.15 shows flux reconstructions in 3D following co-registration of all surfaces back to a common orientation for clarity. The reconstructed flux distributions appear to be consistent in terms of the spatial distribution and values. Figure 7.16 shows circumferential slice data

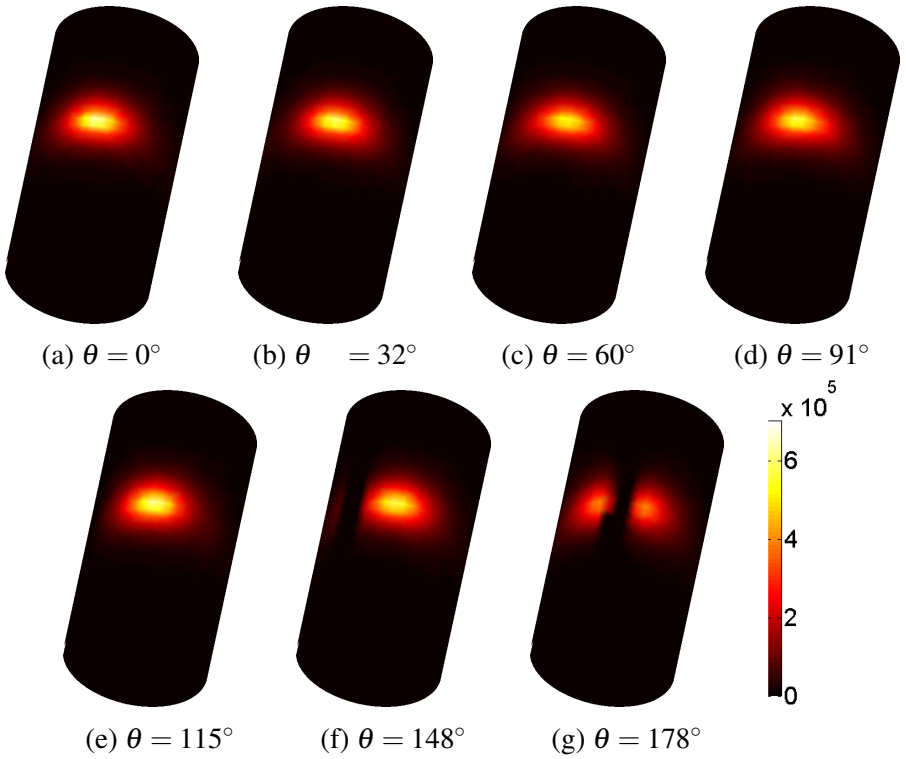

(a) $\theta = 0°$ (b) $\theta = 32°$ (c) $\theta = 60°$ (d) $\theta = 91°$

(e) $\theta = 115°$ (f) $\theta = 148°$ (g) $\theta = 178°$

Figure 7.15: Practical experiment reconstructed fluxes in consistent arbitrary units, co-registered to a common orientation. These figures are reproduced from Guggenheim et al.[2].

extracted as for the simulation data (fig. 7.12). These show the mapped practical data is indeed very consistent. It can be seen in some cases (most notably fig. 7.16(f) and (g)) that there are gaps where mapped data was not obtained, this is owing to the total invisibility of the relevant part of the surface in the corresponding scenario. These results are similar to those of the simulation though there is a little more quantitative variability between scenarios most likely owing to modelling error. Figure 7.17a shows a comparison of summed reconstructed signal based on the results of backprojection, of the new method, and of naive quantification direct from single-view images (i.e. summing background-subtracted top-view image values). The graph only covers data sets one to five because other cases suffer gaps in the flux due to the surface invisibility and comparison with these would be unfair. The graph shows that the naive method produces values that vary 10-fold. The backprojection method offers little improvement on this. Conversely, the new method produces a value that changes by less than $\pm 3\%$.

A final practical experiment was undertaken in which the cylinder was prepared with the source 10mm deep (fig. 7.4c). It was placed on the sample stage and imaged without mirrors

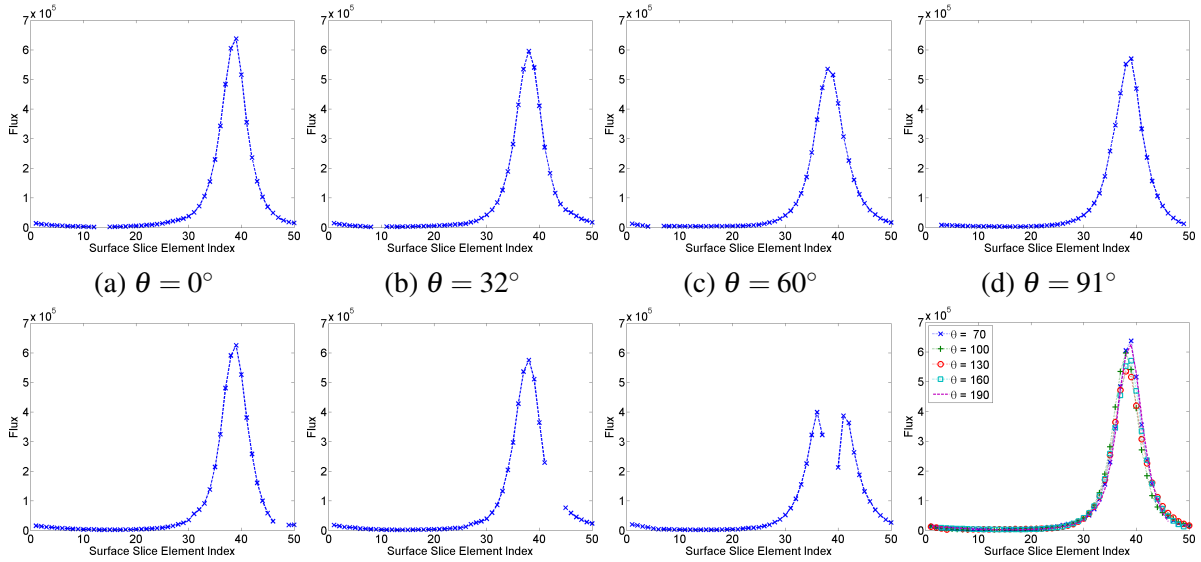

(a) $\theta = 0°$ (b) $\theta = 32°$ (c) $\theta = 60°$ (d) $\theta = 91°$

(e) $\theta = 115°$ (f) $\theta = 148°$ (g) $\theta = 178°$ (h) Overlay

Figure 7.16: Slices through reconstructions using real data with (a-g) individual data sets and (h) sets 1-5 shown overlaid together.

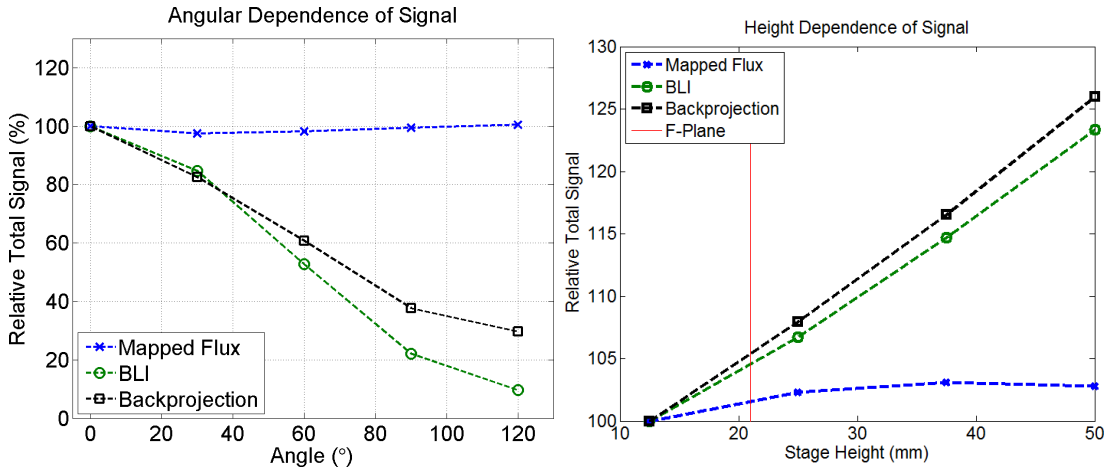

(a) (b)

Figure 7.17: Plots of total reconstructed flux as a function of (a) rotation and (b) height of imaged surfaces in the two practical experimental data sets obtained. These figures are reproduced from Guggenheim et al.[2].

at four different stage heights cumulatively separated by 10mm. It was in effect moved through the focal plane and was in different states of focus for the experiments. Figure 7.17b shows the total flux reconstructed in this experiment. The values obtained using the new method appear to be height (therefore focus) independent varying again by less than 3%. The naive and backprojection approaches both show large changes by up to 25%.

7.5 Discussion

Four experiments have been presented, two simulation and two practical, testing the new method for mapping surface fluxes from images[2].

The first simulation involved a 2D slice through the system and cylindrical surface. Across 50 repeats, of several experiments where the surface was at different vertical positions, the new method accurately reconstructed target fluxes with errors as low as 3% and small standard deviations. Conversely, the backprojection method was found to be quantitatively poor. As such, the presented method represents a substantial improvement.

Note that the maximum SNR in the best-case in this simulation was approximately 95:1 which is a realistically achievable SNR value; a 16-bit detector capturing 50% of its full-well capacity will read $2^{16}/2 = 32768$ counts, therefore will have a shot-noise-only SNR of 181:1. Furthermore, in section 3.3.3 it was shown that values of surface radiance up to 10^{10}p/s/cm^2/sr are reported in BLI literature whereas the maximum signal considered here was approximately an order of magnitude lower. The experiment has therefore shown that, for moderate SNR levels, the method is robust to noise.

The second simulation demonstrated a full imaging system simulation as has been shown previously[194, 190] and showed that the new flux mapping method reconstructed accurate fluxes robust to the orientation of the imaged 3D cylinder. This is comparable to the scenario in which a mouse is imaged from slightly different perspectives in longitudinal studies. Results showed that in cases of a visible surface, flux was mapped effectively and from multiple views simultaneously which has not been previously shown. The simultaneous use of multiple views

might be a useful approach to overcoming noise by considering different observations together.

This experiment highlighted the importance of being able to model a focal plane positioned arbitrarily because it involved multiple focal planes which, owing to the different path lengths, fell both in front of and behind the imaged surface. The new method is robust to the placement of the focal plane as shown in the experiment and discussed in the theory section which is an advancement over previous techniques.

The practical experiments were in support of the simulation results. Acquired images of the cylinder phantom had maximum values that decreased substantially and non-linearly as a function of the turning angle of the cylinder; a 10-fold drop in total intensity was seen across the range of perspectives for which a comparable part of the surface was visible. This is striking because this approach is presently used for quantification in biomedical studies[195, 10, 12] and the real flux was known to be in fact identical. Backprojection crucially was no more robust to changing perspectives. The introduced new method on-the-other-hand produced reconstructed fluxes that were stable across the scenarios. This is very important as it suggests that for an animal imaged in an arbitrary position in the system under study (or another similar system) the method could reliably reconstruct the flux either as an improved BLI/FLI end-point or as input for BLT/DOT.

In the second practical experiment the cylinder was moved through a range of heights and imaged, this tested the robustness of the presented method with respect to another set of changing surface properties. Again, the method worked effectively whereas other approaches did not.

Together the presented experiments suggest that the new technique is the best available; being robust to noise, and to changing system and surface perspectives. The stability of results suggest that with appropriate calibration the system might also be able to image arbitrary surfaces reliably in real world units.

A current limitation is that the optimal level, and to some extent the precise impact of the level, of discretisation used for surfaces and detectors is unknown. In this chapter the values were chosen empirically based on comparison of performance in forward models (qualitative

judgement of simulated image quality). These may have been higher than required costing unnecessary resources, or possibly lower than required to gain even better results. Further investigation into the impact of discretisation level is therefore a recommended topic for future work. An initial observation is that finer discretisation appears to be required when the surface is near to the focal plane, this is most likely due to the mechanism for determining the collection or non-collection of rays by the lens[2]. Approaches involving variable discretisation levels, based on the distance between the surface and focal plane, could therefore be worth exploring. In later chapters (chapters 8 and 9) examples are shown where discretisation levels are increased until the change in the matrix is stabilised to a certain degree; further characterisation of this approach would also be useful.

The major weakness of the new method is the need to select the regularisation parameter, λ, manually. Changing this value has sometimes been seen to affect results significantly, on other occasions it appears to have little affect but this is problem-dependent and selection is an ad-hoc and empirical process at present which is not ideal. There are automated approaches available for choosing such parameters automatically such as the L-Curve[196], maximum likelihood (ML) and generalized cross-validation (GCV)[197], and U-Curve[198] methods which could be investigated for the automatic, ideally optimal determination of the parameter. Equally, novel or existing optimisation methods with fewer free variables might eliminate the issue. This highlights the need to explore further approaches to fully automate the mapping accurately and in general.

7.6 Conclusion

In the imaging system presented in this thesis, there is a requirement to relate, via some *free-space* model, the output light flux leaving a surface (the small animal skin) and the values subsequently measured by the optical system. This is an important part of the full system model required to facilitate the reconstruction of 3D volumetric data from 2D images in BLT and DOT. In principle there are two ways to do this. One is to perform inversion of the free-

space model, thereby obtaining boundary data, and then perform BLT/DOT reconstructions as in the paradigms introduced in sections 2.5 and 2.6, this is the approach used in this work. The other is to combine the internal and free-space models and invert the composite model; the approach used by Ripoll et al.[187] for FMT. In both cases the model must be quantitatively accurate and in the first case an inversion approach is required explicitly to relate images back to surface fluxes.

This chapter has focused on an enhanced method for building and inverting the required relationship between surface and detector which is particularly tailored towards the current system, and which represents an extension of the previously available methods in several ways. Arbitrary placement of the focal plane is handled which is important for the current multi-view system as a good balance in focusing between the mirrors and direct views is likely to place the focal planes in front and ahead of the surface. Multiple views have been dealt with simultaneously for the first time. Quantitative stability for equivalent fluxes has also been achieved with respect to changing many imaging and surface parameters, this is important to the current system because the ultimate goal is performance of accurate 3D *in vivo* imaging with meaningful real world units delivered.

It has been demonstrated in this chapter with explicit experiments that relative surface measurements generally vary significantly even for slightly different geometries when 2D image data is used directly for quantification. These are important demonstrations given that this approach is used in biomedicine currently. The new mapping method goes some way to overcome this problem by reconstructing flux distributions accurately using the images, system geometry, and surface capture data; obtained by methods described elsewhere in this thesis.

Specifically in the tests performed, the method produced a stable result (to within 3%) for total light output at the surface independent of the surface position in the system and the orientation. This was supported by simulation studies which showed the same trends and additionally showed good robustness to noise. This improves significantly in quantitative terms on existing methods such as the discussed backprojection approach[193].

The method therefore represents a novel, better way (than direct quantification from images)

of interrogating an animal surface. As part of the DOT/BLT multi-modal system work-flow it is suggested that the work of this chapter also makes significant strides towards a general, effective and quantitative 3D imaging solution for small animal studies.

This being said, the method should be tested with real animal subjects so that the performance on more realistic samples can be evaluated. For diffused light within nude or shaved mice, the cylinder phantom used in the experiments is expected to be an effective analogue but this needs to be validated, and an animal study is needed to test performance further in general.

An interesting potential extension of the presented work is the modelling of the non-contact excitation light used for DOT (or equally that used in FMT systems). This is a closely related problem since the light paths considered from the DOT source are essentially identical to those detection-side. The problem is reversed because it concerns projection rather than collection of light but the method might be adaptable. This requires further work and experiments but could improve the accuracy and efficiency of DOT which in the current work (chapter 8) relies on simplifying assumptions of the DOT sources which may be inaccurate.

A more elaborate but potentially exciting opportunity for future work involves the use of the full system simulation capability to design new imaging system layouts. For example, mirror designs could be explored that would maximise the utility of the CCD detection area and minimise the cost, in a similar spirit to the use of the conical mirror by Li et al.[146]. This could provide a powerful tool for system design optimisation.

The methods of this chapter are discussed further in context of internal image reconstruction in chapters 8 and 9 in which they are used in experiments to facilitate BLT and DOT image reconstructions.

CHAPTER 8

DIFFUSE OPTICAL TOMOGRAPHY

8.1 Introduction

The basic mechanism underlying diffuse optical tomography (DOT) has been outlined in section 2.6 in which the standard method, involving reconstruction of μ_a and/or μ_s' throughout a volume, and the spectrally-constrained method, involving reconstruction of absorber/chromophore concentrations and scatter power and amplitude throughout a volume, were introduced.

To re-iterate, the goal of diffuse optical tomography in the presented system is to reconstruct absorbing features and properties at the level of molecular agents and scattering properties throughout a mouse given knowledge of the sources and detectors and measurements of the light transmittance between them. The results will be used for their direct physiological value and also as prior information to inform the light propagation modelling used when reconstructing bioluminescence tomography (BLT) images.

The aim is to do DOT in non-contact mode. As such, the many, typically surface-bound, DOT 'detectors' will correspond to surface flux values at many positions sampled from visible locations on the camera-imaged subject whilst the sources correspond to light introduced at the object boundary supplied by the DOT light source.

This chapter evaluates methods used to acquire calibrated multi-spectral data and perform DOT image reconstruction with physical phantoms (appendix A.1.1). Results are presented for cases where bulk properties are reconstructed (spectroscopy), and cases where spatially

heterogeneous distributions are reconstructed (tomography).

As has been discussed in chapter 4 there has been relatively little work in DOT for small animal imaging, much less than fluorescence or bioluminescence, and there exists even less using multi-spectral or CW data.

The chapter first discusses methodology, which is standardised for the results presented in this chapter following a great many rounds of refinement in the course of the project. The protocol for DOT data acquisition is discussed including the chosen illumination scheme. This is followed by a demonstration of the data acquisition, mapping and processing used to produce calibrated boundary data. The specific image reconstruction method is then laid out for clarity, building on the introduction to DOT reconstruction in section 2.6. Data sampling and parameter selection for image reconstruction are then discussed. Three experiments are then described involving spectroscopy of the mouse-shaped phantom, and tomography in the cylindrical and cuboid phantoms. Results are presented for these experiments and all are discussed.

The reasons for using simple phantoms in these experiments is discussed in appendix A.1.1. Summarising briefly, there are two key reasons. Firstly that the phantoms are very well understood and characterised as compared to real tissues and animals thus results can be quantitatively evaluated. Secondly, by starting as simple as possible it may be possible to identify, as complexity is increased, the breaking point for the system and methods, this can provide valuable insight into the system properties and lead improved development. In the presented experiments the homogeneous phantom can be seen as the simplest case in terms of optical properties, and the heterogeneous phantoms as slightly more complex and slightly more realistic. The heterogeneous inclusions used can be thought of as simple approximations to the organ structures expected in mice.

Together the presented experiments represent a preliminary study set illustrating the type of experiment that can be carried out with the system, providing some early insights into the capabilities of the DOT system. The next chapter discusses BLT imaging and experiments including one that takes a DOT result (for the cylinder phantom experiment) presented in this chapter and uses it as prior information, thus illustrating the desired final work-flow for the

system.

8.2 Methods

8.2.1 Transillumination Imaging

The image acquisition performed to provide data for DOT is referred to as transillumination imaging because the imaged subject is illuminated by the DOT source and the transmitted light is imaged. The subject is placed on the imaging platform directly above the source projector and many images are acquired sequentially for different projected source patterns. This is repeated for each required wavelength.

In contrast to direct laser-based DOT systems (for example that of Yan et al.[166]), which illuminate single points on an object and require mechanical changes to move the source, the projector can produce arbitrary patterns and therefore illumination at arbitrary locations[1]. DOT has been performed by others using a similar light source[167] using wide-field illumination patterns. This approach has the benefit that a large amount of (spatially distributed) light can be used and a large amount of the volume can be interrogated simultaneously, potentially reducing experimental time. Patterns can also be optimised on-the-fly for uniform volume coverage[199]. In the current system, a simple approach is taken in which images containing Gaussians drawn from a regular grid are projected sequentially. The effect is to mimic a raster scan of a single point-like source such as from a laser or fiber-bundle illuminating the phantom from underneath. This approach was chosen because the modelling of point-like sources is well-established in DOT (section 2.6) and as such this is the most appropriate starting point. The use of wide-field source distributions remains a topic for future work. This is expected to be worthy of investigation because of the potential for making imaging more efficient and possibly even more effective[168].

Under the current scheme, illustrated in fig. 8.1, each pattern contains a 2D Gaussian, drawn

[1] within the field of view

from a 6×6 grid, with a maximum intensity equal to the maximum projectable intensity and a standard deviation of approximately $2mm$ on the sample (40 pixels in the original projection image).

The maximum projectable intensity is used in all cases to maximise image SNR. The intensity is non-uniform due to variations in the projector, and these are accounted for with calibrations (section 8.2.2). Another approach could be to modify the input power of the pattern to adjust for the irregular spatial output, though this would necessarily reduce SNR.

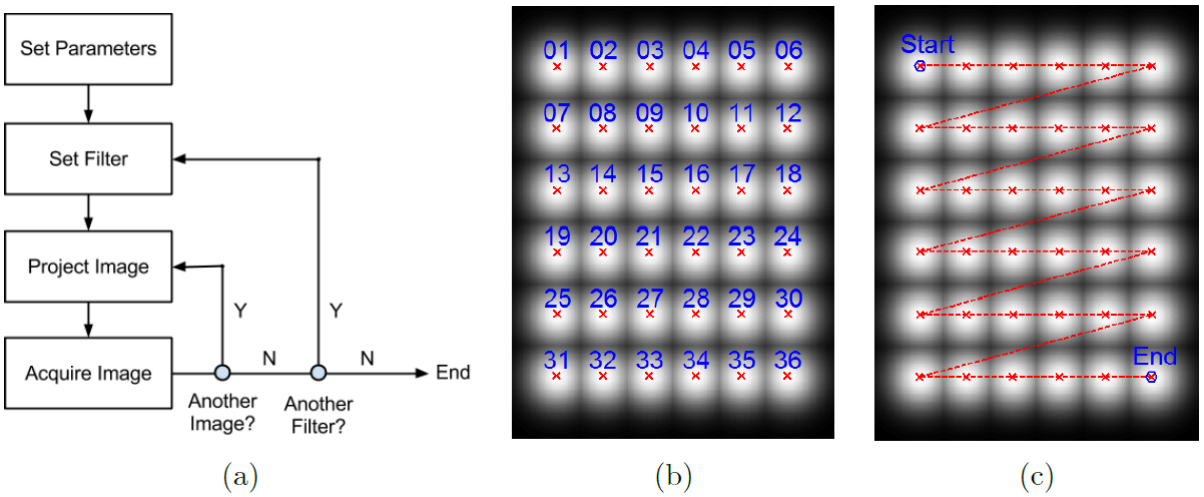

Figure 8.1: (a) DOT imaging protocol; (b) grid of projected sources with sequence order numbers; (c) indication of the effective raster scan performed. This figure is reproduced from Guggenheim et al.[1].

For each wavelength used when imaging patterns, a *dark* image is also acquired where the projected source pattern is uniformly zero; the darkest possible level at all pixels. This is required because the projector always has some non-zero light output due to uncontrolled light leakage which must be subtracted.

Imaging parameters are chosen to maximise SNR (see appendix A.5 and section 5.3.3); hardware binning is 4×4, CCD mode is normal (not EMCCD), read mode is 1. The exposure time is generally selected by first establishing manually the brightest source at a single wavelength via a short-exposure image of each transmittance. Then a single image is acquired at each wavelength for the brightest source and an exposure time is calculated that would be expected to deliver a high maximum count (typically 40000) without saturating the camera by

scaling the maximum intensity (in units of counts per second) measured at each wavelength. The subsequently calculated filter-resolved exposure times are used for all sources. With the assumption that the sources are similar in strength this approach is effective, however it would be better to image each transmittance (i.e. of every source through every filter) and calculate a time precisely for each case. In the current methodology this would be a very lengthy process owing to the large number of patterns and wavelengths used, thus the current method is a compromise between optimal imaging conditions and experimental effort and time. Usually, there is a maximum exposure time also enforced to limit experimental time. In all cases presented in this thesis, the maximum exposure time allowed for trans-illumination imaging was 30s.

Acquired data sets comprise 37 images for each wavelength (transillumination images with 36 source patterns and 1 *dark* image). All images are first processed and converted into units of $e^{-}/s/pixel$ as detailed in section 5.4.1, the dark image is then subtracted for each wavelength.

DOT imaging is accompanied by a surface capture to acquire a representation and mesh of the imaged surface (chapter 6). The acquired multi-source, multi-spectral image set is mapped onto the surface according to the mapping method described in chapter 7. The mapping is performed for every image but the same free-space transfer matrix is used in all cases. As such the mapping scales well as a function of the number of images assuming there is no significant movement of the object or change in the system.

Note that the free-space mesh and mapping defines the locations of the FEM detectors used in DOT. The FEM sources are placed in a more simplistic fashion by taking the known projection locations of the sources (centroids from the spatial calibration method described in appendix A.6.3) and moving these to the nearest point on the mesh[200]. In most cases this does not involve moving them, as the mesh is created such that it sits flush with the stage (chapter 6), though in one case presented below, involving imaging a cylinder for which a pre-defined mesh is used, this does involve moving the sources a small distance in some cases.

8.2.2 Data Calibration

Rather than using explicitly measured values for the system response and the source output, mapped DOT data is calibrated using a reference data set in which the imaged reference phantom is well-known. It is assumed that a single (linear) scaling factor per-wavelength and per-source can adequately model the mismatch between mapped data and simulated data for the known source positions (appendix A.6.3), detector positions and reference phantom properties. An advantage of this approach is that systematic errors can be cancelled out[1]. Another advantage is that any parameter variations with time (such as reduction in the transmittance of the filters which is known to occur) will be accounted for because the reference data is acquired directly before or after the test data.

An example is now presented to illustrate the calibration process in which a homogeneous block phantom (appendix A.1.1), acting as the reference phantom, was imaged at 650, 720, 750, 810, and 850nm. Figure 8.2 shows three sample images from the acquired transillumination dataset.

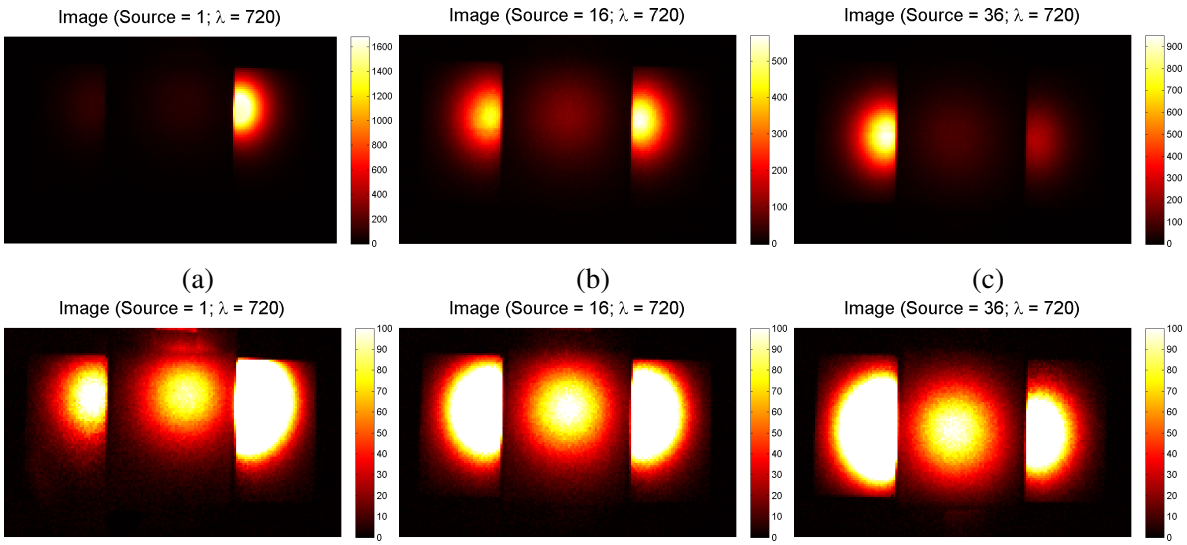

Figure 8.2: Three sample transillumination images for the homogeneous block phantom acting as a reference phantom. First images (a-c) show auto-scaled data whilst (d-f) are the same images on a common re-scaled colormap so that the direct-view data can be seen easily. The images are at 720nm and in units of electrons per second.

[1]This can also be seen as a disadvantage given that this may make errors harder to find, analyse and understand. This should be kept in mind when planning further studies.

The images were mapped onto a mesh generated by the automated surface capture and meshing method (chapter 6) using the free-space mapping method (chapter 7). Examples of the mapped data are shown in fig. 8.3. Four mappings were performed; three in which data was mapped independently for each view (mirror 1, mirror 2, direct) and one where all data was combined.

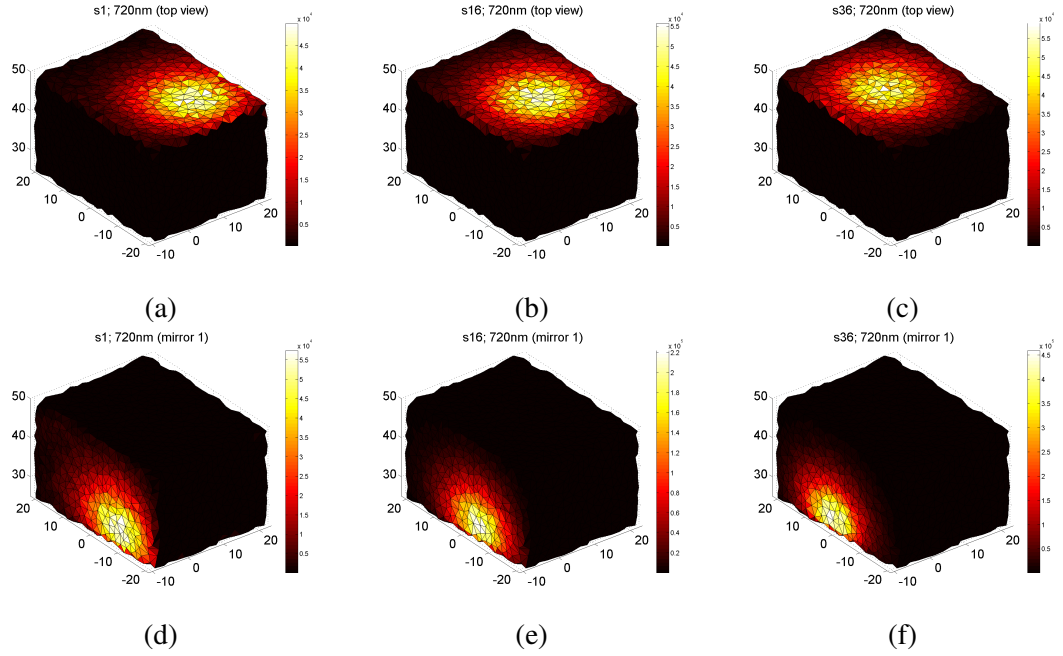

Figure 8.3: Mapped data following application of the inverse free-space transfer method to the reference data. Independently mapped data is shown for the data presented in fig. 8.2 for the top-view data and for a single mirror view.

Using independently established spectral scattering and absorption properties of the homogeneous block (appendix A.7.4), simulated measurements were calculated using NIRFAST for the FEM source positions associated with each projected source and the FEM detector positions corresponding to the mapped data positions. Thus for every practical measurement (per-wavelength, per-source) a corresponding observation was obtained in simulation.

Figure 8.4 shows a direct comparison between the simulated and measured flux for one source at 720nm taken from the reference data set. It can be seen that qualitatively the appearance of the source is very similar, whilst quantitatively there is a large difference in the actual values evident in the colour bars. The large difference is due to the fact that up until this point the values are not corrected for the system response or source power. There is also

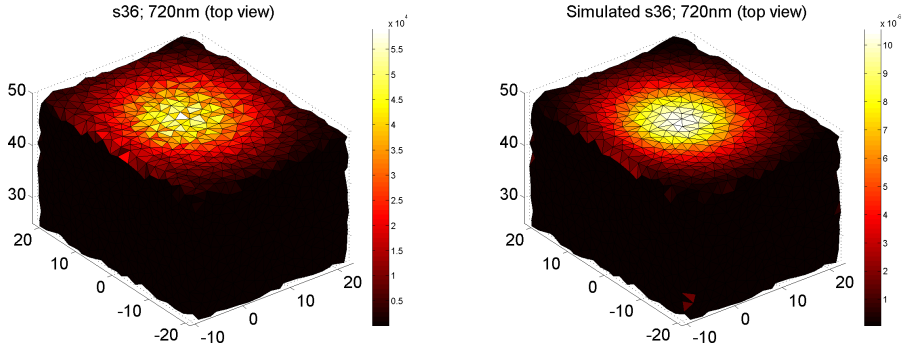

(a) Practical Data (b) Simulated Data

Figure 8.4: Mapped data for the reference phantom for a single source at a single wavelength and corresponding simulated data from which a scaling factor is derived.

some smaller quantitative variation which may be due to a combination of noise and free-space mapping errors due to mesh coarseness, under-discretisation, or poor regularisation parameter choice.[1]

Scaling factors were then found to match the practical data to the modelled data. Note that not all data points were used, instead only points mapped from image data within a selected region on the CCD were used (see section 8.2.4). This is to reduce the problem size and fitting time, to select only a signal area (because fitting noise and zeros is no use), and to avoid any errors introduced by meshing errors on the sharp corners of the cuboid mesh.

The fitting was performed using the *fminunc* function in MATLAB with an L2-norm objective function. The data for all sources for a single wavelength is plotted before and after subsequent scaling along with the directly compared reference data in fig. 8.5.

It can be seen that following the scaling, the match appears to be good at least for relatively high values. There is an apparent *noise-floor* below which the values do not match well, this is assumed to be due to the smaller dynamic range available in the practical measurements effectively raising the lower limit of detection.

Finally, the full set of per-wavelength, per-source factors derived from the presented calibration experiment are shown in fig. 8.6. It can be seen that the plots appear similar in shape at each wavelength, indicating a degree of independence of spatial and spectral output variations.

[1]The results presented in chapter 7 suggest this could be done more accurately.

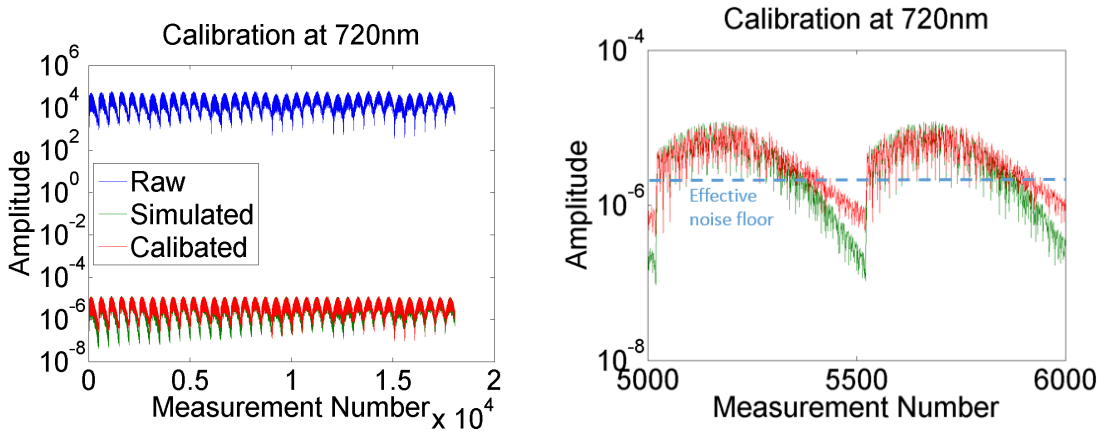

(a) (b)

Figure 8.5: Mapped measured data for a single wavelength compared to corresponding simulated data before and after fitted per-source, per-wavelength calibration factors were applied.

This suggests that calibration factors could be measured separately as a function of wavelength and source number, reducing experimental time with little if any loss in efficacy. This idea may be worth exploring further by acquiring many data sets and explicitly comparing these two approaches.

8.2.3 Image Reconstruction

Recall that DOT image reconstruction is carried out using NIRFAST in which the formulation of the light-transport problem within the volume is based on the diffusion approximation (section 2.3). Recall also that the absorption and scattering coefficients can be respectively expressed in terms of chromophore abundances and extinction coefficients (section 2.1.1), and an amplitude-power law representation approximating Mie theory (section 2.1.2), i.e.

$$\mu_a(\lambda) = \sum_{i=1}^{n} \varepsilon_i(\lambda) C_i, \tag{8.1}$$

and

$$\mu_s' = a\lambda^{-b}. \tag{8.2}$$

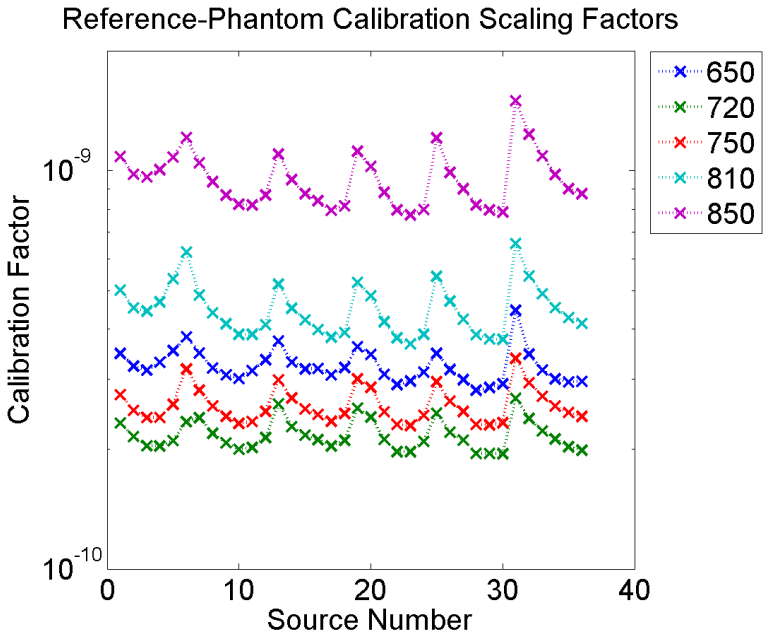

Figure 8.6: Calibration factors plotted per-source for each wavelength at which they were measured for the presented homogeneous block reference-phantom calibration experiment.

Under the spectral regime, the Jacobian matrix J is calculated, which relates the entities C_i (absorber concentration for chromophore i), a (scatter amplitude), and b (scatter power) to DOT boundary measurements ϕ_{dot} given an initial guess of $\mu = [C, a, b]$; a vector representing node-wise chromophore concentrations and scattering parameters throughout the FEM mesh. The reconstruction is undertaken by use of an iterative Tikhonov-regularised Levenberg-Marquardt type update term[25].

$$J^T (JJ^T + m\rho I)^{-1} \delta\phi_{dot} = \delta\mu \tag{8.3}$$

Where ρ is a regularisation parameter and m is the maximum of the diagonal of the matrix JJ^T. Reconstruction is terminated by stopping the algorithm after a certain number of iterations or awaiting convergence (judged as the point at which data re-projection error change is less than 2%).

By utilising spectral DOT, the concentration of absorbers and the scattering properties within the medium are computed directly rather than solving for μ at multiple wavelengths and then curve fitting. This serves to constrain the solution space given the spectral characteristics of the finite set of known absorbers assumed to be present.

8.2.4 Data Selection

The number of data points acquired by the CCD is very large compared to classical DOT experiments that used a few optical fibre based detectors. This is beneficial in that many useful observations are made but detrimental in that it increases computational complexity in flux and image reconstructions; both time and memory requirements can become practically prohibitive. Steps can be taken to remove unnecessary or unwanted data points to reduce complexity, ideally without losing useful information (though this is not generally the case[201]).

Uniform Down-Sampling

Random or uniform down-sampling is the simplest way to reduce the problem complexity. It has the disadvantage of being indiscriminate with respect to measurement quality. However, there is generally a lack of independence between adjacent measurements in DOT because the sensitive area within the volume for a single FEM detector is very similar to that of neighbouring detectors (due to the diffuse nature of the light). For this reason there is little new information providing by detectors that are very close together.

It is therefore assumed that a certain amount of data can be discarded with little loss in information or reconstruction accuracy. However, quantitative validation of this assumption is challenging and has not been carried out. It is worth noting that even repeated observations can be useful to overcome noise and data compression methods that preserve the information whilst reducing the complexity are highly attractive such as that presented by Styles et al for BLT[8].

Region-of-Interest Propagation

It was discussed in section 8.2.2 that a region-of-interest (ROI) was selected in images to select data. This was done by selecting a box in images around an area of signal and then propagating the region to the mapped flux space. The propagation is done by including in the final selection any data point that has a corresponding image point in the image ROI (one that is sensitive to

it with a value greater than zero[1] in the transfer matrix). This data selection method has the advantage that it is intuitive and easily controlled to define signal or interest area. It has the disadvantage that it involves manual intervention which slows the imaging process and moves away from the aim of full automation.

Thresholding

Figure 8.5 showed that accurate measured data is not obtainable below a noise floor. The bad measurements may have a negative impact on reconstructions as well as adding to complexity, as such it is desirable to remove these points in particular.

Following the mapping procedure it is not possible to estimate the noise-floor relative to the measurements owing to the surface flux having been estimated using an optimisation algorithm (chapter 7). However, the noise can be estimated according to the system model on the image side (appendix A.5) and it may be possible to use the noise properties of pixels that are *sensitive* (linked via ROI propagation) to each surface element (with a mapped value) to somehow estimate the reliability of the resulting value. Such an approach is not currently implemented but is a suggestion for future investigation.

8.2.5 Selection of Reconstruction Parameters

There are a number of free variables and operations associated with DOT reconstruction.

Regularisation Parameter

The regularisation parameter ρ has an impact on the reconstruction process; it smooths the solution search-space and makes the system matrix better conditioned. This means that the data is in practice fitted less exactly and reconstructed images appear smoother. The choice of lambda is empirically determined which can be seen as a weakness as it moves away from full automation of the reconstruction, and introduces some subjectivity.

[1] or some small value.

NIRFAST, by default, initialises ρ to 100 and reduces it at each iteration; there is a high level of regularisation to guide the first iterations, but as the reconstruction proceeds finer changes are made to the reconstructed image. This approach has developed following several years of empirical refinement and is not modified here. A more objective assignment of the regularisation parameter is a potential topic for future work.

Reconstruction Pixel Basis

Reconstruction is performed using a pixel-basis; a spatial grid placed over the original mesh that then defines a different (coarser resolution) mesh on which the reconstruction is actually performed, with recovered values then being interpolated back onto the original mesh[25]. There is a trade-off introduced between final image resolution, which is limited to the resolution of the basis in the best case, and computational complexity which is reduced by using a coarser mesh. Additionally, the coarser the mesh the fewer unknowns in the reconstruction and hence the less ill-conditioned the inverse problem is and this can lead to better performance.

Number of Iterations

Choice of number of iterations is typically a heuristic process that currently has no well-defined, accepted method available for its choosing. This is a clear limitation, particularly as it has been observed throughout this project that the reconstructed image often looks best (most like the target in well-understood situations) prior to convergence meaning that it may be worth pursuing a better stopping criterion. On the other hand as the system is applied to more practical problems it may be possible to identify by rigorous testing a value that is practically most useful as a fact of experience.

8.3 Phantom Experiments

8.3.1 XPM-2 Phantom Spectroscopy Experiment

The XPM-2 phantom (appendix A.1.1) was imaged at 640, 720, 770, 810, and 850nm. Sample images are shown in figure 8.7. In all cases the CCD mode was normal (not EMCCD), the read mode was 1, the binning was 4×4, and the exposure time was 1s. The exposure time was chosen to make imaging fast; it could have been made longer to improve SNR (the camera would not have saturated imminently) but at the expense of added experimental time.

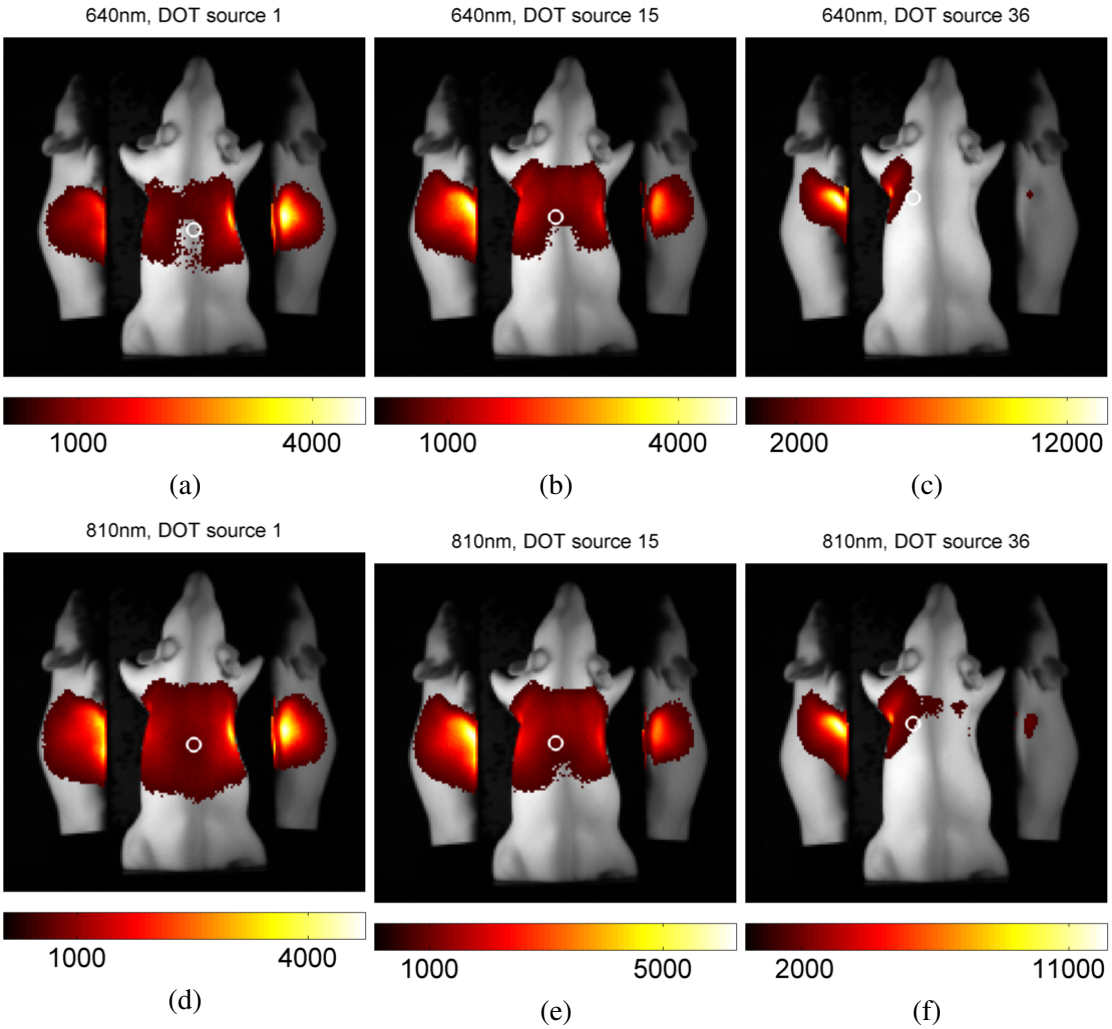

Figure 8.7: Background-subtracted transillumination images for the XPM-2 phantom under illumination by DOT sources (section 8.2.1). Data is shown for three sources and for two wavelengths. The units are electrons per second on the CCD. Source positions are indicated by white circles.

The image processing, mapping and calibration were performed as described above. The homogeneous block phantom was used as the reference phantom (as in the example of section 8.2.2).

In reconstruction, the FEM model was assigned homogeneous spectral optical properties based on the best estimate established in appendix A.7.4,

$$\begin{cases} C := 1 \\ a := 1.28 \\ b := 1.79 \end{cases} \tag{8.4}$$

in which the single concentration value, C pertains to the invented absorber *Composite Dye 1* (appendix A.7.4) which is the only absorber considered in this experiment.

These were then intentionally corrupted in experiments by modifying the initial estimate of the absorber concentration; a perfect reconstruction would be expected to recover the target values of eq. (8.4). The values used for the initial guess of C were 0.5, 0.8, 1, 1.2 and 1.5. The setting of 1 can be considered a control. Thus there were a total of five reconstructions performed.

They were performed with top-view data only. The main reason for this choice is that the data from the top-view necessarily represents light that has transmitted through the whole thickness of the phantom and therefore has sampled a larger part of the volume. Other reasons are that it reduces the problem complexity, that it is the simplest possible starting point to be built upon, and that there is some light leakage at the base of the phantom (see for example the right-view in fig. 8.7d) due to the fact that it is suspended slightly above the stage owing to its shape and rigid nature. Top-view-only (cropped) images are shown in figure 8.8. Corresponding free-space mapped data is shown in figure 8.9.

Reconstructions were performed with a hard-prior and a single region[25] (forcing homogeneity as a constraint on the FEM). The concentration and scattering properties were solved for. A pixel basis with approximately 12000 nodes was used. No down-sampling was per-

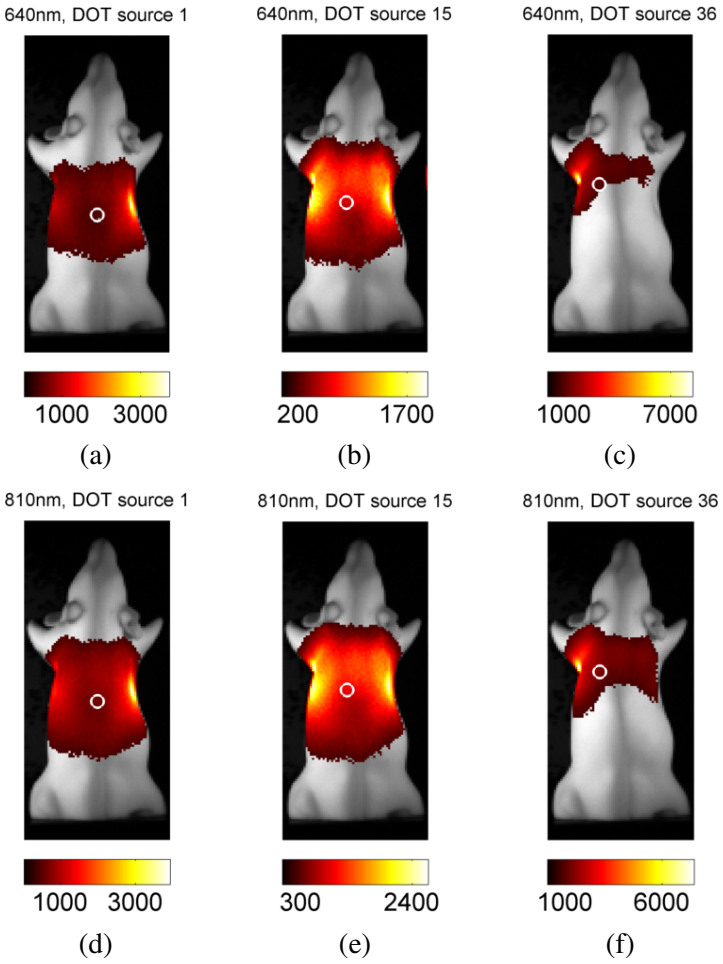

Figure 8.8: Top-view only version of figure 8.7 with cropped images to visualise the top-view data with higher contrast.

formed on the data. Regularisation was set to automatic (section 8.2.5) and an upper limit on the number of iterations was set at 30 though in practice the reconstruction always terminated due to having converged before 30 iterations.

Results of the five reconstructions are shown in figures 8.10, 8.11, 8.12, 8.13, and 8.14. The results show that in both cases where the absorber concentration was initialised to a value lower than the target value, the reconstructed absorber concentration converged to the expected value. This was also the case when the absorption was initialised at the target value. However, it can also be seen that in both cases where the concentration was over-estimated, the reconstruction converged with the value still over-estimated, though there is also a downward turn in both time traces after the second iteration suggesting this may have improved if more iterations were carried out. This could be tested by modifying the convergence criteria. The final result for

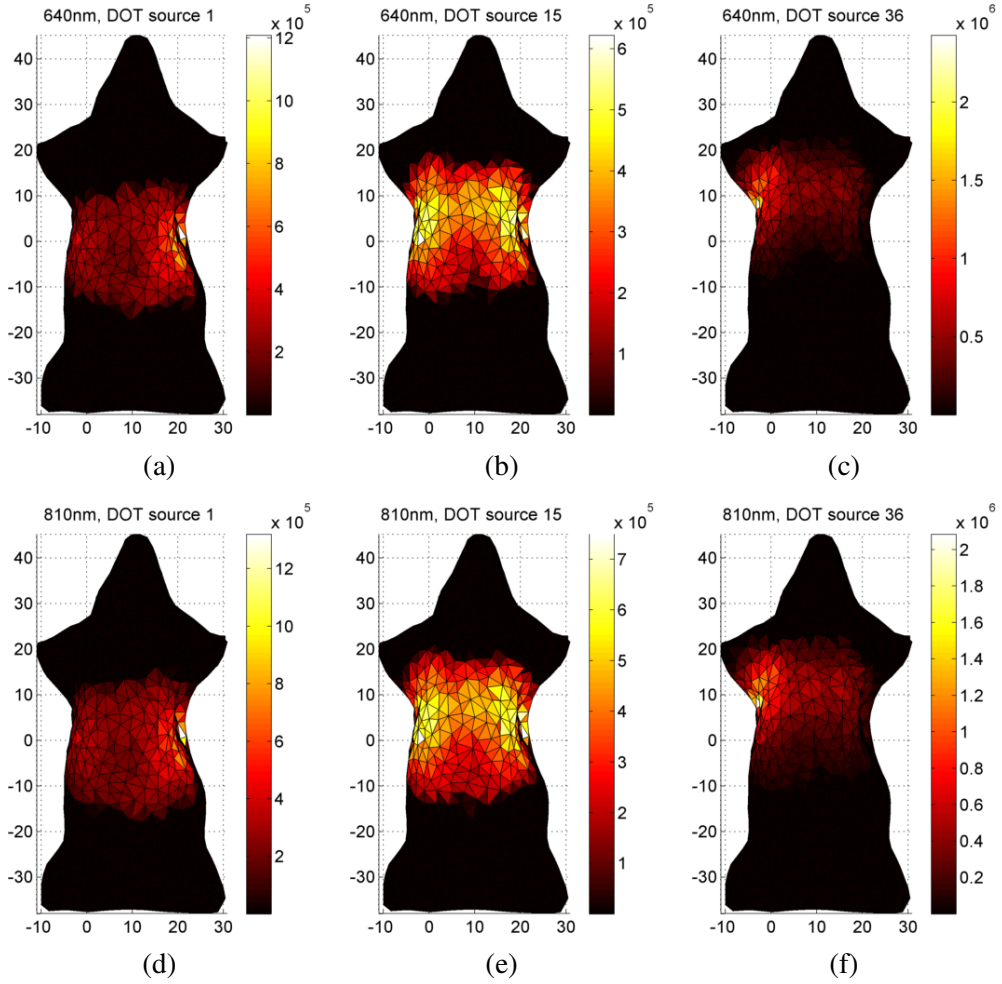

Figure 8.9: Free-space mapped data for the images shown in fig. 8.8 onto a mesh produced using the automated meshing method from surface capture (chapter 6). Values are pre-calibration.

the absorption was therefore correct for three of the reconstructions and over-estimated in the remaining two.

In terms of scattering, it can be seen that the scatter power converges on a value of 2 in all cases. This is surprising because it is somewhat higher than the expected value. The scatter amplitude shows a tendency to be initially increased in the course of the reconstruction but then returns to approximately the expected value in all cases. The final result is a similar looking scattering spectrum in all cases with a steeper incline and generally higher scattering values than expected, particularly at lower wavelengths. The final scatter amplitude is slightly lower in the latter two reconstructions compared to the first three, possibly this is compensating for the over-estimated absorption.

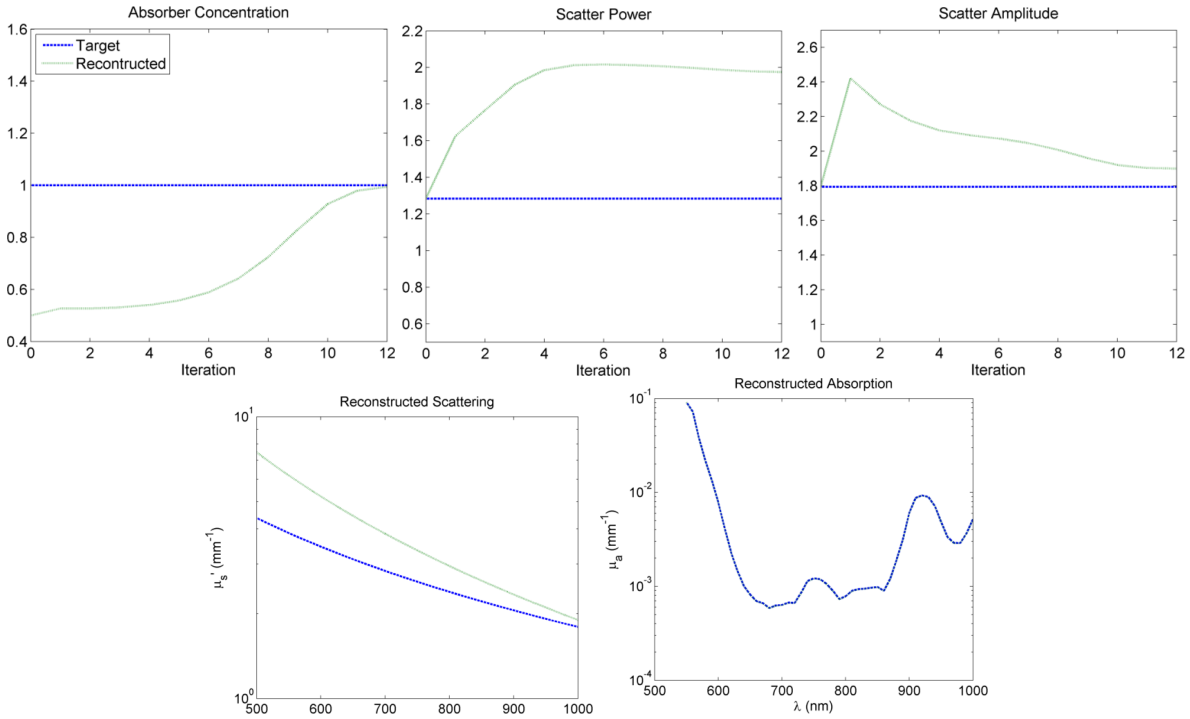

Figure 8.10: Results of spectroscopic reconstruction of XPM2 properties with initial concentration estimate at 50% of the expected value. The upper three subfigures show reconstructed parameter values as a function of iteration. The lower two subfigures show final scattering and absorption values as a function of wavelength (calculated from the above values at the final iteration). In all cases the calculated and expected (ground truth established in appendix A.7.4) are shown for comparison.

8.3.2 Cylinder Phantom Tomography Experiment

A DOT experiment was performed with a cylindrical phantom (appendix A.1.1). The phantom was made optically heterogeneous by the insertion of a double-absorbing rod into the shallow inclusion tunnel resulting in the configuration shown in fig. 8.15[1]. The phantom was positioned in the system in the orientation shown and images were acquired at 650, 700, 750, 810 and 850nm. Exposure time and other parameters were selected as described in section 8.2.1. Experimental data was processed as described above with the exception that rather than creating a mesh on-the-fly from surface capture data, a pre-made cylinder mesh was registered to the surface capture points as done previously in chapter 7, and discussed in section 7.3.3.

In this experiment, the reference data used for calibration was also acquired with the cylin-

[1]Note that this imaging run formed part of a multi-modal experiment which is described further in section 9.3.4, for this a light source was also placed in the other cylinder tunnel,

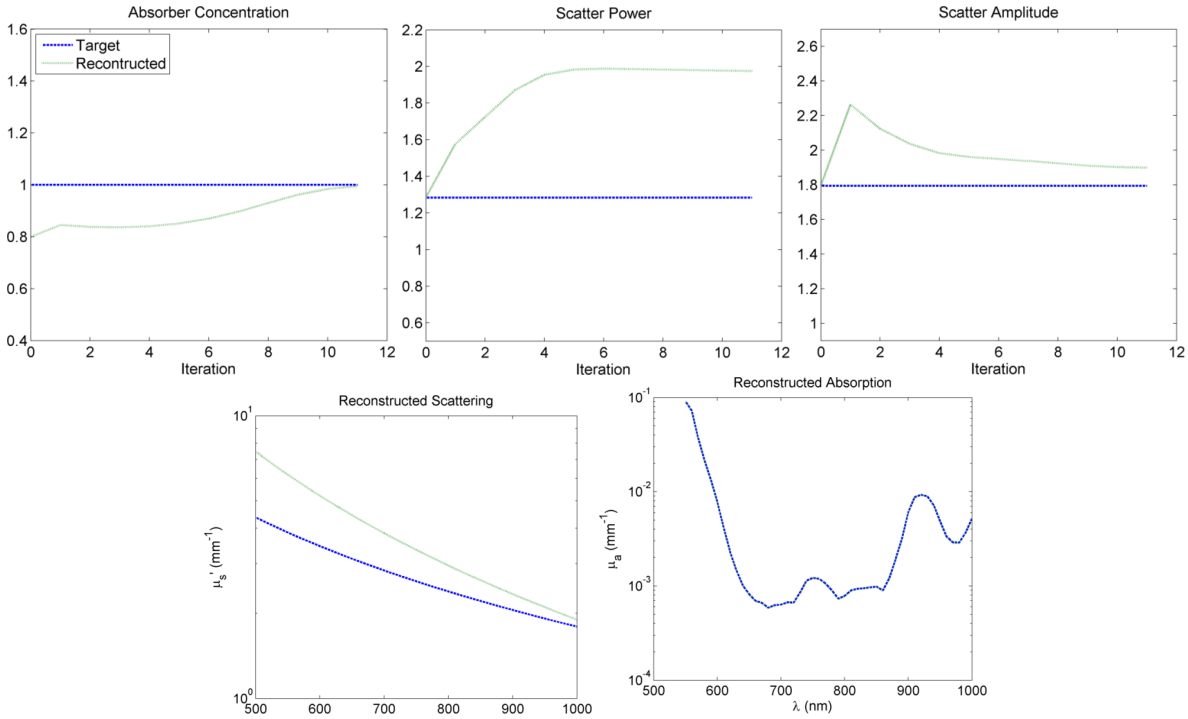

Figure 8.11: Results of spectroscopic reconstruction of XPM2 properties with initial concentration estimate at 80% of the expected value.

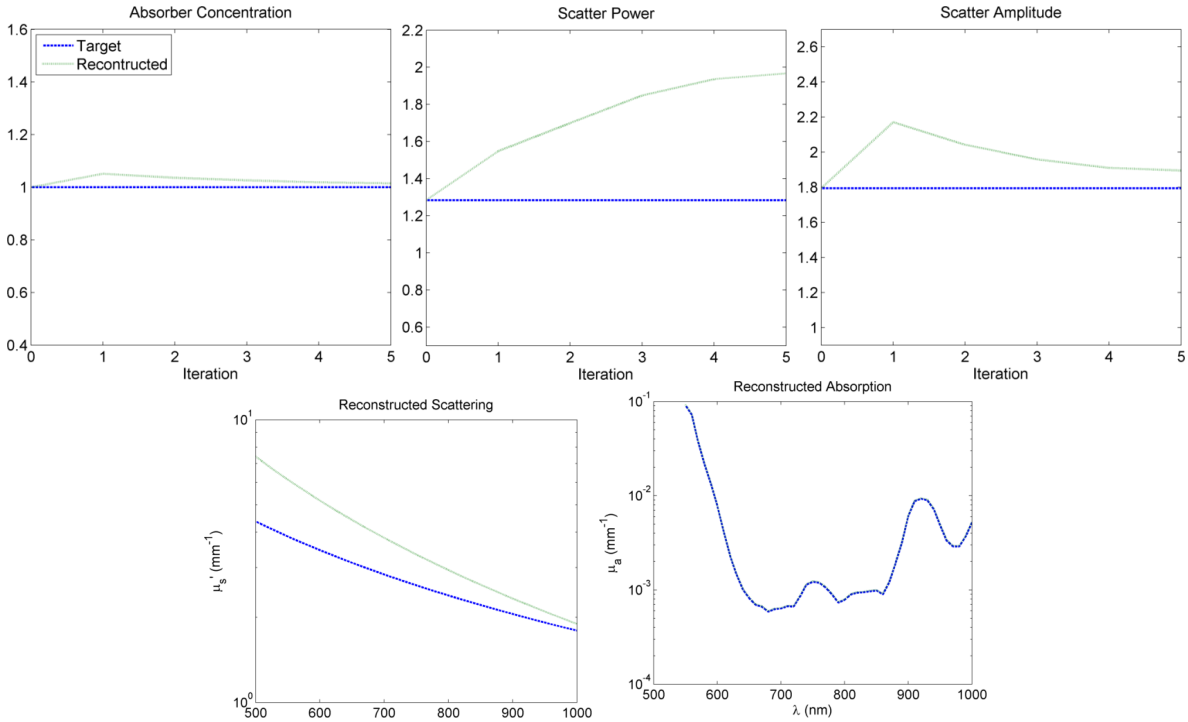

Figure 8.12: Results of spectroscopic reconstruction of XPM2 properties with initial concentration estimate equal to the expected value.

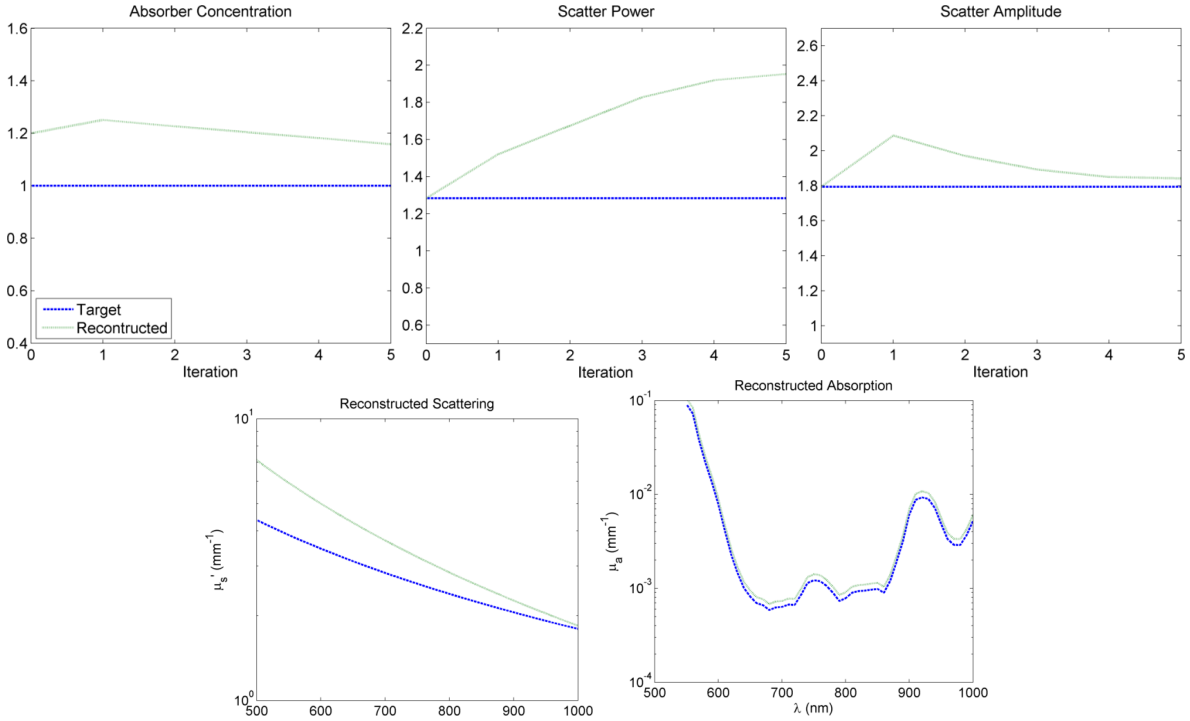

Figure 8.13: Results of spectroscopic reconstruction of XPM2 properties with initial concentration estimate at 120% of the expected value.

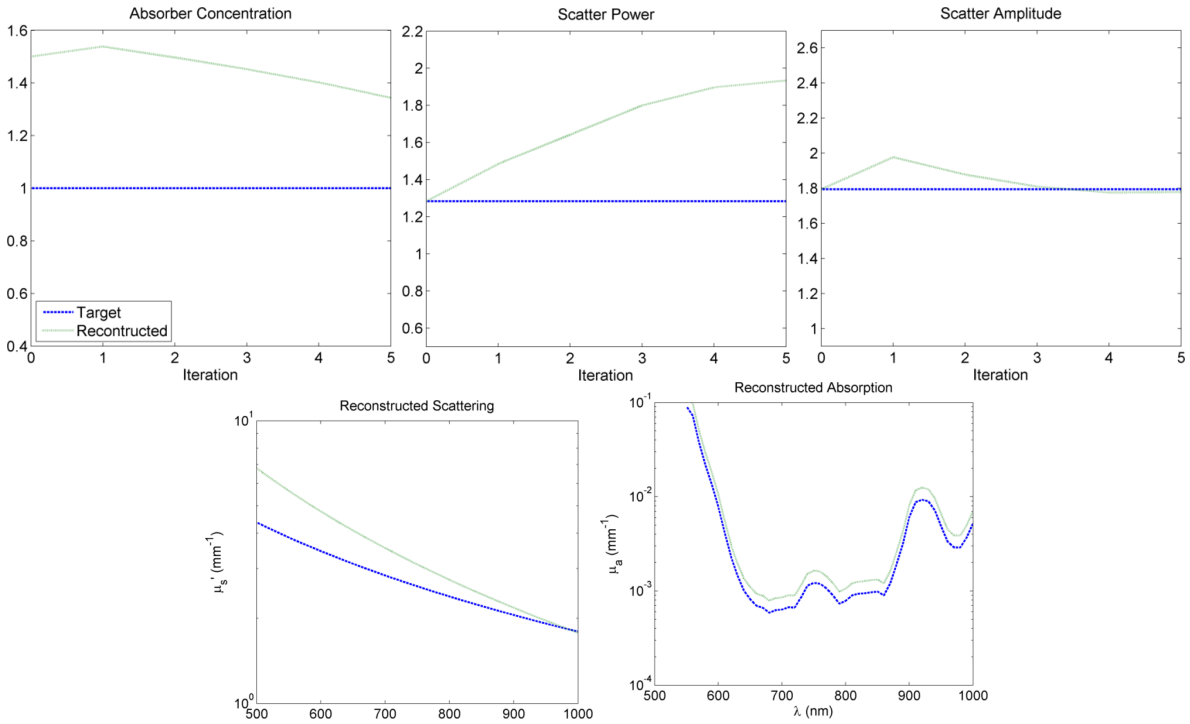

Figure 8.14: Results of spectroscopic reconstruction of XPM2 properties with initial concentration estimate at 150% of the expected value.

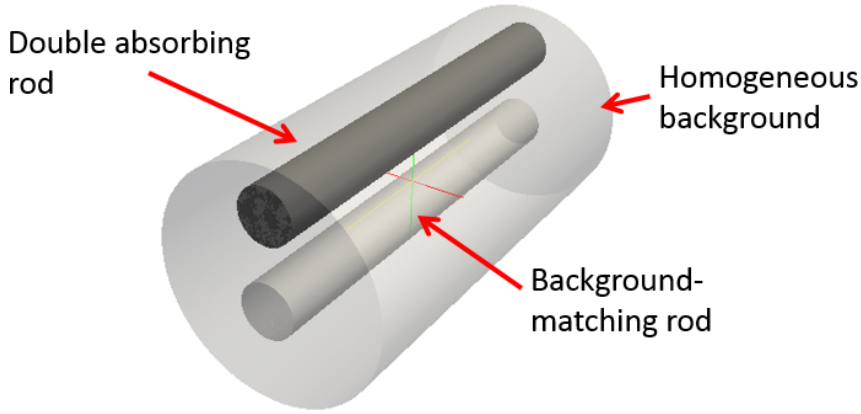

Figure 8.15: Schematic of the heterogeneous cylinder configuration used in the cylinder DOT experiment.

der phantom, but with background-matching rods inserted in both tunnels so as to render it homogeneous. The cylindrical reference was chosen in this case because of the unique shape of the underside of the phantom compared to the others available; it was judged that this would affect source input and that the only reliable reference phantom was therefore the homogeneous cylinder.

Spectral DOT reconstruction used default regularisation and a maximum number of iterations of 7. For simplicity, it was assumed that scattering was known; it was fixed at the expected value (appendix A.7.4) and not solved for. As in the XPM2 spectroscopy experiment, a single dye was considered that was responsible for the absorption. To model the extinction coefficient in this case, a curve was fit to the manufacturer data for the range 500 to 850nm. This is shown in figure 8.16, along with sample images. The initial guess for the absorber concentration was set to that of the background, i.e. 1, whilst the target value for the anomaly region was 2.

It can be seen that the images of diffuse flux change in an intuitive way as a function of the source position. It can be seen that, owing to the shape of the cylinder, the signal is weakest in the middle as compared to at the sides. In the case of the cylinder, owing to the geometry, direct reflections of the source would be visible in the mirrors if they were used, for this reason and those already discussed in the XPM2 spectroscopy case, top-view only data was once again used.

Figures 8.17 and 8.18 show the results of performing spectral DOT reconstruction. The

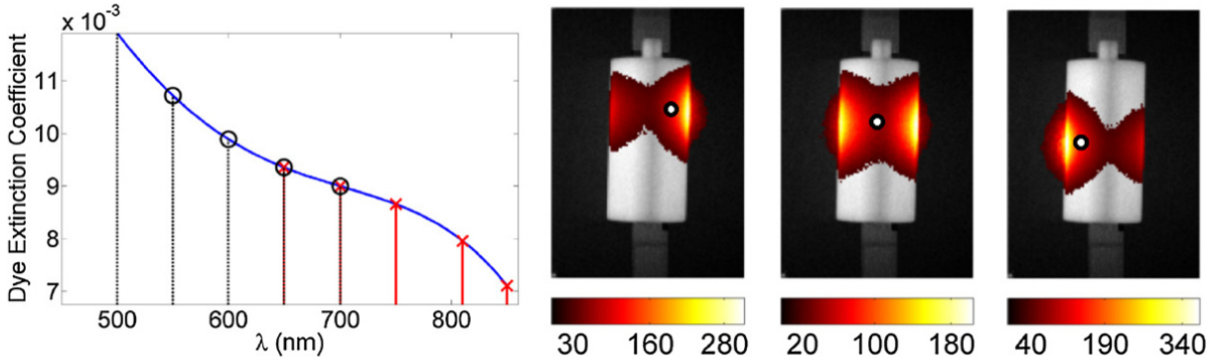

Figure 8.16: Spectral dye extinction (left) with wavelengths at which DOT data was acquired and at which properties were read off for use later in BLT (see section 9.3.4) indicated by vertical lines (red lines and crosses for DOT and black dashed lines and circles for BLT), and sample transillumination images at 750nm (right) with source positions indicated by white circles.

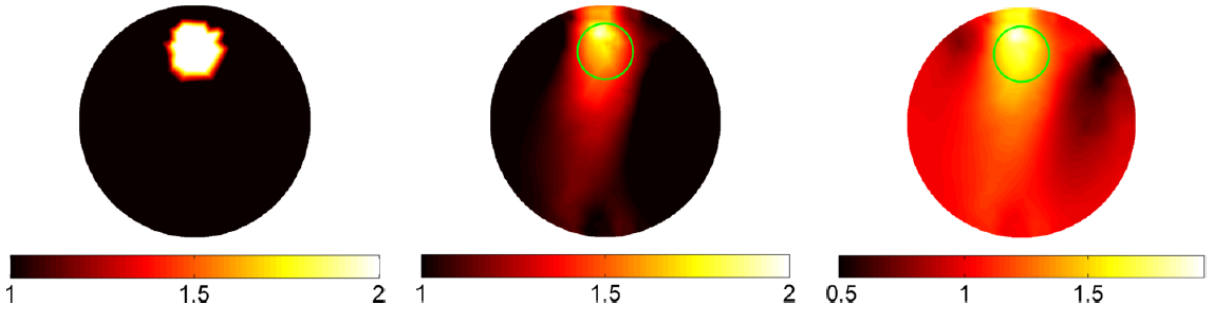

Figure 8.17: Result of spectral DOT reconstruction with (a) target absorber concentration, (b) reconstructed absorber concentration scaled to the same color-scale as the target, and (c) reconstructed absorber concentration auto-scaled. The slices shown are taken through the fully 3D reconstruction at the axial location of the centre of the DOT source grid. In the reconstruction slices, the green circle indicates the bounds of the true anomaly for reference.

image was reconstructed containing a blurred anomaly with approximately the correct spatial location and quantitative concentration. It can be seen that there are some image artefacts around the anomaly where the background concentration is underestimated.

8.3.3 Block Phantom Tomography Experiment

Another DOT test was carried out using the drilled block phantom (appendix A.1.1). In this experiment, 'anomaly' and 'blank' rods made for use with the cylinder phantom (and therefore not matching the block in either case, see appendix A.1.1) were placed in the drilled tunnels. The effect was to create a heterogeneous phantom with background optical properties of the

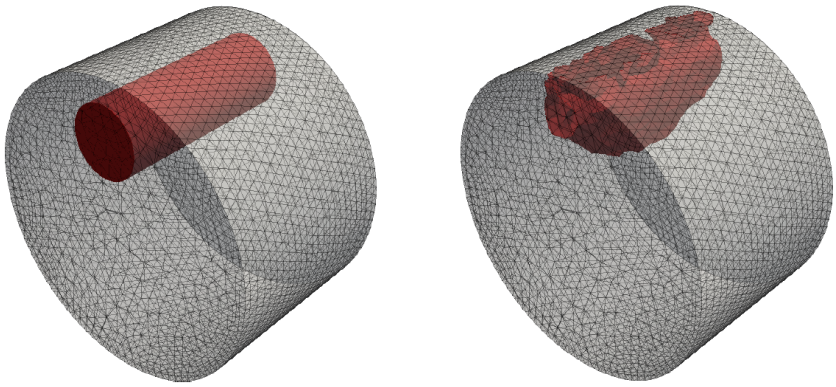

(a) (b)

Figure 8.18: 3D visualisation of (a) the target dye concentration and (b) the reconstructed concentration. The displayed contour plots are produced by thresholding the data at 1.5 (the mid-way point between the background and anomaly ground truth values). The axial range corresponds to a cropped section, approximately 20mm long, centred over the source grid.

block and two anomalies with cylinder phantom properties (fig. A.46), the upper anomaly having double absorption as compared to the lower anomaly. The resultant set-up is shown in fig. 8.19.

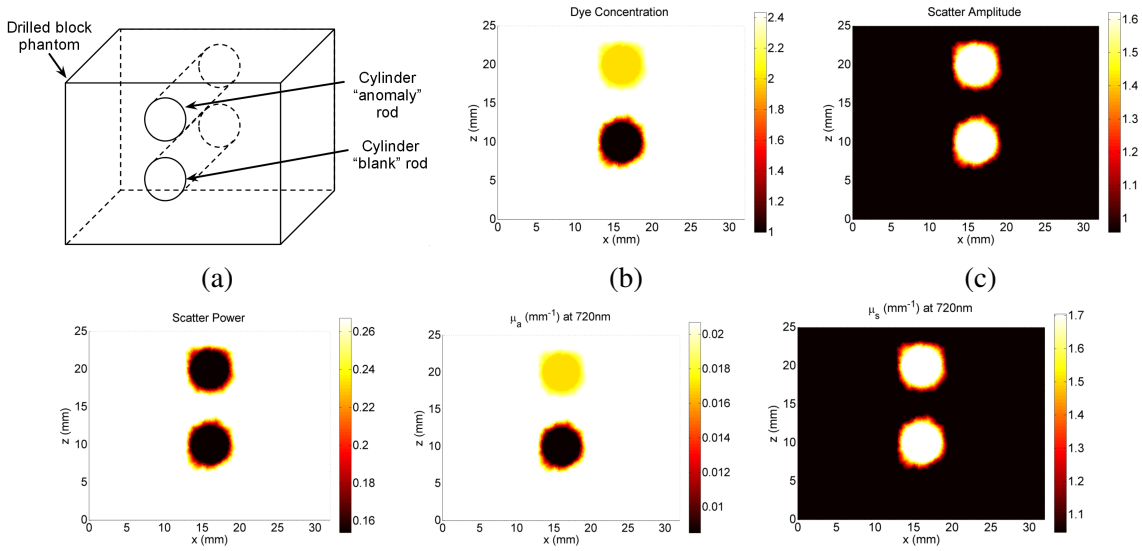

Figure 8.19: (a) Schematic of set-up and (b-f) target optical property cross-sections of the heterogeneous block phantom used in this experiment.

Diffuse imaging was performed as outlined in section 8.2.1 for the wavelengths 650, 720, 750, 810, and 850nm. The homogeneous block phantom (appendix A.1.1) was used as the reference phantom.

Reconstruction was performed with initial optical properties assumed to be those of the homogeneous block (shown in A.46), using the same single-dye representation method as before. In this case both scattering and absorption properties were reconstructed for. The data selected was once again top-view only, with regular down-sampling to 10%. The reconstruction was allowed to run for a maximum of 100 iterations but in practice it terminated after 13, default regularisation was used. Figure 8.20 shows the results at the final iteration.

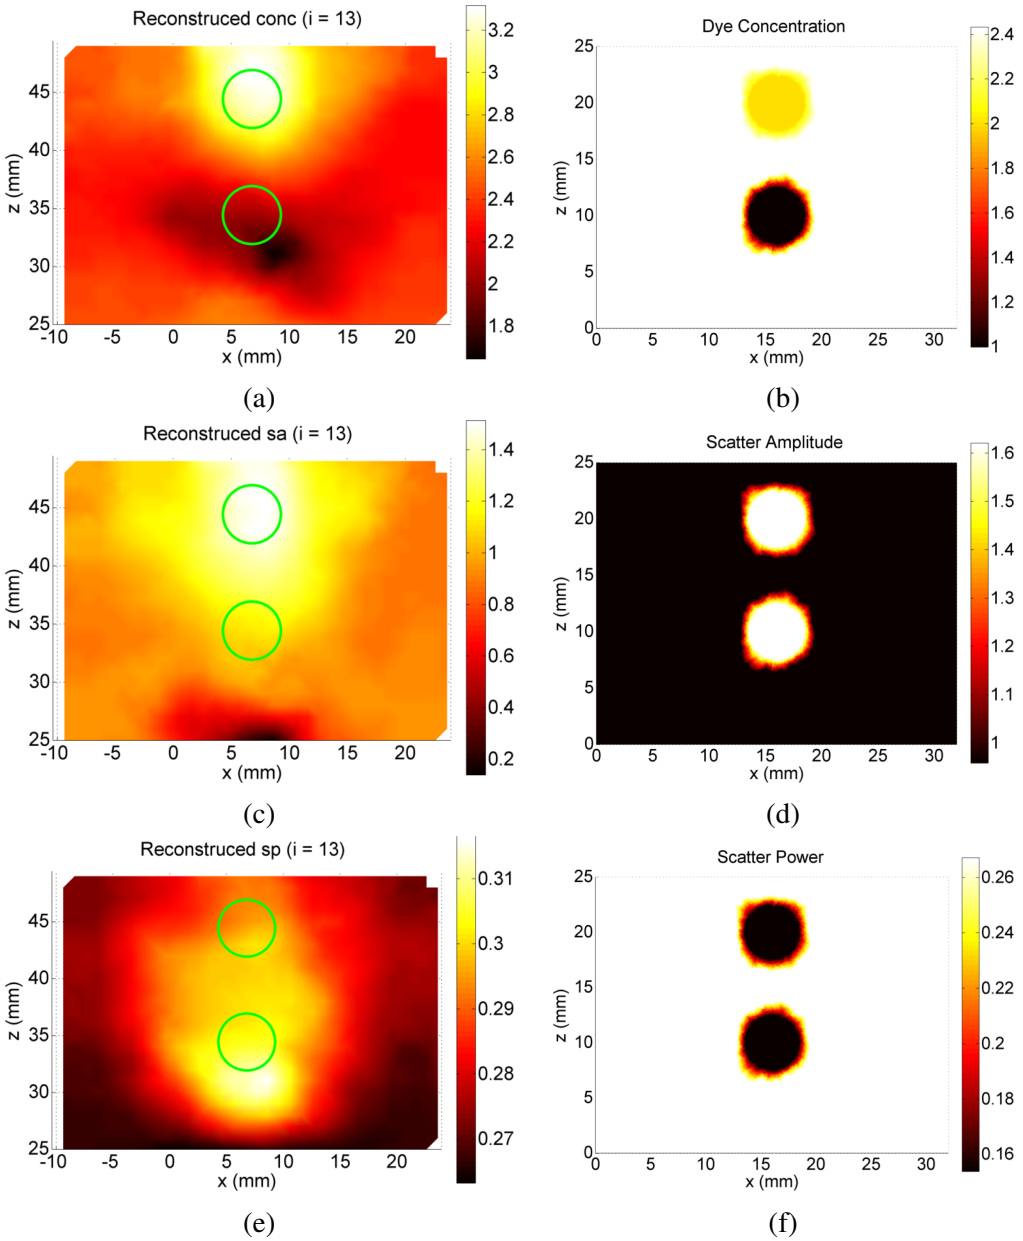

Figure 8.20: Cross-sections through reconstructions of spectral properties of the heterogeneous block phantom with corresponding target values.

It can be seen that the values are not quantitatively accurate. Though qualitatively it can be seen that some of the expected structure is visible and some of the values are roughly correct. In concentration, the target upper and lower anomalies should have slightly and significantly less than the background material and in practice there is a region of low concentration found in the area of the lower anomaly which is encouraging. There is also a structure in the concentration image that looks like the upper anomaly but with contrast in the wrong direction (i.e. the reconstructed concentration is higher than the background whereas it should be slightly lower). As compared to the concentration image, the scatter amplitude image appears to show one large distribution covering both the rods which might indicate it is somewhat recognising the increased scattering in the rods, whilst for the scatter power the result is similar in shape with somewhat more distribution around the lower rod. Though the scatter power values show contrast in the wrong direction, the range in the values is notably small.

In order to consider the result in terms of the impact on the optical properties μ_a and μ_s' which are particularly important for the dual-modality system, fig. 8.21 shows the μ_a and μ_s' reconstructed and target values for a single wavelength.

It can be seen that the image structure is qualitatively similar to the expected structure, the absorption properties being essentially the same as the concentration ones discussed above, in terms of quantitative impact on the concentration, it is shown that at this wavelength the absorption is generally over-estimated. In terms of the scattering it is apparent that there is an unwanted low scattering region at the base of the image which may be compensating for inaccuracies in the source position, representation or intensity calibration being as this is the region close to the sources. Additionally, the peak scattering value is lower in reconstruction than in the target case which perhaps is expected in conjunction with the generally over-estimated absorption. This and particularly the inverted upper concentration anomaly suggest that the system is experiencing cross-talk between the absorption and scattering, compensating for the higher scattering to some extent with absorption in this case, which is to be expected using continuous-wave data.

In general, qualitatively the images are useful, showing anomalies to an extent where they are. Though there is some quantitative inaccuracy there is also a tendency in many cases towards

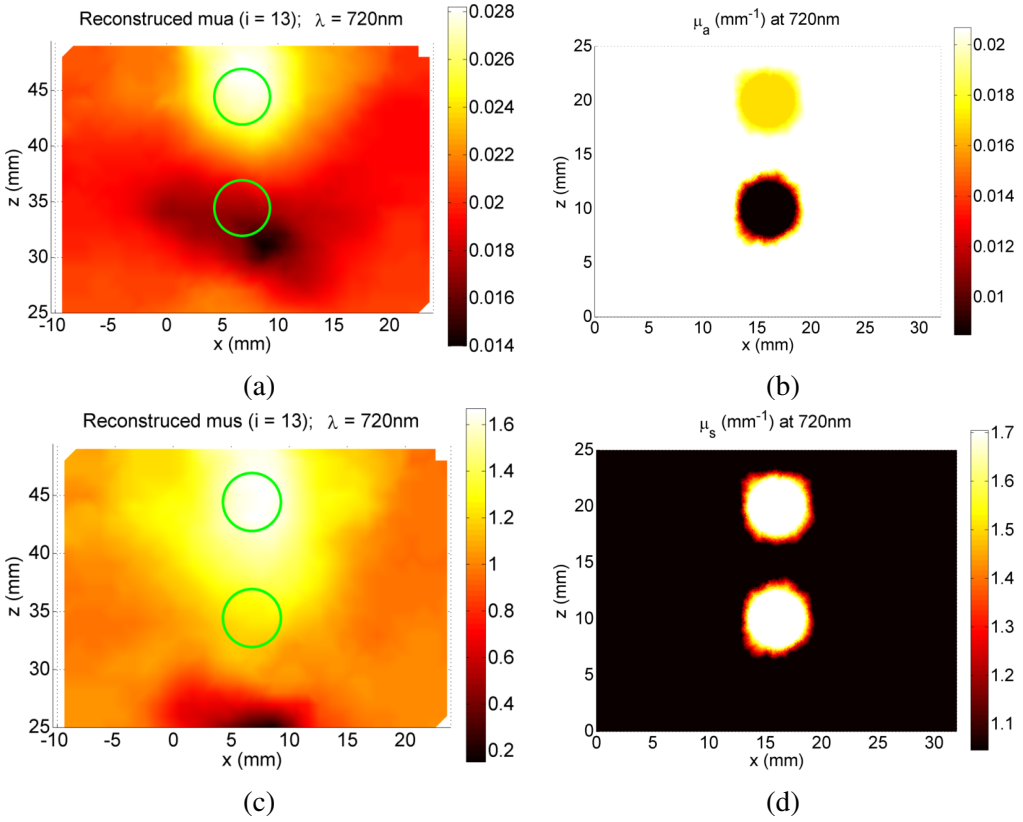

Figure 8.21: Cross-sections through reconstructions of optical properties of the heterogeneous block phantom with corresponding target values.

the correct result.

8.4 Discussion

Three experiments have been reported involving DOT of different phantoms, demonstrating the imaging system and work-flow in DOT mode. The reconstruction parameters and phantoms both varied in complexity to provide an overview of system capabilities.

In the first experiment, a set of spectroscopy (single-pixel tomography) results were obtained for the mouse-shaped XPM2 phantom. The phantom is homogeneous but realistically shaped. The initial guess for scattering was set at the ground truth and the amount of error in the initial guess of the absorption was varied. Concentration and scatter properties were reconstructed for.

The experiment demonstrates the use of the developed system for performing diffuse optical spectroscopy (DOS), the simplest form of DOT. Results showed that with under-estimated or correctly estimated absorber concentration[1], the expected absorption was recovered/maintained in reconstruction. However, further experiments where the absorption was over-estimated showed that the absorption remained over-estimated in these cases. The scattering appeared to converge on the same unexpected spectrum in all cases, which could suggest that the *ground truth* values obtained independently could be wrong, or might indicate some systematic issue with the system or the single experimental data set used in all reconstructions.

The experiment demonstrates a method that can and should be repeated (many times with rigorous attention to the variables and isolating these where necessary) before reliable conclusions can be drawn regarding the phantom properties or the system performance. It would be interesting to investigate spectroscopy problems of increasing difficulty, perhaps containing more known or partially known absorbers, making use of the multi-view data, or using real tissues. Where other well-known phantoms are used it may be possible to isolate any systematic issues that might be giving rise to errors. It would also be interesting to corrupt the scattering, or both the scattering and absorption together, and reconstruct.

[1] to the best of available knowledge.

The second experiment was DOT of the heterogeneous cylinder phantom. The phantom contained a single anomaly which could be seen as the simplest possible approximation to an absorbing organ such as the liver or spleen. The result of the experiment was an image which showed clearly an anomaly with approximately the correct characteristics of the known inclusion. This is an encouraging result as it illustrates the capability to image heterogeneous phantoms and produce qualitatively and quantitatively meaningful images.

The images did however contain some artefacts. Isolation of the cause of the artefacts and improvement of image quality are subjects for future work; it is clear that a great deal further data would be required to draw accurate conclusions. One potential weakness is the DOT source placement in this case which, owing to the curvature of the cylinder and the fact that source locations were measured in the stage plane and merely moved to the nearest point on the mesh, is likely to have introduced some error. This is discussed further below.

In the final experiment, the block phantom with two heterogeneities was imaged, in this experiment the scattering and absorber concentration parameters were all solved for in tomography-mode. Resultant images showed evidence of detection of both anomalies but with more accurate results in the case of the shallower inclusion. This being said the absorption was significantly over-estimated in the shallower inclusion. There was also some evidence of cross-talk between the scattering and absorption parameters.

The rendered image showing scattering and absorption coefficients at 720nm suggested that the recovered parameters in practice would result in an over-estimated absorption and an under-estimated scatter, though qualitatively all of the images were encouraging owing to the similarity of spatial structures in the target and reconstruction.

This experiment represents a proof-of-concept for imaging organ structures or otherwise spatially-resolved optical properties with both scattering and *absorption*, or scattering and *absorber* based contrast. Nevertheless, the lack of quantitative accuracy in general demonstrates the need for a great deal of further investigations with more phantoms, tailored to isolating variables and rigorously analysing system deficiencies so as to arrive at improved results. A possible reason for the poorer performance in the block experiment as compared to the cylinder

experiment quantitatively is the unrealistically large thickness of the block.

A current limitation, and potential cause of errors in the XPM and cylinder experiments, is the method of modelling DOT sources. To re-iterate, projected DOT source centre-points in the stage plane were established by experiment (appendix A.6.3) and these are placed at the nearest point on the mesh when building the model for reconstruction.

In the case of the cylinder experiment this was identified as one potential cause for image artefacts because these were seen towards the sides of the cylinder where the nearest point on the mesh would represents a larger distance from the measured location and therefore would be expected to be less accurate (consider for example the further perspective effects on the projected Gaussian travelling further to hit the curved surface at the sides).

The model is also likely to be inaccurate when the object is not flush with the stage in general, and this is compounded by inaccuracies that will be introduced into the meshing. It is known that the XPM sits a small distance ($\approx 1mm$) above the stage owing to its rigidity and shape but this is not known the system and therefore the mesh is created assuming that it is flush with the stage. As such, the mesh will be thicker than it should be and source positions will be slightly wrong .

Moreover, when using the block phantom as the reference, as in the XPM case, the model-matching calibration may be inaccurate because the true source position and distribution not the same in the reference and test cases (because the block *is* flush with the stage). Note that this aspect of the problem was bypassed in the cylinder experiment by using the cylindrical reference phantom also to match the shape and any systematic source error. The extent of this issue could be studied further, though because it is not expected to be an issue for flat surfaces (e.g. the block, resting mice) this perhaps should just be kept in mind when performing further phantom studies.

It is possible that these matters could be improved by the modelling of the projection of the NIR source in a more thorough physical way e.g. by utilising the free-space propagation model in conjunction with an additional projector calibration for the DOT source projector to map fluxes from the projector onto the object. This could be challenging particularly because

of the small working distance which might invalidate the pin-hole camera (or thin lens) approximation of the optics for this unit and a more thorough model might be necessary. Alternatively, a small camera such as a web-cam placed below the object and the stage could image directly the projected source pattern and with co-calibration could report the spatial distribution as visualised directly on the object, this could make for a more thorough analysis and more accurate treatment.

Another approach might be to solve for the source location in the reference phantom using BLT reconstruction (section 2.5), this might allow for a more accurate representation of a projected source from the perspective of the model although this would not help with changing geometries between reference and test subjects.

An elaborate method that could allow for the calibration phantom to match the test subject directly is using the surface capture generated mesh as input to a 3D printer to produce a replica of an imaged animal in a plastic with known optical properties, this could potentially allow for a calibration on a per-detector (as well as per-source and per-wavelength) basis to empirically determine the free-space function for all points on the surface, though it would be costly at present in time and materials.

An alternative approach to the presented calibration is to compare collected images of the reference phantom with projected (free-space forward modelled) images of the NIRFAST simulated version. This might be more accurate than the demonstrated method because model inversion errors associated with the flux reconstruction for the reference images would not be present (the forward model is expected to be more accurate because it is not an inverse problem with regularisation etc.), otherwise scaling factors could be derived on the same basis only image-side rather than flux-side.

8.5 Conclusion

In this chapter, the method for using the developed system in DOT mode was introduced. This included a description and justification for the transillumination imaging protocol and the data-

model matching calibration method, as well as the non-contact reconstruction scheme.

Details and results were then presented for three separate experiments with three phantoms highlighting and preliminarily investigating different areas of system usage and performance. Between the experiments, it was shown that the work-flow could be applied with some success to arbitrarily-shaped, in one case realistically mouse-shaped phantoms. It was also demonstrated that single-pixel reconstruction (spectroscopy) as well as spatially-resolved tomography could be performed solving for scattering and absorption parameters.

Together, the results provide a thorough initial demonstration of the usage of the presented methods; the novel point-by-point wide-field illumination scheme, the novel combination of surface capture, free-space mapping and DOT, and the final spectral reconstruction, image production and parameter extraction. The preliminary investigations have also highlighted several important directions for future study. These include many further experiments with phantoms in spectroscopy and tomography mode to isolate the discussed free-variables, optimise imaging parameters and explore the capabilities and limitations of the system in a thorough scientific fashion. The key aim being to improve image quality in quantitative terms and understand why high quantitative accuracy is not achieved in all cases at present.

Another highlighted direction for future work is more advanced modelling of the DOT sources, though in practice this may not be necessary dependant upon the impact of modelling errors on other flat-bottomed phantoms and particularly on real animals.

Regularisation has been discussed but not interfered with in terms of changing the default NIRFAST behaviour. Assuming that the spatial changes in bulk optical properties are expected to be smooth this is highly appropriate; there is a large body of work demonstrating this where this is typically the case for example in breast cancer imaging problems where DOT has been widely applied[65, 66, 61]. The validity of this in small animal imaging is less clear, and other methods of regularisation or indeed entirely other ways of including prior knowledge and perhaps foregoing the need for regularisation are worthy of future study.

Further experiments are required to understand the system. If possible better phantoms are also needed that could close the gap in terms of complexity between those used and real animals.

Either this or some other imaging technique capable of measuring optical properties would be needed as a gold standard. The aim should be in either case to image something perfectly like a living mouse but with perfectly well-known properties for validation.

Ultimately, blind studies where the real properties are not known should be carried out from a position of naivety to investigate the potential to characterise new samples using the presented techniques. This is important in the work-flow of the long term aim of the system because in general animal properties will not be known in advance and this is the key reason for using DOT. A hard-prior DOS experiment, as demonstrated, may prove to be a useful step prior to spatially resolved DOT to obtain background bulk property estimate to use as an initial guess and this work-flow could be further explored.

For now the presented DOT method has been demonstrated to some extent and approaches to imaging have been shown using the novel method and system. In the following chapter one of the results is taken and used to provide prior knowledge in a BLT reconstruction experiment. This draws together the final anticipated work-flow of the system, demonstrated for the first time with a simple phantom.

CHAPTER 9

BIOLUMINESCENCE TOMOGRAPHY

9.1 Introduction

A primary goal of the developed system is the performance of accurate bioluminescence tomography (BLT), that is the production of quantitatively accurate 3D images of reconstructed bioluminescent source distributions throughout a living mouse. Such images could for example be obtained for luciferase-tagged cancer cells such that tumour growth could be monitored and visualised in 3D. This would provide a powerful tool for monitoring a great many biomedical problems.

The general outline for BLT has already been introduced in section 2.5. In this chapter details are provided of how BLT is done in the presented system and results are presented for practical experiments demonstrating accurate BLT imaging in phantoms. The whole BLT work-flow is demonstrated as was the diffuse optical tomography (DOT) work-flow in the previous chapter. Quantitative results are promising and a single experiment utilising a DOT reconstruction from the previous chapter ties together the full work-flow intended for the final developed system.

The chapter first provides an overview of the BLI image acquisition and processing methods and the BLT reconstruction method used in the system. Results are then presented for three sets of experiments. The first of these is a study on reconstructing single and multiple sources inside the XPM2 mouse-shaped phantom (appendix A.1.1), thereby representing the study of

a realistically shaped object with simple known properties and varied source distributions. Repeats of this study are also reported so as to shed light on the reproducibility of the system. The second study uses the cylindrical phantom (appendix A.1.1) with light sources placed at 8 different internal positions. This allows an analysis of the consistency of quantitative source measurements as a function of internal source position and a comparison to standard BLI in the same situation. Finally, the last study locks together the intended work-flow for the system by reporting DOT-BLT of the heterogeneous cylinder phantom. The chapter is then concluded with discussions.

9.2 Methods

9.2.1 Luminescence Imaging

The bioluminescence image acquisition work-flow is identical to that presented for DOT (section 8.2.1), except that whereas in DOT there was an image acquired for each illuminating pattern and one without illumination, there is just a single image acquired per wavelength in luminescence mode because the source is unchanging and uncontrollable. The major difference between BLT and DOT imaging, other than the smaller number of images acquired, is the longer exposure times typically used to acquire high signal.

For the DOT experiments thus far carried out, only single-view data is used (i.e. mirrors are not used). This is not the case in BLT mode; mirrors are placed around the imaged subject following its positioning in the stage and their locations are established in the course of surface capture (section 6.2.6).

The image processing is identical to the DOT case other than there being no dark subtraction. Images acquired are first converted into units of electrons per second and then mapped onto a surface captured mesh using the free-space mapping method. As before the FEM model detector locations are defined by the locations of the mapped data. This is all that is needed to then reconstruct because there is no external source to model. The exception to this is the use

of pre-made meshes in the case of the cylinder phantom as already discussed in section 7.3.3.

Note that in contrast to DOT where calibration drew the values into real-world units, no such calibration has been established or used here with BLT. As such, all results are only compared and evaluated in a relative sense. An experiment could (and should) be designed in the future to achieve imaging in real-world units though there is still a great deal of scientific value in the relative data presented here.

In all cases presented, as in the DOT case, imaging parameters were chosen to optimise SNR based on the analysis of appendix A.5. Thus the read mode was 1, the binning was 4×4, and the CCD mode was normal.

Exposure times were selected by acquiring images at short exposures and extrapolating the values to longer times. Times were selected that maximised the signal received up to a target value of 60000 counts, thus allowing some uncertainty to avoid saturating the 16-bit detector. A maximum exposure time of 10 minutes was set as a cut-off point to avoid infeasibly long experimental time.

9.2.2 Image Reconstruction

As in the case of DOT, all light modelling is performed with NIRFAST software tools and is based on the diffusion equation (section 2.3). Recall from section 2.5 that the FEM-based luminescence system is described by

$$\mathbf{y} = \mathbf{W}\mathbf{b}, \tag{9.1}$$

where $\mathbf{y}$ is the vector of boundary measurements, $\mathbf{W}$ is the spectral sensitivity matrix. and $\mathbf{b}$ is the source vector (a value per-node in the FEM mesh). This system cannot be inverted directly because $\mathbf{W}$ is non-square and the inversion problem is ill-posed and under-determined, thus in general it is necessary to solve for the source vector $\mathbf{b}$, i.e. some variant of

$$\mathbf{b} = \min_{\mathbf{x}} ||\mathbf{y} - \mathbf{W}\mathbf{x}||_2^2. \tag{9.2}$$

In the course of this project, a novel algorithm was developed to solve this problem specifically; compressed-sensing based conjugate gradient (CSCG). This reconstruction algorithm was described in detail by Basevi et al.[4]. It combines two objective functions. The first is the least-squares fit with data as shown in eq. (9.2) and the second is the L1-norm of the reconstructed values which is used to approximate the L0-norm. The objective function is

$$\mathbf{b} = \min_{\mathbf{b}} ||\mathbf{y} - \mathbf{W}\mathbf{x}||_2^2 + \gamma||\mathbf{x}||_1, \tag{9.3}$$

where γ is a parameter controlling the relative weighting of the two objectives, the second of which favours spatial sparsity in the reconstructed image. In the presented work, the objective is initially weighted heavily in favour of sparsity and then reduced throughout the course of the algorithm (in the same manner described in Basevi et al.[4]) such that the overall effect is to have a guided least-squares search whose starting condition is influenced by sparsity but which nevertheless is not "regularised" in the same way as the Tikhonov method that is often applied to BLT[55, 58][1]. As such, final images will not necessarily be sparse but will be a good fit to the data.

The algorithm uses a conjugate-gradient optimisation method[4] which proceeds until it converges. Convergence is assumed when there is an absolute percentage change in the objective function value between adjacent iterations of less than 0.001% (chosen empirically).

CSCG was developed with a particular intended application in imaging scenarios where sparse, compact source distributions are expected such as in the case of a cancer study with metastases or a macroscopic cell tracking study. It was shown that CSCG improved robustness to noise when using a different imaging system and performing BLT on the XPM2 phantom (appendix A.1.1) over the Tikhonov approach and one other non-negative least squares algorithm[4]. CSCG is used in all cases in the presented work.

[1] Incidentally this is the method applied to surface flux mapping in chapter 7.

9.3 Results

9.3.1 BLT of XPM2

The XPM2 mouse-shaped phantom (appendix A.1.1) was imaged in luminescence mode. The phantom contains two light sources, A and B, the positions of which have been established (appendix A.8), that can be switched on and off independently. This feature was exploited to facilitate imaging each source individually and both together. The result was three different experimental scenarios used in three separate experiments. Imaging was performed at 560, 580, 590, 600, 620, and 640nm.

The results for the source A experiment are shown in fig. 9.1. It can be seen that the re-

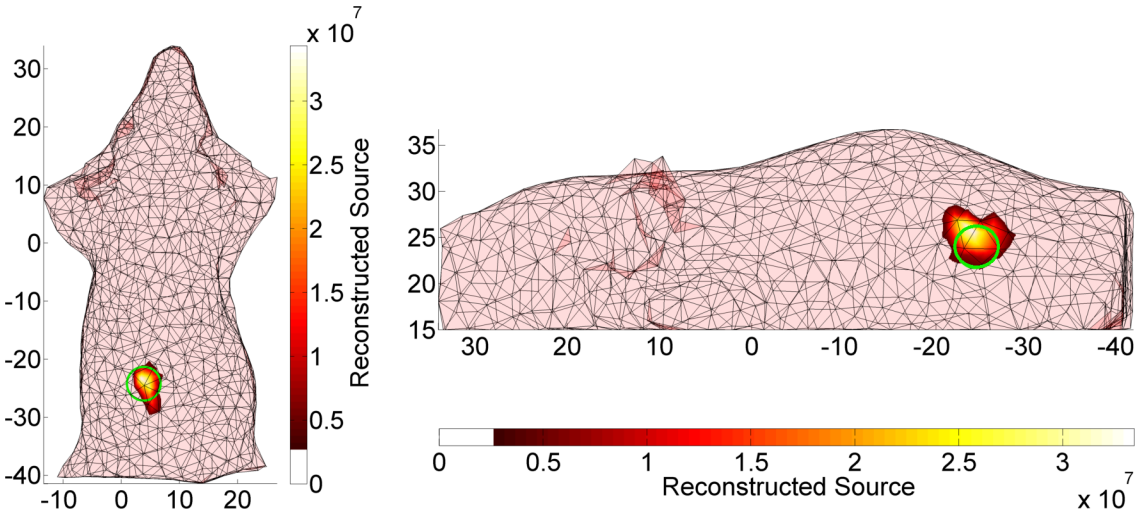

Figure 9.1: XPM-2 BLT Reconstruction with only source A switched on. Reconstructions are shown (a) from above and (b) from one side. The source is shown as interpolated slices through the centre-of-mass of the source distribution. The green circle indicates the expected location for the source.

constructed source distribution is well co-localised with the target region, the source is fairly compact and has a peak intensity of 30 million (a.u.). The results for the source B experiment are shown in fig. 9.2.

The source B reconstruction is also well co-localised with the target region. The source distribution is less compact as in the source A case, and the peak value is lower at 7 million

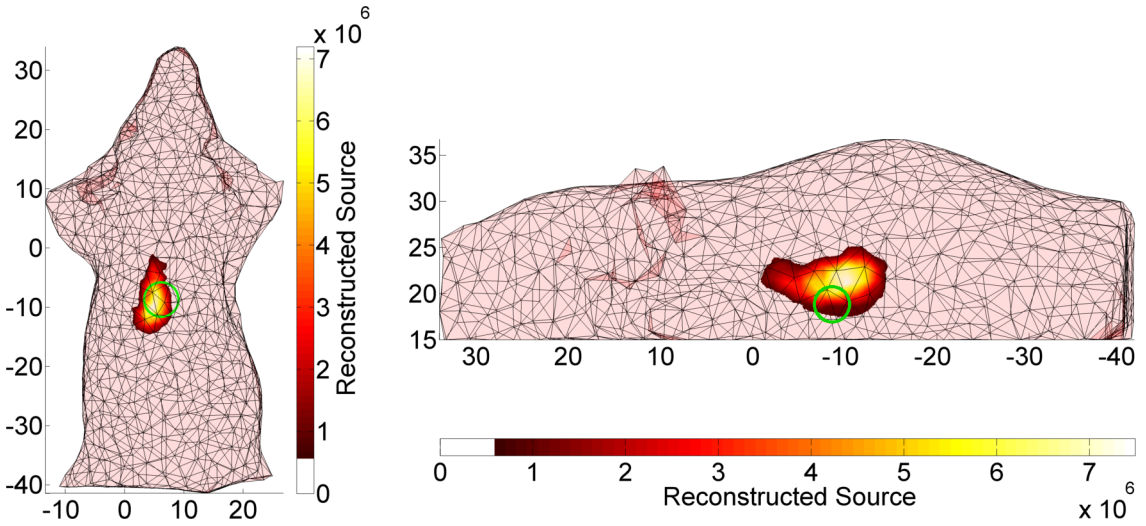

(a) (b)

Figure 9.2: XPM-2 BLT Reconstruction with only source B switched on.

(a.u.). The results for the simultaneous source A and B experiment are shown in fig. 9.3.

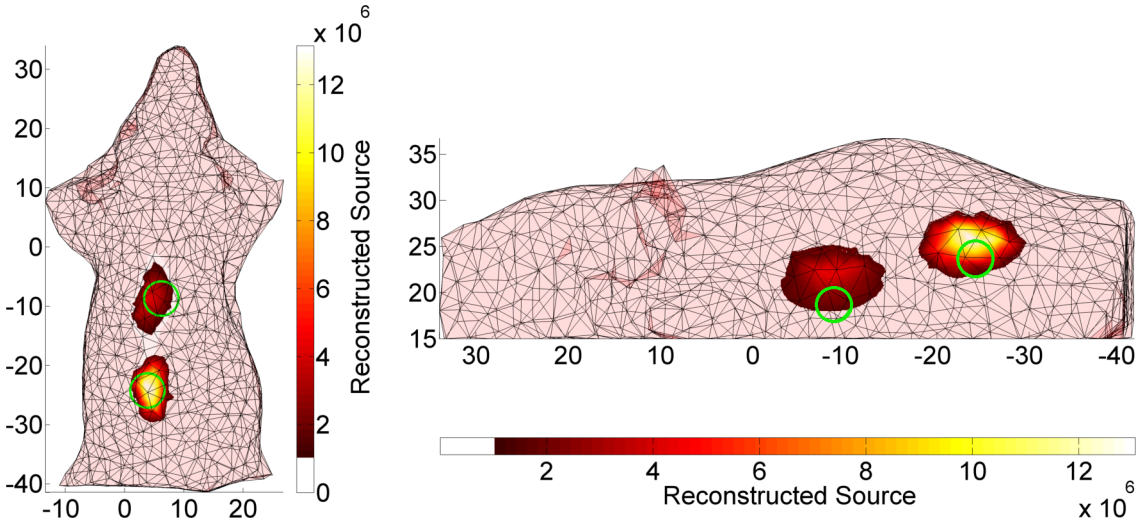

(a) (b)

Figure 9.3: XPM-2 BLT Reconstruction with both sources A and B switched on simultaneously.

The multiple-source reconstruction shows both of the expected source distributions co-localised with the target positions, they are about the same size, and clearly separable with high contrast. The peak value of the whole distribution is 12 million (a.u.) which indicates that the part of the distribution corresponding to source A (reconstructed still with a higher peak intensity) has weakened, this is in-line with it also spreading out somewhat as is also observed. The peak of the reconstructed source B is around 5 million (a.u.). Table 9.1 shows quantitative data

for the reconstructions including the reconstructed centres-of-mass of the source distributions calculated as

$$c_b = \frac{\sum_{i=1}^{n} \mathbf{x}_i b_i}{\sum_{i=1}^{n} b_i}, \tag{9.4}$$

where b_i is the source value at the i^{th} FEM node at position $\mathbf{x}_i$. This is compared to the target source location (appendix A.8) to establish a localisation error metric. The table also shows target and reconstructed total source intensity calculated as

$$s = \sum_{i=1}^{n} b_i. \tag{9.5}$$

The target source value has real-world physical units whereas the reconstructed values do not owing to the system calibration being only relative, however, the relative strength of the sources is still meaningful and comparable to that of other sources.

The relative total source reconstructed in the source A only experiment is 93% of that found for source B. This is an encouraging result because the expected relative intensity is 96%. This indicates that the source intensity may be being measured accurately. In the multi-source experiment the total reconstructed source was 28% lower (as a proportion of source B) than expected which is also encouraging but not as good as in the single-source cases.

The localisation error for sources A only and B only were 1.83mm and 2.59mm respectively (± 2 for the uncertainty established in appendix A.8), as such the error estimates are within and almost within the estimated range of the uncertainty on the target location. Both of these localisation errors are increased in the multi-source case which is indicative of the added complexity of the problem, it is possible that there is some cross-talk between sources and the amount of data in terms of both independent observations and the dynamic range available for both sources is less which may have an impact. Nevertheless the error is only just over 1mm from the expected region in the worst case. A better understanding and knowledge of the internal expected source might allow stronger conclusions to be drawn. It can be seen that most of the localisation error is in the z component, i.e. the depth w.r.t. the top of the phantom. This could indicate that there is an effect due to errors in the wavelength that is measured - as it is the spectral

Source Configuration	**A only**	**B only**	**A and B**
Expected total source (photons/s)	*8.090e+010*	*8.400e+010*	*1.649e+011*
Expected total source (percentage of B)	*96.31*	*100.00*	*196.31*
Reconstructed total source (a.u.)	*3.674e+008*	*3.950e+008*	*6.636e+008*
Reconstructed total source (percentage of B)	*93.01*	*100.00*	*167.99*
Expected centre-of-mass of source A (mm)	*(3.85,-24.24,23.54)*	-	*(3.85,-24.24,23.54)*
Reconstructed centre-of-mass of source A (mm)	*(4.36,-24.48,25.28)*	-	*(4.31,-23.55,25.53)*
Components of localisation error for source A (mm)	*(0.50,-0.25,1.75)*	-	*(0.46,0.69,1.99)*
Localisation error for source A (mm)	*1.83*	-	*2.16*
Reconstructed volume at half-max of source A (mm^3)	*3.94*	-	*11.55*
Radius of sphere of equivalent volume to source A (mm)	*0.98*	-	*1.40*
Expected centre-of-mass of source B (mm)	-	*(6.19,-8.73,18.78)*	*(6.19,-8.73,18.78)*
Reconstructed centre-of-mass of source B (mm)	-	*(4.86,-8.00,20.89)*	*(4.84,-8.65,21.66)*
Components of localisation error for source B (mm)	-	*(-1.32,0.73,2.11)*	*(-1.35,0.08,2.88)*
Localisation error for source B (mm)	-	*2.59*	*3.18*
Reconstructed volume at half-max of source B (mm^3)	-	*14.92*	*32.19*
Radius of sphere of equivalent volume to source B (mm)	-	*1.53*	*1.97*

Table 9.1: Quantitative analysis of XPM-2 BLT results. Expected total source is taken from manufacturer documents (appendix A.1.1), the measured is calculated by eq. (9.5). Expected centre-of-mass is found in appendix A.8, the measured is calculated using eq. (9.4). The error is the Euclidean distance between them. Measured volume is the summed volumes of all FEM elements in which every nodal source value is higher than 50% of the maximum, the radius of a sphere with equal volume is given to aid interpretation of this value. For the double-source experiment, nodes with source values greater than 10% of the overall maximum value were clustered using K-means with two clusters to identify two regions to treat as sources A and B.

characteristics that provide most of the depth information[56, 44, 59, 60] which might be due to the known issue with the bandpass filters (appendix A.2.5).

The reconstructed volumes of sources as-reconstructed varies by a factor of approximately 2 or 3. The genuine target volume is not known but it is expected that whatever the true effective source volume it should (if accurate) be consistent between measurements (e.g. between single and multi-source experiments).

9.3.2 Reproducibility of XPM2 Results

The above experiment was repeated from start to finish; the phantom was removed and re-placed in the imaging system, the mirrors were removed and re-placed, the same imaging protocol was then repeated from surface capture to luminescence imaging in the same three source scenarios, the system was switched off between experiments. A further repeat was then performed without the phantom being removed and replaced (i.e. only the imaging part was repeated). The same reconstruction method was then used (with identical parameters to those used above) to reconstruct luminescence images.

Due to the low number of samples for each scenario, and the variable change between experiments not being identical this is a *preliminary* investigation into reproducibility and accuracy of the BLT part of the system and further repeats of each part of the imaging protocol are required in order to analyse reproducibility in quantitatively accurate terms.

Figure 9.4 shows reconstructed images acquired from each of the experiments, visualised as before[1]. It can be seen immediately that in all cases the correct qualitative appearance is achieved, either a single or both sources are seen located in approximately the correct position.

Table 9.2 shows quantitative analysis of the source A only experiments. It can be seen that across these sets there was a drop in the total source intensity reconstructed as compared to the first experiment of around 10%. As a fraction of source B for the corresponding experiment

[1]Note that, in the figure, the expected region (green circle) and source appear in what looks like different positions owing to the different amounts of the phantom that happened to be masked into the surface capture and therefore were meshed, in terms of the absolute positions in the physical phantom the green circles are in the same place.

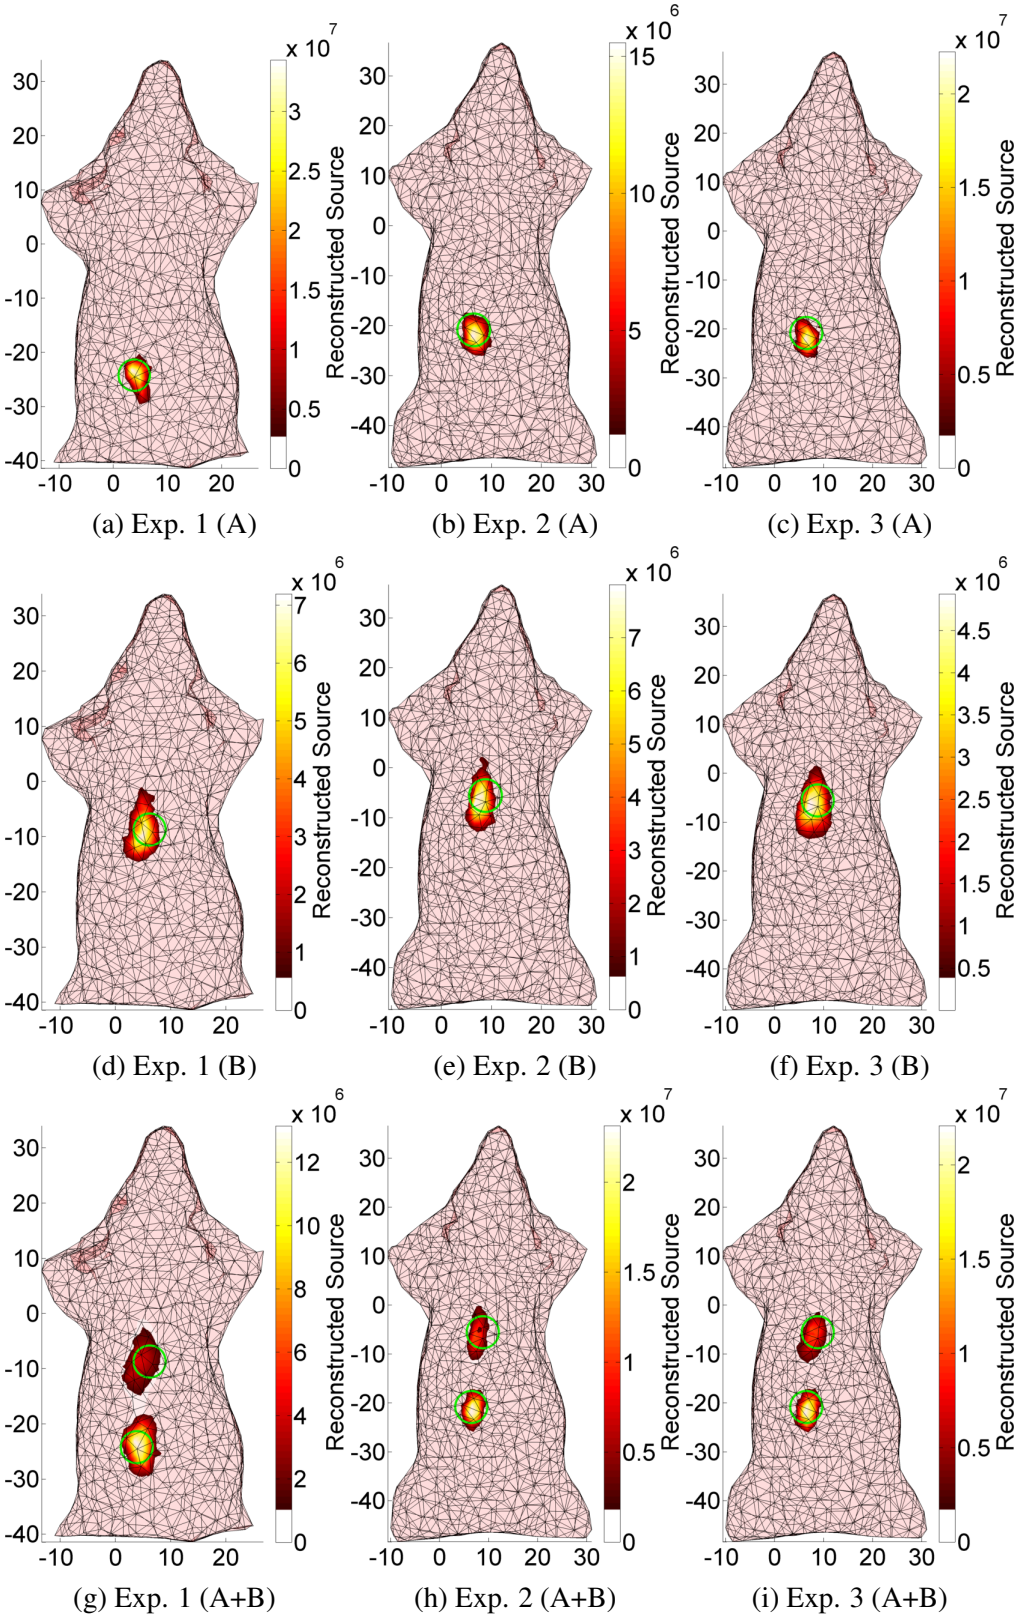

(a) Exp. 1 (A) (b) Exp. 2 (A) (c) Exp. 3 (A)

(d) Exp. 1 (B) (e) Exp. 2 (B) (f) Exp. 3 (B)

(g) Exp. 1 (A+B) (h) Exp. 2 (A+B) (i) Exp. 3 (A+B)

Figure 9.4: Reconstructed source images for three repeats of the XPM BLT experiments, shown from a single perspective only.

number (below) there was also a change of around 10%. This makes the worst case estimate across the set (compared to the expected relative value) that of set 2 which is approximately 14% below the target value. However this is a relatively small error still (compared to BLI for example) and it should be noted that because it is a relative measurement there is the combined error of the source B and A only reconstructions. The extremal values being within 12% of one another suggests good reproducibility of this quantitative measurement.

The error in source localisation across the sets ranges from 1.49 to 1.83mm which is a change on the order of 25%, which is a small change and might indicate reproducibility in the reconstructed location. Particularly it should be noted that the components of the localisation are similar in each case (i.e. the relative position of the reconstructed source is the same with respect to the target, as opposed to being off by the same distance but on the other side of the source for example) with the position in z being under-estimated consistently by around a millimetre, the y value always underestimated and the x value always over-estimated. Given the level of uncertainty on the source location (2mm) this is especially encouraging. The reconstructed volume varies approximately 3-fold.

Table 9.3 shows the same analysis for the source B only experiments. In this case the reconstructed source as a percentage is not meaningful because this was merely used as the benchmark for comparison and therefore is 100. The variation in absolute reconstructed intensity between the repeats was around 13% which is slightly higher than the change seen for source A though still low. It is seen that the source B total intensity values rise then fall across the experiments which is unlike the behaviour seen in source A.

The total localisation error appears to be quite consistent, being practically identical in experiments 1 and 2 and and increased by around 12% in 3. As was seen in the case of source A, the error characteristics in terms of the x, y, z components are the same across the sets with the majority of the error being in the vertical direction (depth being under-estimated as seen in the original images), approximately half as much error in the x component and very little error in the y component. This is interesting as the ground truth position in the y direction of the source is considered the most likely to be accurate given how it was found (appendix A.8). The

Experiment Number	1.00	2.00	3.00
Expected total source (photons/s)	*8.090e+010*	*8.090e+010*	*8.090e+010*
Expected total source (percentage of B)	*96.31*	*96.31*	*96.31*
Reconstructed total source (a.u.)	*3.674e+008*	*3.525e+008*	*3.341e+008*
Reconstructed total source (percentage of B)	*93.01*	*82.10*	*87.41*
Expected centre-of-mass of source A (mm)	*(3.85,-24.24,23.54)*	*(6.46,-20.86,24.31)*	*(6.46,-20.86,24.31)*
Reconstructed centre-of-mass of source A (mm)	*(4.36,-24.48,25.28)*	*(6.70,-21.99,25.25)*	*(6.52,-22.09,25.49)*
Components of localisation error for source A (mm)	*(0.50,-0.25,1.75)*	*(0.24,-1.13,0.94)*	*(0.06,-1.23,1.18)*
Localisation error for source A (mm)	*1.83*	*1.49*	*1.70*
Reconstructed volume at half-max of source A (mm^3)	*3.94*	*5.80*	*1.96*
Radius of sphere of equivalent volume to source A (mm)	*0.98*	*1.11*	*0.78*

Table 9.2: Repeated results for source A experiments.

reconstructed volume varies by around 2-3 times, similarly to source A.

Table 9.4 shows the analysis for the double-source experiments. It can be seen that the total reconstructed source varies by around 11% with is consistent with the single-source sets. It can also be seen that with this variation the two repeats of the experiment are closer to the expected relative intensity (as a proportion of B only intensity) being 25% and 11% away rather than the 28% of the first set.

More striking is the difference in the localisation error for source A, in the first set this was around 2mm whereas it is around 1mm in the repeats. This is interesting as this is also the first observation of different characteristics in terms of the x, y, z errors which show the centre-of-mass of the reconstruction to be a different direction from the target location. However, the differences are small and the majority of the error is dominated by the top-relative depth error as before. The localisation error for source B is more consistent and shows a consistent relative position from the expected source location. Interestingly, the position of the sources (i.e. the positive/negative characteristics of the components of localisation) are very similar to those

Experiment Number	**1.00**	**2.00**	**3.00**
Expected total source (photons/s)	*8.400e+010*	*8.400e+010*	*8.400e+010*
Expected total source (percentage of B)	*100.00*	*100.00*	*100.00*
Reconstructed total source (a.u.)	*3.950e+008*	*4.293e+008*	*3.823e+008*
Reconstructed total source (percentage of B)	*100.00*	*100.00*	*100.00*
Expected centre-of-mass of source B (mm)	*(6.19,-8.73,18.78)*	*(8.80,-5.58,18.87)*	*(8.80,-5.58,18.87)*
Reconstructed centre-of-mass of source B (mm)	*(4.86,-8.00,20.89)*	*(7.57,-5.44,21.13)*	*(7.63,-6.04,21.49)*
Components of localisation error for source B (mm)	*(-1.32,0.73,2.11)*	*(-1.23,0.14,2.26)*	*(-1.17,-0.46,2.62)*
Localisation error for source B (mm)	*2.59*	*2.58*	*2.91*
Reconstructed volume at half-max of source B (mm^3)	*14.92*	*19.58*	*35.14*
Radius of sphere of equivalent volume to source B (mm)	*1.53*	*1.67*	*2.03*

Table 9.3: Repeated results for source B experiments.

observed in the A and B only experiments indicating a systematic difference. The reconstructed volumes vary in a similar manner to the variations observed in previous data sets.

Across the experimental repeats there is good quantitative stability in the results for total source, varying by only around $10-13\%$ for all sets. The error in localisation is also generally consistent, and in most cases the observed offset shows that the reconstructed source is in a similar position relative to the source as opposed to just being a similar distance away. The volume is observed to be more variable on the order of 2-3 fold variations across the sets.

Overall the results appear from this preliminary study to be quite reproducible. The consistent position of the reconstructed sources relative to the expected centre-of-mass could indicate that these are the true positions of the source (because there was uncertainty of estimated 2mm on the positions) or that there is a systematic error being introduced such as might be expected due to the known issues with the bandpass filters (namely the dependence on angle of the filtering which is not compensated for in this study). Another potential source of error is the source spectrum of the XPM-2 LED outputs, this was obtained from a tiny graph provided by the man-

Experiment Number	**1.00**	**2.00**	**3.00**
Expected total source (photons/s)	*1.649e+011*	*1.649e+011*	*1.649e+011*
Expected total source (percentage of B)	*196.31*	*196.31*	*196.31*
Reconstructed total source (a.u.)	*6.636e+008*	*7.377e+008*	*7.100e+008*
Reconstructed total source (percentage of B)	*167.99*	*171.84*	*185.72*
Expected centre-of-mass of source A (mm)	*(3.85,-24.24,23.54)*	*(6.46,-20.86,24.31)*	*(6.46,-20.86,24.31)*
Reconstructed centre-of-mass of source A (mm)	*(4.31,-23.55,25.53)*	*(6.81,-21.46,24.98)*	*(6.72,-21.50,25.16)*
Components of localisation error for source A (mm)	*(0.46,0.69,1.99)*	*(0.35,-0.61,0.67)*	*(0.26,-0.64,0.85)*
Localisation error for source A (mm)	*2.16*	*0.97*	*1.09*
Reconstructed volume at half-max of source A (mm^3)	*11.55*	*5.36*	*4.45*
Radius of sphere of equivalent volume to source A (mm)	*1.40*	*1.09*	*1.02*
Expected centre-of-mass of source B (mm)	*(6.19,-8.73,18.78)*	*(8.80,-5.58,18.87)*	*(8.80,-5.58,18.87)*
Reconstructed centre-of-mass of source B (mm)	*(4.84,-8.65,21.66)*	*(7.44,-5.77,21.48)*	*(7.72,-6.55,22.09)*
Components of localisation error for source B (mm)	*(-1.35,0.08,2.88)*	*(-1.36,-0.19,2.61)*	*(-1.09,-0.97,3.22)*
Localisation error for source B (mm)	*3.18*	*2.95*	*3.53*
Reconstructed volume at half-max of source B (mm^3)	*32.19*	*18.57*	*12.37*
Radius of sphere of equivalent volume to source B (mm)	*1.97*	*1.64*	*1.43*

Table 9.4: Repeated results for simultaneous source A and B experiments.

ufacturer and as such is likely to have substantial error (appendix A.1.1), given that this impacts upon the interpretation of the spectral model it could well cause a depth misjudgement. As has been stated more experiments are needed to clarify these suggestions, but the preliminary results suggest that the system performance in BLT is likely to be reproducible to within around 15% for total source and to within around 1*mm* in terms of localisation error.

In terms of reconstructed volume, results are less consistent. Coupled to the fact that the other quantitative results seem quite stable this might indicate a variability owing to the stochastic nature of the reconstruction algorithm, which simply may have found a different local minima, the fitness landscape of the reconstruction problem could be analysed to shed light on the validity of this idea. Alternatively owing to stochastic nature of the mesh generation there could also be variability in the precise internal node structure which might result in a different distribution of values geometrically. As indicated, further study is required to establish the source of variability with isolation of the many variables in the complex work-flow being necessary to draw firm conclusions.

9.3.3 BLT of Cylinder Phantom

BLT experiments were performed using the cylinder phantom that took advantage of the cylindrical geometry, specifically its rotational symmetry, to image multiple light source positions within an otherwise identical phantom. Recall figure fig. 7.4 and the description of the cylinder phantom set-up in section 7.3.2; the phantom was prepared in exactly the same way for this experiment; it was effectively homogeneous, with known properties, and with an orange trigalight source placed halfway along one of the tunnels. The source was therefore located at one of two depths (5 or 15mm) inside the cylinder phantom.

In a first experiment, an imaging run was performed with the source in each of the tunnels. The cylinder was then rotated by 180 degrees and this was repeated. Across four data sets the target source location was therefore one of four depths along the central axis of the cylinder as shown in fig. 9.5 (a-d). In a second experiment, this was repeated with the addition of a further rotation of 45 degrees applied at the beginning. As such in the second experiment the

appearance of the target source positions is of an off-axis version of those in the first set (fig. 9.6 (a-d)).

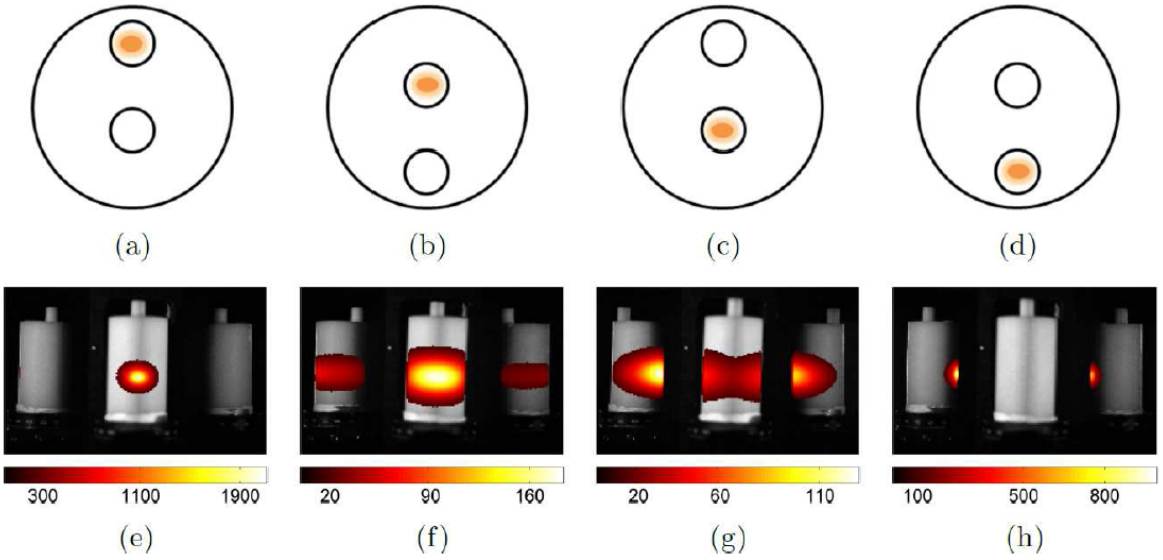

Figure 9.5: Sample BLI images for the on-axis data set: (a-d) phantom schematics; and (e-h) luminescence images at $\lambda = 600nm$ shown overlaid on bright images from surface capture image sets. Overlaid images are set to transparent for values less than 10% of the image maximum.

As in the cylinder phantom experiments in chapter 7, it was fixed in a rotation mount that was mounted directly to the sample stage to provide accuracy and repeatability when turning and re-positioning the phantom[1].

BLI was performed at 500, 550, 600, 650, and 700nm. Sample images are shown in figures 9.5 and 9.6[1]. The figures also indicate the source location diagrammatically.

The total acquisition time for BLI was dependent on the level of signal available, and subsequently upon source depth and external perspective, but as an example the total imaging time for the case presented in figure 9.5(a) was approximately 8 minutes.

It is worth noting the striking diversity of apparent image features across these scenarios given that the actual source, though situated at different positions, is known to be identical in each case. For example a tightly packed source is evident in the shallowest case (figure 9.5(e)) yet two intuitively separable blurred surface flux structures exist in a deeper case (figure 9.5(g)). Quantitatively, it can be seen that the signal is lower by a factor of ≈ 10 between

[1]Note that image presentation is deliberately designed to mimic the format often used to present results in biomedical studies (see for example Contag[202])

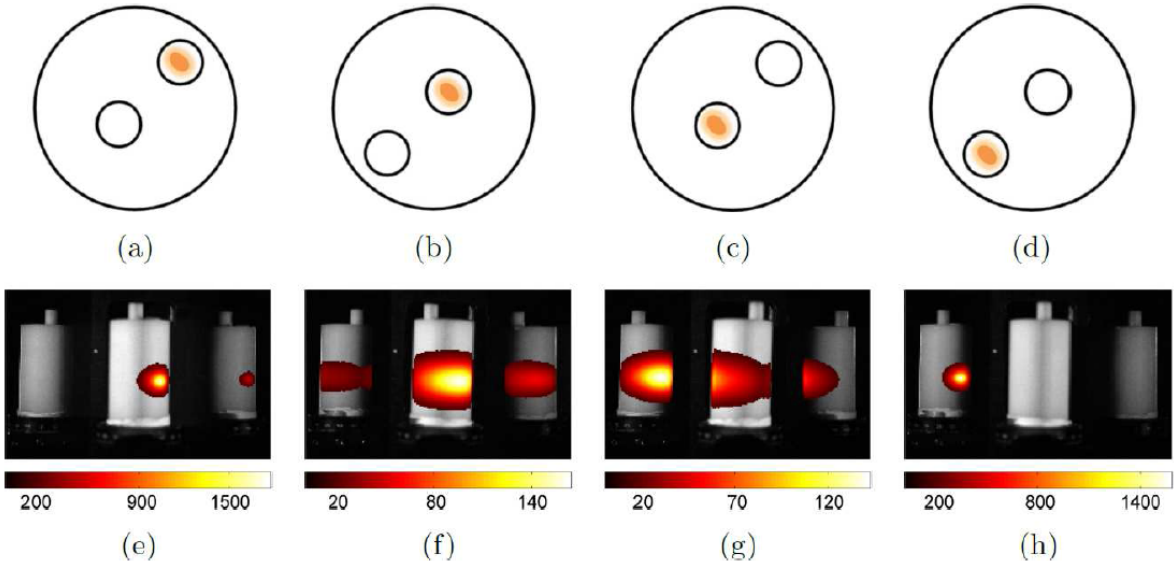

Figure 9.6: Sample BLI images for the off-axis data set at $\lambda = 600nm$.

the first and second scenarios in both the on and off-axis data sets. Whilst the mirrors allow access to otherwise invisible signals (e.g. figure 9.5(h)) this also confuses matters under naive interpretation in that a deeper source with respect to the camera view-point can be seen to appear quantitatively more intense in a mirror view due it being shallow with respect to the nearest visible surface point; in this case a point visible through a mirror. It is therefore not possible to accurately deduce a photon count (or, therefore, cell-count in a biological study) directly from this type of data when the source is a collection of bioluminescent or fluorescently-labelled markers. These results are a clear indication of the need for more advanced tomographic image recovery.

The image sets were used for BLT with a pre-made FEM mesh containing around $11,000$ nodes. Figure 9.7 shows a single luminescence reconstruction (corresponding to the example in figure 9.5(a)), rendered in 3D using Paraview (Kitware, NY, USA) and illustrating a method of visualisation of slices through the volume at the centre-of-mass of the reconstructed source used to aid comparison with the 2D cross-sectional target diagrams. Remaining BLT results are shown in this way in figures 9.8 and 9.9. These figures also show the BLI data again for direct comparison; recall that a major aim of the BLT approach is to improve on the accuracy and interpretability of BLI data.

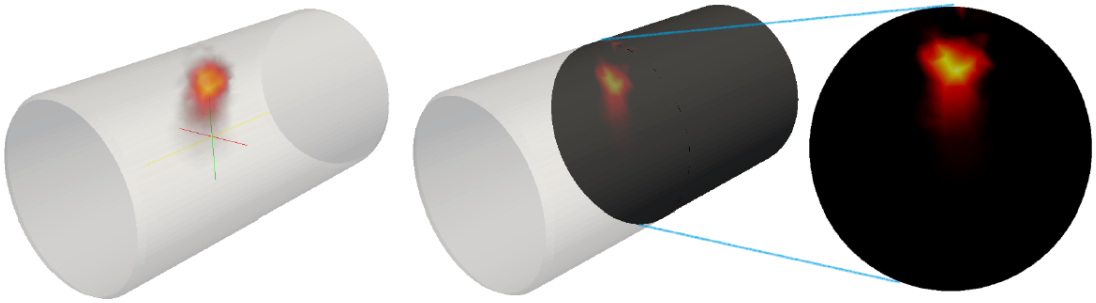

(a) (b)

Figure 9.7: Visualisation of (a) 3D reconstruction for the data set with the most shallow source (see figure 9.8(a)) along with (b) indication of how the following 2D slice representations of results are obtained

In the on-axis data set (figure 9.8), it can be seen that qualitatively the reconstructed images are accurate and clear; in contrast to the results of BLI, it is easy to interpret the images as showing a single source in the correct location. It can also be seen that in the cases where the source is farthest from the surface (figure 9.8(j) and figure 9.8(k)) the reconstructed distributions are slightly broader than in other cases and in the case where the source is farthest from the detector and least visible in recorded images (figure 9.8(l)) the recovered distribution appears qualitatively less well-localised.

In the off-axis data set (figure 9.9) the results are very similar, with the images once again qualitatively clear and accurate in terms of showing a single source in the correct location in each case. There is a qualitative improvement in the image of the deepest source (figure 9.9(l)) compared to the on-axis set primarily due to the improved signal level and surface flux visibility in the rotated case; it is expected that the deepest case in the on-axis experiment would be the most challenging problem since the source is most difficult to see from all three views and is therefore a worst case scenario.

Figure 9.10 shows a quantitative analysis of the BLT reconstructions. 2D and 3D localisation errors are shown, calculated by two different metrics. The first metric assumes the position of the reconstructed source to be the centre-of-mass (eq. (9.4)) and the second assumes it to be the position of the maximum-valued node. The target position is assumed to be the centre of the tunnel containing the source, the Euclidean distance from which is the error. The reason for

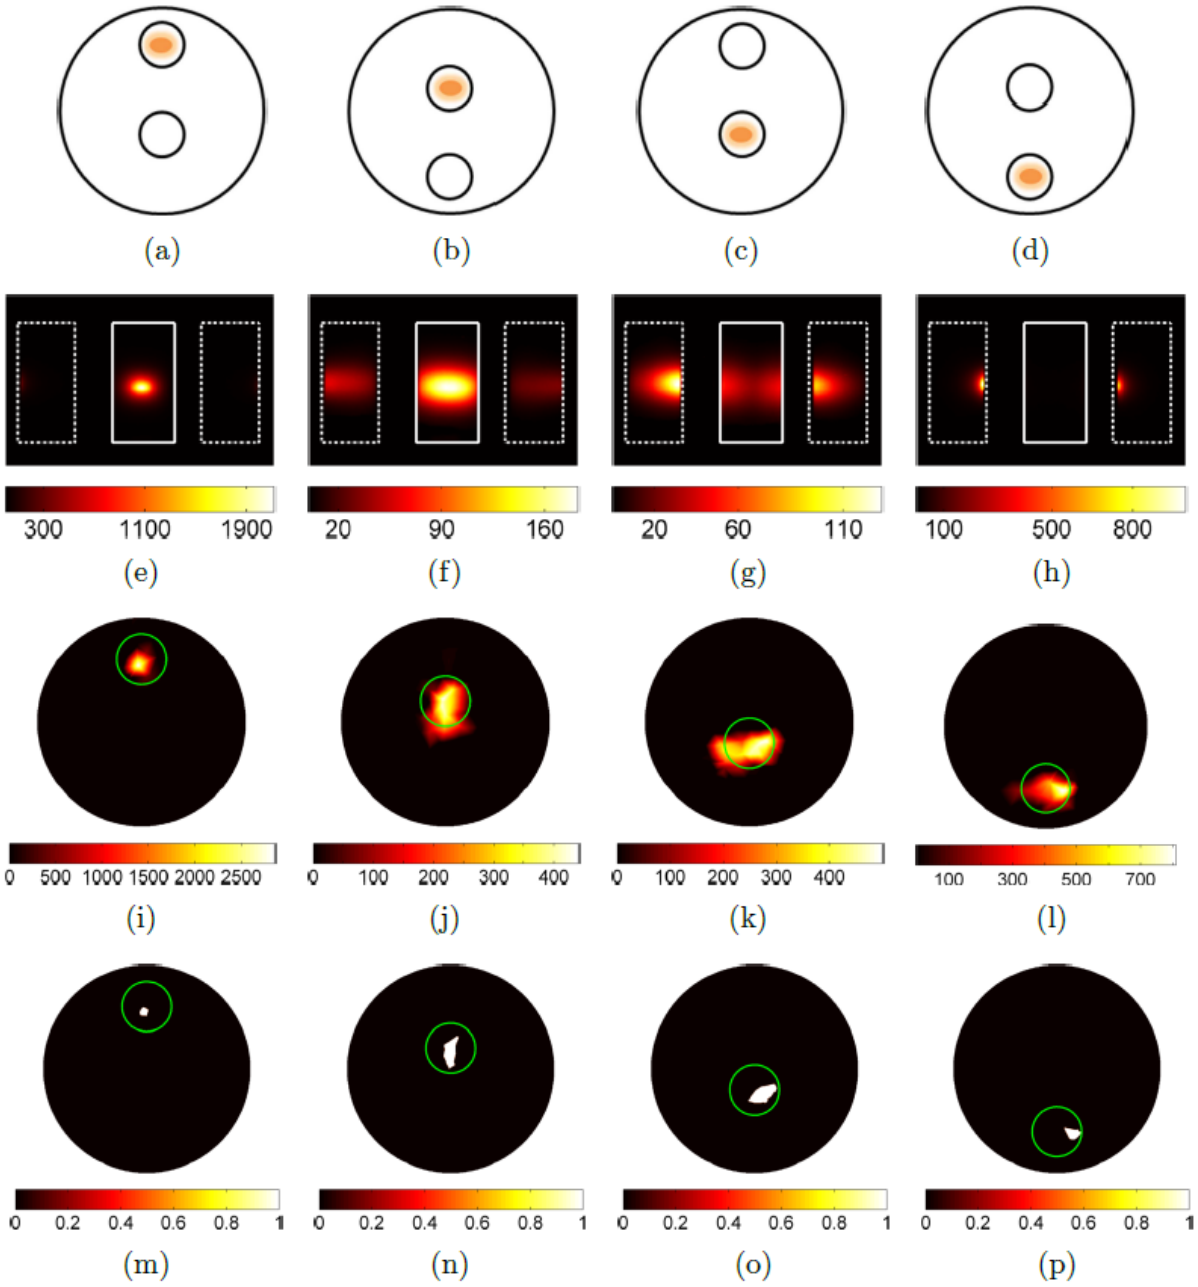

Figure 9.8: Summary of on-axis, homogeneous BLT experiment results showing: (a-d) schematics of source experimental locations in 2D projection; (e-h) BLI images (CCD measurement e^{-}/s) of the phantom at $\lambda = 600nm$ with approximate phantom outlines; (i-l) slices through corresponding BLT reconstructions at the axial offset corresponding to the centre-of-mass of the reconstruction; and (m-p) the slice images thresholded at 75% of the maximum value.

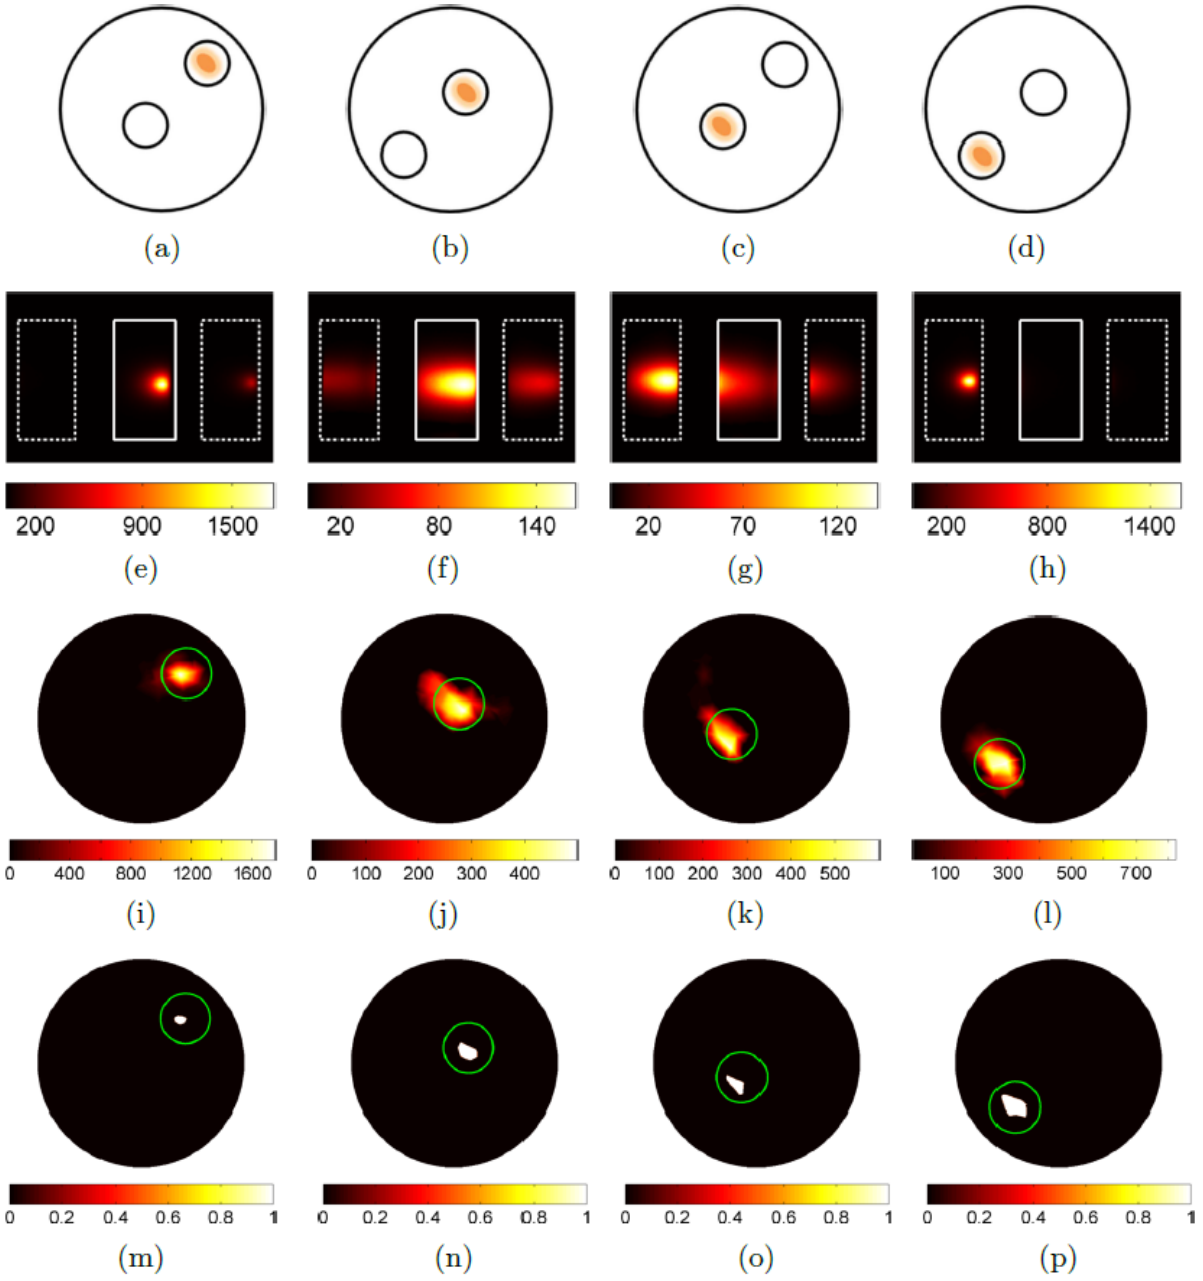

Figure 9.9: Summary of off-axis, homogeneous BLT experiment results showing: (a-d) schematics of source experimental locations in 2D projection; (e-h) BLI images (CCD measurement e^{-}/s) of the phantom at $\lambda = 600nm$ with approximate phantom outlines; (i-l) slices through corresponding BLT reconstructions at the axial offset corresponding to the centre-of-mass of the reconstruction; and (m-p) the slice images thresholded at 75% of the maximum value.

showing both 2D and 3D metrics is that the axial depth of the source was difficult to control and as such the 3D target position has some added (unknown) uncertainty. Total reconstructed source is also shown in arbitrary but consistent units and as a percentage of the overall mean value for clear comparison.

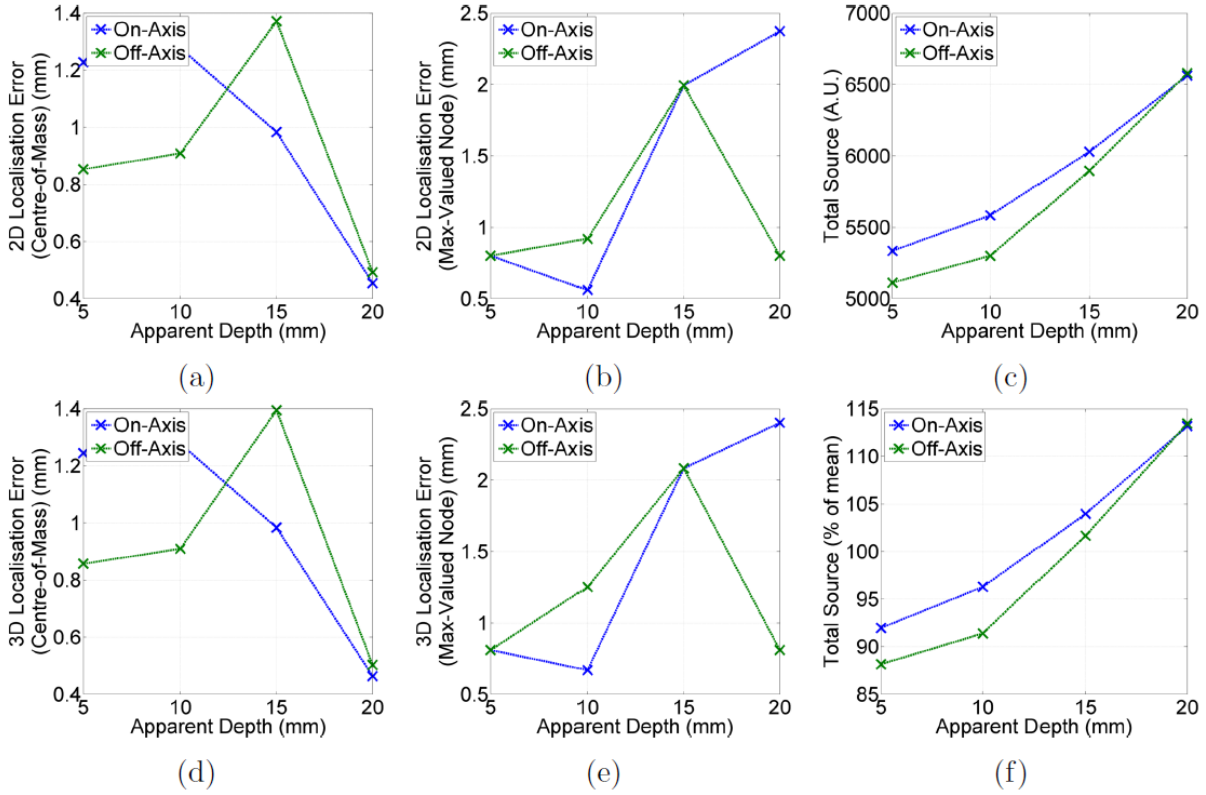

Figure 9.10: Summary of BLT results in quantitative terms for both the on-axis and off-axis (45° rotated) experiments ordered versus effective depth w.r.t. the top of the phantom: (a) the 2D error in reconstructed source position based on the centre-of-mass metric; (b) the 2D position error based on the max-valued node metric; (c) the total (summed) reconstructed source shown in arbitrary units; (d) and (e) the 3D versions of (a) and (b) respectively; and (f) the values of (c) shown as a percentage of the mean value across both sets.

It can be seen that by both metrics the localisation error is less than 2.5mm in all cases; less than 1.5mm in all cases using the centre-of-mass metric which is consistent with the best previously published BLT results[203, 204, 56, 192]. Using the max-valued node error metric it can be seen that there is less accuracy as the depth increases in the on-axis set, which is indicative of the source becoming less well focused (i.e. more diffuse) whilst remaining centred in approximately the correct location. It can be seen that the reconstructed source intensity is consistent across both sets, within 15% either side of the mean value. This is very encouraging as it shows

quantitative stability in the reconstruction with respect to a diverse variety of source locations and depths. This is a dramatic improvement on the quantitative variation within bioluminescence images (section 9.3.3) and suggests that in the case where optical properties are known, and assuming the diffusion equation holds sufficiently well, the presented BLT approach could be effectively applied to cell-counting applications and would offer a substantial improvement upon BLI.

9.3.4 Multi-modal DOT-BLT of Cylinder Phantom

To test and demonstrate the approach of using DOT results as prior information to BLT image reconstruction, in-keeping with ultimate multi-modal imaging aim of the system, a luminescence imaging run was performed with the light source in one cylinder tunnel and a double-absorbing rod in the other. The source position was as in the third on-axis scenario presented above (fig. 9.5(c)). The DOT part of this experiment has already been presented in section 8.3.2.

Two BLT reconstructions were performed, one in which homogeneous phantom properties were assumed, and one where the spatially varying DOT-derived properties were used. The same FEM mesh and reconstruction parameters were used as throughout section 9.3.3. Results are shown in figure 9.11.

The reconstructed images are different, the image based on the DOT data is qualitatively superior as it shows the single source clearly visible as has been seen before. The other reconstruction (with homogeneous properties assumed) by contrast shows a divided feature that appears to be weighted outside of the expected region. Table 9.5 shows a quantitative analysis of the results.

3-View Reconstructions	**Homogeneous**	**DOT Prior**
2D Position Error (mm; centre-of-mass)	1.23	1.20
3D Position Error (mm; centre-of-mass)	1.23	1.20
2D Position Error (mm; max-valued node)	1.99	1.99
3D Position Error (mm; max-valued node)	2.08	2.08
Total Source (A.U.)	5791	6677
Total Source (% of BLT experiment mean)	100%	115%

Table 9.5: Quantitative results of BLT-DOT experiment.

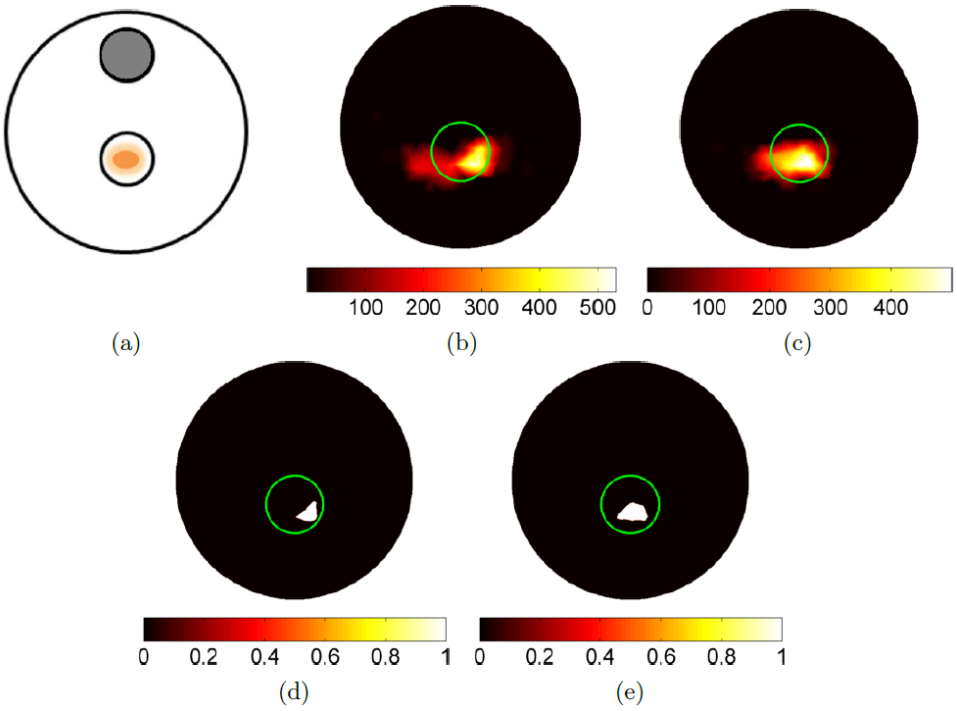

Figure 9.11: Reconstructions of luminescence source distribution in the case where the target is in the lower tunnel with an anomaly in the upper tunnel: (a) schematic of target slice showing anomaly and source positions; (b) BLT reconstruction where the absorber concentration is assumed to be the background level throughout the volume; (c) BLT reconstruction where the absorber concentration is assumed to be that obtained via DOT; (d) and (e) 75% thresholded versions of (b) and (c).

The quantitative results are very similar for both cases; although localisation error judged by the centre-of-mass approach is slightly lower in the prior case this is a tiny (2.5%) difference. The equal max-valued localisation error reveals that despite visually apparent differences, the max-valued node was the same in each case. The total source reconstructed is closer to the mean value in the previous experimental results (figure 9.10) in the homogeneous-assumption case although it is perhaps worth noting that in the equivalent homogeneous experiment in terms of source location (the on-axis experiment at depth 15mm; figure 9.10(f)), the reconstructed source was above 100% and the reconstruced source generally increased with depth indicating that possibly the mean across all depths might not be the most appropriate bench-mark for evaluating the single BLT-DOT result.

Despite the lack of substantial quantitative improvement judged by these metrics, the qualitative image improvement when using the prior information provided by DOT suggests that this is a useful technique to pursue. It is additionally not surprising that in this experiment the homogeneous reconstruction worked reasonably well as the multiple-view approach combined with the source position means that the highest signal for this data set was obtained in the side-views and therefore would be affected relatively little by the presence of the anomaly positioned between the source and the already less sensitive top-view.

To investigate further whether the good quality of the homogeneous reconstruction was due to the enhanced coverage made available by the multi-view set-up and whether it was this that was overcoming the lack of anomaly knowledge in this experiment, a final pair of reconstructions were performed that used only the top-view (non-mirror) data from the above experiment. Again one reconstruction assumed background, homogeneous properties and one utilised the DOT prior information. The results are shown in figure 9.12 and quantitative results are summarised in table 9.6.

It can be seen that in the case where side-view data is not used the effects of making the homogeneous assumption are far more significant. Qualitatively the DOT-prior reconstruction looks quite similar to the multi-view version albeit a little less well-focused whilst the homogeneous reconstruction is no longer recognisable as a small source cluster in the right location,

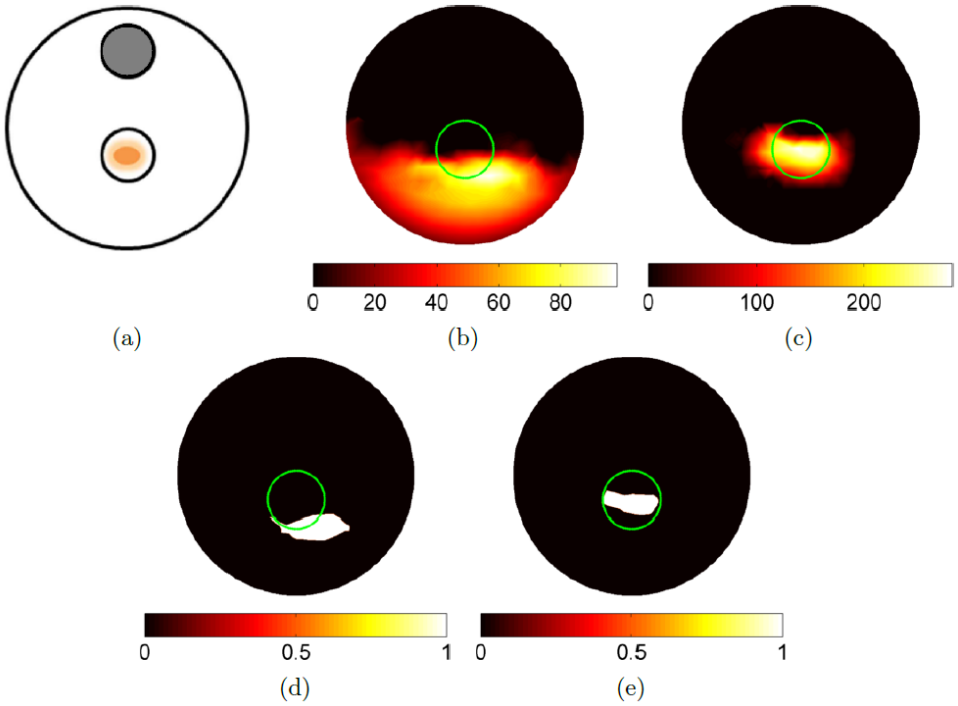

Figure 9.12: Reconstructions of luminescence source distribution in the case where the target is in the lower tunnel with an anomaly in the upper tunnel, with reconstruction performed using top-view (direct) data only: (a) schematic of target slice showing anomaly and source positions; (b) BLT reconstruction where the absorber concentration is assumed to be the background level throughout the volume; (c) BLT reconstruction where the absorber concentration is assumed to be that obtained via DOT; (d) and (e) 75% thresholded versions of (b) and (c).

rather it is very broad and blurred, and somewhat deeper than the true source location.

Top-Only (1-View) Reconstructions	**Homogeneous**	**DOT Prior**
2D Position Error (mm; centre-of-mass)	5.17	0.78
3D Position Error (mm; centre-of-mass)	5.21	0.81
2D Position Error (mm; max-valued node)	3.39	0.56
3D Position Error (mm; max-valued node)	3.57	0.67
Total Source (A.U.)	21525	5476
Total Source (% of BLT experiment mean)	371%	94%

Table 9.6: Quantitative results of BLT-DOT without mirrors.

The quantitative metrics now show that the homogeneous reconstruction also contains around 3 to 4 times the expected total source as well as being clearly in the wrong position. The quantitative accuracy of the DOT-prior top-only reconstruction is actually better than in the 3-view case which could be due to a diminished influence of the DOT artefacts (located towards the sides) although this may merely be due to limitations in the metrics used.

9.4 Discussion

Three sets of experiments have been presented, one where the XPM2 was imaged repeatedly and with multiple sources simultaneously present, one where the homogeneous cylinder was imaged with varied source positions, and one where the heterogeneous cylinder was imaged with combined DOT-BLT data used.

In the XPM2 experiments, the results were mostly positive with images being rendered that were clearly accurately interpretable in qualitative terms showing, as they did, the correct number of sources in roughly the correct positions and with similar sizes in each case. This is an important set of results because it demonstrates the system can perform BLT imaging of a mouse-shaped subject and provide clear, meaningful images.

In terms of quantitative accuracy, the experiment also had some positive results with the relative source strengths between experiments being approximately as expected for the single-source experiments, albeit less so in the multi-source experiment. This suggests the system may perform with some level of quantitative accuracy which could be useful for functional biomedical imaging and should be studied further. Localisation of the sources was also quantitatively good with errors calculated that were all within or near to (up to 1mm away from) the estimated uncertainty on the expected source position. This is encouraging and highlights the need for futher experiments where the source position is more well-known.

Repeats of the XPM2 experiments showed a level of consistency in that, again initially in qualitative terms, the reconstructed images all looked similar and showed the expected sources in the expected positions. In quantitative terms the differences in the total reconstructed source was on the order of 10-15%, providing a preliminary indication of reasonably quantitative reproducibility. Nevertheless, it is unknown where this variability has come from precisely and the variation highlights the need for further studies.

In terms of localisation errors, all sets were quite consistent, with differences observed of less than 1mm. As noted, the spatial characteristics of the source errors in terms of the amount contributed by x, y and z components was quite consistent indicating perhaps a systematic error, possibly in the expected source position. All of the measured locations were very close to or

within the expected uncertainty on the source position. Consistently the largest error was in the z direction (vertical depth) which, as discussed, could be due to some spectral inaccuracy such as is expected with the known issue with the filters (appendix A.2.5). This reinforces the importance of modelling or otherwise fixing this filter issue. It was also noted that there is some unknown uncertainty in the XPM2 emission spectrum which might also cause the depth errors. This can be addressed by imaging more phantoms with better known source characteristics.

The cylinder experiment addressed some of these needs; the source was better understood and its expected position was more precisely known, also a large number of scenarios were shown with the same source in different positions such that the many reconstructed intensities of the source could be meaningfully compared for a better idea of quantitative stability as well as spatial sensitivity and robustness to source position.

This experiment again demonstrated excellent qualitative image production with the single source being clearly visible in all cases and visually in the correct position. The importance of this was accentuated by the comparison with BLI data for the same problems which showed a great deal of variability and was visually far less clear. In some cases with the BLI data it would have appeared for example that there was more than one source where there was not, this is just one demonstration of the danger of interpreting BLI data directly and of the importance and future role of BLT.

Quantitatively, results were also good, with the centre-of-mass based localisation error being less than 1.5mm for all cases considered. This is a lower error than was seen in general in the XPM case which could be due to the more well-known target source position. The total source was within 15% either side of the mean reconstructed value indicating a good level of quantitative stability considering the range of depths, positions, and therefore dominant views in terms of locations of high values in images and on mirrors. One possible reason for the better result here than in the XPM2 experiment is that the wavelengths used were more spaced out and the cylinder optical properties are flatter than those of the XPM2 as a function of wavelength, as such any systematic errors in spectral terms such as those conjectured above owing to the filter issues for example, might have less negative impact. Further studies are required to clarify this

and in any case the filter issue and any others should be addressed. The quantitative results are also shown to be especially good when compared to the range of values visible in BLI images. It is clear that the potential for BLT in the presented system as an improvement over BLI is very high from these results.

The final cylinder experiment in DOT-BLT mode demonstrated the use of DOT data from the system to improve BLT reconstruction. This was clearly achieved in the second reconstruction shown where mirror data was not used. It was less clear whether or not the DOT prior information was useful in the case where mirrors were used, which may be due to the particular scenario, nevertheless qualitatively the produced image was clearer where the prior information was used. As a single experiment this final test can be seen as a proof-of-concept that draws together the work-flow intended for the system and provides a preliminary indication of the value of using DOT to inform BLT.

9.5 Conclusion

This chapter has introduced the methods used in conjunction with the developed imaging system to perform BLT and demonstrated this in phantoms. As in the DOT chapter, the chosen selection of experiments are designed to illustrate the type of experiment that can be carried out, gather insights into the system performance, and demonstrate the developed work-flow in application in the system.

The combination of presented phantom experiments serve to illustrate several achievements. Firstly, it has been possible reconstruct single and multiple sources accurately, with a degree of reproducibility in a mouse-shaped phantom with qualitative clarity and greater quantitative accuracy than is typical with BLI. The cylinder experiments demonstrated sensitivity throughout the mouse-sized volume, showing the importance of the mirror system and multi-view mapping and illustrating robustness of quantitative reconstruction accuracy to source position.

These experiments together provide a strong case to say that when the properties are known in advance, the imaging system can perform accurate BLT and this is far superior as an imaging

technique to BLI. The results are also consistent in terms of localisation errors with other BLT systems under development which typically quote errors in the range of 1-3mm[58, 44, 60, 205].

The final experiment presented showed that DOT could provide useful prior information to BLT. This is a very important proof-of-concept because the aim of the imaging system is to perform multi-modal imaging used in just this way. The experiment showed that DOT could identify departures in terms of optical properties from the background values and provide this data to improve BLT. There is clearly a great deal of future work necessary to push this approach further in other known phantoms starting from a position of total naivety and seeing if the system can cope. Many studies of this type should be carried out to analyse and refine the system using the presented experiment as the starting point.

CHAPTER 10

SMALL ANIMAL IMAGING

10.1 Introduction

A single feasibility study was carried out in which mice were imaged *ex vivo*, with implanted luminescent light sources imitating bioluminescence.

The major achievements of this study were to show that it was possible to image through an in-tact mouse using the DOT trans-illumination scheme (chapter 8), that the investigated protocol for imaging implanted sources (chapter 9) could provide BLI-like data and to image in BLT and DOT modes without apparent cross-talk. The study was carried out early in the course of the project and parameters and methods were not at that stage well-optimised for tomographic reconstruction, as such the results are limited to camera images and corresponding discussions.

10.2 Experiments

10.2.1 Sample Preparation

Four partially shaved mice were sacrificed using increasing levels of CO_2 in accordance with applicable guidelines. A single trigalight source (appendix A.1.2) was then surgically implanted in each of the mice. Each of the trigalights was a different colour and implanted at a different

Mouse	trigalight	Implant location
1	Green	Subcutaneous, back
2	Yellow	Subcutaneous, left arm-pit
3	Orange	Deep, under liver
4	Red	Subcutaneous, near left kidney

Table 10.1: Summary of implant locations for light sources in *ex-vivo* mouse experiment. For trigalight emission spectra refer to appendix A.1.2

anatomical location as indicated in table table 10.1. The mice were identified by ear clipping. A photograph of the mice prior to implantation is shown in fig. 10.1a.

(a) (b)

Figure 10.1: (a) Mice before light-source implantation, (b) mouse 4 after experiment with skin peeled away, showing location of implanted light source.

It was actually intended that the source in mouse 4 be placed deeper than it was. It was surgically implanted above the kidney but imaging suggested it was nearer to the skin surface (sections 10.2.2 and 10.2.3). To clarify the location, the skin of the mouse was peeled back after imaging revealing that the source had indeed *fallen* to the surface (fig. 10.1b); it is assumed that it was not in a well-enough supported position initially and effectively immediately moved to a subcutaneous position following the initial implantation.

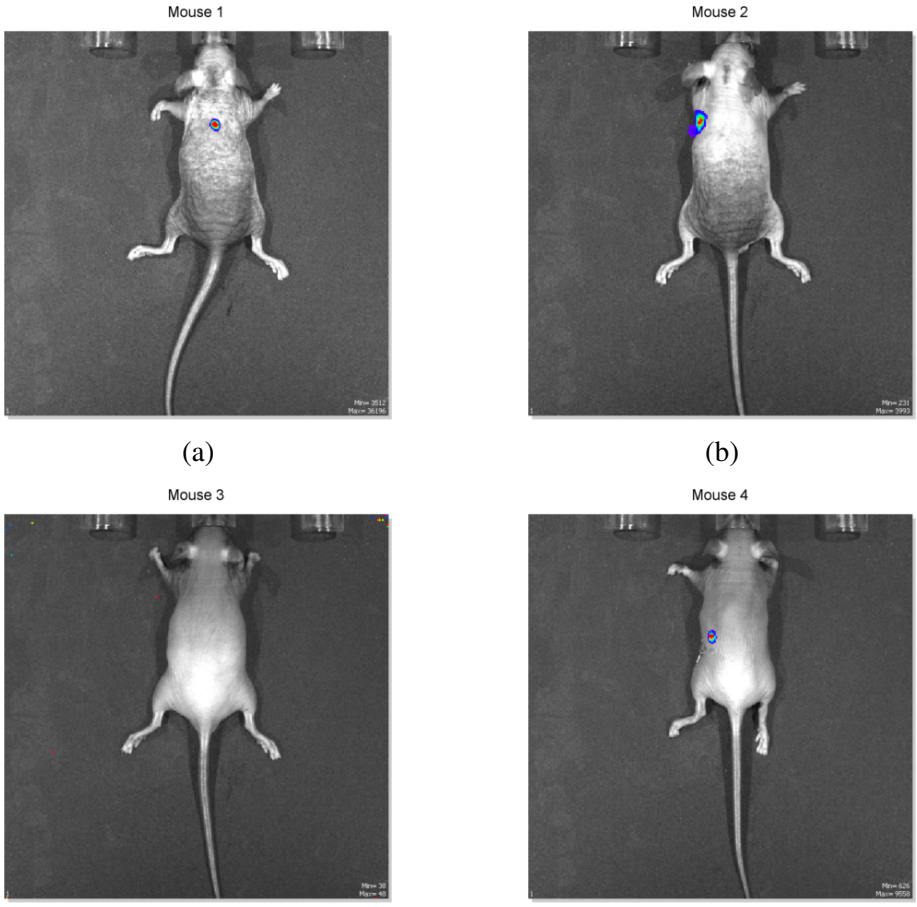

Figure 10.2: Images acquired using the IVIS system. The imaging parameters were (a) 520nm, 1s exposure, 4×4 binning (b) 580nm, 10s exposure, 8×8 binning (c) 600nm, 60s exposure, 8×8 binning (d) 580nm, 2s exposure, 8×8 binning. The maximum and minimum image values are shown in the bottom corner, these read: (a) 3512, 36196, (b) 231, 3993, (c) 38, 48, (d) 626, 9558.

10.2.2 BLI in the IVIS System

Directly following their preparation, each of the mice was imaged in an IVIS Spectrum imaging system (Caliper Life Sciences, now part of Perkin Elmer, the system is described in the paper by Kuo et al.[44]). The purpose of this was to establish signal in a standard commercial imaging system. Both the presence and approximate strength of the signals was of interest. Images were acquired for wavelengths near the peak of each source emission[1] with automatic binning and exposure settings. The images produced are shown in fig. 10.2.

[1] except for mouse 4 for which the trigalight emission spectrum (of the red trigalight; see appendix A.1.2) has several peaks; in this case a peak with similar nominal light output to the other lights was chosen.

Being a commercial system, the precise image processing performed by the system to produce the images in fig. 10.2 is unknown, therefore it is not clear exactly what the numbers in these images correspond to. However, the system does allow access to raw luminescence, background and photograph (back-lit) images. These were processed such that a background-subtracted, exposure corrected, and binning corrected image was obtained for each case. These were then adjusted in software such that the effective binning-level was equivalent to 4×4, and cropped to remove uninteresting data. The images are shown in fig. 10.3.

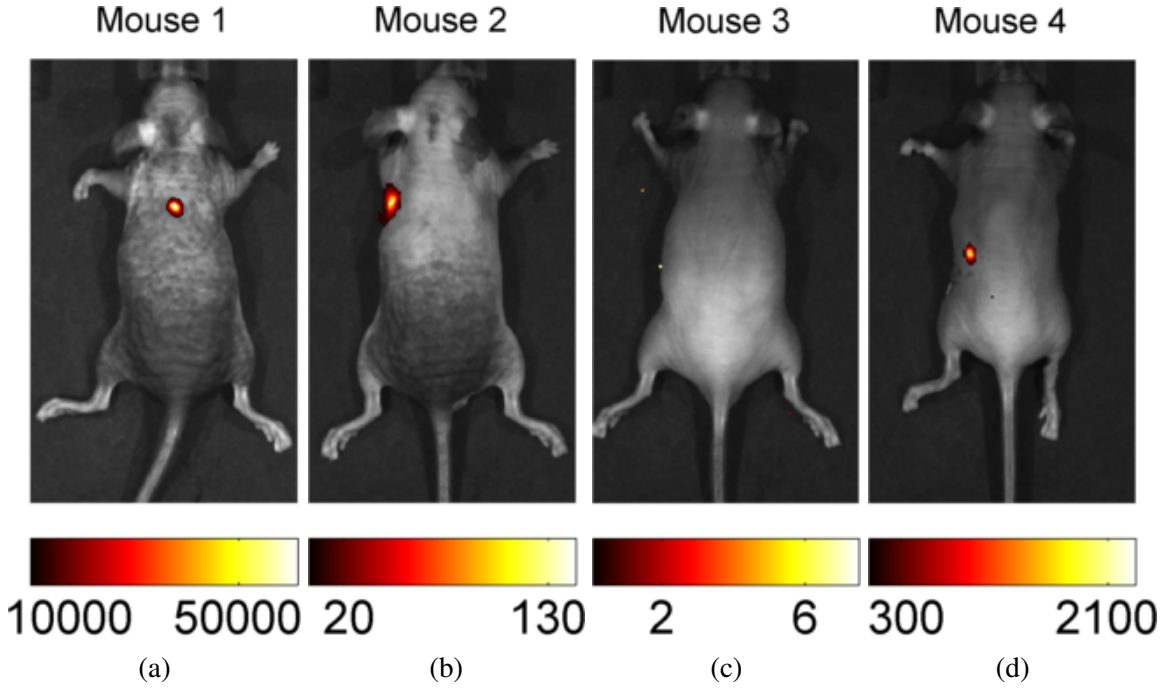

Figure 10.3: Post-processed IVIS images. The effective binning level is now 4×4, with units of counts per second per (binned) pixel. Without greater knowledge of the imaging system hardware it is not possible to process further whilst maintaining certain quantitative meaning. The overlay images are rendered in the same way as those in chapter 9 (for example fig. 9.5) with a ten percent signal threshold for the signal image over the photograph.

It can be seen that the signal is easily identifiable in the case of mice 1, 2 and 4 but absent in the case of mouse 3; it appears that the depth and position of this source render it below the limit of detection of the system. The signal received is strongest in the case of mouse 1, then 4, then 2. These cannot be compared absolutely given the different sources and wavelengths used; the signal difference is due to a combination of this as well as perspective (the emission location in mouse 1 is *square-on* with the camera whereas the others are off to one side), slight

differences in depth, variations in the mouse skin-tones (therefore optical properties), and the filter transmittance and response characteristics of the system.

10.2.3 BLI in the Developed System

The mice were transported to the new imaging system and imaged in luminescence mode. Images were recorded at 6 wavelengths. Resultant images for three sample wavelengths are shown in fig. 10.4. Images were not recorded for mouse 3 because preliminary images indicated that, as expected based on the IVIS image, there was no detectable signal.

Note that read mode 3 and EMCCD mode were used in all cases, with an EM gain index of 0 (equivalent to 10$\times$ (appendix A.2.4)) despite the fact that the noise properties are typically poorer in this mode (appendix A.5). The reason for this is that this work was carried out before the noise characteristics were well studied and understood.

Whilst in the case of mouse 1 the wavelength shown for the IVIS was not used for the presented system, if it is assumed that the nearest wavelength to that shown (580nm) is comparable, it can be seen that for mouse 1 and 4 the signal level received in the presented system was similar to that measured by the IVIS, that is to say, the same order of magnitude. Specific differences will be due to a combination of system hardware, response characteristics and perspective differences. It can be seen that for mouse 2 the signal is relatively a lot higher, this is due to the mirrors providing in this case a more direct view of the light source; it can be seen that the direct view contains no signal due to the overwhelming signal seen on the side view (causing that that is present in the direct view to fall below the 10% threshold for display).

There is clearly a spectrally varying signal, which, given that the images are not corrected for system spectral response or source characteristics, and given the shallow source positions, is clearly strongly due to a combination of the spectral source emission and system response variation, as well as tissue interactions.

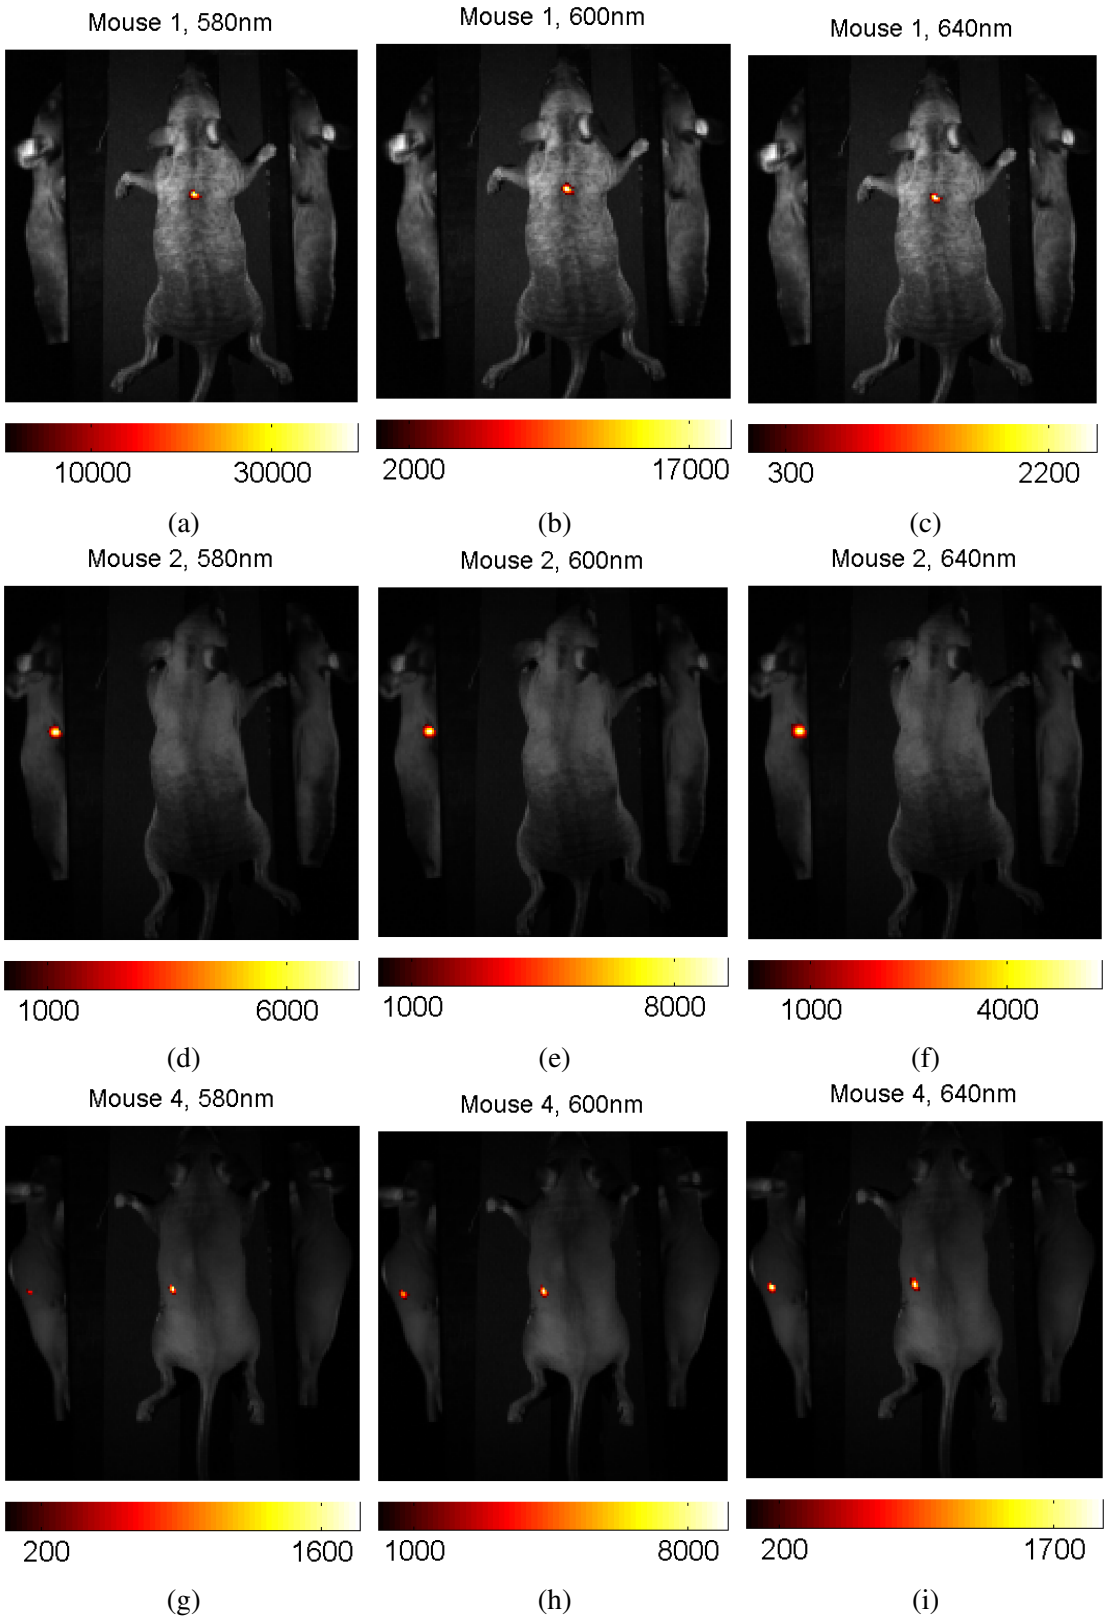

Figure 10.4: Luminescence images acquired using the presented system. A sample of three images is shown for each case where there was signal with wavelengths as indicated. 4×4 binning was used in all cases. The exposures used for all wavelengths were 5, 30, and 10 seconds for mice 1, 2, and 4 respectively.

10.3 NIR Transillumination Imaging

Mouse 4 only was imaged in NIR transillumination mode. This involved using the 6×6 DOT source grid described in chapter 8. Images were acquired for six wavelengths; 680, 700, 750, 770, 810, and 850nm. In all cases the exposure time was 5 seconds, read mode was 3, CCD mode was EMCCD, binning was 4×4, and EM gain index was 100 indicating a multiplication of approximately 37 times (appendix A.2.4). Sample resultant images are shown, in units of multiplied and corrected electrons per second, in figures 10.5 and 10.6 respectively.

As in the case of the presented BLI, the choice of parameters is sub-optimal, in light of the characterisations presented in appendix A, in terms of expected SNR. Indeed this work, which was carried out early in the project, drew attention to the need for these analyses. Accordingly, it can be seen that the images appear noisy by visual inspection with the appearance of a 'salt and pepper' noise. Given that the exposure time was short and the dark noise is expected to be very low for short exposures, and given that the read noise was effectively minimised by the use of EM multiplication, it can be assumed that this is a combination of multiplied shot noise and excess (multiplication-related) noise (see appendix A.5).

The major positive result of this preliminary study is that there is a clearly detectable transmitted signal, which through parameter optimisation might be obtainable with good SNR. Furthermore, there is evidence of the impact of optical heterogeneities. Particularly in the case of images shown for sources 15 and 36, it can be seen that the imaged surface flux is shifted so as to avoid the centre of the mouse body. This may in part be due to an increased thickness of the body at this point, but it may also be due to the presence of the liver which represents a large high absorbing anomaly in this approximate position (chapter 2).

In total, with 36 sources, 6 wavelengths and 5 second exposures, the NIR transillumination part of the mouse 4 experiment took approximately 20 minutes.

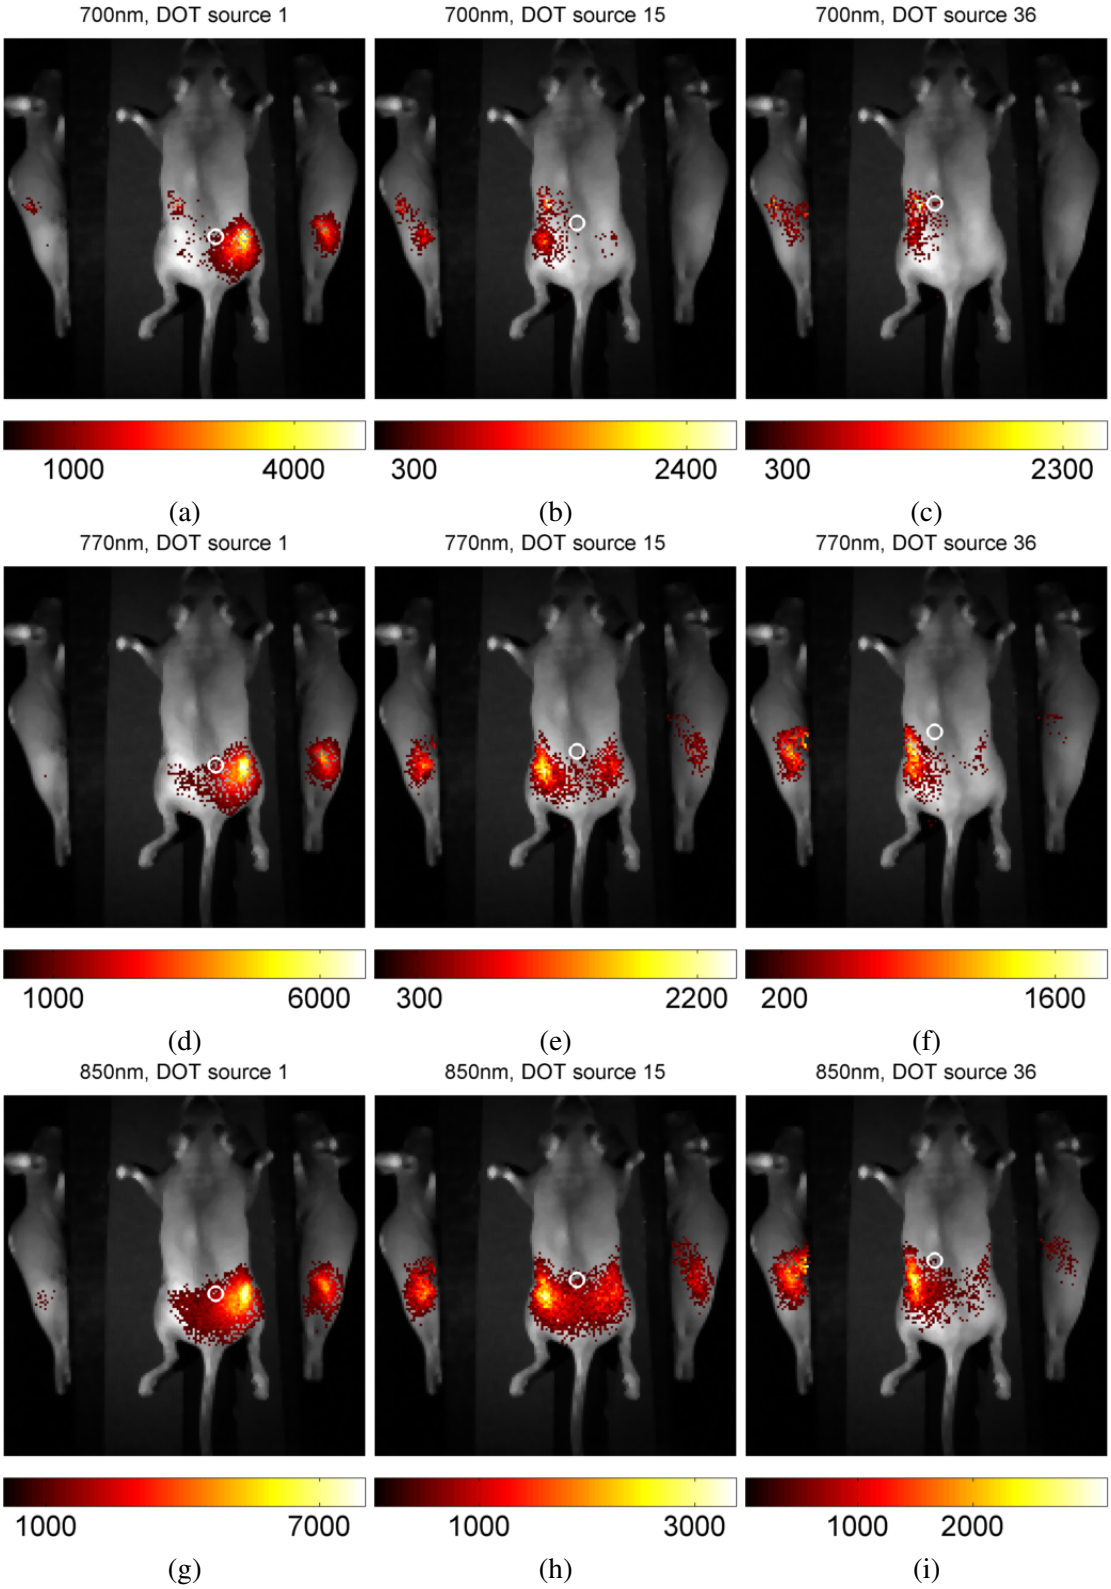

Figure 10.5: NIR Transillumination images for mouse 4. Images are shown for three sources taken from the source grid, with source locations indicated by white circles. Images are shown for 3 wavelengths as indicated. These images are in (multiplied) electrons per second on the CCD and are not corrected for system response or source-to-source variations.

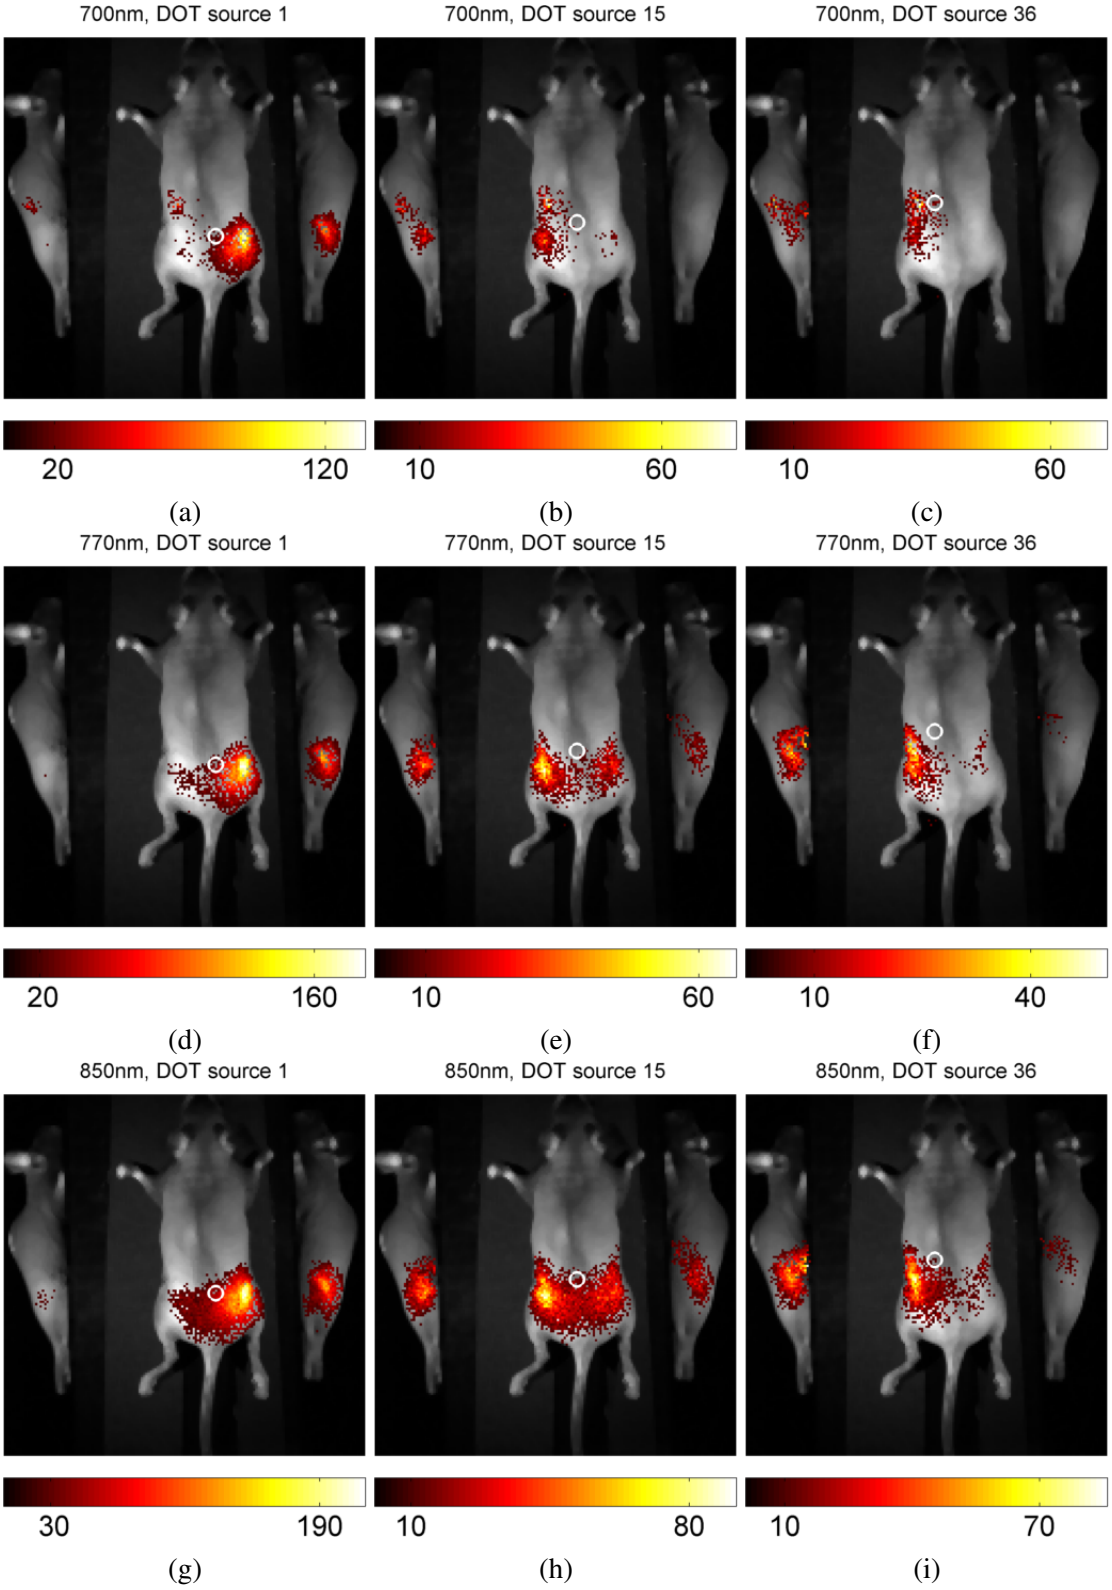

Figure 10.6: NIR Transillumination images for mouse 4, those of fig. 10.5, corrected for EM multiplication; units of electrons per second.

10.4 Surface Capture

Each of the mice was also surface captured, this appeared to work very effectively and samples of the results have already been shown in chapter 6.

10.5 Discussion

Results have been shown from the first and only mouse experiment set carried out in the course of the presented project. This was the only such experiment set due primarily to practical constraints. It was carried out early in the project and as such the work is preliminary with known limitations. Particularly, as stated, the parameters for imaging were not well-optimised. This prompted the work presented in appendix A which provides a deeper understanding of the system, parameter impact and noise characteristics. Despite this there are several useful outcomes from the experiments.

Firstly, luminescence images from the presented system were comparable to those from the IVIS, both quantitatively and qualitatively. This is confounded to some extent by uncertain knowledge of the IVIS hardware but it indicates a broadly similar level of performance to a commercial-grade system which is a positive result.

It can be seen from the images that there is substantial benefit gained by using mirrors for BLI. For instance in the case of mouse 2, the signal seen was dramatically stronger in the mirror view. In the case of mouse 3, the mirror also showed a usable signal. This is an encouraging result as it justifies the use of mirrors and demonstrates the benefit of expanded field-of-view and the ability to obtain scenario-dependent higher signal.

The light source implantation method was successful to some extent but also appears to have been imprecise. In one case, that of the fourth mouse, the target implantation location was not obtained and this resulted in a subcutaneous source rather than a deeper one. There is ongoing work to establish a better implantation protocol and one suggested approach is the use of a large syringe to deposit the source at a known depth from the skin. Despite improvements being

necessary to the protocol, the light sources themselves appeared to be suitable for imitating bioluminescence and relatively easy to work with.

The NIR transillumination experiment showed that it is possible to obtain signal through the whole in-tact mouse body using the developed DOT illumination scheme and provided preliminary evidence for the detection of heterogeneities (superficially; by visual inspection). This is an important proof-of-concept for the system. The signal however was rather weak which highlighted the need for the characterisations and optimisations of the system which have now been performed.

10.6 Conclusion

Together, the animal imaging in this chapter has provided an important preliminary insight into the performance in imaging real animals. It highlighted important areas for improvement which have now been addressed by experiment and research presented in other parts of the thesis, or which have alternatively been indicated as important future directions.

Key achievements are that a similar level of signal to a standard commercial grade system has been demonstrated in luminescence mode and the ability to image, so far with sub-optimal imaging parameters, through a whole animal in DOT mode has been demonstrated.

A further positive observation is that, unlike in the XPM2 DOT study, there is no visible light leakage around the underside of the mouse as it is flush with the stage. Note that this is due to the mouse resting unsupported on the stage in a natural position, representative of how it might rest *in vivo*. This behaviour was anticipated but nevertheless the observation represents an important validation of this idea and it suggests that side-view data is more easily usable without risk of light leakage causing misleading results as was observed in certain phantoms.

Due to a combination of incomplete methods being available at the time of performing this study, sub-optimal parameter selection leading to low signal-to-noise ratio in DOT mode, and imprecision in the implantation method leading to only scenarios where there was no signal or only subcutaneous (i.e. unsuitable for tomography) sources, the study was terminated at the

point of images. In other words, tomography was not possible with this data. However, by returning the results discussed the study is a valuable addition to the thesis.

Future directions, as indicated, now include the development of a better implantation method for the sources. Then it is clear that a large amount of data should be collected with real animals using the newer methods presented in this thesis (the tomographic reconstruction methods, the surface flux reconstruction method, and the automated meshing and surface processing methods) to obtain DOT and BLT data sets and analyse the performance of the system and go on to improve this in this scenario. The current study has provided a foundation for this work and illustrated that such studies are certainly feasible.

CHAPTER 11

CONCLUSION

A novel multi-modal optical imaging system has been developed for performing bioluminescence tomography (BLT) and diffuse optical tomography (DOT) along with optical surface capture of small animals. The aim of the system is to provide a way of making BLT feasible and accurate without the need for other more expensive or time-consuming imaging methods such as CT, PAT or MRI to inform the methods.

In the conceived work-flow, the system is designed to capture the surface of the imaged mouse, build a geometric (FEM mesh) representation of it with which to model light transport inside and outside the volume in association with a model of the imaging system. With this approach the system will capture surface fluxes for transillumination imaging and reconstruct optical properties, i.e. do DOT, and map fluxes for emitted luminescence and perform BLT with this data, using the DOT result as the basis for accurate light modelling as well as to inform on chromophore concentrations that might have biological relevance.

In this way the system would offer a cheaper, therefore more widely attainable and useful alternative to other molecular imaging methods such as nuclear techniques, with all the other advantages of the already widely popular bioluminescence imaging (BLI) but with the addition of quantitative accuracy and 3D spatially-resolved information. Studies could then be performed on animal disease models such as monitoring cancer models with bioluminescently-labelled tumours so as to watch the tumours grow over time or respond to therapy in 3D. In this scenario it can be envisaged that functional DOT could also provide information on the tumour

oxygenation status thus providing a novel source of added useful information content.

At several points the presented work has demonstrated clear justification for the concept of the system. In chapter 7 for example it was shown that BLI will often fail to accurately measure the light leaving the animal surface due to a lack of accounting for perspective and other imaging system image formation considerations. In chapter 9 it was similarly shown that BLI is quantitatively and qualitatively very poor as an indicator of underlying source position and intensity because images show very large qualitative and quantitative variation as a function of source position. Both of these issues have been accounted for within the present system.

In some respects the demonstrable lack of quality of BLI in general, when combined with its wide usage (chapter 3), is an indicator of the wide-held desire in the scientific community to use BLI for its other advantages, showing the importance of these. This is further justification of the importance of developing technology to do more accurate, more meaningful and useful BLT.

The major achievements of the presented work are that the system has been designed and constructed (chapter 5), that it has been refined and understood by a great deal of rigorous characterisation and testing (appendix A), and that methods have been implemented to make the desired imaging work-flow possible.

Several innovations have been made in order to make this possible with good results. Firstly, the surface capture technique introduced in chapter 6 was designed specifically for the presented application and has been shown to work very well and make imaging simultaneously in multiple views and through mirrors possible for the first time with such a technique[3]. This has been shown to capture real animal and phantom surfaces with high accuracy (hundreds of microns of error) and facilitate an, also bespoke, automated meshing routine to provide FEM models in the system work-flow. The surface capture method was developed in conjunction with a specific design of the hardware involving the two projectors pointed at the subject from either side of the system camera, making use of existing resources in the camera and the mirrors and thereby maximising the scope of the surface covered and acquired in this novel design.

Another such innovation is the presented surface flux mapping method for reconstructing

maps of the diffuse light leaving the surface from capture images using a system model (chapter 7). It was shown that this was necessary and that current techniques had limitations which were overcome by the developed technique[2]. There is still important future work to be done in optimising this technique and making it more objective with respect to regularisation. Assuming this can be done effectively this might represent a powerful way of measuring surface fluxes in other domains as well as being integral in the work-flow presented.

Beyond experiments for characterisation and testing of these novel techniques, the majority of the presented work has been demonstrative of the utility of the system in either BLT or DOT mode and has provided preliminary insight into the capabilities of the system in cases of studying well-understood phantoms and mice *ex vivo*.

In chapter 8, experiments were presented where the system was used for spectral DOT of phantoms. Little work has been done on spectral, CW, intrinsic contrast (as opposed to fluorescence-enhanced) DOT in the small animal imaging domain and as such this work represents a novel application of the established imaging method, the expected performance therefore could not have been anticipated without experiment. The preliminary experiments showed that DOT was possible in the designed work-flow using surface capture, automated meshing and the novel illumination scheme proposed with the system (chapter 5). Encouraging results were found when imaging a homogeneous mouse-shaped phantom where quantitative recovery of absorber parameters was achieved in some cases, though this was also shown to vary with the initial guess of the properties and requires further studies. In spatial DOT studies of cylindrical and cuboid phantoms including anomalies there was evidence that the system could identify the heterogeneous inclusions but quantitative accuracy was mixed. Establishing the reasons for the mixed performance and improving the system accuracy is now a topic for future work where experiments can be carried out that build on the presented demonstrations.

In chapter 9 results of BLT using the system were presented for cases of imaging the mouse-shaped and cylindrical phantoms. This chapter provided a demonstration of the BLT work-flow in action and showed a great many positive results. In all cases qualitative accuracy was obtained in imaging in that luminescent sources, in some cases multiple sources simultaneously, in other

cases identical sources at many different positions, could be identified clearly in reconstructed images. Quantitative stability as a function of source position in the cylindrical experiment and as a function of repeated runs in the XPM2 experiment was shown to be good; a huge improvement over equivalent BLI quantification. This chapter also demonstrated a positive impact of using DOT as a precursor to BLT. In the experiment reported, the cylinder phantom with an absorbing inclusion was shown to be more accurately imaged in BLT mode when the DOT prior was used.

Chapter 10 showed the application of some aspects of the system to small animal imaging, though due to practical constraints the amount of data obtained and the extent of the studies was limited. Nevertheless, it was demonstrated that the transillumination method used to obtain DOT data, as used in the phantom case for this, was feasible in the animal imaging case. That is to say, it was shown that signal could be obtained from the DOT sources transmitted through the real animal. It was also demonstrated in this chapter that the level of signal measurable was comparable to that obtained on a commercial system in luminescence mode.

Now that the system and general work-flow has been established and demonstrated, there are a great many opportunities for future study in either improving, testing, or utilising the new technique. These have been highlighted as important owing to particular results that have been obtained or limitations or opportunities that have otherwise been recognised in the course of the presented work.

A current limitation to the studies presented is that they have mostly been limited to simple phantoms. This was by design for the reasons laid out in appendix A and was a logical and important starting point. It is also the case that on many occasions it has been indicated that further similar phantom studies are still required particularly in larger studies designed to isolate variables and thereby better understand and optimise performance of aspects of the system. The limitation associated with this approach is that the phantoms are unrealistically simple. To gain better insight into how the system will behave in the scenario of imaging a real mouse with real bioluminescence, phantom complexity should be increased in a controlled fashion such that the problem becomes more realistic but is still well-understood. Only in this way can the system be

fully understood, evaluated and validated.

Examples of future directions with more complex phantom studies include those where the diffusion equation may hold less well, as might be expected in some parts of the small animal (chapter 2). Phantoms containing more realistic heterogeneity indicative of organ structures and real mouse anatomy are also important. From the point of view of spectral DOT evaluation it is important that phantoms also contain multiple known chromophores so that the proposed method of spectral quantification, and effective un-mixing, of multiple contributors can be tested. For both BLT and DOT evaluation in a more realistic context it would also be better to use real chromophores and real bioluminescent reporters. From the results presented here no direct data has yet been obtained for these cases. This is an important future step.

Assuming that the system has been further characterised with more advanced phantoms. It is also important that it be applied blindly to samples for which there is no other source of prior knowledge (other than for validation) otherwise the starting condition is always subject to bias.

In general there are still several opportunities for parameter optimisation in the system. For example there is an open question of how much signal is actually needed (therefore what exposure and experimental time is required to obtain it) in order for accurate DOT or BLT to be performed. This is a question that can be answered in further phantom studies. Another similar question is concerned with the optimal wavelengths for imaging. These are complex problems because the answer will depend to some extent on the particular imaging conditions, thus there is great scope for future investigation.

The system geometry may also be optimised further. It has been shown that the system works in the present configuration, but it may not be optimal in terms of mirror placement, camera and projector positions for surface capture and imaging data acquisition. A benefit of having a full system model by virtue of having the FEM code (NIRFAST) in conjunction with the developed free-space propagation code (chapter 7) is that new configurations could conceivably be tested in simulations. Therefore it is plausible that system designs could be trialled and optimised in an automated, fast and objective fashion. This might prove useful.

A further opportunity presented by the full system model is the solving for DOT and BLT

reconstructions directly from image data. This would follow the work of Ripoll et al.[189] applied to fluorescence tomography. This would involve re-formulating the inverse problem with a combined sensitivity matrix (the free-space **T** matrix multiplied by the FEM Jacobian). This is conceptually a trivial extension to the BLT problem because it would remain linear. It may be more of a challenge for DOT because it would remain non-linear. In both cases it is worthy of study because it would reduce the number of optimisation problems requiring regularisation and introducing errors from 2 to 1.

A difficulty that could arise with using the full system model is that the sensitivity matrix could be larger, this is already a problem to some extent because reconstruction time and memory complexity is high. Styles et al.[8] have recently presented the first application (in simulation) of a matrix-compression technique involving random projection of data onto a new basis to reduce the BLT problem complexity by reducing the size of the multi-spectral Jacobian and data set without loss of information. This is a promising approach and deserves attention. Importantly, it can be validated with the presented system.

Another area with great scope for further investigation is the DOT source. The current arrangement has been shown to work quite effectively but is also relatively slow and does not make the best use of the wide-field illumination potential of the projector system[199]. There is an opportunity for further work in optimising the use of input source patterns for signal and reconstruction quality.

A limitation of the currently demonstrated DOT method is that it neglects the mirror data available from the system. The main reason for this is the difficulty of having such a large dynamic range owing to varying path lengths between mirror and main views. One possible solution to this might be to include large planar neutral density filters or spatially variable filters at either side of the animal in a similar design to the spectral filter design of Wang[147]. This simple hardware addition might lower the dynamic range between views and allow them to be used together more easily in DOT mode.

The small animal imaging experiment, as well as general consideration of the system requirements in the long-term, show the need for several additions to the system to practically

accomadate mice in the system. One such addition is some form of anaesthesia system, such as that used in the IVIS system[44], this is needed to anaesthetise living animals to keep them still during imaging. Similarly, some form of thermal management system is needed to maintain and regulate the animal temperature, this is important to avoid unwanted alteration of the animal physiology which could harm the animal or confound imaging results, and also to avoid modification of the bioluminescent output of the reporter which will change with temperature[206]. These features must be considered before the system can be used in *in vivo* imaging.

When considering system components, it was noted that better hardware could be bought for the system at higher cost. The presented work has demonstrated the development of novel and effective methods that are independent of the particular hardware used to some extent but the system hardware could be upgraded to improve SNR and this could improve efficiency and accuracy. For instance, the latest version of the IVIS imaging system[173] is approximately 16 times more sensitive than the presented system by virtue of some simple hardware features alone; the camera sensor area is four times larger than the one used here, the pass-bands of the bandpass filters used have approximately double the FWHM and double the peak-efficiency. Thus exposure times can be reduced sixteen-fold without loss in performance. The reason that more advanced hardware was not utilised here is the high cost, but it is worth consideration when moving towards imaging *in vivo* where time is more important.

Also considering experimental efficiency, a future direction is the pursuit of imaging in DOT and BLT modes simultaneously. This would save time because imaging could be parallelised and filter changing would not be required. This could be performed by the addition of another camera and filter wheel into the system. It has been found in preliminary DOT acquisitions that any residual luminescence, so long as it is constant, visible through DOT-wavelength filters can be effectively removed via the dark subtraction already included. Furthermore this is not such an issue given the low-light in luminescence. The reverse problem may present a challenge, i.e. when projecting DOT sources some visible light may be seen owing to the source through the luminescence filters even though the DOT light is filtered at source to remove most of the visible light emitted. This could be more of an issue owing to the low-light in the luminescence

mode. As to whether this practically will cause problems and as to how to address any that do arise, this is subject for future experiments and investigations.

When this project began there was a camera in a box. In three years, this has become the first experimental bioluminescence tomography system in the UK and the first combined bioluminescence tomography, spectral diffuse optical tomography and optical surface capture system in the world. The opportunities for improvement and future work are many as will be the opportunities for performing novel science as the system develops further. It is hoped that work on the system and the technique will continue and that biomedical science will benefit from it.

APPENDIX A

CHARACTERISATIONS AND CALIBRATIONS

This appendix reports upon investigations that were carried out in order to characterise and calibrate the presented imaging system and acquired data. These investigations are supported by thoroughly-researched explanations of the workings of various system components and materials.

First, several tools are introduced that are used for system tests and evaluations. These are physical phantoms, reflectance standards, and self-sustained light sources. A proper understanding of these components is important to the accurate interpretation of the results of experiments in which they feature.

Technical parameters of the system are then introduced and measured including the detection system spectral response, image digitization functions, gain mechanisms and noise characteristics. These all must be taken into account when working towards obtaining accurate quantitative results from the system.

Geometric calibration methods are discussed in detail which support the full system model needed for surface capture (chapter 6) and freespace light modelling (chapter 7), again integral to accurate and/or quantitative reconstructions; the goal of the system.

A.1 Introduction to Materials for Characterisation

A.1.1 Solid Physical Tissue Phantoms

An important tool in characterising the system, particularly in evaluating and understanding imaging performance, are tissue phantoms. These are objects that in some relevant ways are representative of the real biological entity of interest, i.e. small animals in this case. The role of phantoms is to act as well-understood imaging subjects for which the target image can therefore be well-defined, thus allowing a comparison between the image that is really obtained and the ideal image. This then provides an opportunity to derive quantitative metrics of system performance. Ideally, the phantom should be perfectly tissue-like, though in practice abstractions have to be made in order to obtain a well-known phantom. It also is important as part of the evaluation process to start with very simple phantoms and then step up levels of complexity, in this way variables can be isolated and it can become clear what aspect of complexity underlies what level of performance, as such the system limits can be obtained and directions for improvement can be identified.

Phantoms can by physical objects, used in practical experiments, or virtual phantoms used in simulations. In the presented work the emphasis is on practical experimentation and as such presented results mostly feature physical phantoms.

Several different materials have been investigated in the field of biomedical optics for the creation of phantoms. Types of phantom include liquid phantoms which often comprise some source of image contrast, such as a tube containing blood or ink, suspended in a bath of background liquid with tissue-like properties. Liquids such as milk or solutions of intralipid (a commercially available, well-characterised lipid formulation), titanium dioxide particles and polymer microspheres (effective scatterers) are commonly used. Another type of phantom is semi-solid phantoms which can be made along the same lines but using a base material of agar, gelatin or cast silicone rubber. In some cases these can have tissue-like mechanical, electrical or magnetic properties in addition to similar optical properties.

The types of phantom discussed so far are very useful but tend to degrade over time and

are not easily stored or transported. They also must generally be made in-house and therefore require some expertise and need to be characterised after making. The optical properties can also be difficult to predict.

In the presented work another type of phantom, those made of solid plastics, are used. The choice to use solid phantoms was based on their long-lasting robustness (important for repeat studies and developmental testing) their fixed and potentially complex geometry (being more realistic for surface capture and practical considerations) and their availability from third party manufacturers that provide certain guarantees on their properties and a level of geometric precision aiding the well-known nature and numerical modelling of the phantoms.

The three phantoms used in the majority of the presented work are introduced in the remainder of this section. These are a cylindrical and a cuboid phantom each with optional absorbing and luminescence inclusions to create optical heterogeneity, and a mouse shaped phantom containing fixed sources. Later in this appendix (section A.7) characterisation experiments are described in which the precise properties of the phantoms were investigated. This work was particularly important because, as indicated in the following subsections, the properties were not all well-known from previously available data.

Cylindrical Phantom

The cylindrical phantom (INO, Quebec, Canada), shown in fig. A.1, is 25mm in diameter and 50mm in length; approximately the same size as a mouse. The material is a solid plastic with spatially homogeneous but spectrally varying absorption and scattering properties that were characterised by the manufacturer for a range of visible and NIR wavelengths (appendix A.7).

The phantom body contains two 6mm diameter tunnels at depths of 5mm and 15mm from the surface. Rods with different optical properties can be inserted into the tunnels to create optical anomalies, representing simplified, well-known organs or tumours for example, or background-matching rods can be inserted to make the cylinder homogeneous. Bioluminescence can be modelled by placing a light source some way along a tunnel enclosed between rods.

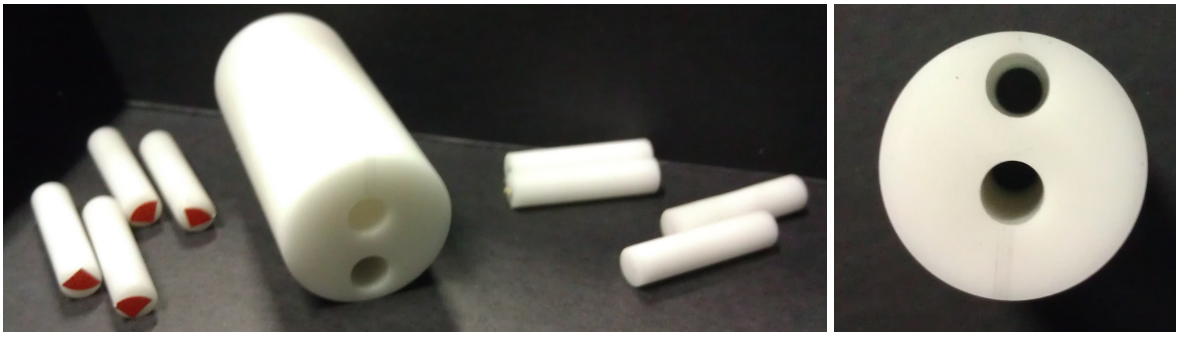

(a) (b)

Figure A.1: Photos of 25mm diameter by 50mm length cylinder phantom; (a) the phantom with the half-length, 6mm diameter inclusion rods, those with red stickers on the end are double-absorption compared to the background whilst those without match the background (dubbed "blanks"); (b) the cylinder seen on end showing the tunnels clearly.

This phantom is used because it has a number of advantages. Firstly it is flexible in that optical anomalies and/or light sources can be included at different positions making it suitable as a simple test subject for DOT and BLT and combined DOT-BLT imaging. Secondly, because inclusion positions can be varied by rotation of the whole cylinder which can result in many anomaly or source positions without effectively changing the bulk position of the phantom (as is done in chapter 7 and chapter 9).

XPM-2 Mouse Phantom

The XPM-2 (Caliper Life Sciences), shown in fig. A.2, is a nominally homogeneous bioluminescence phantom that is approximately $40 \times 20 \times 75$mm. The mould used to create the phantom was cast using a real mouse so the shape of the phantom is highly realistic[44]. Two light sources exist inside in fixed positions. These are created by the tips of optical fibres that are coupled to external LEDs (in a box attached to the phantom). The manufacturer provides a graph of the phantom optical properties, which are similar in spectral shape to tissue spectra but quite low in absolute terms (appendix A.7). The spectral emission of the light sources is also provided, shown in Figure A.2, this is similar to the spectral profile of luciferase (section 3.3.1).

The manufacturer also provides some quantitative data regarding the sources that is shown in table A.1.

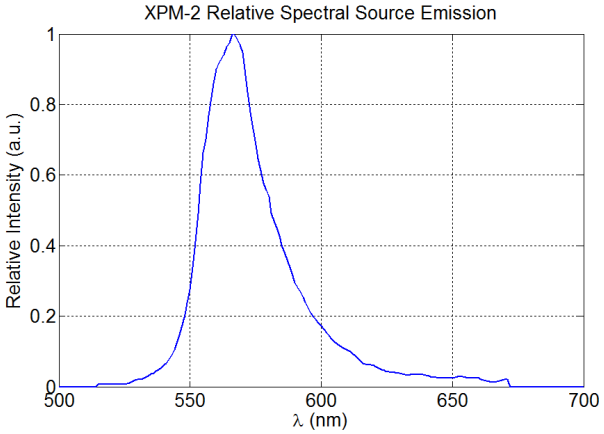

(a) (b)

Figure A.2: Photo of the XPM-2 and spectral emission of internal light sources. Note that this data was only available on a graph that was small in size and therefore difficult to read accurately. To improve this the graph was scanned at high resolution using a scanner and interpreted using MATLAB tools to select points on the curve and re-build it rather than taking values by visual inspection.

Source	Power (photons/s)	D_{Dorsal} (mm)	$D_{Ventral}$ (mm)	D_{Axial} (mm)	$D_{Lateral}$ (mm)
A	8.09e10	8.38	10.96	24.4	18.69
B	8.40e10	16.58	15.47	39.45	10.77

Table A.1: XPM-2 source output data provided by the manufacturer, in terms of total output power and position. D_{Dorsal} is the depth downwards from the top of the phantom $D_{Ventral}$ is the depth upwards from the bottom of the phantom, D_{Axial} is the depth forwards from the back of the phantom, $D_{Lateral}$ is the depth across from the right of the phantom when seen from the back.

The XPM-2 has a number of advantages. Being of highly realistic geometry, it is very appropriate for surface capture evaluation (chapter 6) and investigations of the effects of the complex animal shape on other modalities. The two internal sources can be enabled individually or together providing the opportunity to evaluate the robustness of BLT in scenarios of increasing complexity.

Block Phantom

The block shaped phantom was provided by colleagues at Dartmouth college. The original phantom was cut in half to create two approximately identical phantoms of half the original length. The resulting phantoms are cuboid in shape and measure $33 \times 26 \times 40$mm. One of

the phantoms was drilled to create two cylindrical tunnels, at depths of 5 and 15mm along the central axis (fig. A.3), 6mm in diameter, used to incorporate inclusions as in the cylinder phantom. The other half of the phantom was left in-tact (and homogeneous) so as to serve as a reference. The homogeneous block phantom is discussed in the remainder of this section.

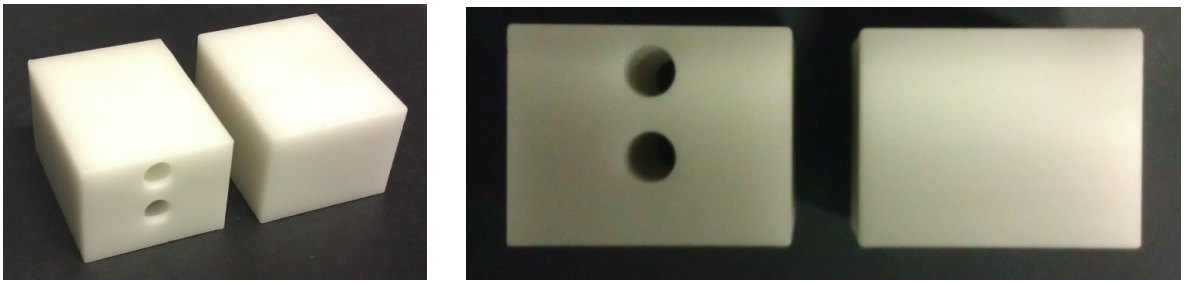

(a) (b)

Figure A.3: Photographs of homogeneous and drilled block phantoms. It can be seen that the cut phantoms are the same size and that one has tunnels passing through it whilst the other is homogeneous.

Optical properties for the phantom measured by the original manufacturer were provided (INO, Quebec, Canada), though it has been several years since the original manufacturing of this phantom.

As with the cylinder, an advantage of the block is that various inclusions can be incorporated at a variety of positions (particularly when rotations are used). It is also useful to have a homogeneous reference version. Furthermore the simple shape of the phantom is easy to model. However, the thickness of the block is unrealistically challenging being almost double the thickness of a normal mouse.

A.1.2 Self-Sustained Light Sources

For characterising aspects of the system involving BLI and BLT, a well-understood analogue to a bioluminescence source is needed. Tritium-based light sources (Trigalight series, mb-microtec, Niederwangen, Switzerland) were found to be appropriate. These have been used in BLT characterisation experiments by others such as Cao et al.[171]. The sources comprise a beta-emitting gas encapsulated within a small ($\approx$ 2.5mm length and 1mm diameter) cylindrical

Name	Peak λ (nm)	μJoule	μWatt	μJ/cm^2	μW/cm^2	Photons/cm^2/s
Orange III	606	9.3097E-5	1.2884E-5	7.4084E-2	1.0252E-2	3.1513E10
Yellow II	572	7.2175E-5	1.2587E-5	5.7435E-2	1.0017E-2	2.9376E10
Green I	528	4.4236E-5	5.3393E-6	3.5202E-2	4.2489E-3	1.1532E10
Red IV	627	1.2246E-5	8.6607E-6	9.7452E-3	6.8920E-3	2.1999E10

Table A.2: Trigalight source data provided by the manufacturer. How these values were measured is unknown, as is therefore the precise physical meaning of some of the values (the units is all that was provided), nevertheless they are included for completeness.

glass vessel with a phosphorescent coating. When bombarded with the beta particles the coating emits light. The sources have a half-life of approximately 10 years meaning that they are highly stable over the duration of imaging experiments and in addition they require no external power source or regulation.

The sources are available in a range of colours, in this work orange, red, yellow and green sources are used. The emission profiles as provided by the manufacturer are shown in fig. A.4. It can be seen that the orange, yellow and green sources have Gaussian-like profiles which are similar to the luciferase emission spectrum. Additional quantitative data is shown in table A.2.

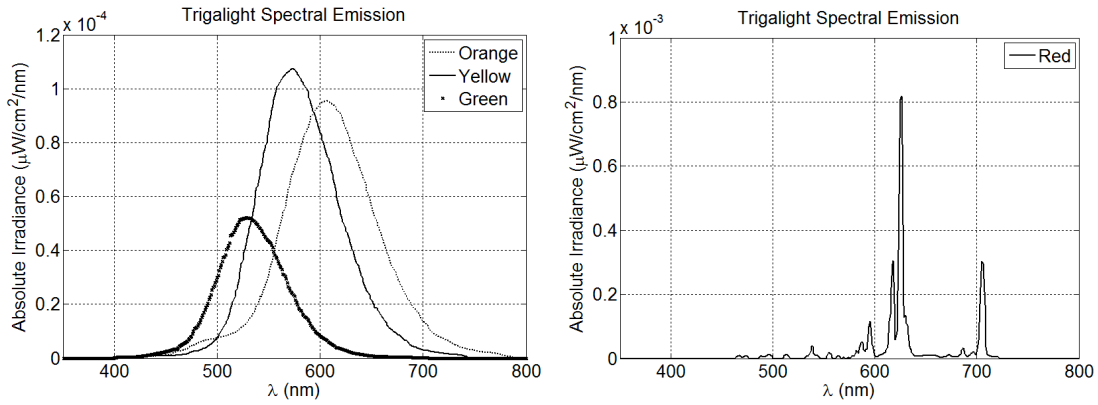

(a) (b)

Figure A.4: Tritium-based trigalight source spectral light emissions with units of absolute irradiance as read from the manufacturer documentation (using the method described in fig. A.2).

In addition to the trigalights described above, another set of self-sustained light sources based on the exactly the same principle, called betalights (Mannings Fishing Tackle, Hampshire, UK) were used in some experiments. The key difference is that the betalights were larger, being 6mm in length and 2mm in diameter.

A.1.3 Reflectance Standards

Spectralon reflectance standards (Labsphere, North Sutton, NH, USA) are used in some calibration experiments. The standards, shown in fig. A.5 are useful because they have a reflectance profile that is very flat[1] and uniform spectrally and also spatially invariant. Furthermore they have negligible specular reflectance compared to diffuse reflectance. As such one of these illuminated by a light source can be expected to create a new light source that is diffuse and has the same spectral profile as the original.

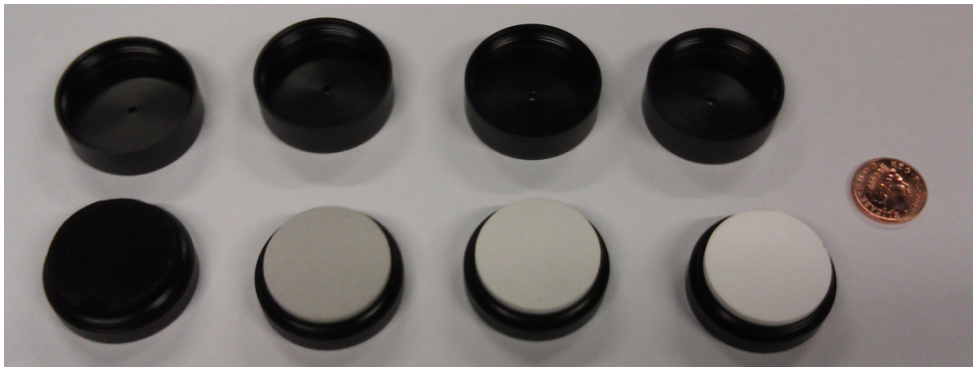

Figure A.5: Photograph of reflectance standards with a coin as a size reference, from left-to-right these are 2, 50, 75, and 99 percent reflective standards.

A.2 Detection System Characterisation

A.2.1 Spectral Instrument Response

The relative spectral response of the whole detection system was measured in a set of experiments in which the (unfiltered) tungsten halogen source was mounted above the stage and set incident upon a 99% reflectance standard on the sample stage. The source was far from the reflectance standard and therefore created an effectively even, diffuse illumination. This was imaged 50 times through each of the bandpass filters and the mean reflected values (in a region of interest containing approximately 1000 pixels) were divided by the known source spectrum

[1] Exactly how flat depends on the particular standard type (i.e. reflectance level) in question, though the characteristics are well-known and therefore can be divided out if they are not flat enough.

to obtain the system response function. To be as representative of the typical use of the system as possible, the filters were loaded into the filter wheel in the standard way. Given that only 6 filters can be loaded into the filter wheel at any one time, the filters were loaded in sets of 5 along with a single common reference filter included to check for (and provide a calibration factor to account for) any inter-experiment drift in the light source intensity. The reference filter was additionally imaged at the start and end of each imaging session (i.e. for each test set in the wheel) to test for any change in the measurement within the time-frame of each experiment- no significant change was observed.

Additionally, the emission of one of the orange trigalight sources multiplied by the system response was measured together as a single quantity by imaging the source directly on the sample stage through each of the filters used in one BLT study (section 9.3.3). Figure A.6 shows the measured spectral response and the measured luminescence source-system spectral response.

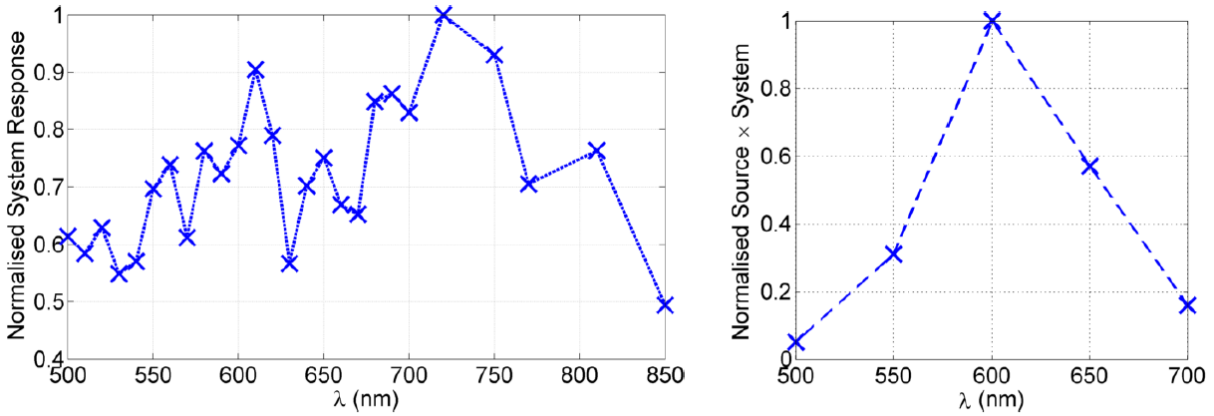

Figure A.6: Normalised spectral system response functions: (a) system response; (b) system response and luminescent source. Standard deviation bars are not shown because they are too small to interpret on the current scale; the maximum standard deviation across all data points was 0.013.

The system response is more variable than might be expected and non-smooth since the dominant factor is the filter transmittance which is somewhat variable between filters in terms of both FWHM and peak height.

A.2.2 Digitizer Conversion Factors

In order to calculate the number of electron-hole pairs created in the detector due to light hitting the CCD, the camera conversion factor must be known that relate this quantity to digitized *counts*. The ImagEM-1K specifications indicate a typical value of 6.3 in EMCCD mode and 0.6 in normal mode. An experiment was performed to test these values for the camera by imaging the same input intensity, provided by the tungsten halogen light source (section 5.3.8) directed onto a reflectance standard, in normal and in EMCCD mode and comparing the resultant measurements. The imaging was performed with an ND4.0 neutral density filter placed in the filter wheel to prevent saturation of the camera, and minimum exposure time was used. The readout mode was 1 and the binning level was 4×4 (the most commonly used imaging mode). Ten images were acquired and averaged through the stack. The averaged images were then processed by offset-subtraction, exposure-correction and the taking of a common region of interest upon the reflectance standard in which the mean and standard deviation of spatial values was recorded. These were then scaled by the gain ($10\times$ in EMCCD mode, 1 in normal mode) and the nominal conversion factors. If the conversion factors are correct then the result would be expected to be the same in both cases.

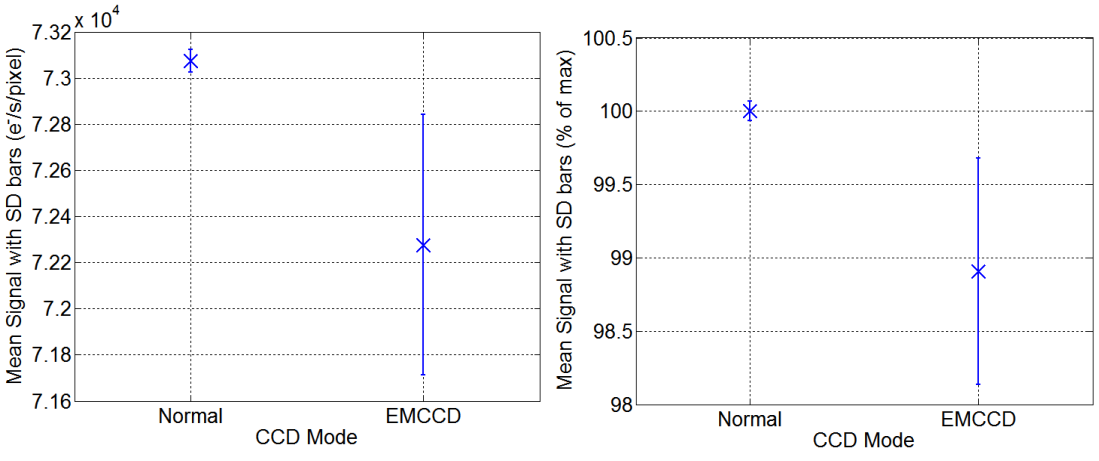

(a) (b)

Figure A.7: Comparison of mode-corrected intensity measured using EMCCD mode and normal CCD mode, serving as a test of the conversion factors.

Figure A.7 shows results. The mean measurements are within just over 1 percent of one

another. The standard deviation is markedly higher in the EMCCD case which is assumed to be owing to the added noise of this mode (appendix A.5). Particularly given the low number of samples it is concluded that there is not currently a reason to suspect that the conversion factors might be incorrect but further testing via repeats of this experiment would be needed to confirm this beyond a doubt.

A.2.3 Camera Digitizer Offsets

When an image is read out from the camera, the retrieved counts are the addition of those due to genuine signal (and noise) and a characteristic digitizer offset i.e. a fixed pattern in the image that is independent of the signal. The camera digitizer offset must be known so that it can be subtracted from images to identify the signal.

An experiment was performed to measure the digitizer offset. The lens cap was placed on the camera so as to remove any external source of counts (other than intrinsic noise) and 100 images were acquired in each combination of CCD, readout and binning modes possible. The different modes were investigated separately because preliminary studies suggested a difference in the offset image between modes. Images were acquired at the minimum exposure time possible in each of the modes to minimise the noise (appendix A.5).

It was found that the mean offset image was clearly different in each of the modes, for example shown in fig. A.8 are two of the images which show very different spatial patterns of background intensity. Owing to the characteristic difference between modes, an offset library was compiled containing the average of each of these images and several more repeated sets for each of the mode combinations. This allows the appropriate library image to be recalled and subtracted at image processing.

Given the proposed method of subtraction of a library offset image, it was important to test that the image was reproducible between different time points. As such, 30 further repeats of each data set were acquired every day for the following 4 days after the original experiment, the order of these was randomized to ensure fairness. The mean offset image (itself a mean of the 30 repeats in each case) of these experiments was then subtracted from each offset image. The

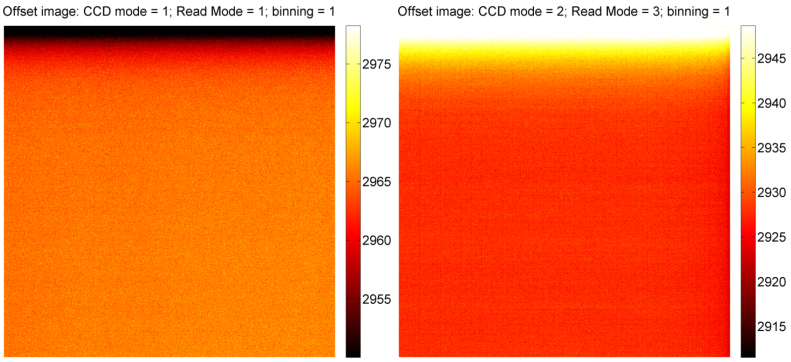

(a) (b)

Figure A.8: Digitizer offset images measured for (a) normal CCD mode, read mode 1 and 1×1 binning, and (b) EMCCD mode, read mode 3, and 1×1 binning. It can be seen that the structure is very different (it actually appears spatially inverted) and the values are quantitatively different.

mean difference across pixels is shown in fig. A.9 along with standard deviation bars.

It can be seen that the mean per-pixel difference values are within 10 counts other than in two specific cases. It can also be seen that the standard deviations amongst the counts is small compared to the mean offset. This is interesting because it suggests that the difference is a systematic offset applied to all pixels approximately evenly. As a tentative possible explanation, it was noted that although the camera was switched on and set to cool and left for an hour before imaging in all cases (such that the temperature was nominally stabilised and appeared stabilised prior to imaging) there was a clear difference in the pattern of changes in temperature over the time period whilst images were acquired. It is possible that this could have given rise to a thermal induced offset; the ImagEM boasts a good thermally-induced charge uniformity[172] and so this might manifest as an approximately flat change in image intensity as observed. The thermal changes also may have been unusually high in these experiments creating an unlikely worst-case because the ordering of the modes used was randomised and it has been observed that changing modes (particularly read mode and CCD mode) appears to cause a small but sudden temperature shift which is less usual in standard imaging runs because the mode is changed relatively rarely.

For the imaging mode that is used in sensitive applications (CCD 1; read mode 1; binning 4; sections 2.5 and 2.6 and chapter 7), the mean difference was within 10 counts either side of

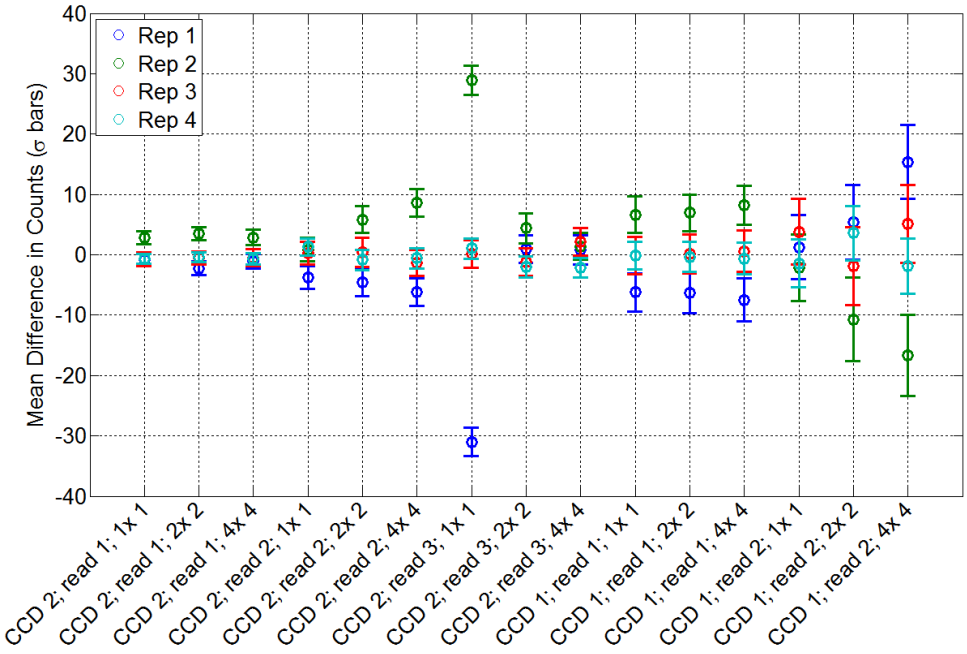

Figure A.9: Analysis of 4 repeats of digitizer offset image estimation (from runs of 30 repeats in each mode acquired on different days). Note that the values for CCD mode 2 (EMCCD) should be divided by 10 to get the non-multiplied value because 10× EM gain (the minimum value) was used in this mode.

the mean and the standard deviation was less than 5 counts, this puts the effect at a similar level to the read noise (approximately 10 counts in this mode; appendix A.5).

Given the generally low variability between recovered offset images, it is concluded that it is appropriate to use an acquired library of offset images to subtract the spatially varying background from images containing signal to obtain the signal alone. The observation that there appears to be a uniform shift in the offset suggests that a per-image background subtraction might also be beneficial but that owing to the apparent uniformity this might be done using combined signal and background images so long as a background region can be defined from which to extract a background level.

It may be beneficial but would be very time-consuming to further categorise the temperature-induced effects and the temperature change-inducing effects of the mode changes in order to validate the current ideas. However, because this would be very costly and the impact appears to be small from the current data this is not currently explored.

A.2.4 Electron-Multiplying Gain Non-Linearity

The camera EM gain is controlled by an input index, M; a value in the range 0 to 255 (8-bit discretisation) which is not directly proportional to the level of gain applied. Furthermore, the gain suffers 'ageing' such that values reported by the manufacturer are likely to be unreliable[172]. Nominally, the camera has a range of EM-gain levels up to $1200\times$. In order to test directly and establish the current relationship between the selected index and the actual gain applied, a set of experiments were done in which the tungsten-halogen source was pointed at a reflectance standard from a distance sufficiently far as to spread a diffuse intensity across the standard. This was then imaged repeatedly through an OD4 neutral density filter for each possible EM gain input index. It was possible to image approximately half of the values before an additional ND filter (OD2) had to be added to reduce the ever-rising signal and avoid camera saturation. The imaging parameters were binning 4×4 and minimum exposure time in addition to EMCCD mode. Several indices were interrogated with and without the additional ND filter in place to use as a reference to scale the results to the same range (i.e. to compensate for the filter changing the input intensity). The experiment was repeated with the same parameters and method and then the first half was repeated twice more using different binning modes to see if this had an effect. Values were obtained from images by averaging a region of interest on the imaged standard and the standard deviation was also recorded.

It was found that all of the measurements acquired with different binning modes fell within one standard deviation of each other (fig. A.10(a)) thus it was concluded that there was no significant impact of binning mode. The data sets were then combined and fitted with an exponential (fig. A.10(b)) which then relates the index to the quantitative gain applied and was used in a calibration file for the imaging system. Given that the measurements were relative, they were also scaled such that the minimum EM gain was $10\times$ which is the nominal minimum value and was found to be consistent with observations of intensities measured between normal and EMCCD mode (appendix A.2.2).

According to the manufacturer documentation, the EM gain is highly dependent on temperature and this is the reason for the high stabilisation of the temperature following cooling. It has

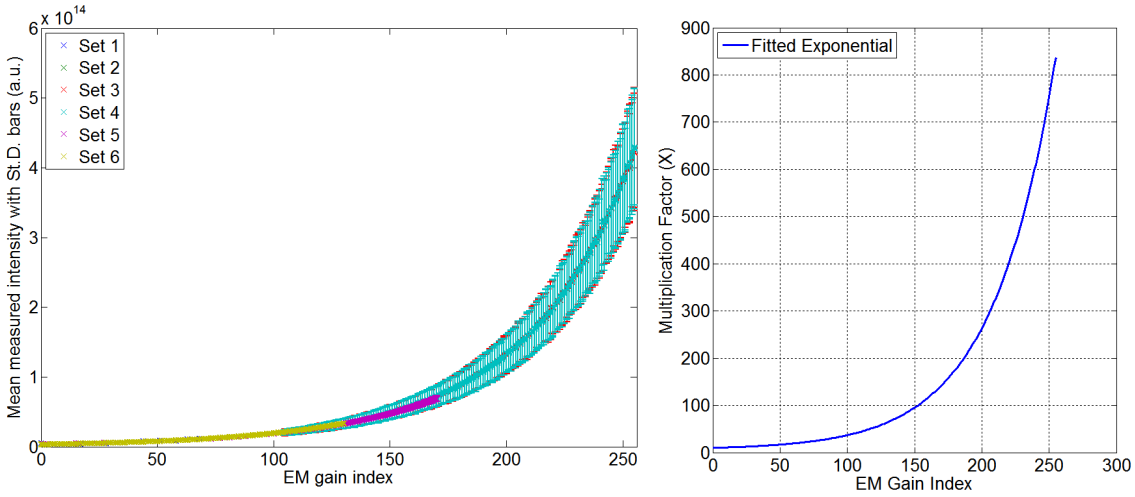

(a) (b)

Figure A.10: Results of EMCCD gain investigation with (a) plot of measured values from 6 data sets (1 and 3 are the original lower-half experiments with 2 and 4 being the original upper-half; 5 and 6 are the repeats with different modes) and (b) resultant fitted curve relating EM gain index to quantitative gain applied.

been observed that particularly following CCD mode and read mode changes there is actually a quick change in the temperature on the order of one or two degrees which might be significant, in the test here the mode was not changed and the temperature was allowed to stabilise following application of the mode settings but before imaging. This is not always feasible, particularly if surface capture and BLT imaging are done together for example in the current workflow it would be unrealistic to wait several minutes to allow re-stabilisation between acquisitions. In the present work EM mode is never used in a sensitive application (it is used in surface capture but only because it offers faster imaging) owing to the unfavourable noise properties in the long-exposure low-light situation (appendix A.5).

A.2.5 Interference Filter Angular Dependence

In this subsection a preliminary investigation of the impact of the angularly-dependent interference filter transmittance characteristics upon the imaging system is presented. It is demonstrated that for the higher angles of incident light in the imaging system (those coming from far into the mirror-views), the interference filter pass-band is changed sufficiently so as to effect

the measurements made by the imaging system. This poses a challenge to better understand the effect and apply principled compensation which is important future work for this project; currently, the effect is assumed to be minor and ignored or corrected for in a simplistic fashion (with a single, per-experiment, empirically derived scaling factor) based on either calibration phantom measurements or simple heuristic corrections which is not satisfactory in the long-term. An alternative to modelling the effect is to use different filters that do not exhibit the same behaviour, such as coloured glass filters, though these are often unavailable for specific desired transmittance bands.

Background

Concerning the transmittance characteristics of interference bandpass filters, it is known that there is a dependence on the angle of incidence of the transmitted light; generally there is a change in central wavelength, as well as size and shape of the pass-band[207]. It is also understood that in most practical circumstances it is not possible to calculate from theory the transmittance characteristics of a filter directly given the angle of incidence because of a lack of availability of knowledge regarding the exact composition of the filter, e.g. number of layers, and the precise properties of each material involved[207].

This fact presents a potential challenge and possible limitation when attempting to use interference filters in an imaging system because the rays of light entering the filter when it is placed, as in the presented system, ahead of the lens, come from a variety of positions and enter at a variety of angles accordingly. In order to examine the impact of the effect on the present imaging system, given the typical geometry, an experiment was carried out as follows.

Experiment

A red laser pointer with central wavelength (CW) of 650 ± 10nm was mounted underneath the sample stage (in the position normally reserved for the NIR projector) pointing upwards through the hole in the stage towards the detection system. The pointer was held in a clamp so as to be constantly switched on. A home-made highly scattering white hemisphere of polyester resin

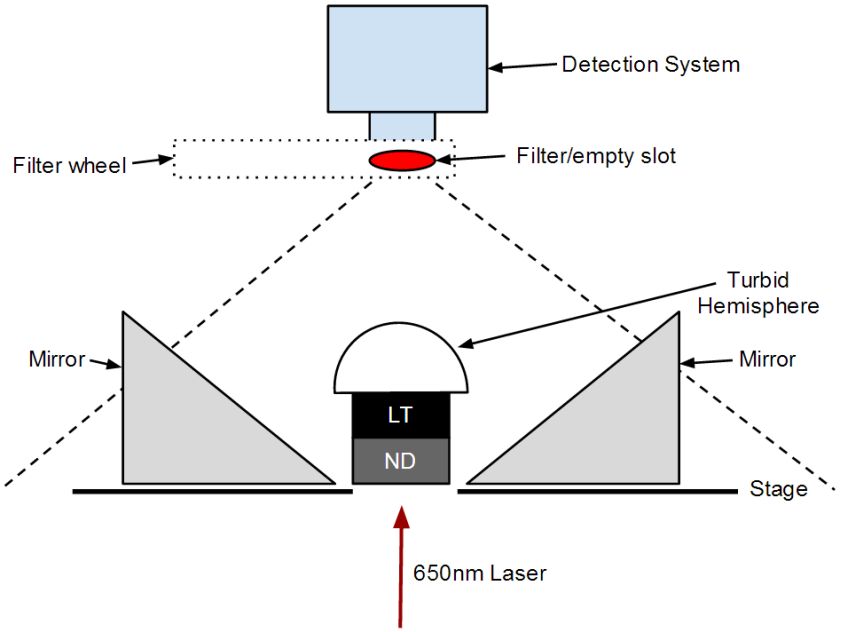

Figure A.11: Diagram of experiment to ascertain impact of angle of incidence of incoming light on filter. Laser light hits first a neutral density filter (ND) housed within a lens tube (LT) then is incident upon the scattering hemisphere. Multiply scattered light leaves the hemisphere and enters the imaging system via the unfiltered/filtered aperture either directly or via a mirror.

was placed on the stage in the path of the laser, supported additionally by a lens tube housing an OD3 neutral density filter (Thorlabs, Cambridgeshire, UK) to reduce the light intensity. In order to investigate the effects of the interference filter on the resultant image, images were acquired with no filter, with a neutral density filter (ND1; Thorlabs) or with one of three interference filters (10nm FWHM, CWs 640nm, 650nm and 660nm; Thorlabs) loaded into the filter wheel. Figure A.11 shows a schematic of the experimental set-up.

The monochromatic light source was used so that there would be no effect of the differing levels of scattering and absorption (due to the wavelength dependence of these qualities) on measurements made with and without filters.

Three different interference filters were selected (640, 650, 660nm) around the expected central wavelength of the laser emission (650nm) in order to capture the light regardless of the actual central wavelength ($\pm$10nm). Five images were taken for each case and averaged.

Results

Images are shown in fig. A.12. The images were dark-subtracted but not scaled by exposure time because the actual inter-image measurement intensity is not important in this experiment, rather the relative intensity of different parts of the image are important. The image at 640nm is not shown as no signal was observed.

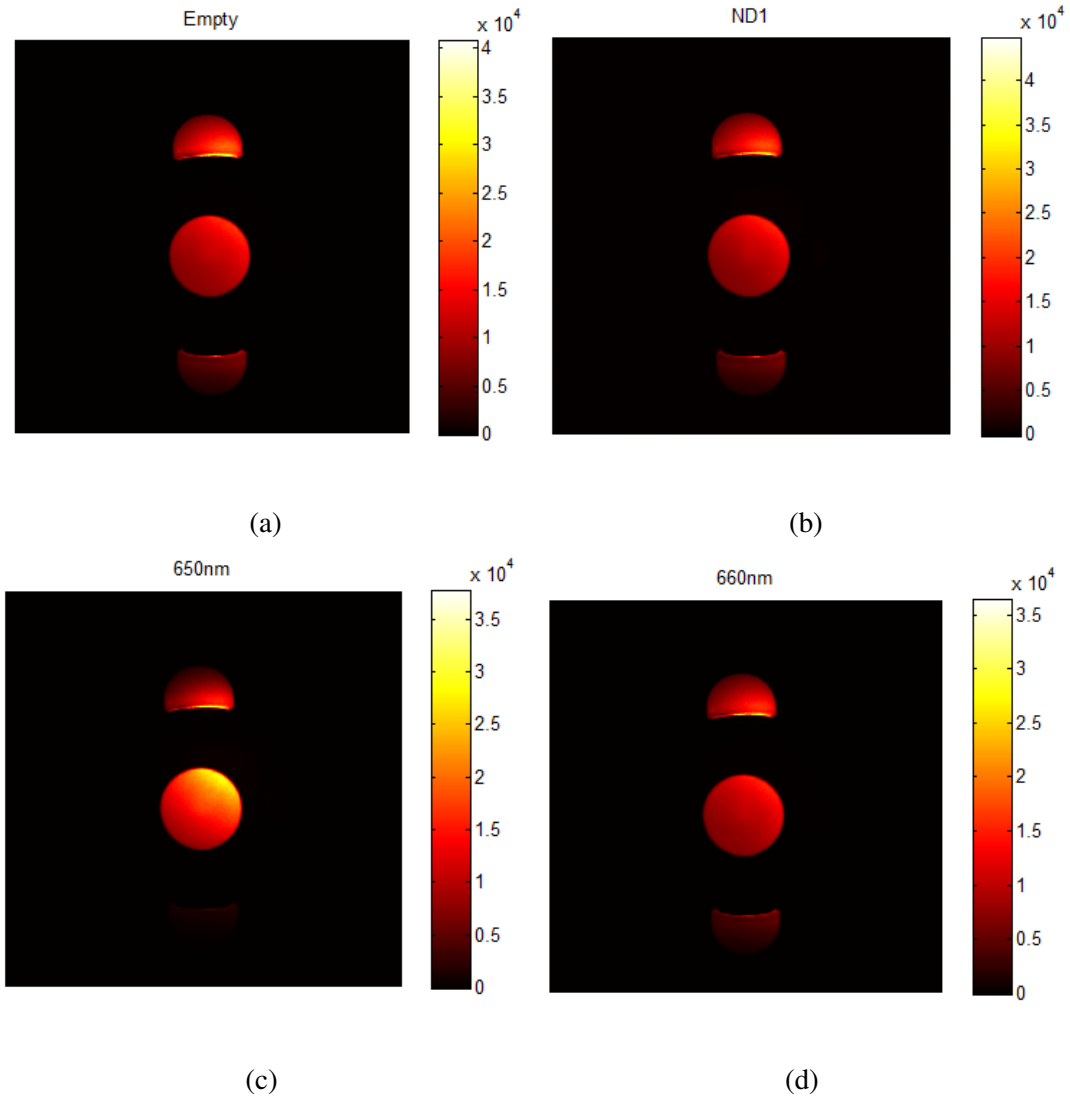

Figure A.12: Multi-view (mirrors/direct) dark-subtracted images of a scattering hemisphere trans-illuminated with 650nm laser, measured (a) without a filter and (b-d) with three different filters; ND1, 650nm and 660nm as indicated.

It can be seen by visual inspection that there is some change in the appearance of the image in terms of the relative intensity of the mirror-view measurements between the images. This

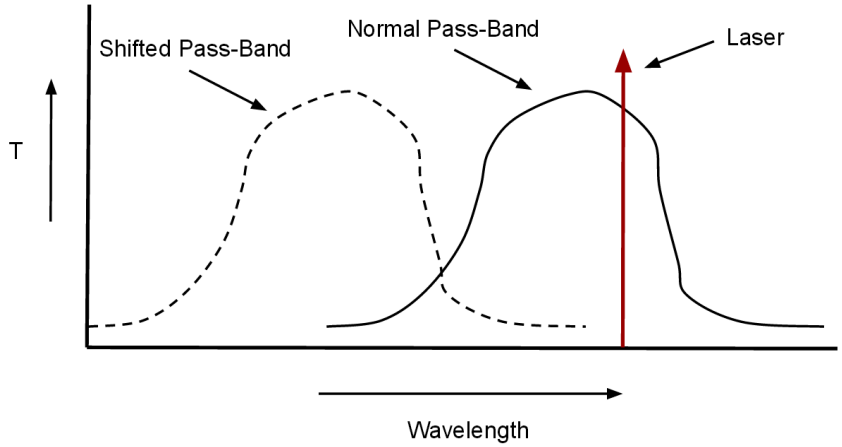

Figure A.13: Sketch illustrating the hypothesised situation in which the shifted filter pass-band no longer admits the laser light and as such image intensity is reduced.

effect is most clearly present in the mirror at the bottom of the presented images, this is assumed to be because this mirror is further from the centre of the field of view (i.e. a higher angle of incidence) than the other mirror. It can be seen that there is little relative difference in the case of the ND filter as compared to the unfiltered measurement. This indicates that it is not the presence of just any filter that causes the effect. The effect is more pronounced in the 650nm case than the 660nm case. This is understandable if it is now assumed that the wavelength of the laser is higher than 650nm but within the pass-band of both filters initially. Then the effect of the high angle of incidence from the mirror view is to shift the filter pass-band to a lower wavelength range (as is typical[207]) so that the shifted 650nm filter pass-band is then beyond the wavelength of laser emission (≈ 650) whilst that of the 660nm filter is not. A sketch illustrating the hypothesis for the 650nm filter case is shown in fig. A.13.

The image data from this experiment is finally summarised by fig. A.14 which shows three regions of interest (ROIs; one for the top-view, one for each of the mirror views) and relative mirror intensity for each mirror (relative to the top-view in each case and all normalised to the unfiltered case) drawn from these ROIs.

The graph shows then the impact of the presence of each of the filters on the relative mirror signals as compared to the direct signal. If there was no effect introduced by the filter one would expect to see a flat line, which is the case for the neutral density filter indicating that this filter does not affect the relative signal and therefore does not have an angular dependence that is

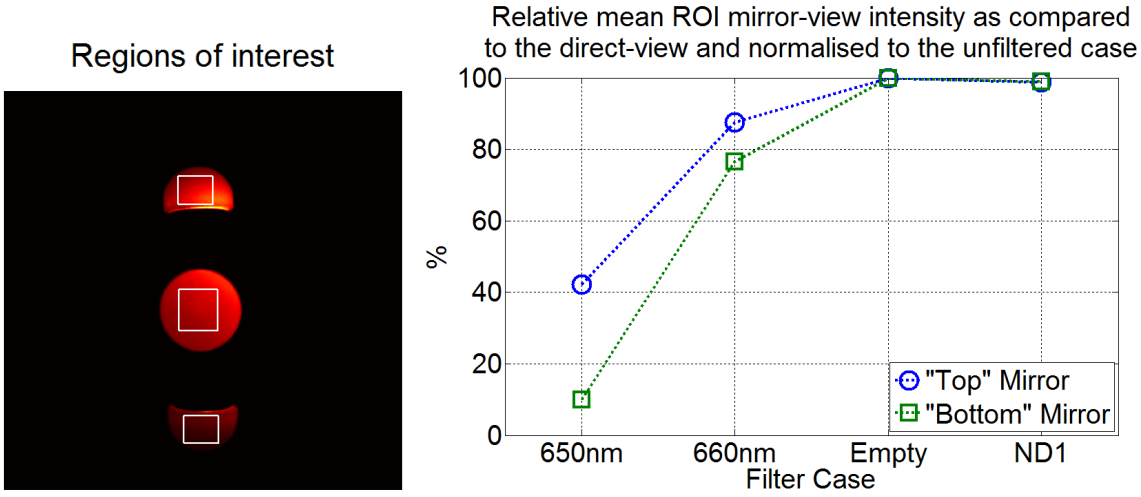

(a) (b)

Figure A.14: Quantitative analysis of the data with (a) indication of three ROIs, one-pre mirror and one direct-view, the values within which are summed to produce single view-resolved measurements. The summed mirror ROI values are expressed as proportions of the top-view values in each case and then these are normalised across the image set to the unfiltered dataset. These values are then plotted (b).

important in the present system. However, in the case of the interference filters, a clear change occurs and it is a stronger change in the case of the mirror that is further from the centre of the image, consistent with the hypothesis that there is an important angular dependence on the signal received through the interference filters.

These results highlight the need for future work in further characterising and accounting for the angular dependence of the bandpass filter behaviour.

A.3 Projector Characterisations

A.3.1 Projector Non-Linearity

The projectors that are used for surface capture do not necessarily have a quantitatively linear output intensity (i.e. number of photons) with respect to the value in the image that is given to be projected. However, the success of the surface capture method relies on being able to project values that are quantitatively linear with respect to the projection image values. Thus

an experiment was done to establish the quantitative relationship between input pixel intensity (from computer to projector) and output light intensity (by the projector).

For the experiment, a 99% diffuse reflectance standard (Labsphere) was imaged under uniform full-field illuminations at every possible input pixel intensities (using 8-bit images this meant 256 different images) supplied by the projector. The normal detection system was used for the imaging with a 580nm filter. There were two reasons for the choice to use this filter: the first is that the intensity of the projector in conjunction with the sensitivity of the EM-CCD camera is such that the camera is saturated when the input light is not filtered, even at the fastest and least sensitive acquisition mode; the second reason was so that no assumption needed to be made that the projector non-linearity would not have a wavelength dependence, as there was no particular reason to make this assumption - by using a single filter and always using the same single filter in future when projected linearity was important (i.e. for surface capture), this ensured that no problems would be encountered due to unknown spectral variations in non-linearity. Images were analysed by taking the mean for a region-of-interest on the reflectance standard. Maximum-normalised values were then taken for the relative intensity (which is the desired quantity). The result is shown in fig. A.15.

It was assumed that the output characteristics of each projector would be identical so this measurement (for one projector) was used to correct images for both. Images were corrected by sampling the curve for the desired intensity to find the appropriate pixel value to project.

A.4 Light Source Characterisation

A.4.1 NIR-Source-Projector Coupling Efficiency

The coupling-efficiency between the NIR light source and the output of the projector used for diffuse transillumination describes the proportion of the light exiting the fibre that emerges from the projector. This is useful in obtaining an estimate of the total amount of power incident upon the sample.

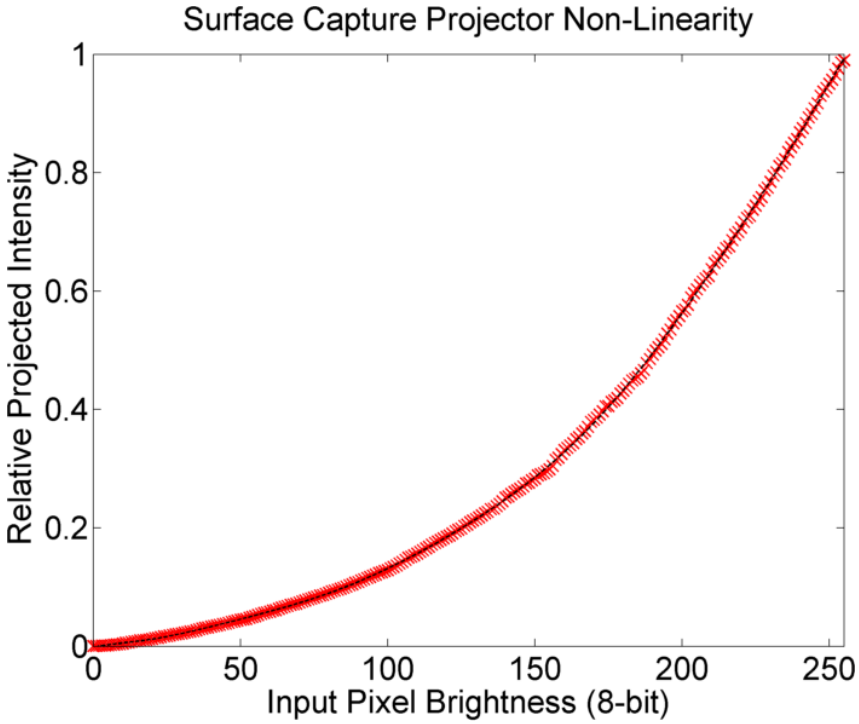

Figure A.15: Measured relative output of one surface capture projector at 580nm. Red crosses indicate measurments whilst the black line shows a fitted curve (fit using the MATLAB polyfit function).

Experiment

The system was set-up in two scenarios shown in fig. A.16 in which a piece of paper was imaged flat on the sample stage being illuminated either by the source fibre directly or by full-field illumination by the fibre-coupled projector. Both the fibre and the projector were positioned so that the whole of the irradiated region of the paper due to the incident light was within the stage hole and therefore the field of view of the detection system. In both cases a neutral density filter (ND30A; Thorlabs; 0.1% transmittance) was positioned ahead of the filter wheel to limit the amount of light entering the system avoiding saturation of the sensitive CCD and a 650nm bandpass filter was set in the filter wheel.

For each experimental set-up above, ten images were acquired sequentially. The images were processed into electrons per second read-out by the CCD and then summed to give the total signal measured. For imaging the source fibre only, the exposure time was 0.42s, the read mode was 2, and the binning was 1×1. For imaging the projector output, the exposure time was 5s, the read mode was 1, and the binning was 4×4. In both cases the CCD mode was

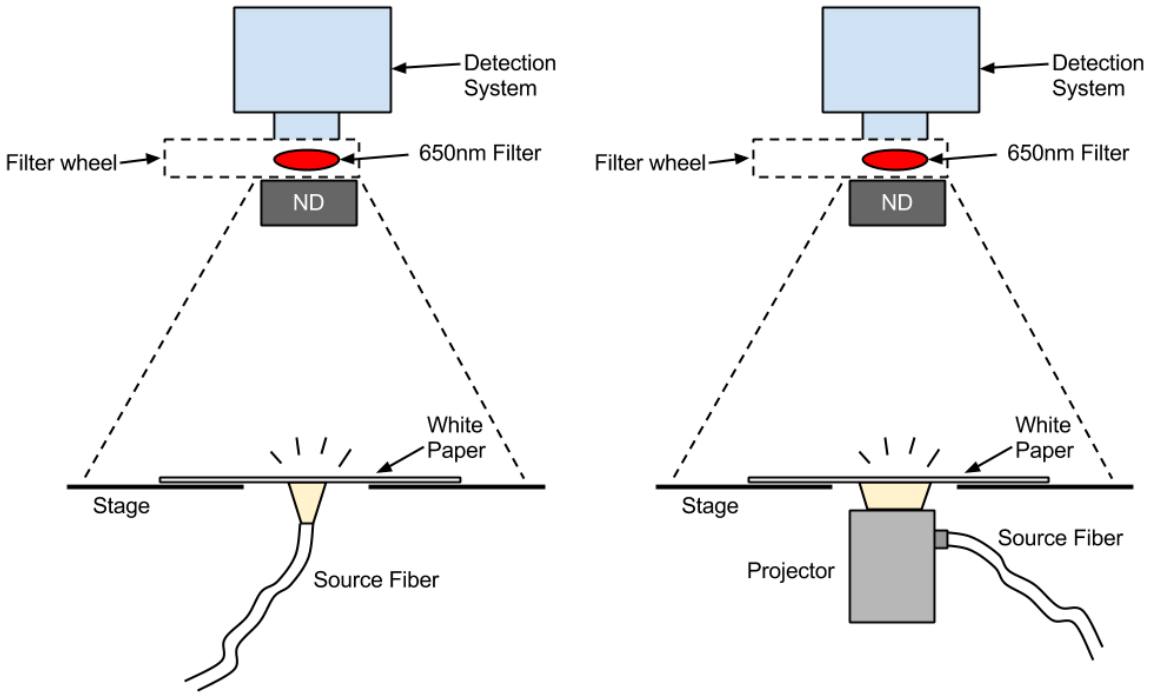

(a) (b)

Figure A.16: Two experimental set-ups used to acquire data to establish the NIR source-projector coupling-efficiency.

normal (non-EM). The difference in the parameters used was due to the need to accommodate differing intensities whilst collecting high signal without saturating the camera.

Results

Figure A.17 shows total signal measurements for the case of the projector-coupled and directly incident fibre source. It can be seen that the signal is very stable across repeats in both cases; specifically the standard deviations in the data sets are 0.11 for projector data and 0.83 for direct data; this is less than 1% of the mean in both cases.

An estimate of the projector transmittance at 650nm was then calculated by dividing the mean projected signal by the mean direct signal. The resultant value was $14.7 \pm 0.2\%$ (where the uncertainty is estimated based on the extremal values calculated assuming uncertainties twice the standard deviation on the original means). This value is low and as such there is considerable scope for improvement.

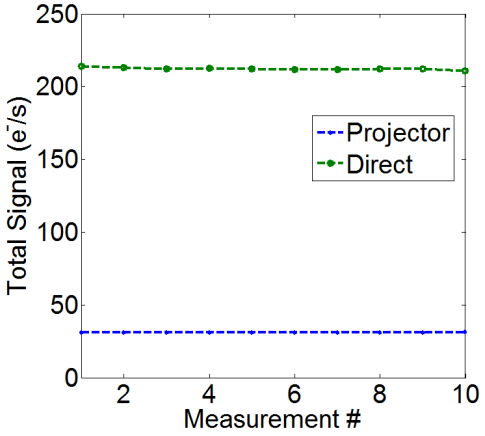

Figure A.17: Total signal transmitted through a piece of paper as-measured by the imaging system for light originating from the fibre-coupled projector (projecting uniform white) and from the fibre alone.

A.5 System Noise Model

Noise describes all contributions to measurements that are not part of the target signal.

A.5.1 Detection System Noise

Detailed understanding of the noise characteristics of the detection system are very important; a detailed understanding allows the accurate prediction of measurement SNR which influences the success of image reconstruction. If the noise characteristics are well-understood then they can also effectively be added to simulations in order to better represent the system in theoretical modelling studies. Noise in the detection system is due only to the camera because the other detection system components have wholly deterministic behaviour on the scale of the measurements considered.

In this section, the main sources of noise in the camera are first introduced and experiments are presented to validate the nominal noise levels stated by the manufacturer where they are available and to measure them where they are not. The noise from all sources is then combined into a single model that is experimentally validated and provides an accurate description of the detection system noise.

A.5.2 Sources of Noise

The sources of noise in images captured by most CCDs can be divided into three categories; read noise, shot noise, and dark noise. In addition to these, the electron-multiplication mechanism of the EMCCD introduces a new type of noise that is referred to as excess noise. Specific formulae and technical details in the following subsections are taken from the ImagEM-1K technical note[172]. The EMCCD is also subject to a 'spurious noise' which is omitted because it is expected to have a negligible effect[172].

Read Noise

Read noise describes measurement variations introduced when charges on the CCD are 'read out' (taken from the CCD and converted into a digitized signal that is presented in the image). For the ImagEM, the readout noise is primarily due to the resetting of the floating diffusion amplifier (FDA) which is the on-chip unit that converts the accumulated charge to a voltage for each pixel serially. The read out noise depends on the pixel clock rate (faster readout being associated with higher noise) which has three possible settings of 11MHz, 2.75MHz and 0.69MHz. According to its specifications the camera normally has a maximum read noise of 19 e^-/pixel when using the fastest readout mode and a minimum level of 10 e^-/pixel when using the slowest. In EMCCD mode the *effective* readout noise (scaled with respect to the signal multiplication) can become 1 e^-/pixel with sufficiently high EM gain though the actual readout noise (equivalent to that encountered with no EM gain applied) in this mode is higher than that in normal mode ranging from 30 to 100 e^-/pixel.

Practically speaking for the ImagEM, the readout noise is dependent on the read mode only, as such it can be expressed as

$$\sigma_r, \tag{A.1}$$

a constant associated with each read mode. A summary of the read mode properties is provided in table A.3 which shows the expected trade-off between imaging speed and read noise.

Read Mode	Pixel Clock Rate (MHz)	EMCCD Read Noise (e^-/pixel)	Normal CCD Read Noise (e^-/pixel)
1	*0.69*	*30*	*10*
2	*2.75*	*80*	*19*
3	*11*	*100*	*N/A*

Table A.3: Nominal ImagEM-1K read mode characteristics showing the related clock speed and noise characteristics. Read mode 3 does not exist in normal CCD mode.

Dark Noise

Dark noise includes all variations that are due to charges appearing on the CCD not due to incident light. These are generally due to thermal effects. Higher temperatures lead to a greater chance of random generation of electron-hole pairs in the CCD. This is the reason for the cooling of the ImagEM camera to $-55°$C at which the manufacturer documentation indicates a dark current, D, of $0.01e^-$/pixel/s. This is per pixel because the process occurs throughout the CCD spatially and per-second because thermal effects are ongoing throughout acquisition. The dark noise is expressed as

$$\sigma_d = \sqrt{(DT)}, \tag{A.2}$$

where T is the exposure time.

Shot Noise

Shot noise is the variation in the signal measured due to the quantized nature of photons; the fact that there are individual packets of energy bombarding the detector such that whilst the signal that is desired is the average photon flux (say for a long period of time), what is measured is subject to statistical variation based on the particular arrival time of individual photons at the detector. For sufficiently high numbers of photons, shot noise is normally distributed and described by

$$\sigma_s = \sqrt{QP}, \tag{A.3}$$

where Q is the camera efficiency and P is the population mean number of photons i.e. the desired signal.

Excess Noise

When using EMCCD mode, 'excess noise' is introduced which is a noise associated with the multiplication process. The extra noise factor is denoted F and influences the signal as it appears on the CCD prior to readout, i.e. including at this stage the shot noise and dark noise which are then amplified along with the signal. The excess noise factor is not provided by the camera specification, although the technical note[172] presents example calculations in which $F = 1.41$ is used.

A.5.3 Calculating Signal, Noise and SNR

In normal CCD mode, the total camera noise is

$$N = \sqrt{\sigma_r^2 + \sigma_d^2 + \sigma_s^2}, \tag{A.4}$$

which is equivalent to

$$N = \sqrt{\sigma_r^2 + DT + QP}. \tag{A.5}$$

In EMCCD mode, with the inclusion of the excess noise factor, F, and the electron multiplication, G_s, the noise is

$$N = \sqrt{\sigma_r^2 + F^2 G_s^2 (\sigma_d^2 + \sigma_s^2)}, \tag{A.6}$$

which is equivalent to

$$N = \sqrt{\sigma_r^2 + F^2 G_s^2 (DT + QP)}. \tag{A.7}$$

Considering the signal ($S = QP$; section 5.4.1) the SNR is

$$\text{SNR} = \frac{SG_s}{\sqrt{\sigma_r^2 + F^2 G_s^2 (DT + S)}}. \tag{A.8}$$

This is valid for both CCD modes (with $F := 1$ and $G_s := 1$ assigned for normal mode).

A.5.4 Nominal Noise Model

The next few sections describe noise modelling results based on the above noise calculations, first with a nominal noise model then with a refined model based on a preliminary fitting of experimental data to gain an estimate of more accurate properties particularly with reference to the excess noise factor.

The technical note from the manufacturer[172] demonstrates situations in which the EM-CCD mode is superior to the normal mode in terms of SNR but these are limited to cases of very short exposure time (tens of milliseconds), which whilst comparable to some biomedical application time-frames such as real-time microscopy applications, are not so relatable to the current application. What is found in practice, is that the EMCCD mode is not generally useful in the current application as it provides measurements with lower SNR than the normal mode except in cases where the incident signal is exceptionally low, i.e. less than 1 electron per second produced due to light on the CCD.

Recall from section 5.4.1, the relationship between CCD electrons and image values:

$$I = \frac{SG_s}{\zeta(\rho)} + I_o(\rho), \tag{A.9}$$

which describes the image digitisation by the camera and can be used to calculate image values read out due to a known signal and camera parameters. For a noisy signal, an estimated upper bound on the value read out can be calculated by:

$$I = \frac{SG_s + 3N(\rho)}{\zeta(\rho)} + I_o(\rho). \tag{A.10}$$

This upper bound calculation is important for understanding the impact of SNR in practice because of camera saturation effects; any value that would be read out that is higher than 2^{16} is useless.

Given the nominal values for read noise in the lowest-noise read mode for each CCD mode (30 e^-/pixel in EMCCD mode and 10 e^-/pixel in normal mode), along with the nominal dark current level of 0.01e$^-$/pixel/s and applying an excess noise parameter F of 1.41 in EMCCD mode (used as an example in the ImagEM technote[172]), predicted SNRs for several constant CCD signals ($S = TS_0$) were calculated to investigate the expected noise characteristics of the camera. An upper bound for the values read-out was then calculated for each case with eq. (A.10).

Figures A.18 to A.20 show the results, the ranges of exposure times that are shown are up to 10 minutes, reflecting the fact that this is used as a cut-off to maintain experimental feasibility in BLT experiments (chapter 9). It can be seen from the first of these figures that when there is a signal of 0.1 electrons-per-second generated in the CCD by incident photons the EMCCD mode offers some advantage over normal mode, most notably at short exposure times. However, it can also be seen that for signal of 1 or more electrons-per-second this is no longer the case; at 1 electron-per-second the SNR is better in normal mode after about 100 second exposure times and very similar beforehand, at 10 electrons-per-second (and higher, though the data is not shown) the SNR is better in normal mode even for short exposure times (e.g. less than 10 seconds).

The EMCCD mode does not perform as well as the normal mode in most cases due to a number of factors. Firstly, the EMCCD mode read noise is higher, secondly, both the dark noise and the shot noise are multiplied by the excess noise factor as well as the EM multiplication.

A.5.5 Read Noise Validation

Experiments were presented in appendix A.2.3 in which the camera digitizer offset was estimated by taking 30 repeated images with minimum exposure times in each mode, it is expected that read noise would dominate in such a data set because there is no signal to give rise to shot noise and a very short exposure time in addition to low temperature giving rise to little dark noise. As such the pixel-wise standard deviations in each of the acquired image sets are expected to mostly be representative of the read noise. These are plotted in fig. A.21.

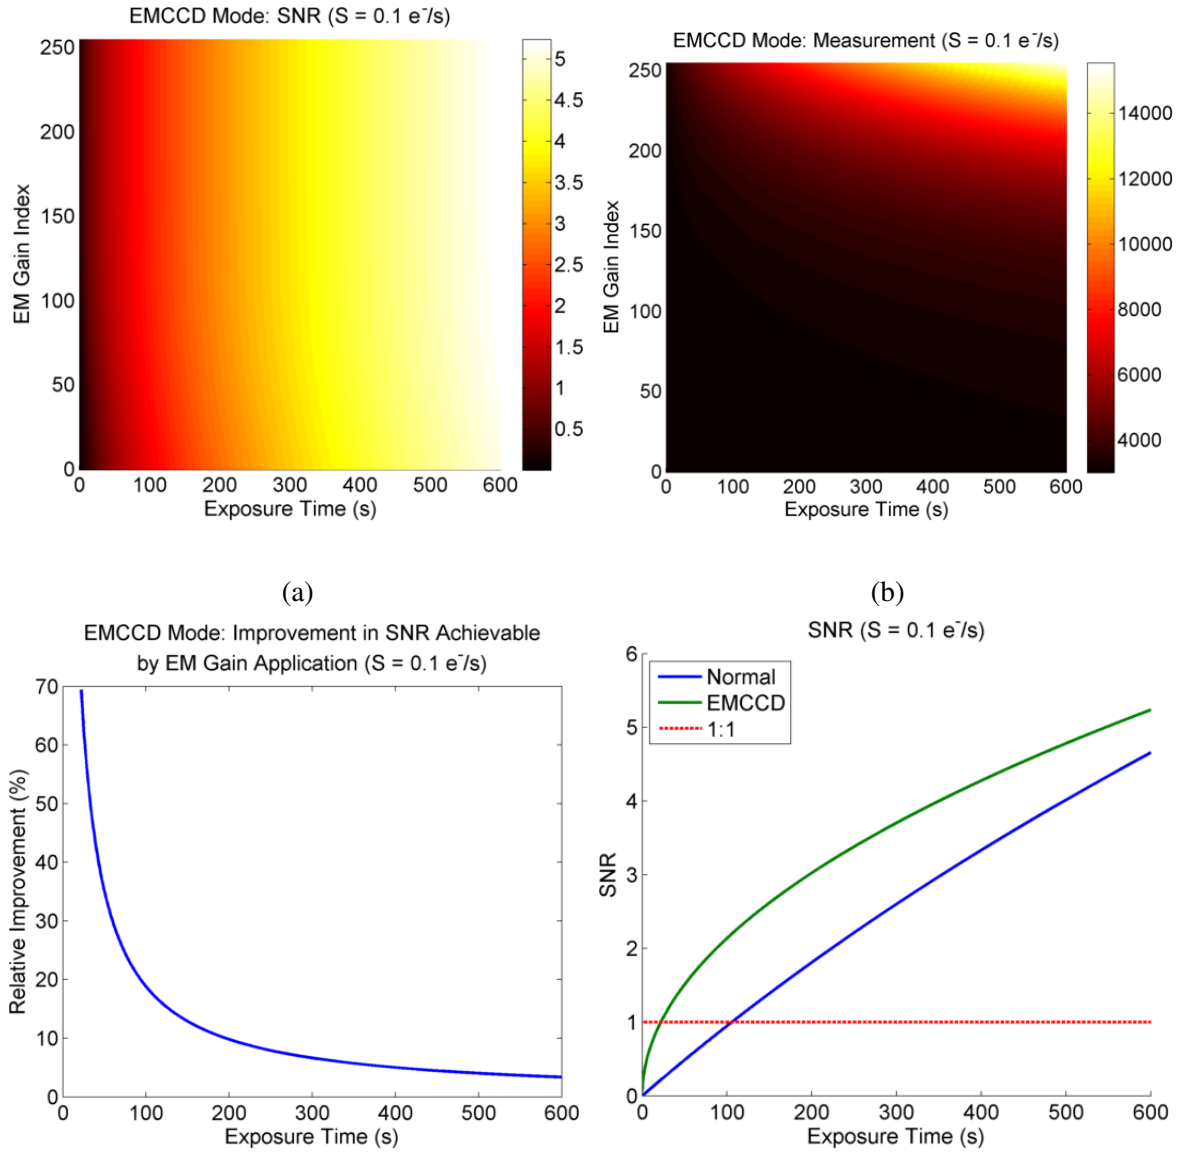

Figure A.18: Nominal noise characteristics (S = $0.1e^{-}/s$) showing (a-b) SNR and measurement (upper bound; eq. (A.10)) as functions of EM gain (appendix A.2.4) and exposure time for EMCCD mode, (c) the maximum percentage increase in SNR achievable across all EM gain levels as a function of exposure time, and (d) the best SNR for EMCCD (for any gain level) compared to the SNR in normal CCD mode, shown against a reference level of 1:1.

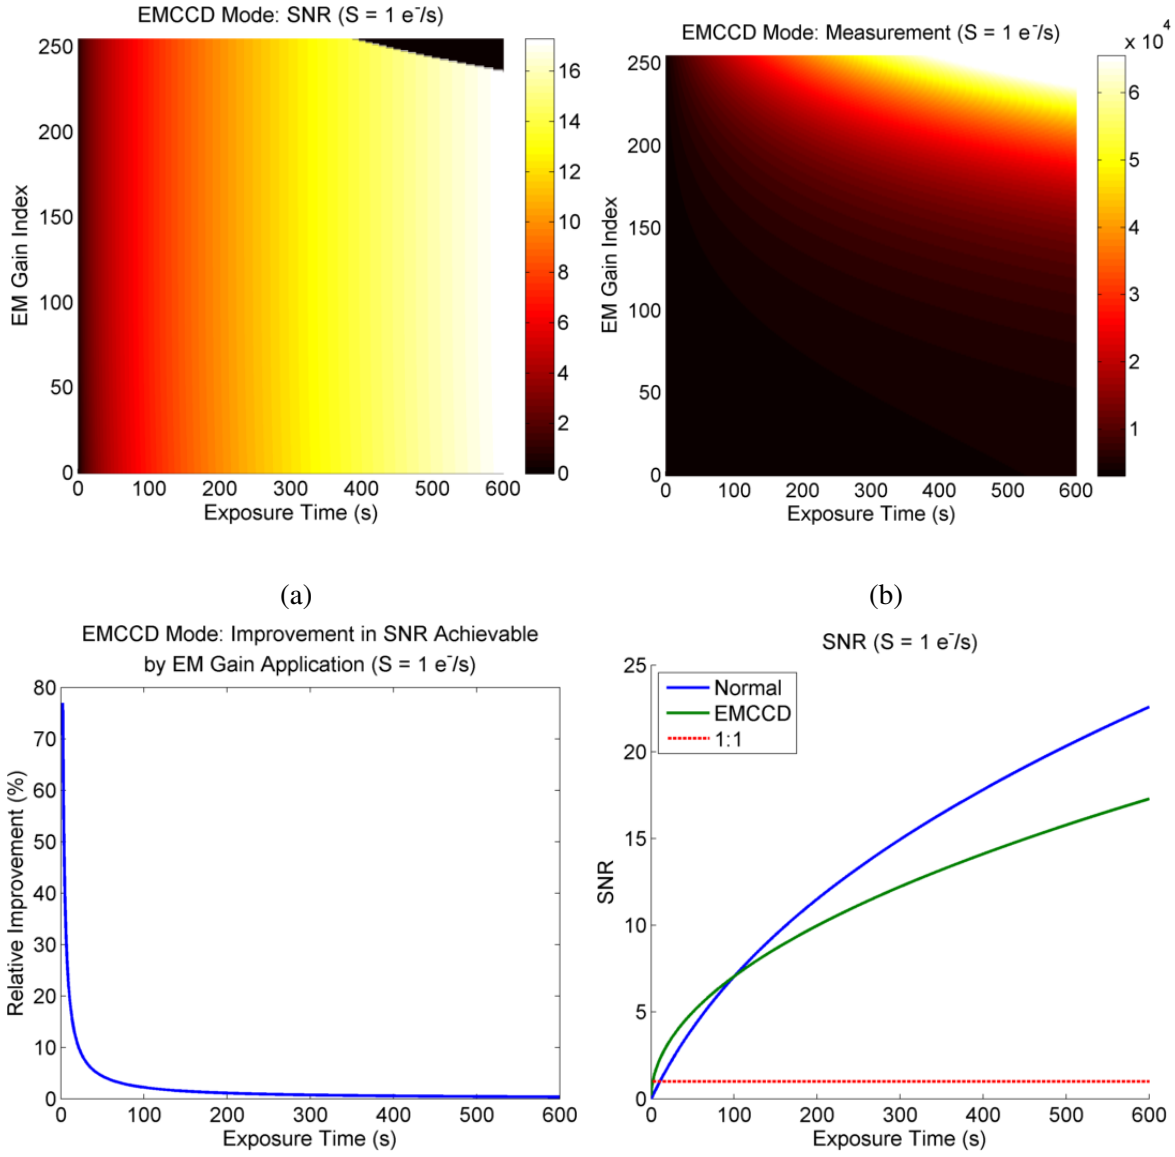

Figure A.19: Nominal noise characteristics (S = $1e^{-}/s$).

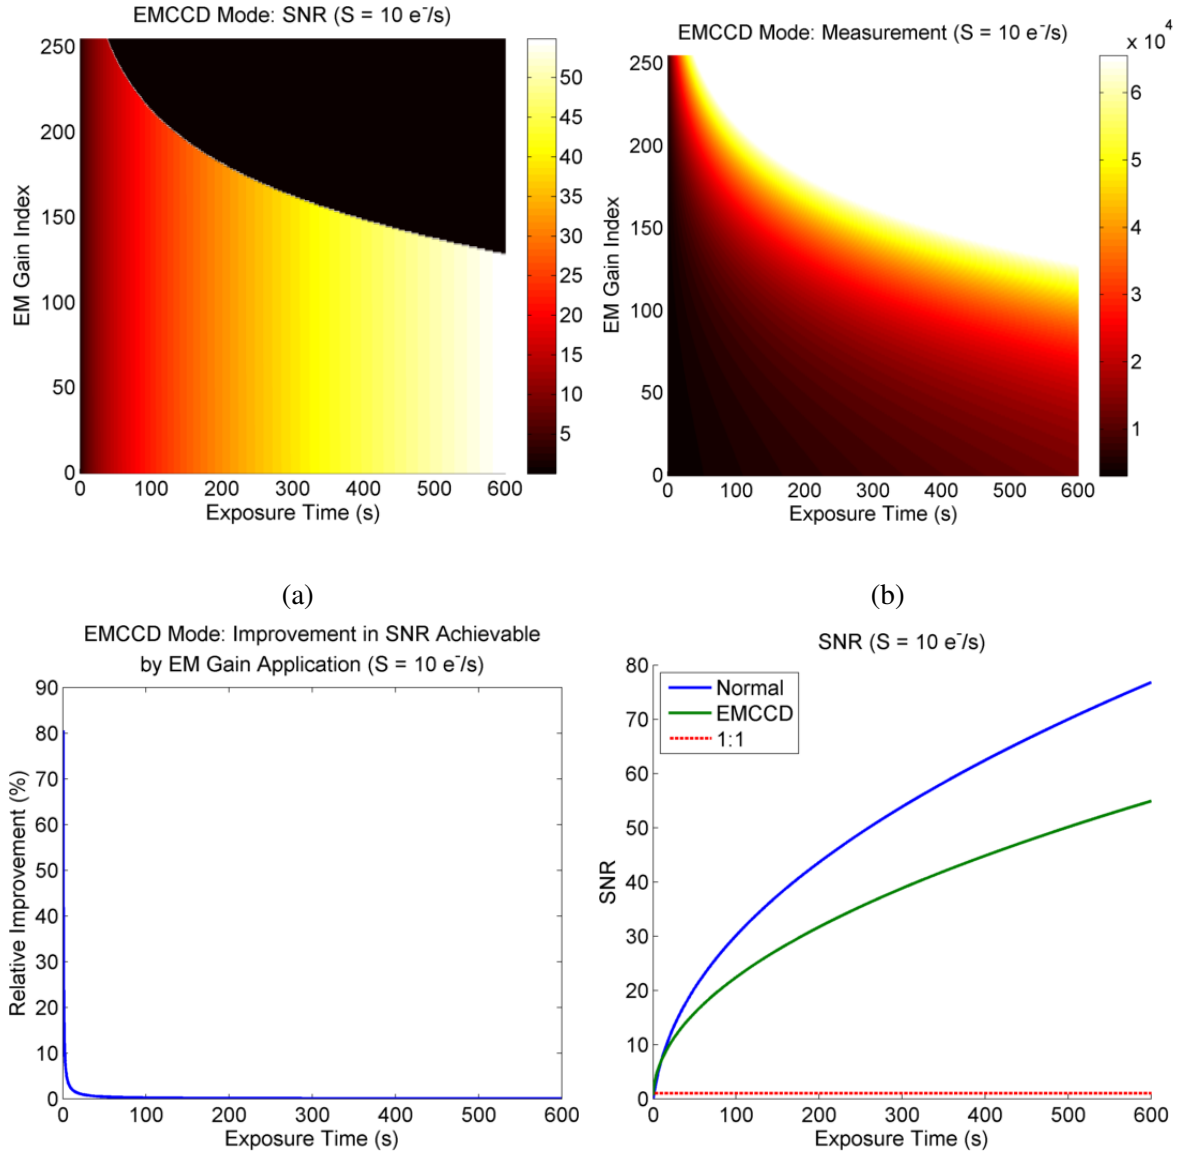

Figure A.20: Nominal noise characteristics (S = $10e^{-}/s$).

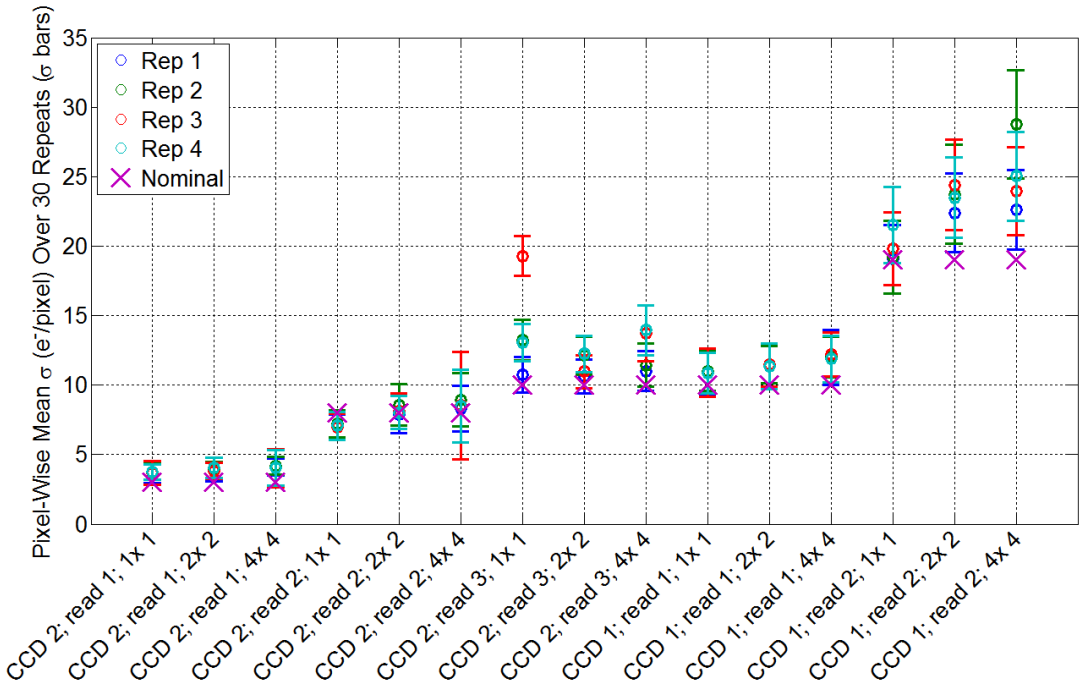

Figure A.21: Plot of the average and standard deviation of standard deviations in pixel intensity through stacks of 30 acquired images obtained on different days with minimum exposure time and no input intensity as compared to the nominal read noise-related σ values. Note that values have been converted into electrons using the camera digitizer conversion factors (appendix A.2.2) and correcting for gain ($10\times$ in EMCCD mode).

It can be seen that the σ values observed are very close to those expected owing to the read noise though in most cases the measured values are slightly higher. A small over-estimate may be expected given to the possible thermal induced variations discussed (appendix A.2.3).

A.5.6 Determination of Excess Noise Factor

Using eq. (A.8) and the nominal noise parameter values provided by the manufacturer, it is possible to calculate the SNR expected for a particular signal level and a set of imaging parameters. However, it is not known how reliable the manufacturer-provided (quoted as typical) values are and a single parameter, the excess noise factor, is not known. Therefore several experiments were undertaken to test the noise characteristics of the detection system empirically. A summary of the methods and findings is now presented.

In order to validate the noise model for normal CCD mode, an experiment was undertaken in which a red betalight (appendix A.1.2) was imaged inside one of the cylinder phantom tunnels (appendix A.1.1). 900 images were acquired; 100 at each of 9 different exposure times ($T \in \{0.5, 0.75, 1, 2, 3, 4, 5, 10, 20\}$s), in normal CCD mode, in read mode 1. 4×4 hardware binning was used meaning that each image contained 2^{16} pixels.

The mean image was calculated for the longest exposure time (20s), this was used to calculate a *best estimate* of the assumed temporally stable input signal[1], this is shown in fig. A.22. This was multiplied by each of the exposure times to give a best estimate of the signal for each exposure. Given the signal and the nominal noise model, the expected noise, v_m, was calculated for each pixel-wise signal and for each exposure time.

Standard deviation images were calculated for each exposure time. The values in these images represent samples of experimental noise, v_e. In total, the values for each pixel for each time represented approximately 5.9×10^5 samples of noise levels spanning many different signal levels.

The mean absolute percentage error in the model as compared to the experimental noise

[1] this approach was used rather than taking the mean of each 100 images for each exposure because in theory the longest exposure images should be the least noisy thereby giving a better estimate of the actual signal when averaged.

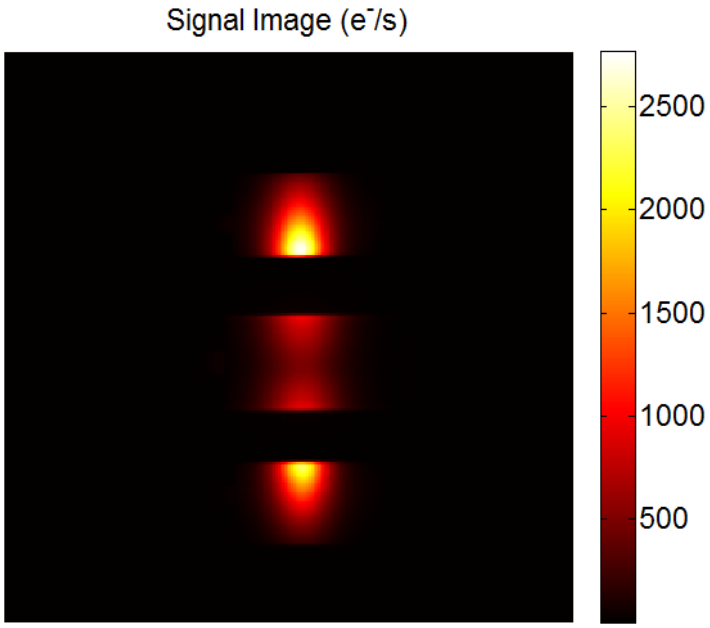

Figure A.22: Signal image (background-subtracted and exposure scaled image for the longest exposure time) used in the noise model evaluation.

values, ε_v, was then calculated by

$$\varepsilon_v = 100 \times \sum_{i=1}^{N} \left| \frac{(v_m^i - v_e^i)}{v_e^i} \right|. \tag{A.11}$$

This value was found to be 5.8%.

A second experiment was performed to investigate the accuracy of the noise model in EM-CCD mode. The experiment was identical except that EMCCD mode was used, the exposure time was fixed at 0.5s, and 20 levels of EM gain were applied ranging from 10× to around 300× with approximately linear increments. The range was chosen to be as broad as possible without causing saturation whilst maintaining a constant exposure time. 2000 images were obtained; 100 for each EM gain level.

The model fit was evaluated as before using the known noise parameters and the assumed excess noise factor. The value of ε_v calculated in this case was 17.33%.

Note that this value was obtained when considering only a subset of the available samples, namely those that had an estimated actual signal value greater than 10% of the maximum value in the whole set. This was deemed a fairer test set than the whole set because for very low-valued

(i.e. ≈ 0) there is relatively high uncertainty in the true signal value. When the calculation above is performed on the whole set the value for ε_v is 36.83%.

The higher error in predicted noise level suggests that the model may be inaccurate. This could be due to an error in the assumed value of F which was, as stated, not actually known. In order to estimate F, the EM-mode data set acquired above was used again. It was first randomly split into two equal-sized subsets, a test and a validation set. The test set was then used to solve for the F value using the MATLAB function *fminunc* to calculate

$$\hat{F} = \min_F \varepsilon_v(F) = \min_F \sum_{i=1}^{N} \left| \frac{(v_m^i(F) - v_e^i)}{v_e^i} \right|. \tag{A.12}$$

Using this method, the value $\hat{F}$ was found to be 1.7, with which the $\varepsilon_v(\hat{F})$ value for the test set was 6.6%. $\varepsilon_v(\hat{F})$ was also calculated for the validation set and found to be 6.6% also.

As with the first test value quoted for the original model, this analysis was carried out following a 10% vale threshold to select a subset of the samples. Using all samples the results were instead $\hat{F}$ = 1.95, test set $\varepsilon_v(\hat{F})$ = 23.5%, and validation set $\varepsilon_v(\hat{F})$ = 23.5% also.

A.5.7 Refined Noise Model

The nominal noise model for normal CCD mode appears to be accurate based on the small error observed when comparing modelled and estimated noise values for a large number of samples. The large error observed in the EM mode noise model inspired the study presented above in which the estimate of F was refined to 1.7. With this updated value of F included, the analysis of appendix A.5.4 was repeated. Results are shown in figs. A.23 to A.25

Whilst the EM mode performed relatively poorly for long exposures in the nominal model, it is still worse in the refined model owing to the increased excess noise. It can be seen that in cases of very low signal there is some small benefit of using EM mode for short exposures, however in general normal CCD mode provides higher SNR. Specifically, if the signal is greater than or equal to $1e^-/\text{s}$ then normal mode provides the best SNR so long as the exposure time is greater than 52.5 seconds.

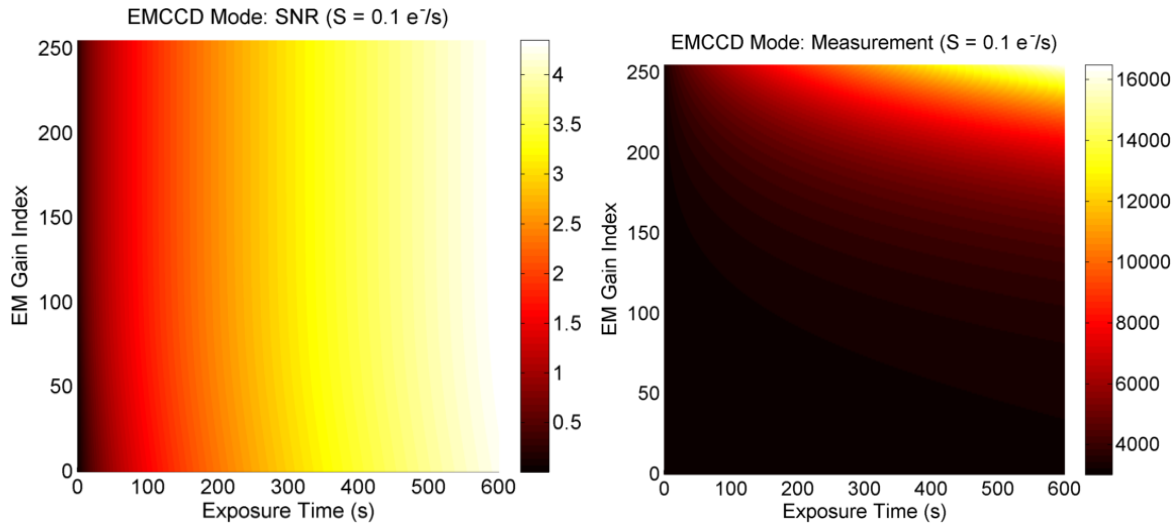

(a) (b)

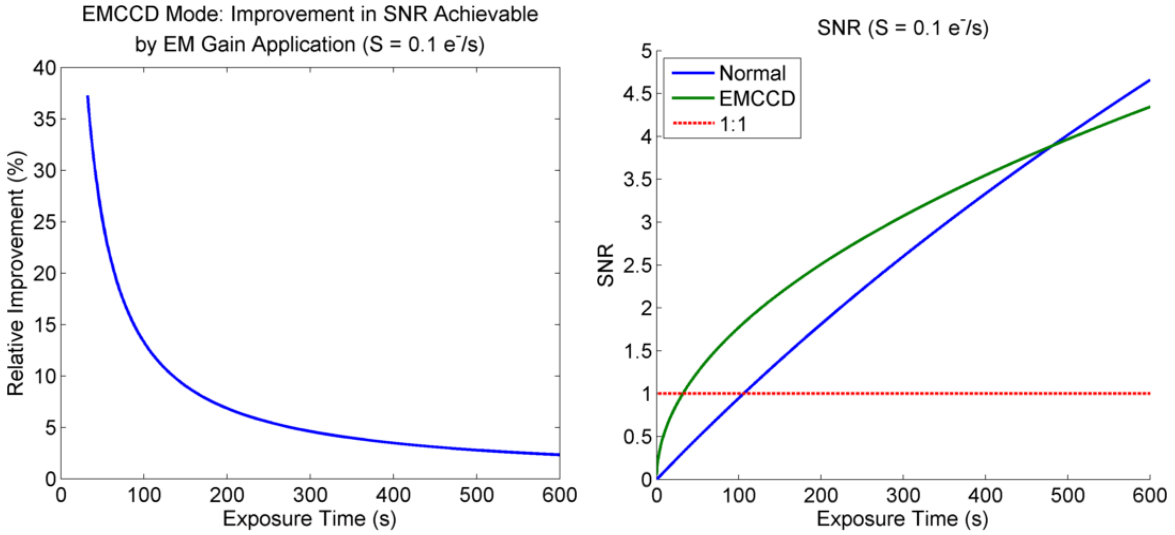

(c) (d)

Figure A.23: Refined noise characteristics (S = $0.1e^{-}/s$).

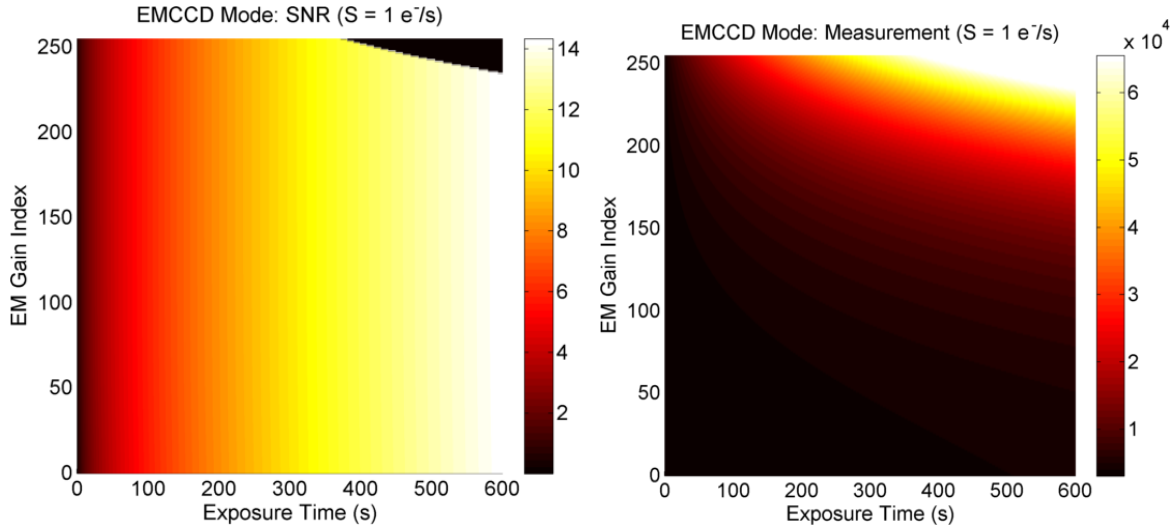

(a) (b)

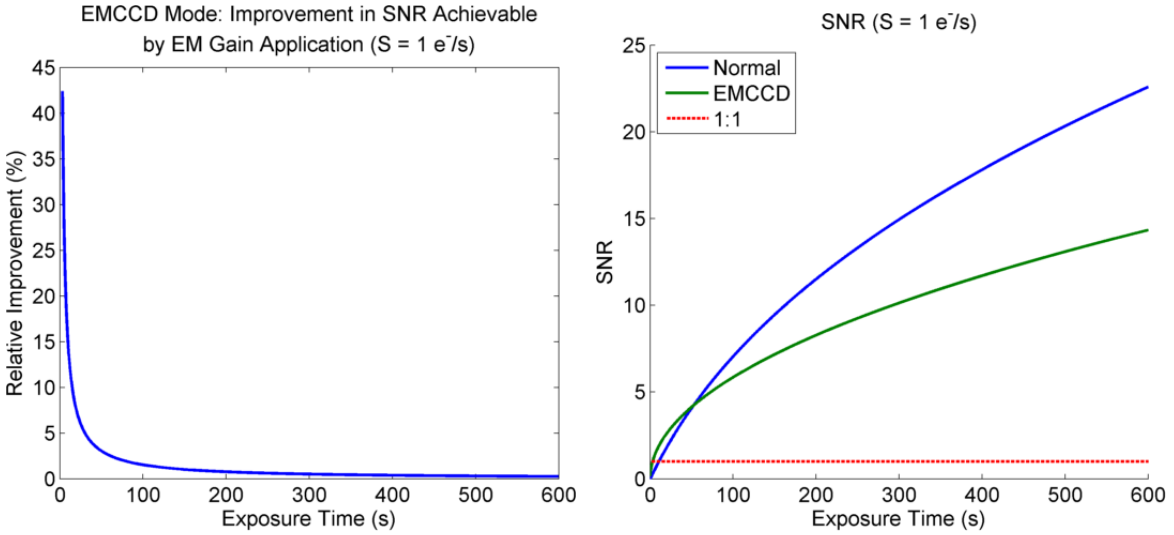

(c) (d)

Figure A.24: Refined noise characteristics (S = $1e^{-}/s$).

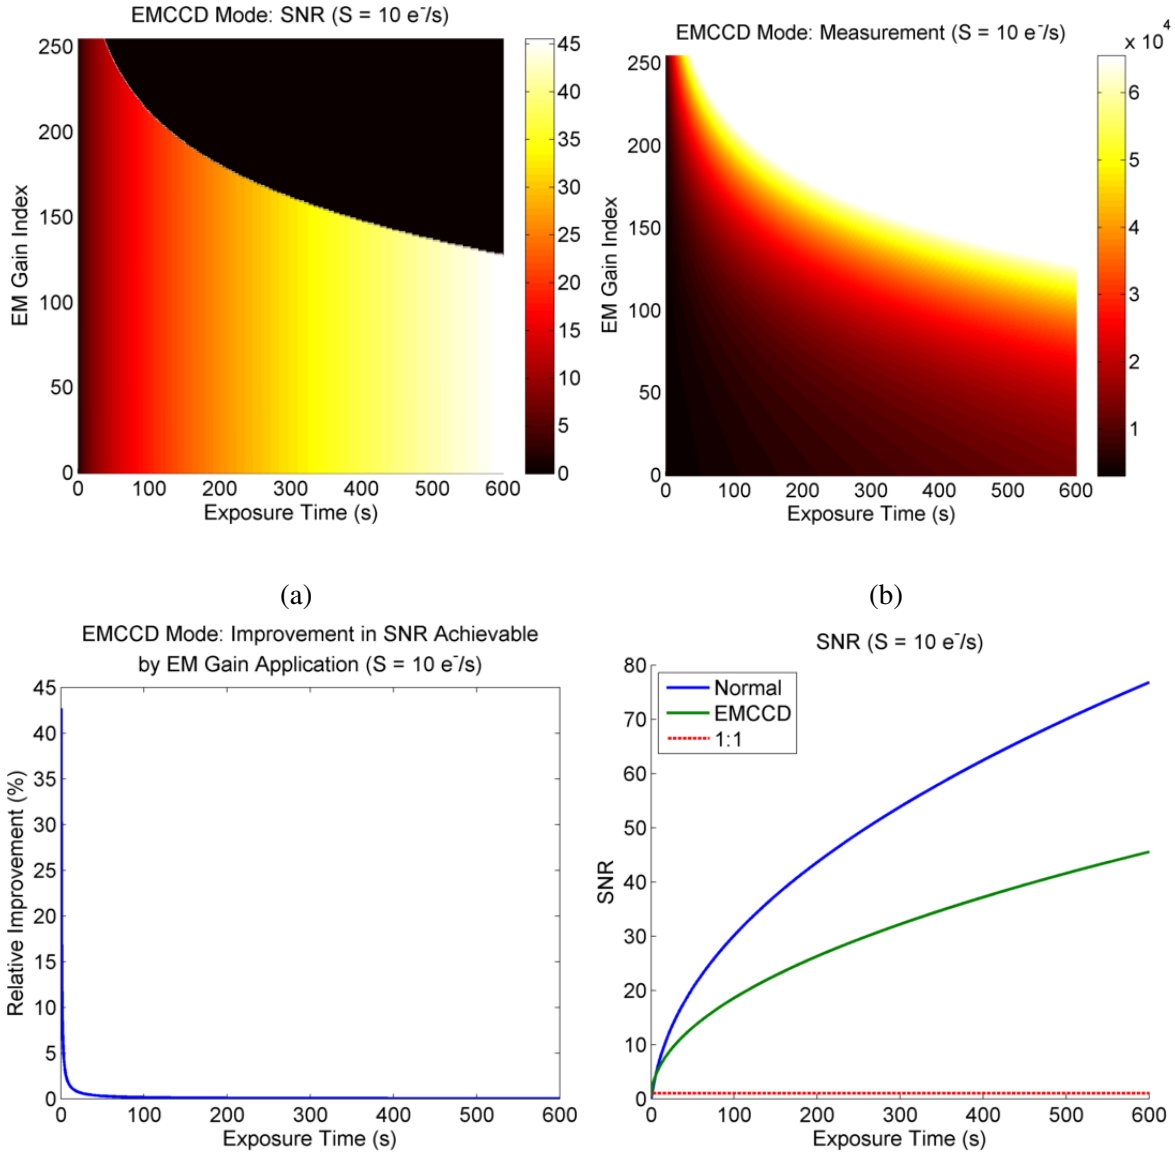

Figure A.25: Refined noise characteristics (S = $10e^{-}/s$).

A.5.8 DOT Source Noise

In the DOT source system (section 5.2.3), the tungsten-halogen lamp is nominally highly stable with a stability of 0.5%. The projector is not optimised for stability and such values are not reported by the manufacturer. The projector has been observed to flicker (emit a flash in spite of the projected pattern) when watched by eye, highlighting the need to characterise its stability.

In order to estimate the stability of the whole source system under typical usage conditions, 30 repeated images were acquired of 4 different DOT source patterns trans-illuminating the block phantom (appendix A.1.1). Images were acquired at 750nm with 30s exposure time, 4×4 binning, in normal CCD mode and read mode 1. For each of the four illumination patterns, a mean (signal) image and a standard deviation (noise) image were calculated. Values were then thrown away that were lower than 10% of the maximum intensity in each set. Approximately 16000 samples remained, for which a noise estimate was calculated using the detection system noise model (appendix A.5.7) based on the signal.

The mean percentage error in the modelled noise as compared to the measured noise, ε_v (eq. (A.11)), was calculated for each of the four patterns. The mean value of ε_v across the four sets was 11.1% ($\sigma = 0.49\%$). Given that ε_v in the noise study experiments with a completely stable source (appendix A.5.6) was around half of this value this indicates that source does introduce some variability, however the fact that the noise model still has good predictive power suggests that the source noise is dominated by the detection system noise in imaging experiments.

A.6 Geometric Calibrations

It is necessary to know the geometry of the imaging system in order to carry out accurate surface reconstruction using the surface capture method (chapter 6) and to accurately calculate the free-space transfer function relating surface light fluxes to CCD measurements in the system (chapter 7). This section presents methods implemented for establishing the geometry of the imaging system in a system co-ordinate space that is defined as being centred around the camera

in x and y with $z = 0$ being co-planar with the top of the sample stage when the stage height is set to 0. A reduced depiction of the resultant complete system geometry model comprising the pupils and principal view-directions of components is shown in figure 6.6 in chapter 6.

A.6.1 Camera Pupil and Pixel Rays

The surface capture method uses a pin-hole camera model of the detection system and the free-space light propagation model uses a thin lens approximation. These models have in common a single effective 'pupil' location which is the pin-hole location in the first case and the centre-point of the thin lens in the second. It is necessary to establish the location of the pupil and beyond this the directions of the (assumed) straight-line rays that pass through it and are incident on the CCD at the pixel centre-points. A calibration routine was devised involving the detection system, the translation stage, and a printed grid. Note that measurement by hand (of the pupil location based on assumption that it was at the lens face) was attempted but found to be insufficiently accurate for surface capture to work effectively. The method that was devised works on the assumption that the camera optical axis is perpendicular to the imaging stage.

Method

A regular grid of dots was printed on standard A4 paper and the total size of the printed grid was measured by hand so that the spacing between dots was known. This was placed on the system stage and imaged at several heights by varying the stage height (specifically $h \in 0, 1..., 50$) with illumination being provided by the projectors. Three images from a sample data set are shown in fig. A.26.

It can be seen that there is a varying illumination but that there is good contrast on many of the dots of the printed grid. Each image was normalised to its maximum and then inverted (i.e. all pixel values are subtracted from 1). The spatially varying background was then isolated by morphological opening of the image (this is an erosion followed by a dilation using the same structuring element and has the effect of eliminating small objects from the image).

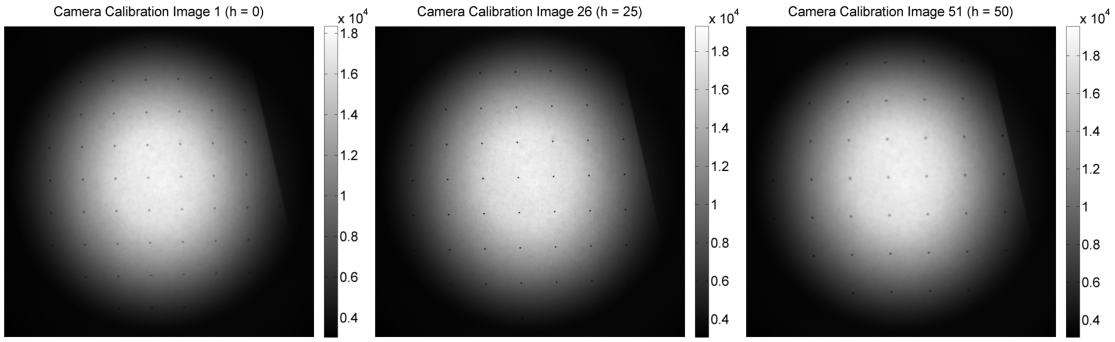

(a) $h = 0$mm (b) $h = 20$mm (c) $h = 50$mm

Figure A.26: Images of a printed regular grid at different stage heights, it can be seen that the perspective is changing and the grid size is different in the different images as is the effective sharpness of focus.

A disk-shaped structuring element was used with a radius of 20 pixels; the grid dots fit inside the structuring element and so are eliminated leaving only a smoothed background. This background was then subtracted from the original normalised image resulting in a high contrast image of the dots that were visible. This image was then thresholded with a threshold at 30% of the maximum intensity. The centroids (centres-of-mass) of all isolated regions were then calculated resulting in estimates of the grid dot centre-points. This work-flow is illustrated in fig. A.27.

A grid-fitting routine was then carried out per-image in which an initial estimate of the size of the grid, s_g, in pixels was made by taking the distance from each region centroid to its nearest neighbour. An initial estimate of the grid orientation, represented by the angle to the horizontal axis, θ_g, was obtained using the directions to the nearest neighbours. The approximate grid centre-point, $\mathbf{r}_g$, was established by taking the point nearest to the mean point.

The position and size of the grid was then refined by solving for

$$\{s_g, \theta_g, \mathbf{r}_g\} = \min_{s_g, \theta_g, \mathbf{r}_g} \sum_{i=1}^{n_r} |c_i - G_i(s_g, \theta_g, \mathbf{r}_g)|, \tag{A.13}$$

where c_i is the centroid of the i^{th} thresholded image (the approximate dot centres) and $G_i(s_g, \theta_g, \mathbf{r}_g)$ is the nearest point to c_i on the grid defined by the size, orientation and position parameters. The L_1 norm is used here because it was found to be more robust to outliers

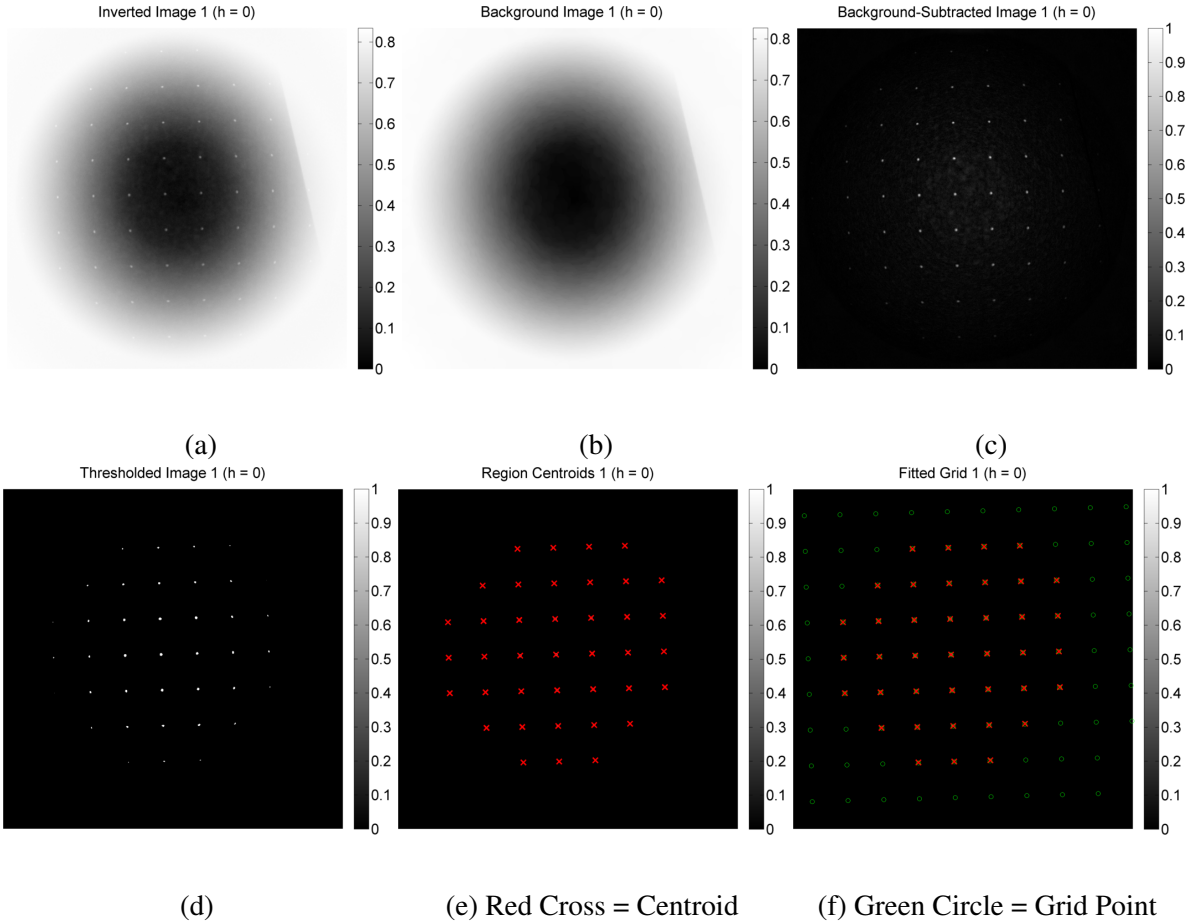

Figure A.27: Step-by-step illustrations of the image processing work-flow. The image is normalised and inverted (a) and the background is isolated by image opening (b), the difference of these images yields (c) the foreground; dots and high-frequency noise only. The image is then thresholded at 30% (d) and centroids of regions are extracted (e), these are then used as the initial guess for a grid-fitting routine to estimate the dot centres (f) with sub-pixel accuracy possible.

(e.g. image artefacts that are above the threshold) in preliminary studies. The fitting step was done for every height-resolved image.

The result was a stack of grid parameters for accompanying stage heights. The height-resolved grid-size parameter was then used in conjunction with the known size of the grid in physical units (mm) to establish a relationship between the stage height and the size of the region on the stage that corresponds to the single image pixels. This relationship is established by linear fitting using the MATLAB *polyfit* function as shown in fig. A.28.

Finally to estimate the height of the camera, h_c from the stage when it is set to height zero, the fitted straight line intercept with the y axis was used. This represents the point at

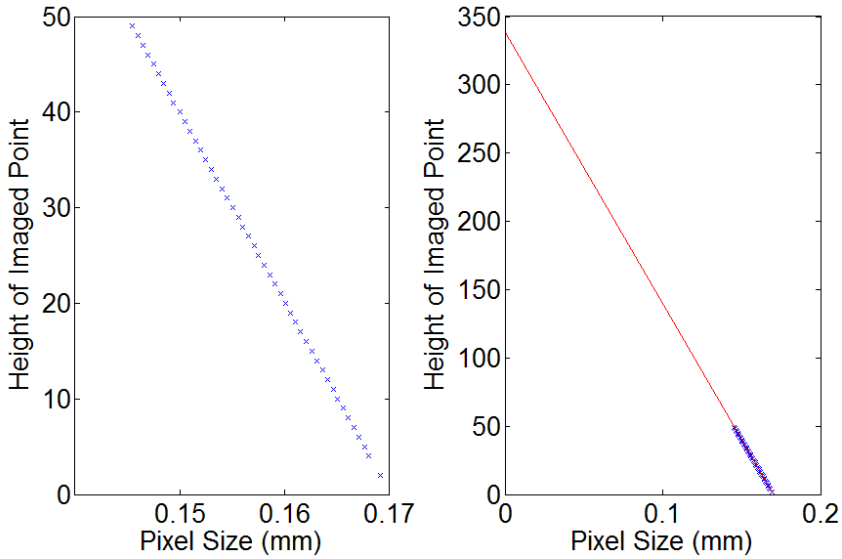

Figure A.28: Illustration of the method used to estimate the height of the camera with (a) a plot of apparent pixel size (the distance on the imaged sample stage corresponding to a single pixel width in each image) against stage height at which the image was acquired, and (b) a straight line fitted to extrapolate this data to the point at which the size of a pixel becomes 0mm which is then taken as an estimate of the camera pupil height.

which the size of a pixel tends towards 0 and therefore is the point of the pupil, in the presented example the extracted height value was 337.9mm. Note that this is only valid under the working assumption that the camera optical axis is perpendicular to the stage. Having obtained the height, the x and y co-ordinates of the camera were declared as 0. Thus the final position, $\mathbf{p_c}$, and direction, $\hat{\mathbf{n}}_\mathbf{c}$, of the camera obtained with the routine were

$$\mathbf{p_c} := (0,0,h_c), \tag{A.14}$$

and

$$\hat{\mathbf{n}}_\mathbf{c} := (0,0,-1). \tag{A.15}$$

The other parameters of the co-ordinate system i.e. the x and y directions are chosen to match the camera image-space x, y directions.

Performance

The benefit of the grid-fitting approach is that it allowed for sub-pixel resolution (because the enforcement of the grid properties coerce the estimated centre-points of the dots towards the correct location) which is necessary because of the large extrapolation needed to establish the camera location (small local errors become large at the extrapolation point).

It is difficult to analyse the performance of the method directly (though the combined performance of this and other methods is evident in the analysis of the surface capture (chapter 6) and the free-space methods (chapter 7). In an effort to at least establish the reproducibility of the result, 10 repeated runs were done from re-placement of the paper grid through analysis for a single system geometric set-up. The resultant height calculations from each set are plotted in fig. A.29.

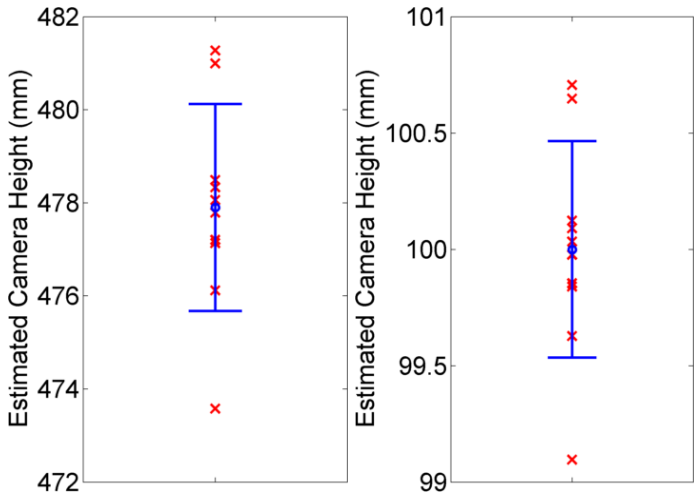

(a) (b)

Figure A.29: Plots of (a) height and (b) set-mean relative heights recovered in each one of ten repeats of the camera calibration routine with re-setting of the apparatus. Red crosses mark the result per-set and the blue circle and bars show the mean with standard deviation bars although with the low number of samples these have limited statistical relevance.

It can be seen that the 10 repeats produce repeatable estimates of the camera height, all 10 being within 1% of the sample mean. It is expected that the performance will be equally good given new system set-ups although it may require supervision in the image-processing phase if the background lighting is changed substantially (e.g. the threshold value may need to be

changed if the image is noisier), further tests with more geometries would be required to test this.

A.6.2 Projector Pupil and Projected Rays

The upper projectors are modelled in the same way as the camera system, with a single pupil representing an effective pin-hole and vectors representing the straight line path of, in this case, the light path *from* each pixel. A similar approach is used for finding the projector information as for the camera case in that a regular grid is used to add robustness to the fitting of dots which are imaged at multiple stage heights, in this case the dots are projected in a regular grid pattern provided to the projector. Image processing is then used to extract the grid co-ordinates in pixel-space as before and the camera model is used to put these into real space, then the projector location is fitted from this information. In the calibration of the system, each of the upper projectors are calibrated one-at-a-time.

Method

Plain white paper was placed on the sample stage to provide a high reflectance and therefore good contrast for imaged projections. For a range of stage heights from 0mm by 1mm to 50mm (inclusive), 4 images were then acquired with the projector projecting a *dark* (uniformly 0), a *bright* (uniformly 255; note 8-bit), a *grid* (a regular grid of 9×9 dots, the central dot of which was in the centre of the projection image), and a *centre-point* (the central dot from the grid projected on its own). Sample images for a single stage height are shown in fig. A.30.

The bright image was thresholded at 20% to provide a *mask* indicating where there was high projected signal in the images. Then the dark image was subtracted from each of the others to remove background illumination. The dark-subtracted images were then divided by the bright image to provide maps of relative intensity projected for the grid and centre-point patterns. These were then median-filtered to remove noise, normalised and thresholded at 10% to isolate the dots. The thresholded images were then masked using the previously constructed image

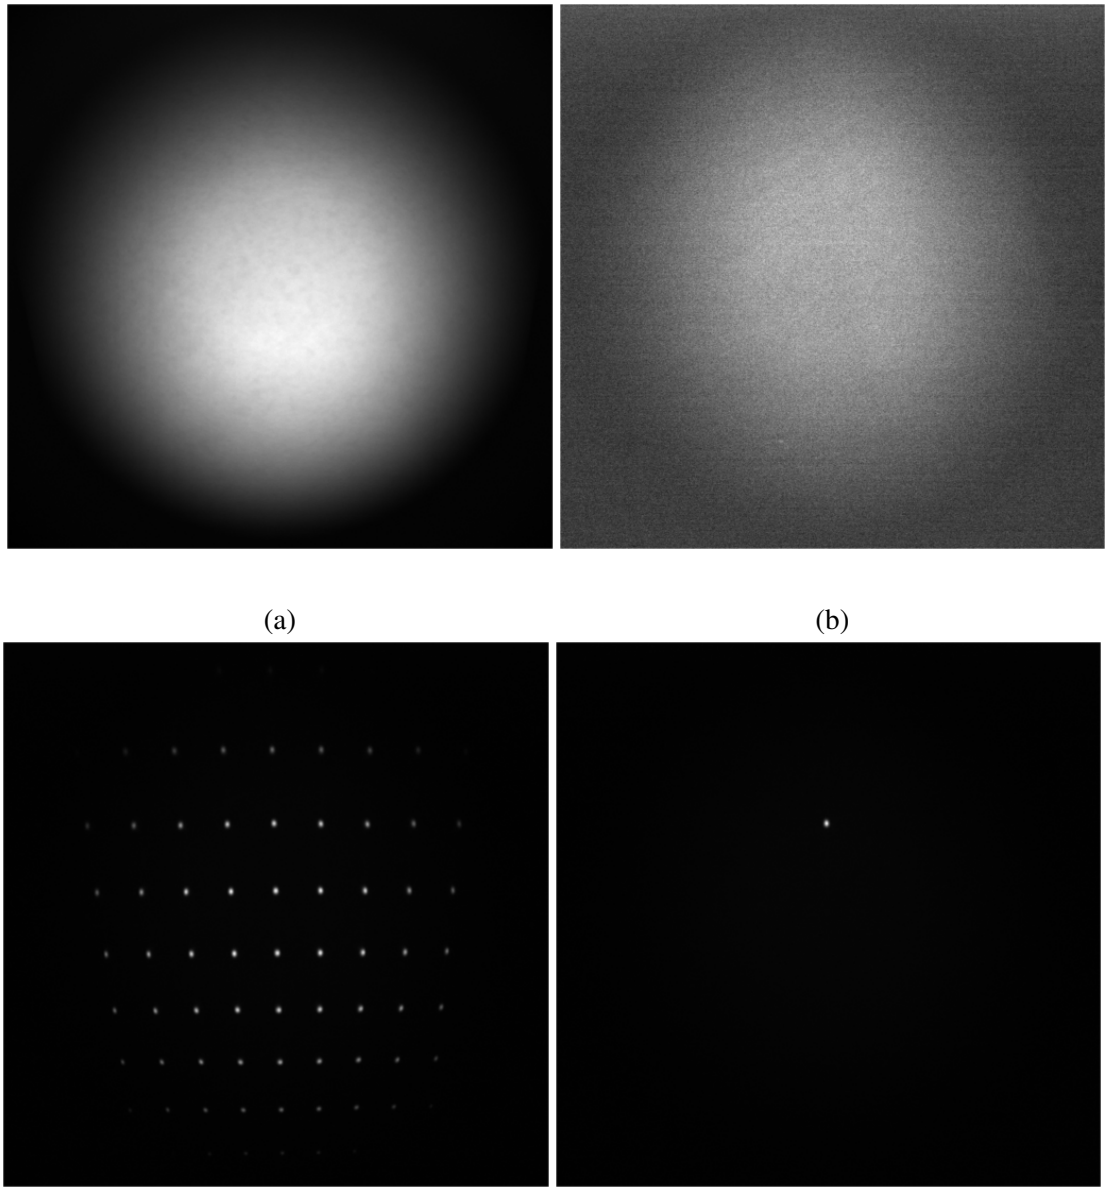

(a) (b)

(c) (d)

Figure A.30: Sample images of projections used in projector geometric calibration routine with (a) bright image projected, (b) dark image projected, (c) grid projected, and (d) grid centre-dot projected.

masks to black-out anything outside the signal area. The centroids of thresholded regions were then established as in the camera calibration. For the centroid image this was then extended such that only the centre-point of the largest region (assumed to be the dot rather than noise) was kept. The region extracted from the grid image whose centroid was closest to the centroid of the centre-point image region was then declared the central dot.

The above image processing produced a set of points in 2-D pixel co-ordinates representing dot centroids at each height with relatively poor accuracy (because no grid constraints or similar have been added at this stage). It was found that some noise would sometimes be present at this stage manifesting as unwanted points. To mitigate the impact of erroneous points, and under the assumption that there would be little change between adjacent heights, any point for which a point at an adjacent height was not within 20 pixels (x-y distance in pixels) was discarded as unreliable. Remaining points were labelled (given a number indicating which dot they correspond to in the original projected image) by initially labelling distinct dots in the lowest-height image (in which the field-of-view is largest on the stage) and propagating these labels upwards through the point stack to the nearest (in x-y) point at each level. The region-labelled points were then discarded if there was not a point visible at all heights for their region. The result was a set of lines traced through camera image pixel-space for projected dots one of which was recognised as the central line (fig. A.31).

Given heights and positions in x-y image space, the camera model (appendix A.6.1) was then used to convert these to absolute co-ordinates in the new co-ordinate system (fig. A.32).

It can be seen that the lines of grouped dots appear curved prior to the conversion into absolute co-ordinates but appear straight afterwards owing to the compensation for the difference in the camera perspective at each height.

At this stage, the points were dealt with directly within a method aiming to fit for the projector geometric parameters directly.

A straight line was then fit to each grouped set of points by a simple least squares fitting in 3D. Intersection points (in practice meaning the closest points) on each line with all others were then calculated and the average of these was taken as the initial guess for the projector location

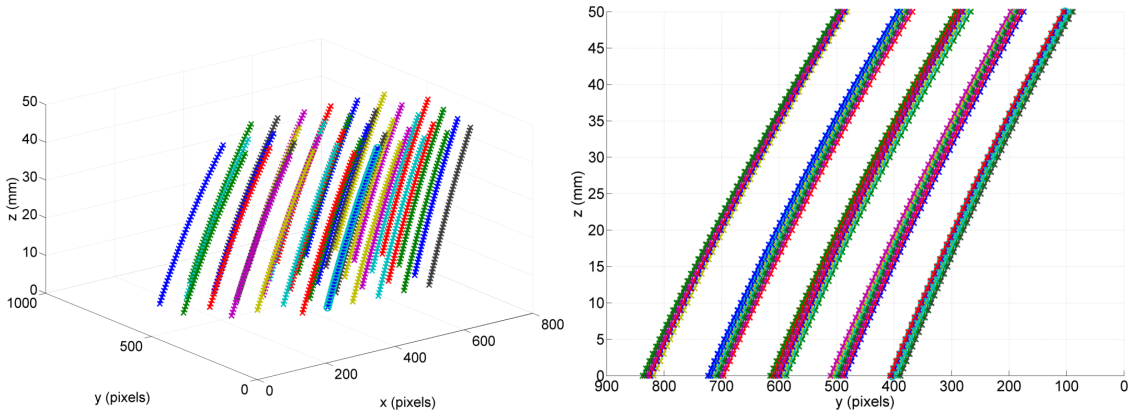

(a) (b)

Figure A.31: Dot centroids extracted from the projector grid analysis, the crosses indicate centroids coloured by their group label (though not uniquely, the same group indicates the same colour but not vice versa) which the central dot group having additional circles shown (visible near the front (towards y = 0) in (a)).

$\mathbf{p_p}$. An estimate for the size of the grid, s_q[1], in the zero-height plane was obtained by taking the minimum-valued second-nearest-neighbour distance for each point in the zero-height point set, the *second-nearest* was used because the first includes those distances artificially shortened by perspective (whereas the second is likely to be in the perpendicular direction to this in the grid in the minimum case). The central direction of projection, $\hat{\mathbf{n}}_\mathbf{p}$ was estimated initially as the direction from the estimated $\mathbf{p_p}$ to the zero-height centre dot. The remaining solved for parameter, a projector grid orientation angle θ_p[2], was initially set at 0.

A constrained optimisation was then performed to solve for

$$\{s_q, \theta_p, \mathbf{p_p}, \hat{\mathbf{n}}_\mathbf{p}\} = \min_{s_q, \theta_p, \mathbf{P_p}, \hat{\mathbf{n}}_\mathbf{p}} \sum_{i=1}^{n_d} |d_i - Q_i(s_q, \theta_p, \mathbf{P_p}, \hat{\mathbf{n}}_\mathbf{p})|, \tag{A.16}$$

where d_i is the estimated centroid of the i^{th} dot in system-space co-ordinates (fig. A.32) and $Q_i(s_q, \theta_p, \mathbf{P_p}, \hat{\mathbf{n}}_\mathbf{p})$ is a function that simulates the intersections of projected dots from the

[1] s_q is the distance between the points in the grid at the plane passing through the intersection of the projection with the $z = 0$ plane that is perpendicular to the projection direction

[2] A single angle is used to parametrise the grid orientation to simplify the optimisation problem. To generate the expected grid points, a co-ordinate basis is calculated using just this value in addition to those already included, a cross-product is taken of the principle projection direction $\hat{\mathbf{n}}_\mathbf{p}$ with the x-axis direction $(1,0,0)$. The result is used as a basis-direction for the projector-centred co-ordinate system, the final basis-direction is found via the cross-product of the first two. Within this basis the grid is rotated around the projection axis by angle θ_p

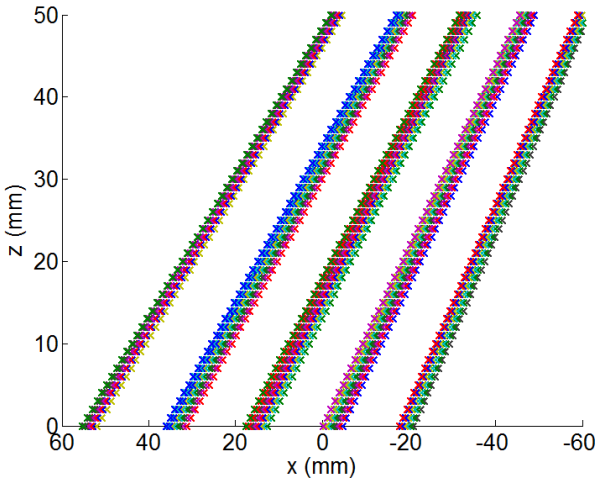

Figure A.32: Projected dot centroids in absolute co-ordinates established using the camera model.

estimated projector point through the parametrised grid with the stage planes at various heights. Thus the projector parameters are obtained, an example of the result is shown in fig. A.33.

Performance

Error in the camera model will propagate to the projector analysis to some extent because the camera model is used to establish the system-space co-ordinates of projected dot centroids, however the effect of this is expected to be small because whilst the camera model has error $\pm 1\%$ in the camera height, this translates to a smaller difference in the height-range of the stage.

Unlike in the case of the camera calibration routine in which it was assumed that the camera pointed perpendicular to the stage, no assumptions need to be made regarding the projector position when using the presented method. However, if the projector is positioned such that the central projection point cannot be seen the method as-is will fail (because it relies on identifying the centre-dot at all heights). It is suggested that this is a reasonable limitation given that the current usage of the projector means it will always be roughly pointing at the same point as the camera through the range of the stage.

In order to estimate the performance of the projector calibration, again without having a ground truth, 5 repeats of the projector calibration data acquisition were performed and then

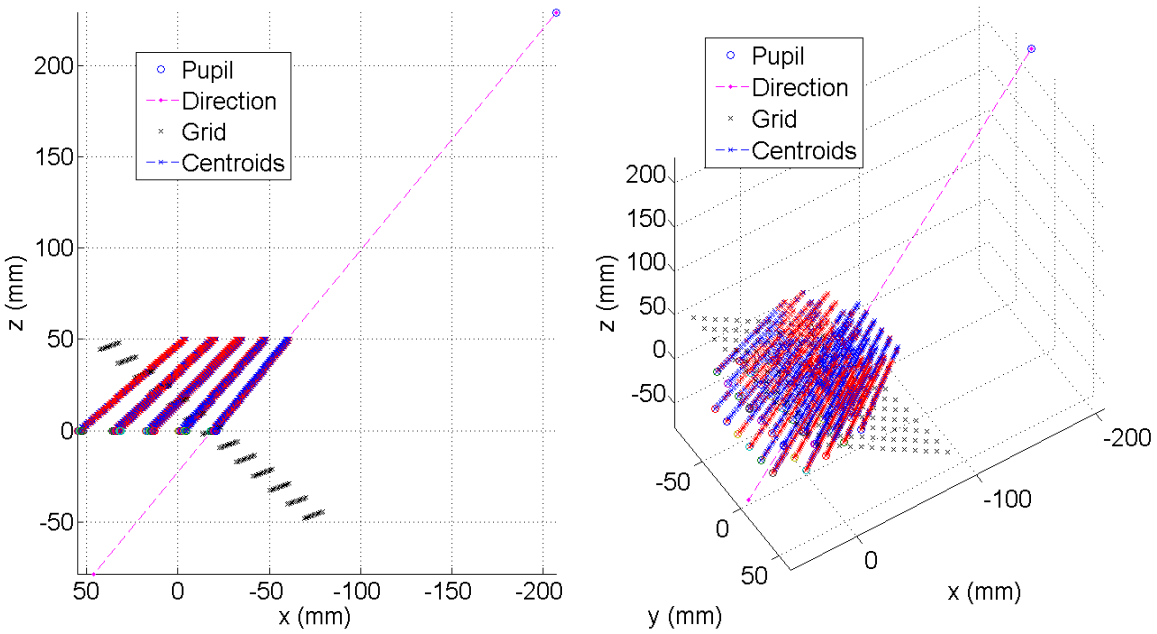

(a) (b)

Figure A.33: Sample result of the projector fitting, in addition to those markers appearing in the legend, the red crosses show the simulated projected grid dots that, given the fitted projector model, the closeness of the red and blue symbols indicates the good fit under visual inspection.

these were processed using each of the camera calibration results in the set of 10 repeats detailed above. This was done for each of the two projectors in the imaging system. The results are summarised in table A.4 and table A.5.

It can be seen that across the repeated runs there is some difference in the recovered pupil location $\mathbf{p_p}$ with most of the difference being in the x and z components. The largest absolute difference from the mean recovered point is 5.9mm indicating a total range of less than 11.8mm. This is similar to the range observed for the camera calibration. At the level of precision shown there is no difference in the calculated direction $\hat{\mathbf{n}}_\mathbf{p}$ for either projector. In order to judge the impact of the variations in $\mathbf{p_p}$ upon the geometry in the height-range of the stage, data for the point t_p is shown which is the point at which the projection line originating at $\mathbf{p_p}$ and in direction $\hat{\mathbf{n}}_\mathbf{p}$ intersects the zero-height stage plane. The standard deviation in the difference from the average t_p position is 0.1 and 0.0 for each projector indicating that there is very little difference in the geometry predicted by the model in the region of the stage. The recovered grid size and angle also vary very little indeed.

Value	Min	Mean	Max	St.D.
Pupil location, $\mathbf{p_p}$ *(mm)*	*(242.9,9.7,365.2)*	*(244.7,9.9,369.1)*	*(246.7,10.1,373.4)*	*(1.2,0.1,2.8)*
Euclidean distance to mean $\mathbf{p_p}$ *(mm)*	*0.8*	*2.9*	*4.7*	*1.1*
Unit direction, $\hat{\mathbf{n}}_\mathbf{p}$ *(mm)*	*(-0.6,-0.0,-0.8)*	*(-0.6,-0.0,-0.8)*	*(-0.6,-0.0,-0.8)*	*(0.0,0.0,0.0)*
Intersection with stage 0, $\mathbf{t_p}$ *(mm)*	*(-47.2,-4.1,0.0)*	*(-47.0,-4.1,0.0)*	*(-46.7,-4.1,0.0)*	*(0.2,0.0,0.0)*
Euclidean distance to mean $\mathbf{t_p}$ *(mm)*	*0.1*	*0.2*	*0.3*	*0.1*
Grid size s_q *(mm)*	*19.5*	*19.5*	*19.6*	*0.0*
Grid orientation $\theta_q(^\circ)$	*5.0*	*5.0*	*5.1*	*0.0*

Table A.4: **Projector 1**: summary of repeated results of projector calibration. Due to the large amount of data it is summarised here by the minimum, mean, maximum and standard deviation of key values across the set and shown to 1 decimal place to save space.

Value	Min	Mean	Max	St.D.
Pupil location, $\mathbf{p_p}$ *(mm)*	*(-212.7,-29.7,308.1)*	*(-208.5,-29.2,312.1)*	*(-205.1,-28.7,316.2)*	*(3.0,0.4,2.9)*
Euclidean distance to mean $\mathbf{p_p}$ *(mm)*	*1.2*	*4.0*	*5.9*	*1.3*
Unit direction, $\hat{\mathbf{n}}_\mathbf{p}$ *(mm)*	*(0.6,0.1,-0.8)*	*(0.6,0.1,-0.8)*	*(0.6,0.1,-0.8)*	*(0.0,0.0,0.0)*
Intersection with stage 0, $\mathbf{t_p}$ *(mm)*	*(4.1,2.4,0.0)*	*(4.2,2.4,0.0)*	*(4.2,2.5,0.0)*	*(0.0,0.0,0.0)*
Euclidean distance to mean $\mathbf{t_p}$ *(mm)*	*0.0*	*0.0*	*0.1*	*0.0*
Grid size s_q *(mm)*	*15.7*	*15.7*	*15.7*	*0.0*
Grid orientation $\theta_q(^\circ)$	*10.3*	*10.3*	*10.4*	*0.0*

Table A.5: **Projector 2**: summary of repeated results of projector calibration. Due to the large amount of data it is summarised here by the minimum, mean, maximum and standard deviation of key values across the set and shown to 1 decimal place to save space.

Figure A.34 shows the recovered projector positions from each analysis for each projector. It can be seen that the positions fall into visiually distinct groups dependent on the projector data set of origin and that there is a characteristic position difference due to the camera set. It can also be seen that all of the positions fall roughly on a diagonal line which is approximately the direction of projection. There is very little ($\approx 1mm$) variation in the y direction in either projector (the direction *into* the system) with rather more ($5-10mm$) in x and z.

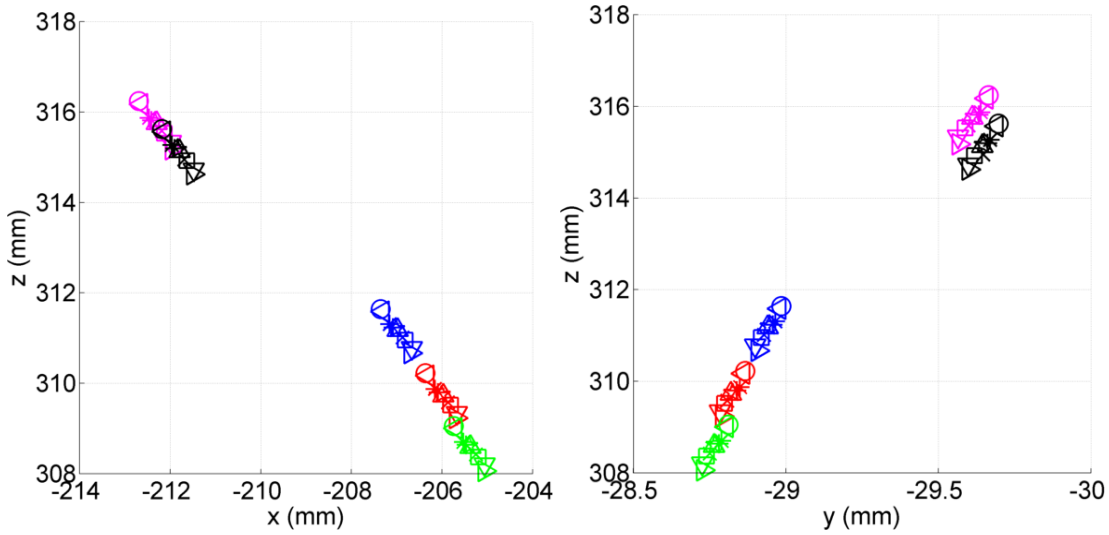

(a) Projector 1 Front View

(b) Projector 1 Side View

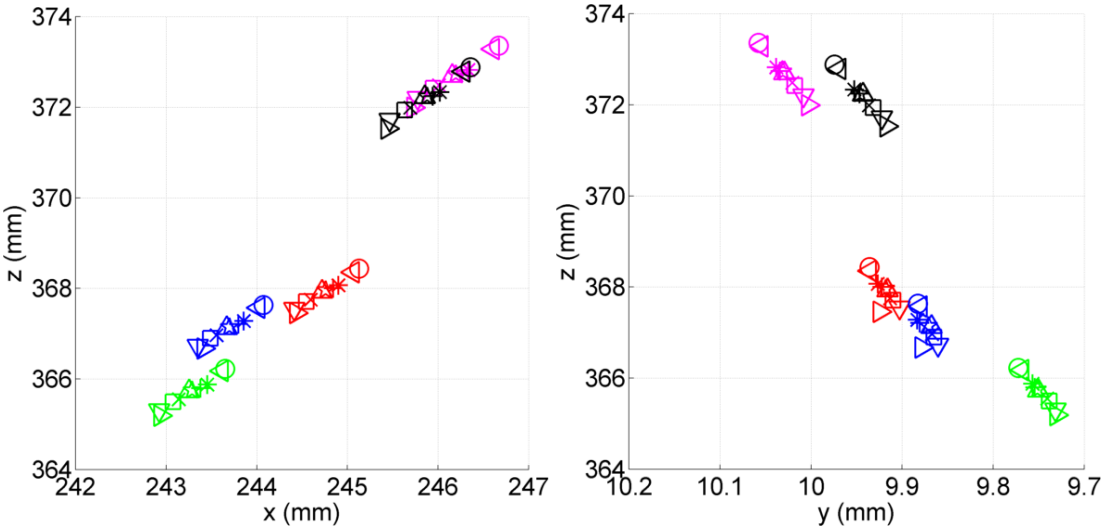

(c) Projector 2 Front View

(d) Projector 2 Side View

Figure A.34: Calculated projector pupil points from 5 repeated projector calibration data sets and 10 repeated camera calibration data sets. Camera sets are indicated by the symbol in the plot whilst projector sets are indicated by colour.

The analyses show that there is some difference in the recovered pupil location but little difference in other parameters. The difference in the pupil position appears to be along the direction of projection having the effect of slightly altering the focal length of the optics in the model and slightly tightening or broadening the field-of-projection, however by looking at the impact of the changes in geometry at the stage region, it appears that the modelled geometry is stable to within a few hundred microns.

A.6.3 DOT Source Model

Whilst Gaussian sources are projected for DOT (section 8.2.1), these are modelled with NIRFAST[25] as point sources[200]. The location of each of the point sources was established using a simple experiment in which each projection was imaged in turn through each filter used for DOT onto white paper. The resultant images were used in conjunction with the system geometric model (section A.6) to establish a 3D source position in the plane of the sample stage. These co-ordinates are used explicitly as the FEM source positions following their movement to the nearest point lying one scattering length inside the surface[200]. This experiment could have also been used to establish the relative brightness of different sources as seen by the camera but in the present work this step has been omitted in favour of calibrating the data directly with a reference data set (section 8.2.2).

A.6.4 Object Distance

The position of the focal plane of the detection system must be known for the free-space light propagation modelling, and is not calculable from the quantities that have thus far been measured/estimated by the methods in this section. Within the thin lens model the focal plane lies at a distance u from the point p in the direction of the optical axis.

In the current model of the imaging system, it is assumed that the sample stage is perpendicular to the optical axis which is also the modelled z-axis. Thus it is assumed that the focal plane lies parallel to the sample stage. Given this assumption and given that an attempt has been

made to manually focus the imaging system on the sample stage when the translation stage is set to 25mm height (approximately half the travel available), it is then assumed that the focal plane lies somewhere within the range travel of the sample stage as the translation stage.

The following experiment was then designed and carried out to measure the distance u under the above assumptions and using the principle that the focal distance u is by definition the distance at which a point source will appear smallest on the sensor (i.e. the point spread function is at its sharpest and narrowest).

Experiment

A piece of black card was pricked with a compass so as to produce a tiny hole in it. This was then placed on the sample stage over the trans-illumination window with the projector source (unfiltered) switched on in full-field illumination mode. This creates the effect of a point-like source near to the height of the sample stage from the perspective of the detection system.

Images were then acquired with the stage height varied from 0mm by 1mm to 50mm inclusive. A symmetric Gaussian was fitted to each image following offset-subtraction (appendix A.2.3) and the stage height that gave rise to the smallest fitted Gaussian size was taken as the height of the focal plane on the understanding that the point-like source would be blurred more giving rise to larger Gaussian images further from the focal plane. By this method the focal plane height, h, was found to be 21mm.

The relationship between the object distance u, the camera pupil $\mathbf{p_c}$ and the stage height h is

$$u = \mathbf{p_c}_z - h \tag{A.17}$$

Thus having estimated h and having established $\mathbf{p_c}$ with the camera calibration (appendix A.6.1), the object distance was calculated.

A.6.5 Focal Length and Image Distance

The size of the image plane (CCD), the size of the focal plane and the distance to the focal plane are known. Given the pin-hole camera approximation, this is sufficient to calculate the image distance v using basic geometry.

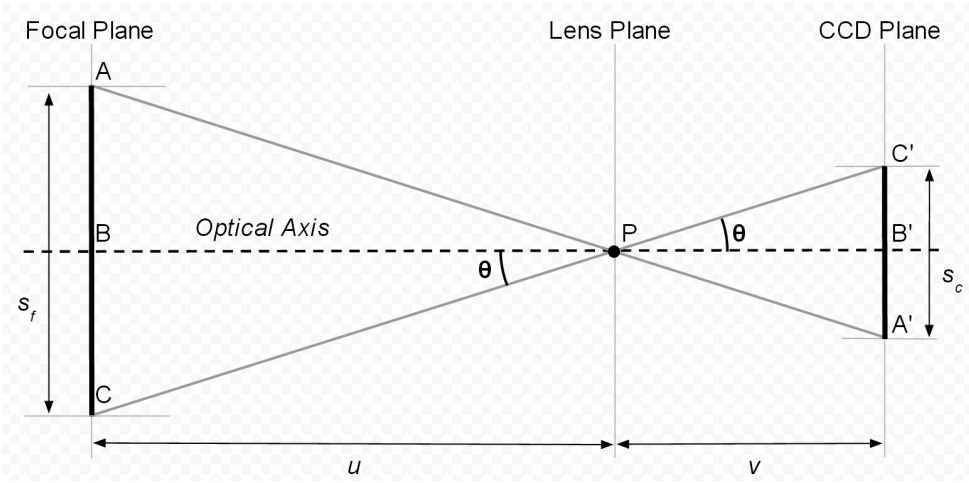

Figure A.35: Schematic of assumed pin-hole camera geometry

As shown in fig. A.35 the triangles $\triangle ABP$ and $\triangle A'B'P$ are similar; the angle θ made with the optical axis is the same in both, as such the image distance v and the object distance u are related by

$$\tan(\theta) = \frac{\overrightarrow{AB}}{u} = \frac{s_f}{2u} = \frac{\overrightarrow{A'B'}}{v} = \frac{s_c}{2v}, \tag{A.18}$$

in which s_f is the width of the focal plane and s_c is the width of the image plane on the sensor. Rearranging gives

$$v = \frac{us_c}{s_f} \tag{A.19}$$

thus v can be calculated directly from quantities already obtained. Given u and v, the focal length can be calculated using the thin lens equation:

$$\frac{1}{f} = \frac{1}{u} + \frac{1}{v}, \tag{A.20}$$

which gives

$$f = \left(\frac{1}{u} + \frac{1}{v}\right)^{-1}, \tag{A.21}$$

Using the equations above, along with the values already obtained, the image distance was calculated as 25.8234mm and the focal length was calculated as 23.8776mm for the present system set-up.

A.7 Phantom Characterisation

A.7.1 Introduction

In order to test and validate the DOT part of the imaging system, it is necessary to have phantoms (tissue mimicking materials) for which the optical and geometrical properties are well-known. This is necessary because in order to evaluate the performance of any DOT results, ground truth or "gold standard" optical properties are first required with which the results can be compared. Note that well-known (in terms of optical properties) phantoms are also required in order to test the BLT component of the system by providing accurate prior knowledge.

The following section describes work that was performed in collaboration with the Centre for Ultrafast Science and Biomedical Optics (CUSBO) at the Politecnico di Milano in order to characterise solid phantoms that are then used for the imaging system validation and testing.

The experiments were performed by Ilaria Bargigia, Andrea Farina, and Antonio Pifferi, and they also performed the processing of the raw system data up to the point of obtaining optical absorption and scattering spectra. Hamid Dehghani calculated libraries of temporal point spread functions for several of the source-detector configurations considered practically.

A.7.2 Time-Resolved Diffuse Optical Spectroscopy

A time-resolved diffuse optical spectroscopy system was used to measure the optical properties of phantoms. The system, which has been described by Bassi et al.[208], utilises ultrafast laser pulses along with single-photon time-resolved detection to measure the temporal point spread function (TPSF) of incident light. A model-based fitting method is then used to fit for the optical properties given the time-resolved measurements.

The system comprises a broadband laser source coupled to a prism which splits the light spectrally, the prism is attached to an automated rotating stage which translates the spectrally separated light, changing the spectral range that passes the pin hole and therefore acting as a variable filter. The filtered light is then put into a sample by a fibre tip in contact-mode and measured by another contact-mode detection fibre which is directed into a detector. The original schematic of the system is shown in figure A.36.

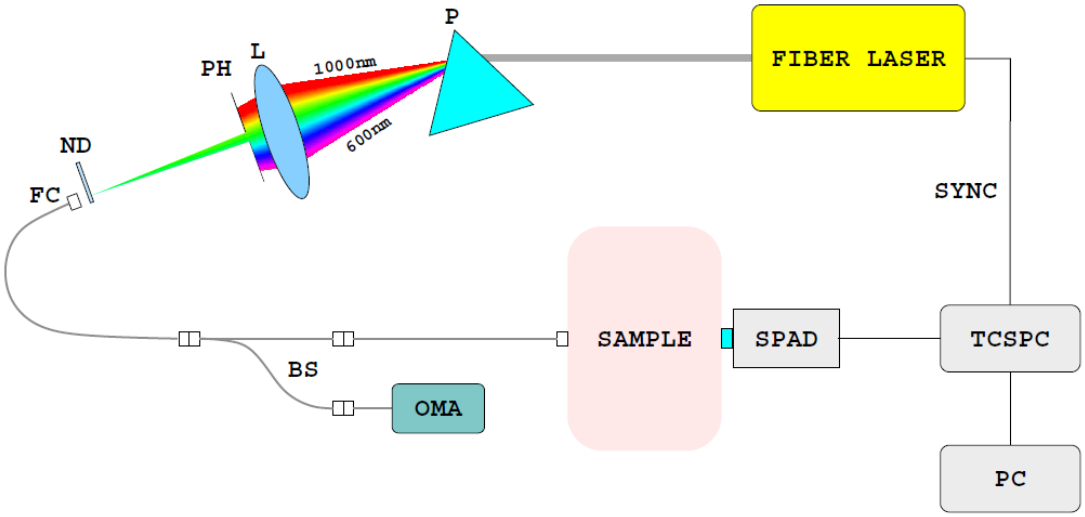

Figure A.36: Time-resolved diffuse optical spectroscopy system layout. This figure is reproduced exactly from Figure 1. in Bassi et al.[208]. Marked components are encoded prism (P), lens (L), pinhole (P), variable neutral density filter (ND), fibre connector (FC), beam splitter (BS), optical spectrum analyzer (OMA), single photon avalanche diode (SPAD), and time-correlated single-photon-counting board (TCSPC).

In the figure the source and detectors are arranged in transmittance mode; they appear on opposite sides of the sample. In the experiments presented here, the particular source-detector configuration was varied, this means that it could sometimes be called reflectance or transmittance measurements based on whether the optodes were on the same side or opposite sides of the phantom. In any case, the optode locations are explained individually for each experiment.

The detector shown in the figure is a single-photon avalanche diode (SPAD). However, in the experiments that follow this detector was not used, rather it was replaced by one of two other detectors, one of which is referred to as ‘S1’ and the other of which is the ‘Hybrid’, the main differences (and the reason for using two different detectors) are their sensitive spectral ranges, their temporal resolution and their sensitivity, using both detectors therefore allowed for

a more comprehensive dynamic range and spectral range as well as a *second opinion* for several results.

A.7.3 Accounting for Arbitrary Irregular Shapes

The optical property fitting method requires a library of expected TPSFs for a range of possible reduced scattering coefficients and no absorption. In the usual approach these are calculated using Monte Carlo with a semi-infinite slab approximation for the geometry of the sample. Libraries also exist for specific common and simple geometries (e.g. a cube or a sphere). However, in the case of the phantoms used here the geometry is in some cases irregular, in all cases far from the semi-infinite slab approximation (partly because of the small size). Errors are found when using the semi-infinite approximation when the geometry is far from this in reality because there are additional unexpected losses at the physical boundaries, so-called "boundary effects". In order to compensate for these effects, the following method was used.

Rather than approximating the shape of the imaged sample using a semi-infinite slab or any other regular shape, the phantom is instead represented as a geometrically accurate finite element model (FEM) mesh. Using time-resolved NIRFAST software[209, 33], the time-resolved propagation of light through the sample is then modelled for the particular source and detector position used in each experiment, for a range of reduced scattering coefficients and with an absorption coefficient of 0. This allows experiment, subject and optode-configuration dependent libraries to be built. The experiment-specific library is then used for the fitting rather than the closest-available standard library. This experiment-specific modelling is theoretical possible using Monte Carlo but, being as it uses the diffusion equation and is somewhat optimised for performance, the NIRFAST approach is much faster allowing for this to be done within a normal experimental time-frame for the system.

	T_{neck}	T_{centre}	$R_{0.5}$	$R_{1.0}$	$R_{1.25}$	$R_{2.0}$	R_neck
XPM2	$*$	$*'$	$**$		$**$	$**$	$*$
Cylinder		$*'*'$					
Block		$*'$			$**$		

Table A.6: Summary of experiments performed with the S1 source. Key: $*$ = completed experiment; $\times$ = failed experiment; $'$ = set-up modelled with NIRFAST.

	T_{neck}	T_{centre}	$R_{0.5}$	$R_{1.0}$	$R_{1.25}$	$R_{2.0}$	R_neck
XPM2		$*'*'$	$*$	$*'$			
Cylinder		$\times*'*'$					
Block		$*'$	$*$	$*'\times$			
Polyurethane		$*$					

Table A.7: Summary of experiments performed with the hybrid source. Key: $*$ = completed experiment; $\times$ = failed experiment; $'$ = set-up modelled with NIRFAST.

A.7.4 Results and Discussions

Experiments were undertaken with the phantoms presented in appendix A.1.1, with different source-detector geometries such as transmittance through the centre of the object, and reflectance on a flat face on the object with different source-detector separations. The aim was to investigate whether the reconstructed optical property estimates were robust to the choice of source-detector geometry. The relevant assumption that was made is that if measurements are stable over different source-detector geometries (being therefore subject to different, if any, boundary effects) then the result is more reliable. The multiple investigations also serve as repeated observations to validate other measurements. Tables A.6 and A.7 summarise the experiments that were carried out using the S1 and the hybrid detectors respectively. A diagram of each of the different source-detector coupling geometries is shown in figure A.37.

It can be seen that in total 14 experiments were carried out using the S1 detector and 10 were completed with the Hybrid making for a total of 24 data sets. The wavelength range measured in each case was dependent only on the detector, the S1 set were done for 500 to 1000nm whilst the Hybrid sets were done for 500 to 900nm. This is due to differing spectral response characteristics of the detectors.

The table also indicates those instances where the geometry was sufficiently well-known or well-controlled as to be representable within NIRFAST or using geometry-specific Monte

id	Experiment No.	Sample	Mode	d (mm)	fig. A.37#	Detector	Fitting
1	0	XPM-2	T	22	1	S1	Semi-Infinite
2	1	XPM-2	R	12.5	2	S1	Semi-Infinite
3	2	XPM-2	R	20	2	S1	Semi-Infinite
4	3	XPM-2	R	20	2	S1	Semi-Infinite
5	4	XPM-2	R	12.5	3	S1	Semi-Infinite
6	5	XPM-2	T	14.7	4	S1	Semi-Infinite
7	6	Block	T	26	8	S1	Semi-Infinite
8	6	Block	T	26	8	S1	Brick
9	7	Block	R	12.5	6	S1	Semi-Infinite
10	8	Block	R	12.5	6	S1	Brick
11	8	Block	R	12.5	6	S1	Semi-Infinite
12	9	Cylinder	T	25	7	S1	Semi-Infinite
13	10	Cylinder	T	25	7	S1	Semi-Infinite
14	11	XPM-2	R	5	2	S1	Semi-Infinite
15	12	XPM-2	R	5	3	S1	Semi-Infinite
16	13	XPM-2	R	5	5	S1	Semi-Infinite
17	14	Block	T	26	8	Hybrid	Semi-Infinite
18	16	Cylinder	T	25	7	Hybrid	Semi-Infinite
19	17	Cylinder	T	25	7	Hybrid	Semi-Infinite
20	18	XPM-2	T	22	1	Hybrid	Semi-Infinite
21	19	XPM-2	R	5	3	Hybrid	Semi-Infinite
22	21	XPM-2	R	10	2	Hybrid	Semi-Infinite
23	16	Cylinder	T	25	7	Hybrid	NIRFAST
24	16	Cylinder	T	25	7	Hybrid	Semi-Infinite
25	17	Cylinder	T	25	7	Hybrid	NIRFAST
26	17	Cylinder	T	25	7	Hybrid	Semi-Infinite
27	20	XPM-2	T	22	1	Hybrid	NIRFAST
28	20	XPM-2	T	22	1	Hybrid	Semi-Infinite
29	21	XPM-2	R	10	2	Hybrid	NIRFAST
30	21	XPM-2	R	10	2	Hybrid	Semi-Infinite
31	23	Block	R	10	6	Hybrid	NIRFAST
32	23	Block	R	10	6	Hybrid	Semi-Infinite
33	24	Block	R	10	6	Hybrid	NIRFAST
34	24	Block	R	10	6	Hybrid	Semi-Infinite
35	25	Block	R	5	6	Hybrid	Semi-Infinite

Table A.8: Processed (optical property) data sets with; (id) a number identifying the processed data set in graphs and discussions of this section, (Experiment No.) a number identifying the experiment that was the source of the data for the fitting, (Sample) the phantom being studied, (Mode) 'T' or 'R' indicating transmittance or reflectance measurement, (d) the Euclidean distance between source and detector, (fig. A.37#) an index to a diagram of the geometric set-up shown in fig. A.37, (Detector) the detector used, and (Fitting) the computational model used to fit.

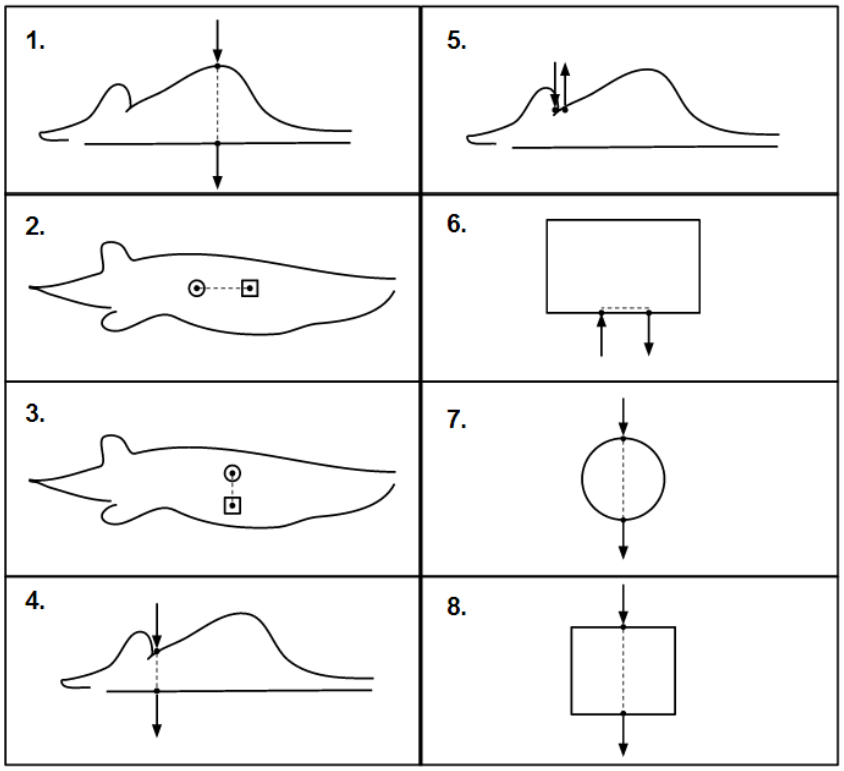

Figure A.37: Source-detector couplings used in experiments.

Carlo models in order to calculate TPSFs which took account of boundary conditions. Note that in cases where it is a reflectance measurement with a source-detector separation of less than 10mm it was found that the NIRFAST model could not produce accurate TPSFs due to the short path-length and the resulting lack of accuracy of the diffusion model.

The experimental data was processed by fitting for absorption and reduced scattering coefficients with the recorded TPSF and a computational model. The usual model is that of a semi-infinite slab (which is valid for large objects with flat surfaces in contact with the optodes and sufficiently high absorption that boundary effects are negligible) but in some cases it was possible to use a theoretically more appropriate model in that a 'brick' Monte Carlo model existed and in other cases the locations of optodes were recorded precisely and FEM meshes were used to create bespoke libraries using NIRFAST. The resultant processed data sets (comprising μ_a and μ_s' values) are itemised in table A.8.

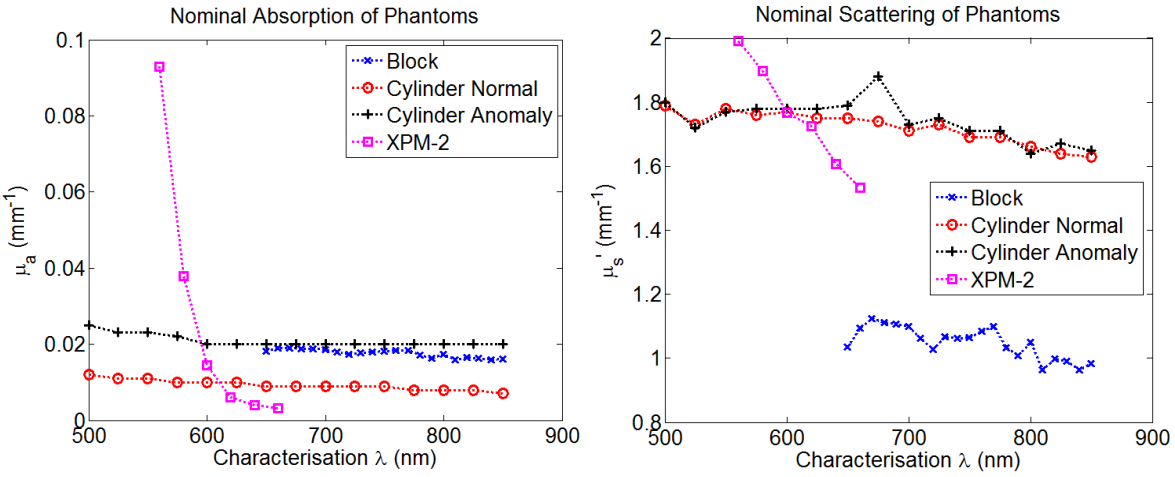

(a) (b)

Figure A.38: Best estimated, measured or assumed optical properties gathered prior to the practical experiment undertaken to measure them.

Previously Assumed Optical Properties

Figure A.38 shows the optical properties that were the estimates prior to the current experiments. The data was different in origin across the phantoms. For the block phantom, it is believed that the data was provided by the manufacturer some years ago but these were received second hand and therefore are somewhat untrusted. The values for the cylinder were received very recently when the cylindrical phantom was manufactured (in 2013), and these are expected to be reliable. The XPM-2 properties were read from the only available resource for it which is a small (10cm by 10cm) booklet provided by the manufacturer with the phantom. The data in this case was presented on a log scale in a very small graph that was therefore difficult to read accurately. As such it is expected that there is some substantial but not precisely known error in the expected values for the XPM-2.

It can be seen that the properties that were available before this experiment were also very different in terms of the spectral ranges for which values were available.

Given these facts, one aim of the present experiment is to validate or otherwise improve upon the estimate of optical properties whilst another is to extend the range of wavelengths for which values are known.

All samples are assumed to have a refractive index of 1.54 because they all used polyurethane

as a base material and this is a typical value for this material.

Fitted Properties: Semi-Infinite Slab Model

Figures A.39, A.41 and A.42 show the first processed results for the fitting of μ_a and μ_s' along with the nominal (fig. A.38) properties for reference. These were attained using a semi-infinite slab Monte Carlo model for the fitting. The legends in the figures indicate the processed data number which indexes the data in table A.8 showing the details of each data set. Across all the semi-infinite slab processed data sets it can be seen that there is substantial variation in the results which it is proposed is most likely due to varying levels of boundary effects - where there is significant mismatch between the semi-infinite model and the practical situation because of the physical boundaries of the phantom that introduce additional losses not due to internal attenuation levels and not represented in the model. This will now be discussed in more detail with reference to each figure.

Figure A.39 shows the semi-infinite slab model fitting results for the XPM-2 mouse phantom. It can be seen (most clearly in the log-plot) that there is substantial difference in the measurements of μ_a in particular where the amount of absorption is lowest. To illustrate this point, the difference between the highest and the lowest measurement of μ_a at 700nm is 0.0028mm^{-1} which represents approximately a five-fold variation. However, it can be seen that the general shape of the curve is the same in each data set. With reference to the specific geometry used in each experiment (table A.8), it is observed that the absorption measurements tend to rise as the distance between the source and detector in the experiment increases, with the lowest (set 15) coming from a reflectance measurement with 5mm optode-separation-distance (which will hereafter be referred to as d in this section). The highest μ_a measurement is from data set 6 which was a transmittance measurement in which d was 22mm, the full thickness of the mouse phantom. It is proposed that the apparent correlation between d and the recovered absorption values is consistent with the notion that the increase in apparent absorption is boundary effects, the amount of which could be expected to increase with d because of the increased opportunity for light losses at boundaries with longer path lengths. Light that is lost at the physical bound-

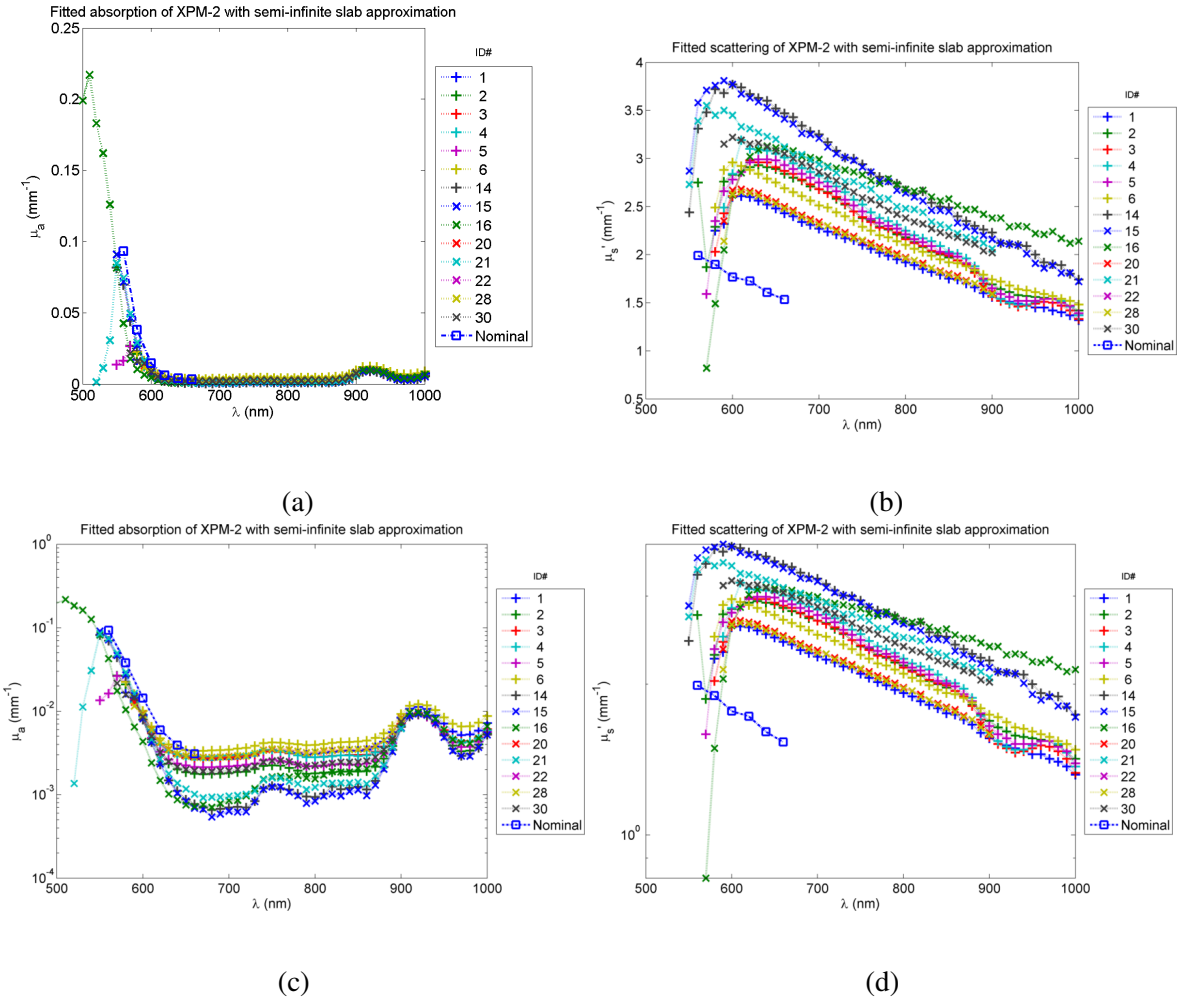

Figure A.39: Fitted (a) absorption and (b) scattering of the XPM-2 phantom with semi-infinite slab approximation, repeated (c-d) on a log-scale plot.

aries (other than the two boundaries of the model) is not accounted for in the semi-infinite model so can only be attributed to absorption. In this line of thinking, it is deemed most likely that the lowest absorption measurement (with the smallest d) is probably most accurate.

Figure A.40 shows an illustration of the difference in a short-distance reflectance measurement and a longer-distance transmittance through a cross-section of the mouse phantom in terms of the present boundary effects that are either known or unknown to the semi-infinite slab model.

Where there is higher absorption there is apparently less difference between the measurements. This is consistent with blaming boundary-effects because higher absorption narrows the detected photon spatial distribution (marked fluence in figure) meaning that less of the modelled

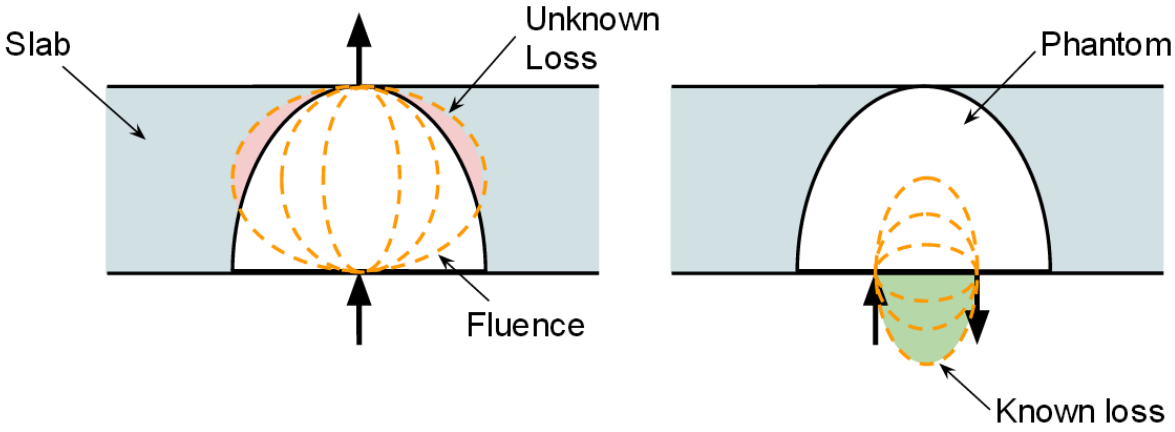

Figure A.40: Illustration of the boundary effects introduced that are not accounted for by the model in some geometries. The illustrative contour lines marked *fluence* indicate the photon path through the slab model (with decreasing number of photons per contour).

light escapes through the physical boundary.

It can also be seen that the nominal μ_a values are similar in shape to the measured ones, and similar in value (obviously more so with some measurements than others due to the range).

For scattering this is much less the case, it can be seen that whilst the scattering varies, the measurements are consistently higher than the nominal values, possibly indicating a poor nominal value set. The scattering values are also substantially more stable across data sets than the absorption ones; whilst it can be seen that there is a great deal of variability near the start of the curve (where many measurements will have had low SNR owing to the high absorption and lower QE of the detectors and leading to a poor result), there is generally the same trend in the curves and a maximum difference in the measurement at 700nm (chosen as an arbitrary point of comparison) of 0.98mm^{-1} (ranging from 2.27mm^{-1} to 3.25mm^{-1}) which in this case only represents a 1.4-fold difference in contrast, much less than is seen in μ_a. This level of variation is not much more than the level expected owing to the repeatability of the instrument (causing reasonable measurement ranges of approximately $\pm 10\%$ in μ_s), but is possibly also influenced to some extent by boundary effects.

Figure A.41 shows the properties fitted for the cylinder using the semi-infinite slab model. It can immediately be seen that the fitted μ_a values are substantially less variable as compared to

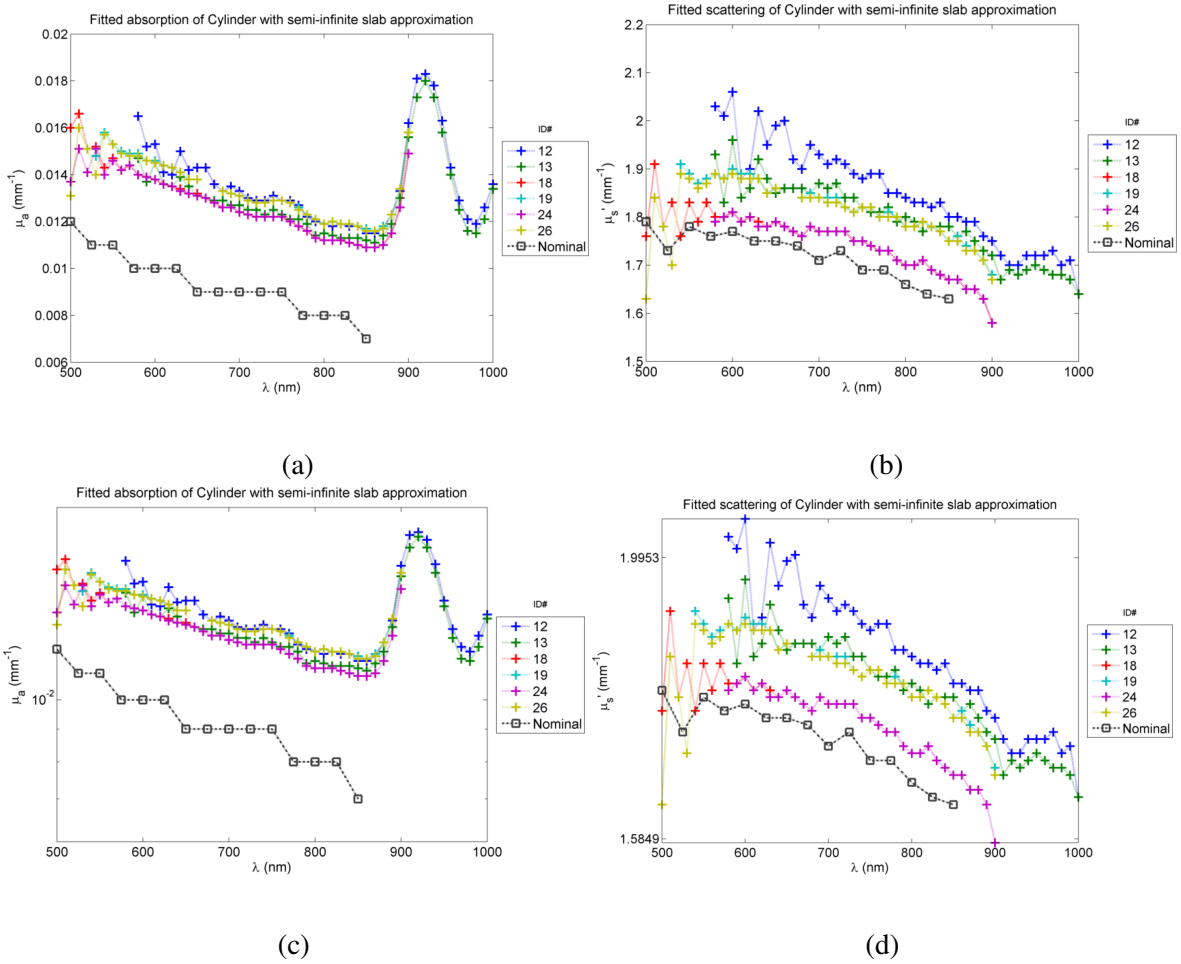

Figure A.41: Fitted (a) absorption and (b) scattering of the Cylinder phantom with semi-infinite slab approximation, repeated (c-d) on a log-scale plot.

those of the mouse phantom being only on the order of $\pm 10\%$ difference. It must be noted that the geometry for the cylinder in all cases was the same but for a rotation of 90 degrees introduced to check whether or not the internal background-matching rods caused any difference. The fact that there is little difference in the values suggests that the rods are not causing a disturbance (which is good because this is their purpose) and that the measured spectrum is also independent of the detector used (because half of these measurements were made with each of the S1 and hybrid detectors). The data tells us nothing about boundary effects because the geometry from this point-of-view was identical in all cases, however, if there *is* some influence as might be expected based on the above discussion, this may explain why the measurements are higher than the nominal values which for this data set are expected to be accurate. The scattering varies in a similar manner, again being on the order of $\pm 10\%$ between sets, this is similar to

that observed on the XPM-2 and to the absorption variation in the cylinder. The scattering values are also close to the nominal values but consistently higher than expected.

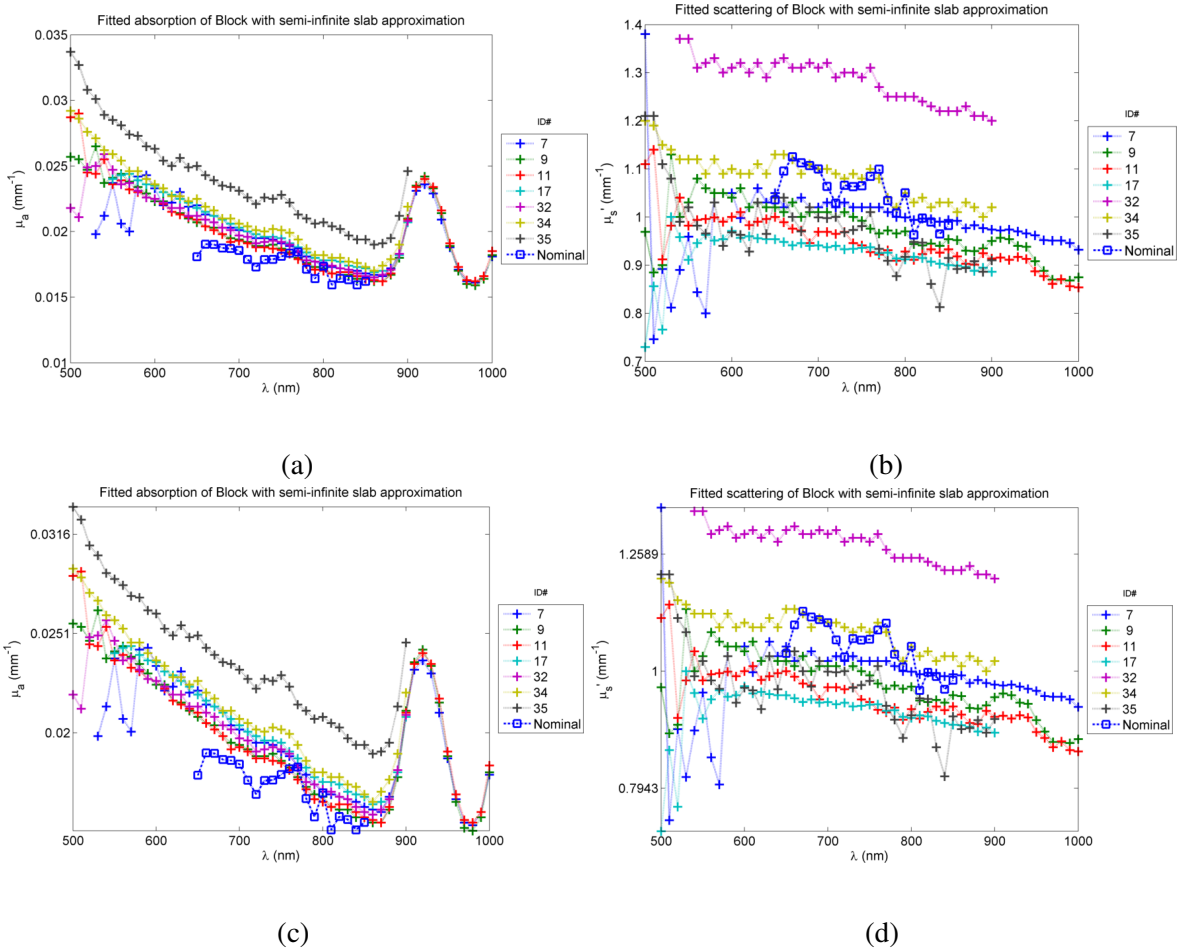

Figure A.42: Fitted (a) absorption and (b) scattering of the Block phantom with semi-infinite slab approximation, repeated (c-d) on a log-scale plot.

Figure A.42 shows fitted optical properties for the block phantom. Unlike the cylinder measurements where the source-detector geometry was not varied, the block experiments were highly variable in terms of set-up; for the 7 results shown, there were a transmittance mode acquisition with $d = 26$ and reflectance mode measurements with $d \in 5, 10, 12.5$. Note that whilst it is still presented, a problem was observed with measurement number 23 (underlying processed set 32 on the graph) in that following the measurement it was found that the black foam protecting the fibre from external light sources had worked loose meaning the measurement was unreliable.

The μ_a values are fairly consistent across data sets varying by around $\pm 20\%$, specifically at

700nm the highest μ_a value found is 0.023mm^{-1} whilst the lowest is 0.019mm^{-1}, though the highest curve clearly stands out on the graph by visual inspection possibly indicating that there was some systematic reason for the measurement being higher in this case. Unlike that that was observed for the mouse phantom, there is no apparent correlation between the geometry or d and the general height of the curve, though consistent with the arguments above this may be due to the higher absorption of the block which might mean there is less boundary effects. In general the shape of the curve is the same in all cases. The μ_a values are also quite a good match to the nominal values, although the nominal values are on the low end of the distribution of those measured.

In the scattering, there is a similar level of variation to that observed in the other experiments (approximately $\pm$10%) when data set 32 is ignored, the fact that this data set was found to be experimentally flawed means that this is reasonable to do. The scattering distribution measured is also a fairly good match to the nominal values, the shape is similar and the values are on the high end of the measured sets.

Fitted Properties: Geometry-Specific Models

Given that the semi-infinite slab fitted results appeared highly variable and that the assumption made is that the variation was due foremostly to the varying boundary effects between data collection geometries, the method described above (appendix A.7.3) was used to create a geometry-specific library based on finite-element model diffusion equation in time-resolved NIRFAST. This was done for one reflectance and one transmittance measurement for the XPM-2 phantom in which the source detector positions were recorded, the model in this case was based on a pre-existing mesh of the phantom made from manufacturer-provided geometry. For the cylinder the geometry was modelled by creating a cylinder mesh of the same dimensions as the practical cylinder. Two measurements were processed with NIRFAST where the measurement geometries were both transmittance (fig. A.37, diagram 7) but the rotation angle of the cylinder was changed by 90 degrees to detect any changes due to the internal rod positions. For the block, a pre-defined Monte Carlo cuboid/Brick model was used (because in the case of this

geometry the model was already available), and a reflectance and a transmittance measurement were taken.

In cases where NIRFAST was used, all measurement d values were at least 10mm so as to ensure the accuracy of the diffusion model. All of the following data was acquired using the hybrid detector. As for all the previous processed data, all the details are shown in table A.8 indexed by the numbers in the following graph legends. The following graphs show data for each phantom in turn as before, with the numbered data being the newly processed geometry-specific data and the range of the previous (semi-infinite model) data being shown for reference as well as the nominal properties.

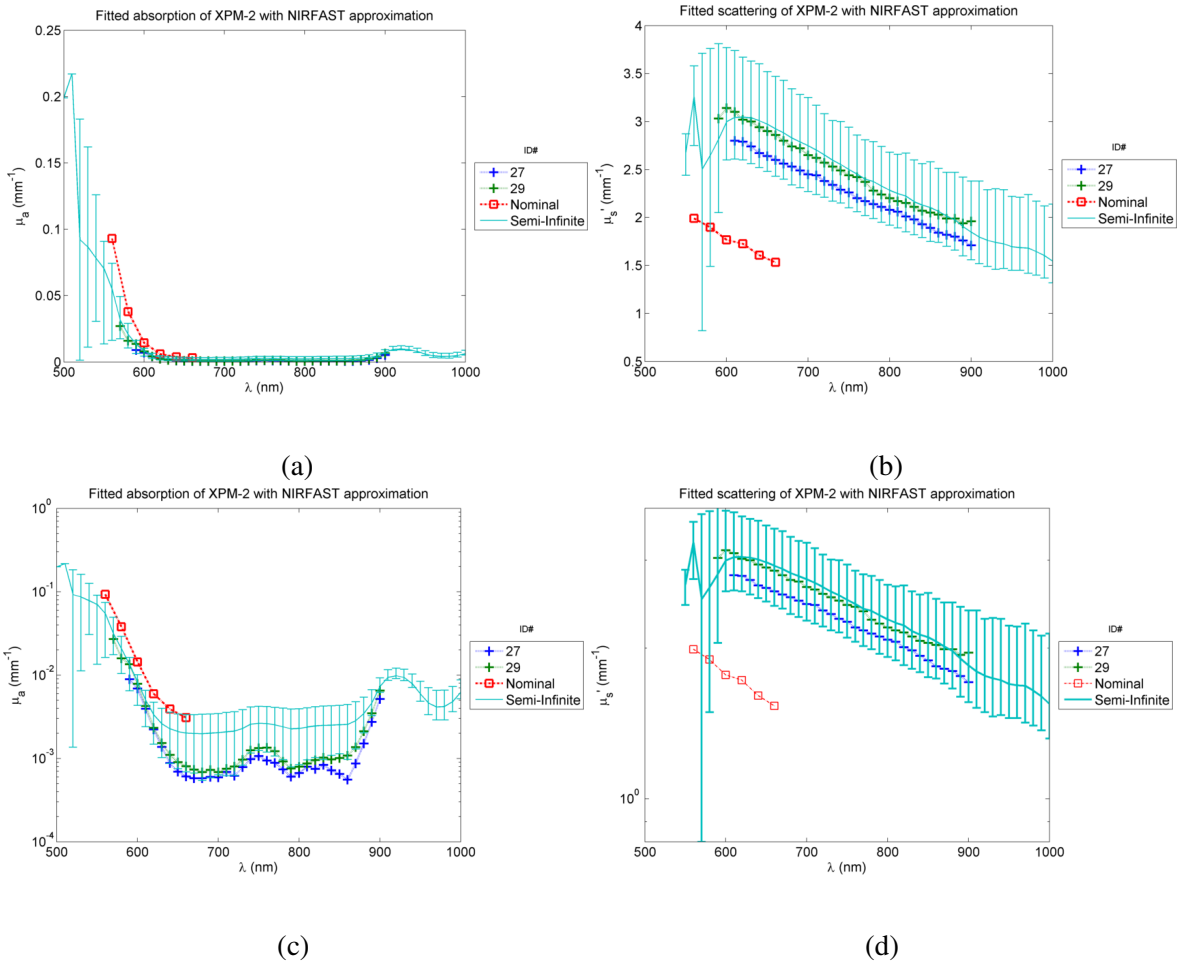

Figure A.43: Fitted (a) absorption and (b) scattering of the XPM-2 phantom with NIRFAST library, repeated (c-d) on a log-scale plot. Additionally for reference, the mean semi-infinite slab values are shown with bars indicating the full range of the previous data (i.e. stretching in total from minimum to maximum per wavelength).

Figure A.43 shows property fitting results provided by the measurement system when the NIRFAST-built library was used to model geometry. It can be seen that for μ_a, the new values appear on the low end of the previous data, taking them further from the nominal data (although this is expected to be unreliable). The fact that μ_a was taken lower is consistent with the argument about boundary effects because these were assumed to increase the absorption previously. It can also be seen that the values are considerably more consistent between the two sets than those previously obtained which is more remarkable given that the new data is taken from the two extremes of the geometries previously used (a full-width transmittance measurement and a short-range reflectance measurement). Together these observations suggest that the new data for μ_a may be more accurate and that there was substantial error in the nominal values. In μ_s', the recovered values are also closer than the range of previous values (although of course there are fewer observations too), and they fall within the range of previous measurements.

Figure A.44 shows the results of optical property fittings for the cylinder using the NIRFAST-library approach. It can be seen that in this case the new data takes the μ_a values closer to the (assumed reliable) nominal values than those previously acquired. If the nominal values are correct then the new measurements are out by approximately 10 and 20 percent. The hypothesis regarding boundary effects would predict this behaviour because the influence of boundary effects is assumed to artificially raise the recovered μ_a. For μ_s', the values are within the range previously found. Both sets (μ_a and μ_s') are still somewhat over-estimates which may indicate some unknown source of error or may indicate that the nominal values are incorrect.

Figure A.45 shows the results for the block phantom when using the block-shape Monte Carlo solution for the fitting. The μ_a values are on the low-end of those acquired with the semi-infinite model which is again consistent with the idea that boundary effects on absorption are being reduced, this is again more remarkable because the data is taken from transmittance and reflectance measurements which are brought into relatively high agreement. The results for μ_a are now also close to the nominal values. For scattering, the values are within the range previously found and near to the nominal values. The results together suggest that the block-model adds more accuracy which is to be expected because it matches the physical shape in this

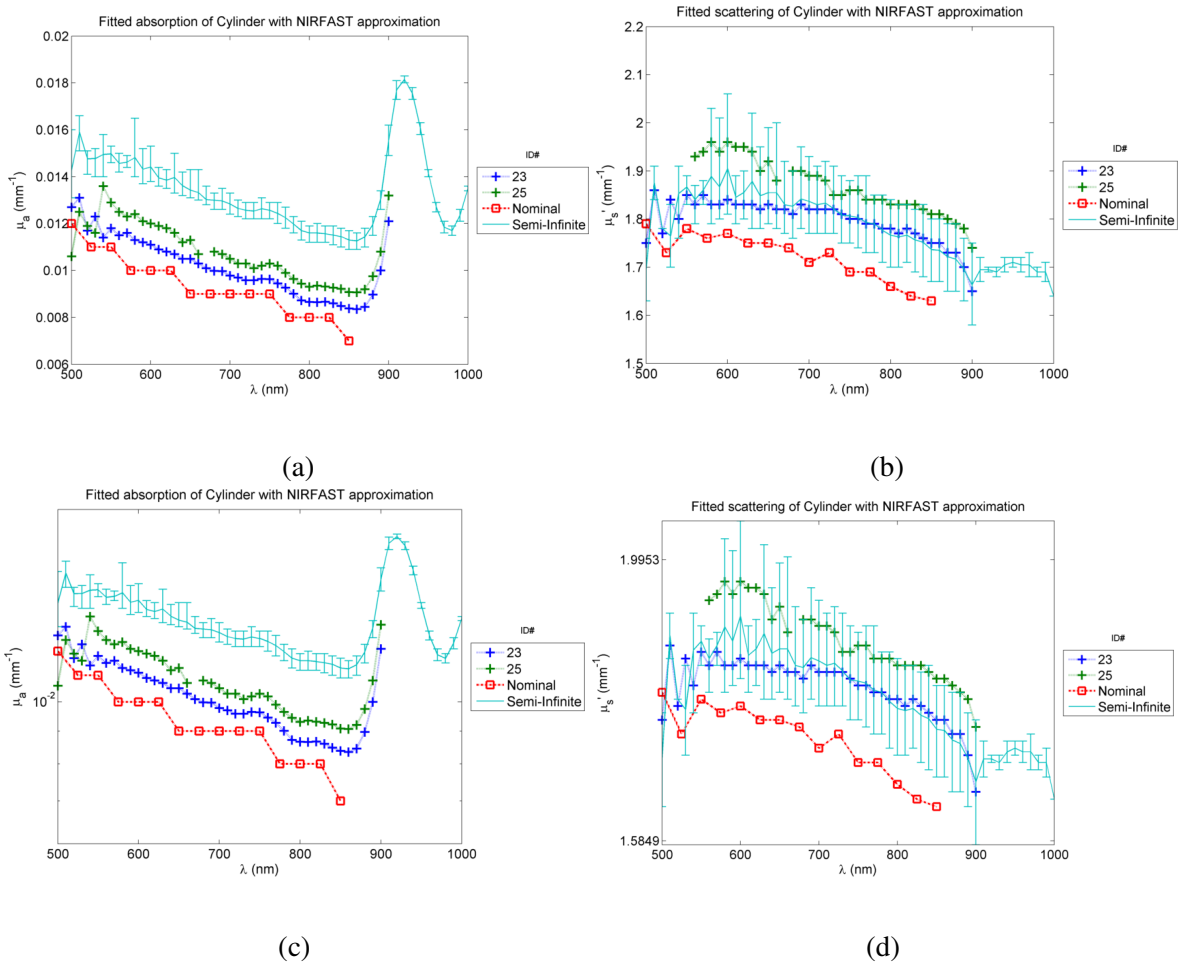

Figure A.44: Fitted (a) absorption and (b) scattering of the Cylinder phantom with NIRFAST library, repeated (c-d) on a log-scale plot. Additionally for reference, the mean semi-infinite slab values are shown with bars indicating the full range of the previous data (i.e. stretching in total from minimum to maximum per wavelength).

case.

Best estimate of optical properties

In accordance with the discussions above, it is now assumed that the results obtained with the NIRFAST model and the block model are more accurate than those acquired using the semi-infinite slab approximation. Given the relatively small number of repeats however, and the experimental variation it is considered that there is not enough evidence to assume an improvement over the nominal values for the cylinder phantom. However, given the low level of confidence in the nominal values for the other phantoms it is assumed that the block model and

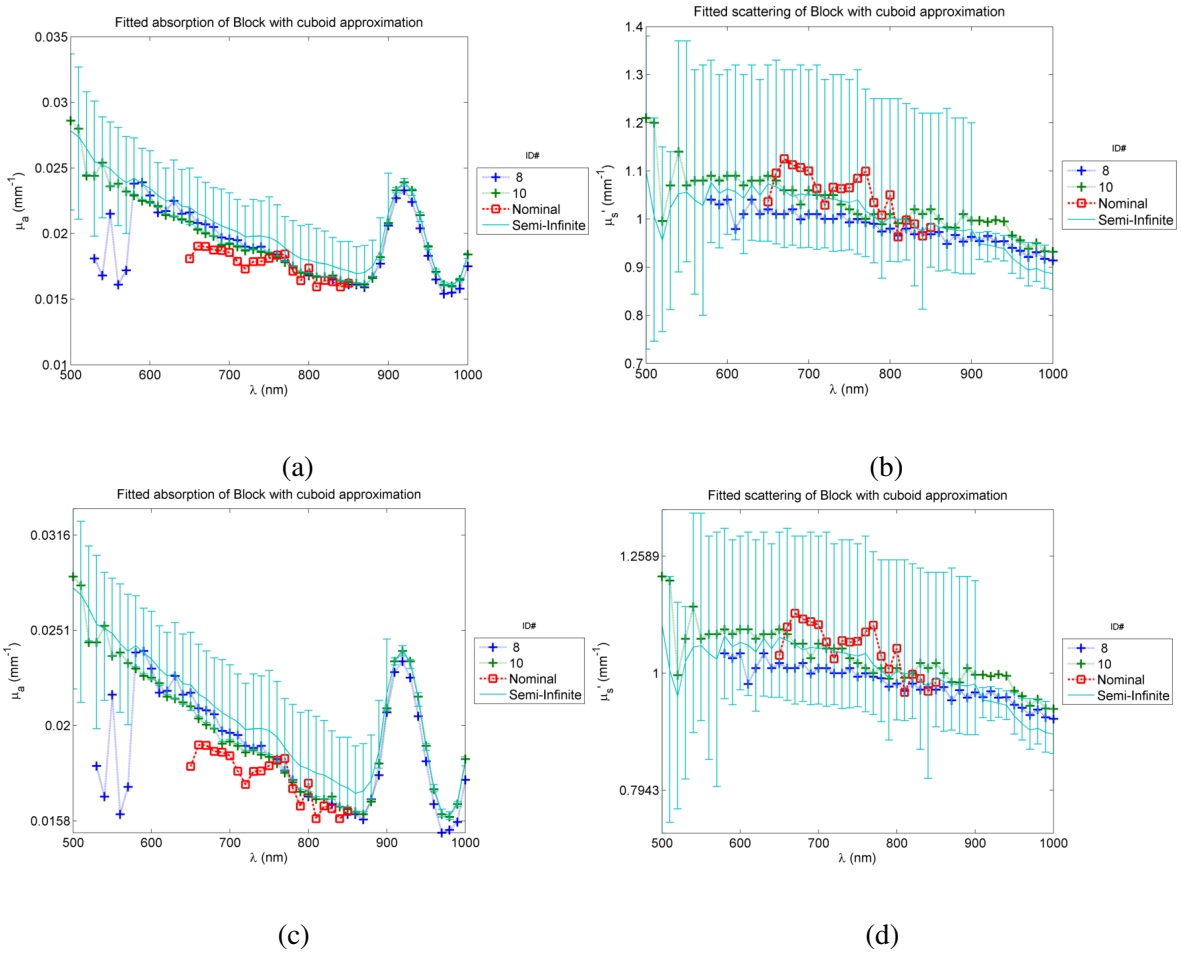

Figure A.45: Fitted (a) absorption and (b) scattering of the Block phantom with Brick-model approximation, repeated (c-d) on a log-scale plot. Additionally for reference, the mean semi-infinite slab values are shown with bars indicating the full range of the previous data (i.e. stretching in total from minimum to maximum per wavelength).

the NIRFAST model versions of each are the best estimate of the optical properties and as such these are used in experiments presented in this thesis.

Note, additionally values from set 16 (fig. A.39) were also used for calculating XPM best-estimate data at low wavelengths where no other recorded data set had a good enough SNR to provide values. This is not ideal as the geometry was based on the semi-infinite slab for this set, however it is the only measurements available for wavelengths less than 570nm. Also, being as the value of d for this set was 5mm (this short distance along with the S1 detector being the reasons for the relatively high SNR) and it was a reflectance measurement, this is likely to be the least influenced by boundary effects (and it measured the lowest absorption in conformance with the expectations this assumption introduces). The same approach was used with set 34 and

the block data. The decision to do this was based on the aim of getting a best estimate for all wavelengths in the range 500 to 1000, and partly on the fact that these spectra appeared close to the NIRFAST and block model fits by visual inspection.

Taking the *best estimate* then as the mean value obtained for all of the chosen sets outlined above for all phantoms, the resultant properties are shown in fig. A.46.

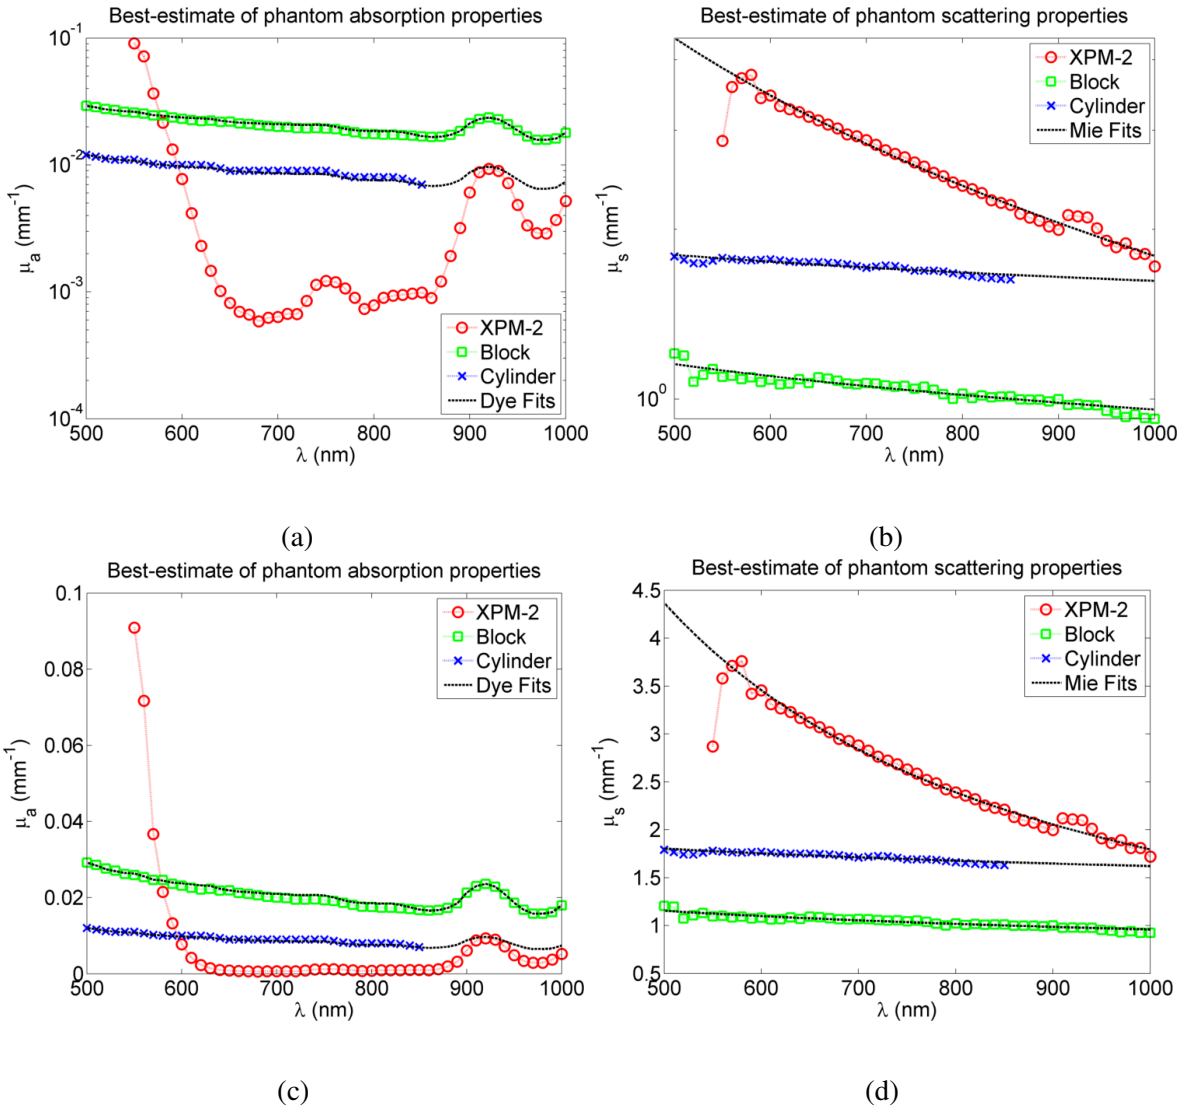

Figure A.46: Best estimate of (a) absorption and (b) scattering of the phantoms, repeated (c-d) on a log-scale plot. Additionally for reference. The dashed black lines indicate fitted spectra when using a single dye to represented absorption in the block and cylinder and a fit to an approximation of Mie scattering (A.9) for all the phantoms.

This figure also shows some additionally fitted lines that illustrate fits used later in the thesis when performing spectrally resolved image reconstruction (chapter 8 section). In the case of

absorption, the fit indicates a derived extinction coefficient spectrum (calculated as the mean of available values across the block and cylinder absorption spectra having first fit them together by a single scaling factor applied to the block data) that along with a concentration term (a single-scaling factor; 1 for the cylinder, 2.43 for the block) now gives rise to the plotted lines. The aim of this is to represent both the block and cylinder absorption as amounts of a single chromophore, being as the manufacturer was the same in both cases and it is assumed that the actual constituents (dye and base material) were the same in both these phantoms this should be valid and it can be seen that the fit is good. Using the same *dye* for representing the XPM-2 absorption would be inappropriate because the composition components are different. In scattering the fit is using an approximation to Mie theory, and represents solutions of best-fits where the searched for quantities were scatter power and scatter amplitude (section 2.1.2), the values for the scatter properties are shown in table A.9.

Sample	Scatter Power	Scatter Amplitude	Chromophore	Concentration
XPM-2	1.79	1.28	Composite Dye 1	1
Block	0.96	0.27	Composite Dye 2	$2.4333^{\cdot}$
Cylinder	1.62	0.15	Composite Dye 2	1

Table A.9: Fitted scattering parameters for each phantom using the model presented in section 2.1.2 (eq. (2.8)) and concentrations of dyes used to represent the absorption spectra: dye 1 represents the XPM absorption characteristic, whilst dye 2 represents both the block and cylinder phantoms.

Interpretation of Spectra w.r.t. Composition

The XPM-2 phantom is made of polyurethane base material, with disperse red 1 dye as an absorber[210, 44]. The disperse red dye has a strong absorption peak at approximately 500nm which tails off as a Gaussian-like curve either side[211]. It is assumed that this is why a large peak is being tended towards at the lower end of the range investigated here and that the remainder of the absorption seen is due to the base material. It is understood that the other two phantoms contain carbon black dye and are again based on polyurethane as a base. This explains the characteristic peak in all three μ_a curves at approximately 930nm because this is a polyurethane peak[212]. The carbon black dye has a slowly descending absorption across the

range considered here which explains the common, slowly descending spectra of both the block and cylinder. The scattering agents are most likely titanium dioxide particles in all cases, it is observed that the resultant spectra for all phantoms fit well with the Mie theory approximation.

Given that the constituent parts are understood, it is possible that these phantoms could be used to do multiple chromophore spectral image reconstruction tests. This is not done in this thesis but could make for interesting future study.

A.7.5 Conclusion

Given the short time-frame available (in a short research visit to the CUSBO lab), a great deal of data was obtained which has been presented above in its entirety. Due to the preliminary nature of the investigation, many variables were changed to get a feel for their impact and therefore all data should be repeated before firm conclusions can be drawn.

There are trends in the data that can be used as a basis for directing future investigation and hypotheses have been formed based on the small number of results already observed. Firstly, there may be a strong influence of boundary effects when measuring phantoms of the shapes and sizes of those considered here, when fitting with a semi-infinite slab approximation based Monte Carlo model, this is already known but is hard to quantify in general so the experiments were worthwhile for gaining phantom-specific insights. It is possible that the presented method for correction based on library building using NIRFAST may overcome this issue and that the use of more model-specific models (such as the brick model used for the block phantom here) also lead to improved accuracy. The fact that the NIRFAST-corrected measurements were closer to one another for strongly varying experimental set-ups than the semi-infinite model counterparts adds weight to this idea, as does the fact that applying the NIRFAST corrections drew the measured properties of the cylinder phantom closer to the nominal values which, of all the phantoms presented, were expected to be most reliable.

It is assumed, based on the increased model appropriateness, that the use of the block Monte Carlo model for the block phantom would add accuracy to the measurements, and the observations support this because the values were drawn closer together between reflectance and

transmittance modes, the values for μ_a reduced towards the nominal values as well. The trend of for the absorption data to reduce with the block model mirrors that seen in the NIRFAST model applied to the other phantoms which may also back up the idea that the approach was valid and helpful.

The estimated XPM properties are rather different to the nominal properties. One possible reason for this could be change in the properties over time which has been noted to occur for polyurethane phantoms in absorption by Moffit et al.[212]. Also the nominal properties contained interpretation error.

A.8 Mouse Phantom Source Identification

The XPM-2 mouse phantom contains two internal light sources at separate locations. The manufacturer's handbook provides details of the expected source locations based on two small diagrams which offer some measurements from particular faces of the phantom (these have been listed in table A.1). However, in order to understand better the composition of the phantom and what to expect from the light sources, further knowledge was required of exactly what it contains. The following is a description of an experiment performed using a (human) X-ray CT instrument and following analysis which allowed better understanding of the phantom and therefore of experiments in which it was later used. It was found that rather than two clear-cut expected positions, the places where the sources are located are better described as expected regions given the current level of knowledge, these areas are now known precisely.

A.8.1 Further Information on the XPM-2 Composition

A patent published by the manufacturer of the XPM-2 describes a method for making such a phantom[210] by using initially empty (i.e. air filled) tunnels in the solid cast plastic phantom, down which small ($\approx 100\mu$m diameter) optical fibres can be fed, then the tunnels back-filled, the fibres then are used to transport light from stabilised LEDs housed in the box at the back of

the phantom to the tips of the fibres now located inside.

It is assumed that this is the technique that was used to make the XPM-2 phantom used here, and the following experiment results back this up.

A.8.2 CT Scanning Experiment

The phantom was scanned using a Philips Brilliance Big Bore CT scanner (designed and usually used for scanning humans). Clearly visible in the CT slices were parts of two tunnels within the phantom that are assumed to house optical fibres.

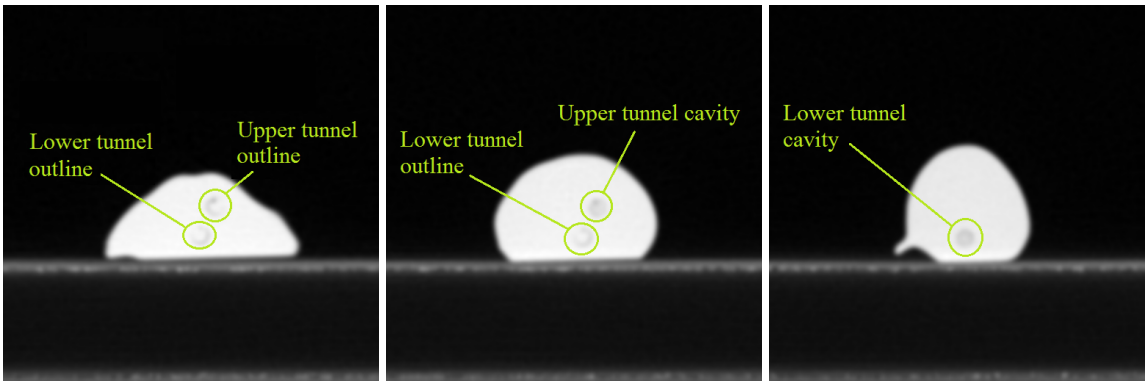

(a) (b) (c)

Figure A.47: Sample cross-sectional X-ray slice Dicoms obtained from the CT scan of the XPM-2 phantom, in particular (a) a slice showing outlines of two tunnels and (b,c) slices showing the last appearance of each of the tunnels.

Figure A.47 shows three of the obtained Dicoms; cross-sectional images showing the X-ray attenuation at different axial offsets. In each image the body of the phantom can be seen clearly, as can the platform on which the phantom was rested. Figure A.47(a) shows a slice at which the outlines of both tunnels are visible in the image, though with poor contrast. Figure A.47(b) shows the last slice (moving through the phantom towards the nose starting at the back) in which traces of the upper tunnel are visible, and the appearance of the tunnel at this slice is particularly clear compared to other slices; the appearance is of an air-gap or cavity rather than the previous inconclusive appearance. Figure A.47(c) shows the final slice in which the lower tunnel can be seen and the same behaviour is seen - whilst the outline of the tunnel is visible in other slices,

the whole tunnel cross-section shows up in this slice. These suggest that the tunnels containing optical fibres have a cavity at the end where they are not filled, this complicates the issue of the precise *location* of the source which now has to be considered differently.

A.8.3 Image Analysis

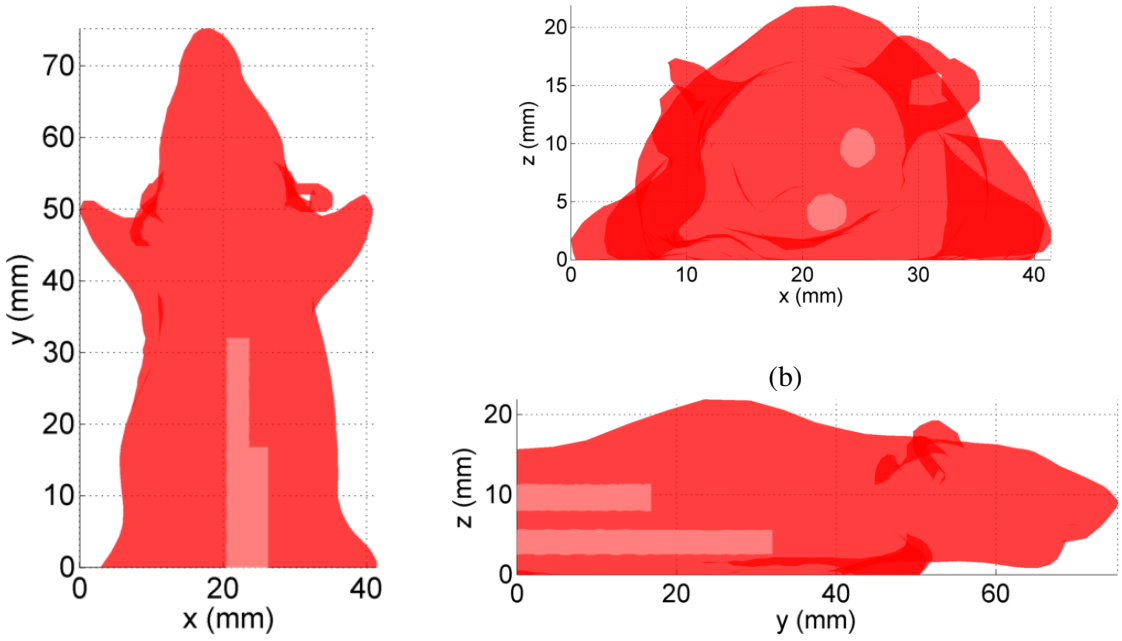

(a) (b) (c)

Figure A.48: Extracted phantom and tunnel outlines for the XPM-2 phantom from CT data show from multiple views.

From the CT data, a 3D point cloud representing the surface of the phantom was extracted by segmenting (via simple thresholding) the phantom and identifying points on its boundary in each slice. The same procedure of segmentation and subsequent edge-point generation was then applied to the ends of each of the light source tunnels (fig. A.47(b) and (c)). A mesh was produced (as above) from each point cloud allowing the visualisation of the tunnels within the phantom in 3D. Note that, whilst the start and end of the tunnels were clearly visible as holes in the CT, it was not clear from intermediate slices whether or not the tunnels had been completely or partially filled in order to enhance the optical homogeneity of the phantom. For simplicity of visualisation it is assumed that the tunnels are continuous but this may not be the case.

A.8.4 Conclusions

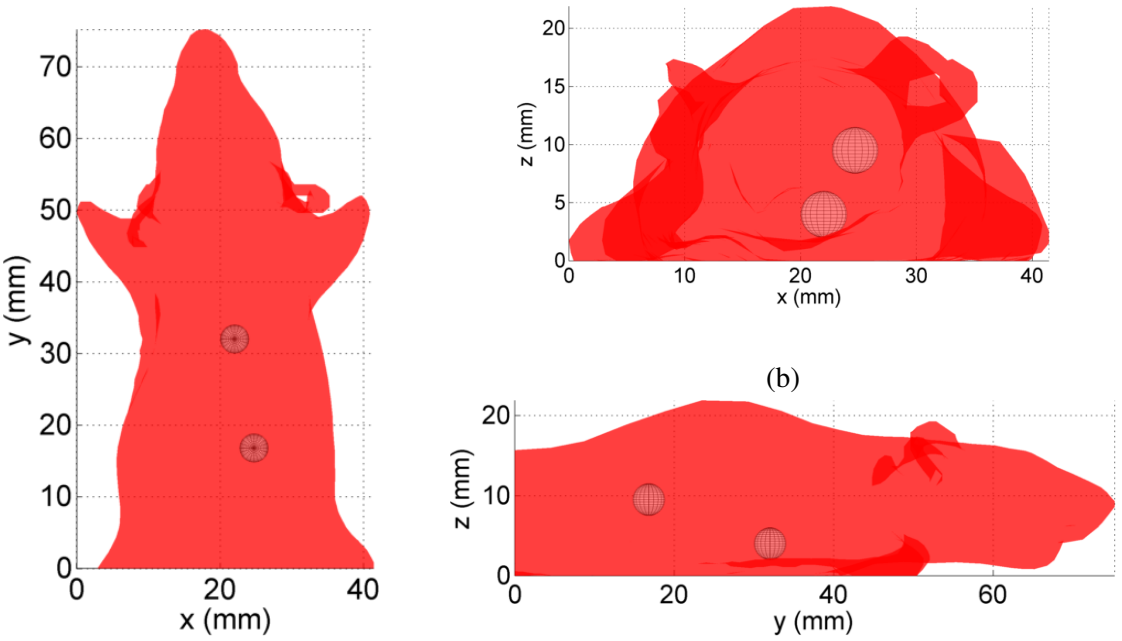

Figure A.49: Extracted phantom source locations given as sphereical regions within which the sources are expected to be found.

It appears that the phantom contains two directed light sources (two instances of light launched directionally from the tip of an optical fibre) that are themselves at an unknown position within small cavities that have been identified on the CT scan. Thus it can be said that the source is somewhere within the cavity but it is likely to be reflected around within and it is unknown what direction the fibre points in. Given that the source is directed, it cannot be reconstructed for directly when using the diffusion approximation therefore the target for reconstructions using this approach must be placed two scattering lengths away from the true locations where an effective diffuse source will be created[20], the target location for the centre of reconstructed sources can therefore be assumed to be anywhere within two scattering lengths of the void region found at the end of the tunnel, such is the uncertainty on the expected location. Because the cavities are only visible in a single slice in each case, and for simplicity, the expected region for each source is now drawn by taking a sphere around the centre-points of the ends of the rods of radius equal to the radius of the tunnel plus a transport scattering length

$(1/\mu_s')$ which is taken as 1/2 mm being a worst case given the spectral optical properties of the XPM-2 (fig. A.46), the final radius of the spherical target regions is 1.975mm, these are shown in fig. A.49. In later analyses, this value is rounded to 2mm for simplicity (chapter 9).

A.9 Appendix Conclusion

This chapter has explored system parameters that are important for obtaining accurate and quantitative data and thereby allowing image reconstruction in BLT and DOT, the primary goals of the system. It has also introduced and rigorously analysed the phantoms and materials used to test aspects of the system. This is vital to a proper interpretation of the results. The chapter has also introduced the methods used for system geometric calibration which are important for the surface capture and free-space light transport modelling performed in the system, as presented in the thesis.

LIST OF REFERENCES

[1] James A Guggenheim, Hector R A Basevi, Jon Frampton, Iain B Styles, and Hamid Dehghani. Multi-modal molecular diffuse optical tomography system for small animal imaging. *Measurement Science and Technology*, 24(10):105405, 2013.

[2] James A Guggenheim, Hector R A Basevi, Iain B Styles, Jon Frampton, and Hamid Dehghani. Quantitative surface radiance mapping using multiview images of light-emitting turbid media. *JOSA A*, 30(12):2572–2584, 2013.

[3] Hector R A Basevi, James A Guggenheim, Hamid Dehghani, and Iain B Styles. Simultaneous multiple view high resolution surface geometry acquisition using structured light and mirrors. *Optics Express*, 21(6):7222–7239, 2013.

[4] Hector R A Basevi, Kenneth M Tichauer, Frederic Leblond, Hamid Dehghani, James A Guggenheim, Robert W Holt, and Iain B Styles. Compressive sensing based reconstruction in bioluminescence tomography improves image resolution and robustness to noise. *Biomedical optics express*, 3(9):2131, 2012.

[5] James A Guggenheim, Hector R A Basevi, Iain B Styles, Jon Frampton, and Hamid Dehghani. Bioluminescence tomography improves quantitative accuracy for pre-clinical imaging. In *European Conferences on Biomedical Optics*, pages 87990G–87990G. International Society for Optics and Photonics, 2013.

[6] James A Guggenheim, Hector R Basevi, Iain B Styles, Jon Frampton, and Hamid Dehghani. Multi-view, multi-spectral bioluminescence tomography. In *Biomedical Optics*. Optical Society of America, 2012.

[7] Hector R A Basevi, James A Guggenheim, Hamid Dehghani, and Iain B Styles. Information-theoretic method for wavelength selection in bioluminescence tomography. In *European Conferences on Biomedical Optics*, pages 879909–879909. International Society for Optics and Photonics, 2013.

[8] Iain B Styles, Hector R A Basevi, James A Guggenheim, and Hamid Dehghani. Random matrix-based dimensionality reduction for bioluminescence tomography reconstruction.

In *European Conferences on Biomedical Optics*, pages 87990J–87990J. International Society for Optics and Photonics, 2013.

[9] James A Guggenheim, Hamid Dehghani, Hector R A Basevi, Iain B Styles, and Jon Frampton. Development of a multi-view multi-spectral bioluminescence tomography small animal imaging system. In *European Conferences on Biomedical Optics*, pages 80881K–80881K. International Society for Optics and Photonics, 2011.

[10] Alnawaz Rehemtulla, Lauren D Stegman, Shaun J Cardozo, Sheila Gupta, Daniel E Hall, Christopher H Contag, and Brian D Ross. Rapid and quantitative assessment of cancer treatment response using in vivo bioluminescence imaging. *Neoplasia (New York, NY)*, 2(6):491, 2000.

[11] Ralph Weissleder. Scaling down imaging: molecular mapping of cancer in mice. *Nature Reviews Cancer*, 2(1):11–18, 2002.

[12] Darlene E Jenkins, Yoko Oei, Yvette S Hornig, Shang-Fan Yu, Joan Dusich, Tony Purchio, and Pamela R Contag. Bioluminescent imaging (bli) to improve and refine traditional murine models of tumor growth and metastasis. *Clinical & experimental metastasis*, 20(8):733–744, 2003.

[13] Shimon Gross and David Piwnica-Worms. Spying on cancer: molecular imaging in vivo with genetically encoded reporters. *Cancer cell*, 7(1):5–15, 2005.

[14] Stefanie Mandl, Christoph Schimmelpfennig, Matthias Edinger, Robert S Negrin, and Christopher H Contag. Understanding immune cell trafficking patterns via in vivo bioluminescence imaging. *Journal of Cellular Biochemistry*, 87(S39):239–248, 2002.

[15] Jonathan Hardy, Matthias Edinger, Michael H Bachmann, Robert S Negrin, C Garrison Fathman, and Christopher H Contag. Bioluminescence imaging of lymphocyte trafficking in vivo. *Experimental hematology*, 29(12):1353–1360, 2001.

[16] Xiuli Wang, Michael Rosol, Shundi Ge, Denise Peterson, George McNamara, Harvey Pollack, Donald B Kohn, Marvin D Nelson, and Gay M Crooks. Dynamic tracking of human hematopoietic stem cell engraftment using in vivo bioluminescence imaging. *Blood*, 102(10):3478–3482, 2003.

[17] Vasilis Ntziachristos, Jorge Ripoll, Lihong V Wang, and Ralph Weissleder. Looking and listening to light: the evolution of whole-body photonic imaging. *Nature biotechnology*, 23(3):313–320, 2005.

[18] CP Klerk, Renée M Overmeer, TM Niers, Henri H Versteeg, Dick J Richel, Tessa Buckle, CJ Van Noorden, and Olaf van Tellingen. Validity of bioluminescence measurements for noninvasive in vivo imaging of tumor load in small animals. *Biotechniques*, 43(1 Suppl):7–13, 2007.

[19] Lihong V Wang and Hsin-i Wu. *Biomedical optics: principles and imaging*. John Wiley & Sons, 2007.

[20] Michael Firbank. *The design, calibration and usage of a solid scattering and absorbing phantom for near infra red spectroscopy.* PhD thesis, University of London, 1994.

[21] Oregon Medical Laser Center (OMLC) "Optical Properties Spectra" compiled by Scott Prahl and Steven Jacques. Web page updated 2001, accessed 2013. http://omlc.ogi.edu/spectra/.

[22] George M Hale and Marvin R Querry. Optical constants of water in the 200-nm to 200-μm wavelength region. *Applied optics*, 12(3):555–563, 1973.

[23] Robert LP van Veen, HJCM Sterenborg, A Pifferi, A Torricelli, and R Cubeddu. Determination of vis-nir absorption coefficients of mammalian fat, with time-and spatially resolved diffuse reflectance and transmission spectroscopy. In *Biomedical Topical Meeting*. Optical Society of America, 2004.

[24] Melanosome absorption coefficient (website by Steven Jacques). http://omlc.ogi.edu/spectra/melanin/mua.html. Accessed: 2013-11-07.

[25] Hamid Dehghani, Matthew E Eames, Phaneendra K Yalavarthy, Scott C Davis, Subhadra Srinivasan, Colin M Carpenter, Brian W Pogue, and Keith D Paulsen. Near infrared optical tomography using nirfast: Algorithm for numerical model and image reconstruction. *Communications in numerical methods in engineering*, 25(6):711–732, 2009.

[26] Joseph P Culver, Turgut Durduran, Daisuke Furuya, Cecil Cheung, Joel H Greenberg, and AG Yodh. Diffuse optical tomography of cerebral blood flow, oxygenation, and metabolism in rat during focal ischemia. *Journal of Cerebral Blood Flow & Metabolism*, 23(8):911–924, 2003.

[27] Andreas H Hielscher et al. Optical tomographic imaging of small animals. *Current opinion in biotechnology*, 16(1):79–88, 2005.

[28] Wai-Fung Cheong, Scott A Prahl, and Ashley J Welch. A review of the optical properties of biological tissues. *Quantum Electronics, IEEE Journal of*, 26(12):2166–2185, 1990.

[29] Hugo J Van Staveren, Christian JM Moes, Jan Van Marle, Scott A Prahl, and Martin JC Van Gemert. Light scattering in intralipid-10% in the wavelength range of 400-1100 nm. *Appl. Opt*, 30(31):4507–4514, 1991.

[30] Judith R Mourant, Tamika Fuselier, James Boyer, Tamara M Johnson, and Irving J Bigio. Predictions and measurements of scattering and absorption over broad wavelength ranges in tissue phantoms. *Applied Optics*, 36(4):949–957, 1997.

[31] George Alexandrakis, Fernando R Rannou, and Arion F Chatziioannou. Tomographic bioluminescence imaging by use of a combined optical-pet (opet) system: a computer simulation feasibility study. *Physics in medicine and biology*, 50(17):4225, 2005.

[32] Frederic Leblond, Scott C Davis, Pablo A Valdés, and Brian W Pogue. Pre-clinical whole-body fluorescence imaging: Review of instruments, methods and applications. *Journal of Photochemistry and Photobiology B: Biology*, 98(1):77–94, 2010.

[33] Kenneth M Tichauer, Robert W Holt, Fadi El-Ghussein, Qun Zhu, Hamid Dehghani, Frederic Leblond, and Brian W Pogue. Imaging workflow and calibration for ct-guided time-domain fluorescence tomography. *Biomedical optics express*, 2(11):3021, 2011.

[34] RC Haskell, LO Svaasand, T Tsay, T Feng, MS McAdams, and BJ Tromberg. Boundary conditions for the diffusion equation in radiative transfer. *Journal of the Optical Society of America, Part A: Optics and Image Science;(United States)*, 11(10), 1994.

[35] Alexander D Klose, Vasilis Ntziachristos, and Andreas H Hielscher. The inverse source problem based on the radiative transfer equation in optical molecular imaging. *Journal of Computational Physics*, 202(1):323–345, 2005.

[36] Alexander D Klose and Edward W Larsen. Light transport in biological tissue based on the simplified spherical harmonics equations. *Journal of Computational Physics*, 220:441–470, 2006.

[37] Lihong Wang, Steven L Jacques, and Liqiong Zheng. McmlâĂŤmonte carlo modeling of light transport in multi-layered tissues. *Computer methods and programs in biomedicine*, 47(2):131–146, 1995.

[38] Jin Chen, Vivek Venugopal, and Xavier Intes. Monte carlo based method for fluorescence tomographic imaging with lifetime multiplexing using time gates. *Biomedical optics express*, 2(4):871, 2011.

[39] Vasilis Ntziachristos. *Concurrent diffuse optical tomography, spectroscopy and magnetic resonance imaging of breast cancer*. PhD thesis, Citeseer, 2000.

[40] Wenxiang Cong, Durairaj Kumar, Yi Liu, Alexander Cong, and Ge Wang. A practical method to determine the light source distribution in bioluminescent imaging. In *Optical Science and Technology, the SPIE 49th Annual Meeting*, pages 679–686. International Society for Optics and Photonics, 2004.

[41] Athanasios D Zacharopoulos, Simon R Arridge, Oliver Dorn, Ville Kolehmainen, and Jan Sikora. Three-dimensional reconstruction of shape and piecewise constant region values for optical tomography using spherical harmonic parametrization and a boundary element method. *Inverse Problems*, 22(5):1509, 2006.

[42] Subhadra Srinivasan, Brian W Pogue, Colin Carpenter, Phaneendra K Yalavarthy, and Keith Paulsen. A boundary element approach for image-guided near-infrared absorption and scatter estimation. *Medical physics*, 34:4545, 2007.

[43] Andreas H Hielscher, Alexander D Klose, and Kenneth M Hanson. Gradient-based iterative image reconstruction scheme for time-resolved optical tomography. *Medical Imaging, IEEE Transactions on*, 18(3):262–271, 1999.

[44] Chaincy Kuo, Olivier Coquoz, Tamara L Troy, Heng Xu, and Brad W Rice. Three-dimensional reconstruction of in vivo bioluminescent sources based on multispectral imaging. *Journal of biomedical optics*, 12(2):024007–024007, 2007.

[45] Belma Dogdas, David Stout, Arion F Chatziioannou, and Richard M Leahy. Digimouse: a 3d whole body mouse atlas from ct and cryosection data. *Physics in medicine and biology*, 52(3):577, 2007.

[46] Digimouse Homepage. http://neuroimage.usc.edu/neuro/Digimouse. Accessed: 2013-11-07.

[47] SR Arridge, M Schweiger, M Hiraoka, and DT Delpy. A finite element approach for modeling photon transport in tissue. *Medical physics*, 20:299, 1993.

[48] Martin Schweiger and Simon R Arridge. Application of temporal filters to time resolved data in optical tomography. *Physics in medicine and biology*, 44(7):1699, 1999.

[49] TOAST Homepage. http://web4.cs.ucl.ac.uk/research/vis/toast/. Accessed: 2013-11-07.

[50] Hui Li, Jie Tian, Fuping Zhu, Wenxiang Cong, Lihong V Wang, Eric A Hoffman, and Ge Wang. A mouse optical simulation environment (mose) to investigate bioluminescent phenomena in the living mouse with the monte carlo method< sup> 1</sup>. *Academic Radiology*, 11(9):1029–1038, 2004.

[51] MOSE Homepage. http://www.imaging.sbes.vt.edu/software/mose/. Accessed: 2013-11-07.

[52] Michael Jermyn, Hamid Ghadyani, Michael A Mastanduno, Wes Turner, Scott C Davis, Hamid Dehghani, and Brian W Pogue. Fast segmentation and high-quality three-dimensional volume mesh creation from medical images for diffuse optical tomography. *Journal of biomedical optics*, 18(8):086007–086007, 2013.

[53] NIRFAST Homepage. http://www.nirfast.org/. Accessed: 2013-11-07.

[54] Ge Wang, EA Hoffman, G McLennan, LV Wang, M Suter, J Meinel, et al. Development of the first bioluminescent ct scanner. *Radiology*, 229:566, 2003.

[55] Ge Wang, Wenxiang Cong, Kumar Durairaj, Xin Qian, Haiou Shen, Patrick Sinn, Eric Hoffman, Geoffrey McLennan, Michael Henry, et al. In vivo mouse studies with bioluminescence tomography. *Optics Express*, 14(17):7801–7809, 2006.

[56] Hamid Dehghani, Scott C Davis, and Brian W Pogue. Spectrally resolved bioluminescence tomography using the reciprocity approach. *Medical physics*, 35:4863, 2008.

[57] Ge Wang, Yi Li, and Ming Jiang. Uniqueness theorems in bioluminescence tomography. *Medical physics*, 31:2289, 2004.

[58] Hamid Dehghani, Scott C Davis, Shudong Jiang, Brian W Pogue, Keith D Paulsen, and Michael S Patterson. Spectrally resolved bioluminescence optical tomography. *Optics letters*, 31(3):365–367, 2006.

[59] Alexander X Cong and Ge Wang. Multispectral bioluminescence tomography: methodology and simulation. *International Journal of Biomedical Imaging*, 2006, 2006.

[60] Abhijit J Chaudhari, Felix Darvas, James R Bading, Rex A Moats, Peter S Conti, Desmond J Smith, Simon R Cherry, and Richard M Leahy. Hyperspectral and multispectral bioluminescence optical tomography for small animal imaging. *Physics in medicine and biology*, 50(23):5421, 2005.

[61] Hamid Dehghani, Subhadra Srinivasan, Brian W Pogue, and Adam Gibson. Numerical modelling and image reconstruction in diffuse optical tomography. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1900):3073–3093, 2009.

[62] David A Boas, Dana H Brooks, Eric L Miller, Charles A DiMarzio, Misha Kilmer, Richard J Gaudette, and Quan Zhang. Imaging the body with diffuse optical tomography. *Signal Processing Magazine, IEEE*, 18(6):57–75, 2001.

[63] Vasilis Ntziachristos, AG Yodh, Mitchell D Schnall, and Britton Chance. Mri-guided diffuse optical spectroscopy of malignant and benign breast lesions. *Neoplasia (New York, NY)*, 4(4):347, 2002.

[64] Albert Cerussi, David Hsiang, Natasha Shah, Rita Mehta, Amanda Durkin, John Butler, and Bruce J Tromberg. Predicting response to breast cancer neoadjuvant chemotherapy using diffuse optical spectroscopy. *Proceedings of the National Academy of Sciences*, 104(10):4014–4019, 2007.

[65] Brian W Pogue, Markus Testorf, Troy McBride, Ulf Osterberg, and Keith Paulsen. Instrumentation and design of a frequency-domain diffuse optical tomography imager for breast cancer detection. *Opt. Express*, 1(13):391–403, 1997.

[66] Troy O McBride, Brian W Pogue, Ellen D Gerety, Steven B Poplack, Ulf L Osterberg, and Keith D Paulsen. Spectroscopic diffuse optical tomography for the quantitative assessment of hemoglobin concentration and oxygen saturation in breast tissue. *Applied optics*, 38(25):5480–5490, 1999.

[67] JP Culver, R Choe, MJ Holboke, L Zubkov, T Durduran, A Slemp, V Ntziachristos, B Chance, and AG Yodh. Three-dimensional diffuse optical tomography in the parallel plane transmission geometry: evaluation of a hybrid frequency domain/continuous wave clinical system for breast imaging. *Medical physics*, 30:235, 2003.

[68] David A Boas, Anders M Dale, and Maria Angela Franceschini. Diffuse optical imaging of brain activation: approaches to optimizing image sensitivity, resolution, and accuracy. *Neuroimage*, 23:S275–S288, 2004.

[69] Joseph P Culver, Andrew M Siegel, Jonathan J Stott, and David A Boas. Volumetric diffuse optical tomography of brain activity. *Optics letters*, 28(21):2061–2063, 2003.

[70] Adam T Eggebrecht, Brian R White, Silvina L Ferradal, Chunxiao Chen, Yuxuan Zhan, Abraham Z Snyder, Hamid Dehghani, and Joseph P Culver. A quantitative spatial comparison of high-density diffuse optical tomography and fmri cortical mapping. *Neuroimage*, 61(4):1120–1128, 2012.

[71] Jeremy C Hebden, Adam Gibson, Topun Austin, Rozarina Md Yusof, Nick Everdell, David T Delpy, Simon R Arridge, Judith H Meek, and John S Wyatt. Imaging changes in blood volume and oxygenation in the newborn infant brain using three-dimensional optical tomography. *Physics in medicine and biology*, 49(7):1117, 2004.

[72] Tuan H Pham, Olivier Coquoz, Joshua B Fishkin, Eric Anderson, and Bruce J Tromberg. Broad bandwidth frequency domain instrument for quantitative tissue optical spectroscopy. *Review of Scientific Instruments*, 71(6):2500–2513, 2000.

[73] Stefan Andersson-Engels, Claes af Klinteberg, K Svanberg, and S Svanberg. In vivo fluorescence imaging for tissue diagnostics. *Physics in medicine and biology*, 42(5):815, 1997.

[74] James R Mansfield, Kirk W Gossage, Clifford C Hoyt, and Richard M Levenson. Autofluorescence removal, multiplexing, and automated analysis methods for in-vivo fluorescence imaging. *Journal of biomedical optics*, 10(4):041207–041207, 2005.

[75] G Weagle, PE Paterson, J Kennedy, and R Pottier. The nature of the chromophore responsible for naturally occurring fluorescence in mouse skin. *Journal of Photochemistry and Photobiology B: Biology*, 2(3):313–320, 1988.

[76] Irene Georgakoudi, Brian C Jacobson, Markus G Müller, Ellen E Sheets, Kamran Badizadegan, David L Carr-Locke, Christopher P Crum, Charles W Boone, Ramachandra R Dasari, Jacques Van Dam, et al. Nad (p) h and collagen as in vivo quantitative fluorescent biomarkers of epithelial precancerous changes. *Cancer research*, 62(3):682–687, 2002.

[77] Kristen E Adams, Shi Ke, Sunkuk Kwon, Feng Liang, Zhen Fan, Yang Lu, Karen Hirschi, Michel E Mawad, Michael A Barry, and Eva M Sevick-Muraca. Comparison of visible and near-infrared wavelength-excitable fluorescent dyes for molecular imaging of cancer. *Journal of biomedical optics*, 12(2):024017–024017, 2007.

[78] Emmet McCormack, David R Micklem, Lars-Erik Pindard, Elisabeth Silden, Pascal Gallant, Alexandre Belenkov, James B Lorens, and Bjorn Tore Gjertsen. In vivo optical imaging of acute myeloid leukemia by green fluorescent protein: time-domain autofluorescence decoupling, fluorophore quantification, and localization. *Molecular Imaging*, 6(3):193–204, 2007.

[79] Douglas J Rowland and Simon R Cherry. Small-animal preclinical nuclear medicine instrumentation and methodology. *Seminars in nuclear medicine*, 38(3):209–222, 2008.

[80] Arion F Chatziioannou. Instrumentation for molecular imaging in preclinical research: Micro-pet and micro-spect. *Proceedings of the american thoracic society*, 2(6):533, 2005.

[81] Tatsuo Torizuka, Nagara Tamaki, Tetsuro Inokuma, Yasuhiro Magata, Satoshi Sasayama, Yoshiharu Yonekura, Akira Tanaka, Yoshio Yamaoka, Kazutaka Yamamoto, and Junji Konishi. In vivo assessment of glucose metabolism in hepatocellular carcinoma with fdg-pet. *The Journal of nuclear medicine*, 36(10):1811–1817, 1995.

[82] Steven M Larson, Yusuf Erdi, Timothy Akhurst, Madhu Mazumdar, Homer A Macapinlac, Ronald D Finn, Cecille Casilla, Melissa Fazzari, Neil Srivastava, Henry WD Yeung, et al. Tumor treatment response based on visual and quantitative changes in global tumor glycolysis using pet-fdg imaging: the visual response score and the change in total lesion glycolysis. *Clinical Positron Imaging*, 2(3):159–171, 1999.

[83] Barbara J Fueger, Johannes Czernin, Isabel Hildebrandt, Chris Tran, Benjamin S Halpern, David Stout, Michael E Phelps, and Wolfgang A Weber. Impact of animal handling on the results of 18f-fdg pet studies in mice. *Journal of Nuclear Medicine*, 47(6):999–1006, 2006.

[84] DJ Wagenaar, J Zhang, T Kazules, T Vandehei, E Bolle, S Chowdhury, K Parnham, and BE Patt. In vivo dual-isotope spect imaging with improved energy resolution. In *Nuclear Science Symposium Conference Record, 2006. IEEE*, volume 6, pages 3821–3826. IEEE, 2006.

[85] Steven R Meikle, Peter Kench, Michael Kassiou, and Richard B Banati. Small animal spect and its place in the matrix of molecular imaging technologies. *Physics in medicine and biology*, 50(22):R45, 2005.

[86] Martin S Judenhofer and Simon R Cherry. Applications for preclinical pet/mri. *Seminars in nuclear medicine*, 43(1):19–29, 2013.

[87] Mark A Brown and Richard C Semelka. *MRI: basic principles and applications*. Wiley. com, 2011.

[88] Thomas E Yankeelov and John C Gore. Dynamic contrast enhanced magnetic resonance imaging in oncology: theory, data acquisition, analysis, and examples. *Current medical imaging reviews*, 3(2):91, 2009.

[89] Chris Albanese, Olga C Rodriguez, John VanMeter, Stanley T Fricke, Brian R Rood, YiChien Lee, Sean S Wang, Subha Madhavan, Yuriy Gusev, Emanuel F Petricoin III, et al. Preclinical magnetic resonance imaging and systems biology in cancer research: current applications and challenges. *The American journal of pathology*, 2012.

[90] Jason A Koutcher, Xiaoyi Hux, Su Xu, Terence PF Gade, Norman Leeds, Xiaohong Joe Zhou, David Zagzag, and Eric C Holland. Mri of mouse models for gliomas shows similarities to humans and can be used to identify mice for preclinical trials. *Neoplasia (New York, NY)*, 4(6):480, 2002.

[91] Martin A Sieber, Hubertus Pietsch, Jakob Walter, Wolfram Haider, Thomas Frenzel, and Hanns-Joachim Weinmann. A preclinical study to investigate the development of nephrogenic systemic fibrosis: a possible role for gadolinium-based contrast media. *Investigative radiology*, 43(1):65–75, 2008.

[92] Tapan K Jain, Susan P Foy, Bernadette Erokwu, Sanja Dimitrijevic, Christopher A Flask, and Vinod Labhasetwar. Magnetic resonance imaging of multifunctional pluronic stabilized iron-oxide nanoparticles in tumor-bearing mice. *Biomaterials*, 30(35):6748–6756, 2009.

[93] Martina A McAteer, Jurgen E Schneider, Ziad A Ali, Nicholas Warrick, Christina A Bursill, Constantin von zur Muhlen, David R Greaves, Stefan Neubauer, Keith M Channon, and Robin P Choudhury. Magnetic resonance imaging of endothelial adhesion molecules in mouse atherosclerosis using dual-targeted microparticles of iron oxide. *Arteriosclerosis, thrombosis, and vascular biology*, 28(1):77–83, 2008.

[94] Jason S Weinstein, Csanad G Varallyay, Edit Dosa, Seymur Gahramanov, Bronwyn Hamilton, William D Rooney, Leslie L Muldoon, and Edward A Neuwelt. Superparamagnetic iron oxide nanoparticles: diagnostic magnetic resonance imaging and potential therapeutic applications in neurooncology and central nervous system inflammatory pathologies, a review. *Journal of Cerebral Blood Flow & Metabolism*, 30(1):15–35, 2009.

[95] Gultekin Gulsen, Ozlem Birgul, Mehmet Burcin Unlu, Roshanak Shafiiha, and Orhan Nalcioglu. Combined diffuse optical tomography (dot) and mri system for cancer imaging in small animals. *Technology in cancer research & treatment*, 5(4), 2006.

[96] David W Holdsworth and Michael M Thornton. Micro-ct in small animal and specimen imaging. *Trends in Biotechnology*, 20(8):S34–S39, 2002.

[97] Erik L Ritman. Micro-computed tomography-current status and developments. *Annu. Rev. Biomed. Eng.*, 6:185–208, 2004.

[98] Connor A Wathen, Nathan Foje, Tony van Avermaete, Bernadette Miramontes, Sarah E Chapaman, Todd A Sasser, Raghuraman Kannan, Steven Gerstler, and W Matthew Leevy. In vivo x-ray computed tomographic imaging of soft tissue with native, intravenous, or oral contrast. *Sensors*, 13(6):6957–6980, 2013.

[99] Michael J Paulus, Shaun S Gleason, Stephen J Kennel, Patricia R Hunsicker, and Dabney K Johnson. High resolution x-ray computed tomography: an emerging tool for small animal cancer research. *Neoplasia (New York, NY)*, 2(1-2):62, 2000.

[100] Srinivasan Mukundan, Ketan B Ghaghada, Cristian T Badea, Chen-Yu Kao, Laurence W Hedlund, James M Provenzale, G Allan Johnson, Emmanuel Chen, Ravi V Bellamkonda, and Ananth Annapragada. A liposomal nanoscale contrast agent for preclinical ct in mice. *American Journal of Roentgenology*, 186(2):300–307, 2006.

[101] Cristian T Badea, Laurence W Hedlund, Ming De Lin, Julie F Boslego Mackel, and G Allan Johnson. Tumor imaging in small animals with a combined micro-ct/micro-dsa system using iodinated conventional and blood pool contrast agents. *Contrast media & molecular imaging*, 1(4):153–164, 2006.

[102] JF Hainfeld, DN Slatkin, TM Focella, and HM Smilowitz. Gold nanoparticles: a new x-ray contrast agent. *British Journal of Radiology*, 79(939):248–253, 2006.

[103] Quan-Yu Cai, Sun Hee Kim, Kyu Sil Choi, Soo Yeon Kim, Seung Jae Byun, Kyoung Woo Kim, Seong Hoon Park, Seon Kwan Juhng, and Kwon-Ha Yoon. Colloidal gold nanoparticles as a blood-pool contrast agent for x-ray computed tomography in mice. *Investigative radiology*, 42(12):797–806, 2007.

[104] Liji Cao and Jörg Peter. Bayesian reconstruction strategy of fluorescence-mediated tomography using an integrated spect-ct-ot system. *Physics in medicine and biology*, 55(9):2693, 2010.

[105] Christophe M Deroose, Abhijit De, Andreas M Loening, Patrick L Chow, Pritha Ray, Arion F Chatziioannou, and Sanjiv S Gambhir. Multimodality imaging of tumor xenografts and metastases in mice with combined small-animal pet, small-animal ct, and bioluminescence imaging. *Journal of Nuclear Medicine*, 48(2):295–303, 2007.

[106] Nirupama Deshpande, Andrew Needles, and Jürgen K Willmann. Molecular ultrasound imaging: current status and future directions. *Clinical radiology*, 65(7):567–581, 2010.

[107] Katherine Ferrara, Rachel Pollard, and Mark Borden. Ultrasound microbubble contrast agents: fundamentals and application to gene and drug delivery. *Annu. Rev. Biomed. Eng.*, 9:415–447, 2007.

[108] A Greco, M Mancini, S Gargiulo, M Gramanzini, PP Claudio, A Brunetti, and M Salvatore. Ultrasound biomicroscopy in small animal research: applications in molecular and preclinical imaging. *Journal of Biomedicine and Biotechnology*, 2012, 2011.

[109] Shashank Sirsi, Jameel Feshitan, James Kwan, Shunichi Homma, and Mark Borden. Effect of microbubble size on fundamental mode high frequency ultrasound imaging in mice. *Ultrasound in medicine & biology*, 36(6):935–948, 2010.

[110] F Stuart Foster, John Hossack, and S Lee Adamson. Micro-ultrasound for preclinical imaging. *Interface focus*, 1(4):576–601, 2011.

[111] Paul Beard. Biomedical photoacoustic imaging. *Interface focus*, 1(4):602–631, 2011.

[112] Vasilis Ntziachristos. Clinical translation of optical and optoacoustic imaging. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1955):4666–4678, 2011.

[113] Lihong V Wang. Multiscale photoacoustic microscopy and computed tomography. *Nature photonics*, 3(9):503–509, 2009.

[114] Rui Ma, Adrian Taruttis, Vasilis Ntziachristos, Daniel Razansky, et al. Multispectral optoacoustic tomography (msot) scanner for whole-body small animal imaging. *Opt. Express*, 17(24):21414–21426, 2009.

[115] A Needles, A Heinmiller, P Ephrat, D Bates, C Bilan-Tracey, C Theodoropoulos, D Hirson, and FS Foster. Development and validation of a combined photoacoustic micro-ultrasound system for in vivo oxygen saturation estimation. In *SPIE BiOS*, pages 78992D–78992D. International Society for Optics and Photonics, 2011.

[116] Bradley E Treeby, Edward Z Zhang, and BT Cox. Photoacoustic tomography in absorbing acoustic media using time reversal. *Inverse Problems*, 26(11):115003, 2010.

[117] Daniel Razansky, Martin Distel, Claudio Vinegoni, Rui Ma, Norbert Perrimon, Reinhard W Köster, and Vasilis Ntziachristos. Multispectral opto-acoustic tomography of deep-seated fluorescent proteins in vivo. *Nature Photonics*, 3(7):412–417, 2009.

[118] Jan Laufer, Peter Johnson, Edward Zhang, Bradley Treeby, Ben Cox, Barbara Pedley, and Paul Beard. In vivo preclinical photoacoustic imaging of tumor vasculature development and therapy. *Journal of biomedical optics*, 17(5):0560161–0560168, 2012.

[119] Andreas Buehler, Eva Herzog, Daniel Razansky, and Vasilis Ntziachristos. Video rate optoacoustic tomography of mouse kidney perfusion. *Optics letters*, 35(14):2475–2477, 2010.

[120] Jan Laufer, Dave Delpy, Clare Elwell, and Paul Beard. Quantitative spatially resolved measurement of tissue chromophore concentrations using photoacoustic spectroscopy: application to the measurement of blood oxygenation and haemoglobin concentration. *Physics in medicine and biology*, 52(1):141, 2007.

[121] Jennifer A Prescher and Christopher H Contag. Guided by the light: visualizing biomolecular processes in living animals with bioluminescence. *Current opinion in chemical biology*, 14(1):80–89, 2010.

[122] Kathryn E Luker and Gary D Luker. Applications of bioluminescence imaging to antiviral research and therapy: multiple luciferase enzymes and quantitation. *Antiviral research*, 78(3):179–187, 2008.

[123] Bruce R Branchini, Danielle M Ablamsky, Audrey L Davis, Tara L Southworth, Braeden Butler, Frank Fan, Amit P Jathoul, and Martin A Pule. Red-emitting luciferases for bioluminescence reporter and imaging applications. *Analytical biochemistry*, 396(2):290–297, 2010.

[124] Hui Zhao, Timothy C Doyle, Olivier Coquoz, Flora Kalish, Bradley W Rice, and Christopher H Contag. Emission spectra of bioluminescent reporters and interaction with mammalian tissue determine the sensitivity of detection in vivo. *Journal of biomedical optics*, 10(4):041210–041210, 2005.

[125] Yoriko Ando, Kazuki Niwa, Nobuyuki Yamada, Toshiteru Enomoto, Tsutomu Irie, Hidehiro Kubota, Yoshihiro Ohmiya, and Hidefumi Akiyama. Firefly bioluminescence quan-

tum yield and colour change by ph-sensitive green emission. *Nature Photonics*, 2(1):44–47, 2007.

[126] Aldo Roda, Massimo Guardigli, Elisa Michelini, and Mara Mirasoli. Bioluminescence in analytical chemistry and in vivo imaging. *Trac Trends in Analytical Chemistry*, 28(3):307–322, 2009.

[127] Tamara Troy, Dragana Jekic-McMullen, Lidia Sambucetti, Brad Rice, et al. Quantitative comparison of the sensitivity of detection of fluorescent and bioluminescent reporters in animal models. *Molecular imaging*, 3(1):9–23, 2004.

[128] David M Barrett, Alix E Seif, Carmine Carpenito, David T Teachey, Jonathan D Fish, Carl H June, Stephan A Grupp, and Gregor SD Reid. Noninvasive bioluminescent imaging of primary patient acute lymphoblastic leukemia: a strategy for preclinical modeling. *Blood*, 118(15):e112–e117, 2011.

[129] S Bhaumik and SS Gambhir. Optical imaging of renilla luciferase reporter gene expression in living mice. *Proceedings of the National Academy of Sciences*, 99(1):377–382, 2002.

[130] Jenny A Hyde, Eric H Weening, MiHee Chang, Jerome P Trzeciakowski, Magnus Höök, Jeffrey D Cirillo, and Jon T Skare. Bioluminescent imaging of borrelia burgdorferi in vivo demonstrates that the fibronectin-binding protein bbk32 is required for optimal infectivity. *Molecular microbiology*, 82(1):99–113, 2011.

[131] Joel C Watts, Kurt Giles, Sunny K Grillo, Azucena Lemus, Stephen J DeArmond, and Stanley B Prusiner. Bioluminescence imaging of aβ deposition in bigenic mouse models of alzheimer's disease. *Proceedings of the National Academy of Sciences*, 108(6):2528–2533, 2011.

[132] John Virostko, Aramandla Radhika, Greg Poffenberger, Zhongyi Chen, Marcela Brissova, Joshua Gilchrist, Brian Coleman, Maureen Gannon, E Duco Jansen, and Alvin C Powers. Bioluminescence imaging in mouse models quantifies β cell mass in the pancreas and after islet transplantation. *Molecular Imaging and Biology*, 12(1):42–53, 2010.

[133] Genevieve C Van de Bittner, Elena A Dubikovskaya, Carolyn R Bertozzi, and Christopher J Chang. In vivo imaging of hydrogen peroxide production in a murine tumor model with a chemoselective bioluminescent reporter. *Proceedings of the National Academy of Sciences*, 107(50):21316–21321, 2010.

[134] Lene Uhrbom, Edward Nerio, and Eric C Holland. Dissecting tumor maintenance requirements using bioluminescence imaging of cell proliferation in a mouse glioma model. *Nature medicine*, 10(11):1257–1260, 2004.

[135] Kelsey L Tinkum, Luciano Marpegan, Lynn S White, Jinwu Sun, Erik D Herzog, David Piwnica-Worms, and Helen Piwnica-Worms. Bioluminescence imaging captures the expression and dynamics of endogenous p21 promoter activity in living mice and intact cells. *Molecular and cellular biology*, 31(18):3759–3772, 2011.

[136] Dirk Stabenow, Marianne Frings, Christina Trück, Katja Gärtner, Irmgard Förster, Christian Kurts, Thomas Tüting, Margarete Odenthal, Hans-Peter Dienes, Karin Cederbrant, et al. Bioluminescence imaging allows measuring cd8 t cell function in the liver. *Hepatology*, 51(4):1430–1437, 2010.

[137] Xuejun Gu, Qizhi Zhang, Lyndon Larcom, and Huabei Jiang. Three-dimensional bioluminescence tomography with model-based reconstruction. *Optics Express*, 12(17):3996–4000, 2004.

[138] Runqiang Han, Jimin Liang, Xiaochao Qu, Yanbin Hou, Nunu Ren, Jingjing Mao, Jie Tian, et al. A source reconstruction algorithm based on adaptive hp-fem for bioluminescence tomography. *Optics Express*, 17(17):14481–14494, 2009.

[139] Jinchao Feng, Kebin Jia, Chenghu Qin, Guorui Yan, Shouping Zhu, Xing Zhang, Junting Liu, Jie Tian, et al. Three-dimensional bioluminescence tomography based on bayesian approach. *Optics Express*, 17(19):16834–16848, 2009.

[140] Vasilis Ntziachristos, Ching-Hsuan Tung, Christoph Bremer, and Ralph Weissleder. Fluorescence molecular tomography resolves protease activity in vivo. *Nature medicine*, 8(7):757–761, 2002.

[141] Nikolaos Deliolanis, Tobias Lasser, Damon Hyde, Antoine Soubret, Jorge Ripoll, Vasilis Ntziachristos, M Lesaffre, F Jean, F Ramaz, AC Boccara, et al. Free-space fluorescence molecular tomography utilizing 360ˆ o geometry projections. *Optics Letters*, 32(4):382, 2007.

[142] L Hervé, A Koenig, A Da Silva, M Berger, J Boutet, JM Dinten, P Peltié, and P Rizo. Noncontact fluorescence diffuse optical tomography of heterogeneous media. *Applied optics*, 46(22):4896–4906, 2007.

[143] Yuting Lin, Han Yan, Orhan Nalcioglu, and Gultekin Gulsen. Quantitative fluorescence tomography with functional and structural< i> a priori</i> information. *Applied optics*, 48(7):1328–1336, 2009.

[144] Fei Liu, Xin Liu, Daifa Wang, Bin Zhang, and Jing Bai. A parallel excitation based fluorescence molecular tomography system for whole-body simultaneous imaging of small animals. *Annals of biomedical engineering*, 38(11):3440–3448, 2010.

[145] Sachin V Patwardhan, Sharon R Bloch, Samuel Achilefu, Joseph P Culver, et al. Time-dependent whole-body fluorescence tomography of probe bio-distributions in mice. *Opt. Express*, 13(7):2564–2577, 2005.

[146] Changqing Li, Gregory S Mitchell, Joyita Dutta, Sangtae Ahn, Richard M Leahy, and Simon R Cherry. A three-dimensional multispectral fluorescence optical tomography imaging system for small animals based on a conical mirror design. *Optics express*, 17(9):7571, 2009.

[147] Ge Wang, Haiou Shen, Kumar Durairaj, Xin Qian, and Wenxiang Cong. The first bioluminescence tomography system for simultaneous acquisition of multiview and multispectral data. *International Journal of Biomedical Imaging*, 2006, 2006.

[148] Ge Wang, Haiou Shen, Ying Liu, Alex Cong, Wenxiang Cong, Yue Wang, Purnima Dubey, et al. Digital spectral separation methods and systems for bioluminescence imaging. *Opt. Express*, 16:1719–1732, 2008.

[149] Alexander D Klose, Bradley J Beattie, Hamid Dehghani, Lena Vider, Carl Le, Vladimir Ponomarev, and Ronald Blasberg. In vivo bioluminescence tomography with a blocking-off finite-difference sp method and mri/ct coregistration. *Medical physics*, 37:329, 2010.

[150] Yujie Lv, Jie Tian, Wenxiang Cong, and Ge Wang. Experimental study on bioluminescence tomography with multimodality fusion. *International journal of biomedical imaging*, 2007, 2007.

[151] Yujie Lu, Hidevaldo B Machado, Qinan Bao, David Stout, Harvey Herschman, and Arion F Chatziioannou. In vivo mouse bioluminescence tomography with radionuclide-based imaging validation. *Molecular Imaging and Biology*, 13(1):53–58, 2010.

[152] Junting Liu, Yabin Wang, Xiaochao Qu, Xiangsi Li, Xiaopeng Ma, Runqiang Han, Zhenhua Hu, Xueli Chen, Dongdong Sun, Rongqing Zhang, et al. In vivo quantitative bioluminescence tomography using heterogeneous and homogeneous mouse models. *Optics express*, 18(12):13102, 2010.

[153] Ralf B Schulz, Angelique Ale, Athanasios Sarantopoulos, Marcus Freyer, Eric Soehngen, Marta Zientkowska, and Vasilis Ntziachristos. Hybrid system for simultaneous fluorescence and x-ray computed tomography. *Medical Imaging, IEEE Transactions on*, 29(2):465–473, 2010.

[154] Xiaoquan Yang, Hui Gong, Guotao Quan, Yong Deng, and Qingming Luo. Combined system of fluorescence diffuse optical tomography and microcomputed tomography for small animal imaging. *Review of Scientific Instruments*, 81(5):054304–054304, 2010.

[155] Dax Kepshire, Niculae Mincu, Michael Hutchins, Josiah Gruber, Hamid Dehghani, Justin Hypnarowski, Frederic Leblond, Mario Khayat, and Brian W Pogue. A microcomputed tomography guided fluorescence tomography system for small animal molecular imaging. *Review of Scientific Instruments*, 80(4):043701–043701, 2009.

[156] Jason Geng. Structured-light 3d surface imaging: a tutorial. *Advances in Optics and Photonics*, 3(2):128–160, 2011.

[157] DL Prout, RW Silverman, and Arion Chatziioannou. Detector concept for opet-a combined pet and optical imaging system. *Nuclear Science, IEEE Transactions on*, 51(3):752–756, 2004.

[158] Ali Douraghy, Fernando R Rannou, Robert W Silverman, and Arion F Chatziioannou. Fpga electronics for opet: a dual-modality optical and positron emission tomograph. *Nuclear Science, IEEE Transactions on*, 55(5):2541–2545, 2008.

[159] George Alexandrakis, Fernando R Rannou, and Arion F Chatziioannou. Effect of optical property estimation accuracy on tomographic bioluminescence imaging: simulation of a combined optical–pet (opet) system. *Physics in medicine and biology*, 51(8):2045, 2006.

[160] Mohamed A Naser and Michael S Patterson. Algorithms for bioluminescence tomography incorporating anatomical information and reconstruction of tissue optical properties. *Biomedical optics express*, 1(2):512, 2010.

[161] Daniel Razansky and Vasilis Ntziachristos. Hybrid photoacoustic fluorescence molecular tomography using finite-element-based inversion. *Medical physics*, 34:4293, 2007.

[162] Weimin Han, Wenxiang Cong, Kamran Kazmi, and Ge Wang. An integrated solution and analysis of bioluminescence tomography and diffuse optical tomography. *Communications in numerical methods in engineering*, 25(6):639–656, 2009.

[163] Qizhi Zhang, Lu Yin, Yiyong Tan, Zhen Yuan, Huabei Jiang, Ge Wang, Haiou Shen, Ying Liu, Alex Cong, Wenxiang Cong, et al. Quantitative bioluminescence tomography guided by diffuse optical tomography. *Optics Express*, 16(3):1481–1486, 2008.

[164] Yiyong Tan and Huabei Jiang. Diffuse optical tomography guided quantitative fluorescence molecular tomography. *Applied optics*, 47(12):2011–2016, 2008.

[165] Julius Pekar. Multispectral bioluminescence tomography with x-ray ct spatial priors. *PhD Thesis*, 2011.

[166] Han Yan, Yuting Lin, William C Barber, Mehmet Burcin Unlu, and Gultekin Gulsen. A gantry-based tri-modality system for bioluminescence tomography. *Review of Scientific Instruments*, 83(4):043708–043708, 2012.

[167] Jin Chen, Vivek Venugopal, Frederic Lesage, and Xavier Intes. Time-resolved diffuse optical tomography with patterned-light illumination and detection. *Optics letters*, 35(13):2121–2123, 2010.

[168] Vivek Venugopal, Jin Chen, and Xavier Intes. Development of an optical imaging platform for functional imaging of small animals using wide-field excitation. *Biomedical optics express*, 1(1):143, 2010.

[169] Yuting Lin, Dave Thayer, Orhan Nalcioglu, and Gultekin Gulsen. Tumor characterization in small animals using magnetic resonance-guided dynamic contrast enhanced diffuse optical tomography. *Journal of biomedical optics*, 16(10):106015–106015, 2011.

[170] Y Lin, MT Ghijsen, H Gao, N Liu, O Nalcioglu, and G Gulsen. A photo-multiplier tube-based hybrid mri and frequency domain fluorescence tomography system for small animal imaging. *Physics in Medicine and Biology*, 56(15):4731, 2011.

[171] Liji Cao, Mathies Breithaupt, and Jörg Peter. Geometrical co-calibration of a tomographic optical system with ct for intrinsically co-registered imaging. *Physics in medicine and biology*, 55(6):1591, 2010.

[172] Hamamatsu. EM-CCD Technical Note, 2009.

[173] Ali Behrooz, Chaincy Kuo, Heng Xu, and Brad Rice. Adaptive row-action inverse solver for fast noise-robust three-dimensional reconstructions in bioluminescence tomography: theory and dual-modality optical/computed tomography in vivo studies. *Journal of biomedical optics*, 18(7):076010–076010, 2013.

[174] Xianyu Su and Wenjing Chen. Fourier transform profilometry:: a review. *Optics and Lasers in Engineering*, 35(5):263–284, 2001.

[175] Mitsuo Takeda and Kazuhiro Mutoh. Fourier transform profilometry for the automatic measurement of 3-d object shape. *Applied optics*, 22(24):3977–3982, 1983.

[176] V Srinivasan, HC Liu, and Maurice Halioua. Automated phase-measuring profilometry of 3-d diffuse objects. *Applied optics*, 23(18):3105–3108, 1984.

[177] Giovanna Sansoni, Matteo Carocci, and Roberto Rodella. Three-dimensional vision based on a combination of gray-code and phase-shift light projection: analysis and compensation of the systematic errors. *Applied Optics*, 38(31):6565–6573, 1999.

[178] Joaquim Salvi, Jordi Pages, and Joan Batlle. Pattern codification strategies in structured light systems. *Pattern Recognition*, 37(4):827–849, 2004.

[179] Henrik O Saldner and Jonathan M Huntley. Temporal phase unwrapping: application to surface profiling of discontinuous objects. *Applied optics*, 36(13):2770–2775, 1997.

[180] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R Marschner, Mark Horowitz, Marc Levoy, and Hendrik Lensch. Dual photography. *ACM Transactions on Graphics (TOG)*, 24(3):745–755, 2005.

[181] P. Cignoni. Meshlab home page. *http://meshlab.sourceforge.net/*, 2013.

[182] Michael Jermyn, Brian Pogue, Hamid R Ghadyani, Scott Davis, Michael A Mastanduno, and Hamid Dehghani. A user-enabling visual workflow for near-infrared light transport modeling in tissue. In *Biomedical Optics*. Optical Society of America, 2012.

[183] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. Meshlab: an open-source mesh processing tool. In *Eurographics Italian Chapter Conference*, pages 129–136. The Eurographics Association, 2008.

[184] Joachim Schöberl. Netgen an advancing front 2d/3d-mesh generator based on abstract rules. *Computing and visualization in science*, 1(1):41–52, 1997.

[185] NETGEN Homepage. http://www.hpfem.jku.at/netgen/. Accessed: 2011-01-01.

[186] Massimiliano Corsini, Paolo Cignoni, and Roberto Scopigno. Efficient and flexible sampling with blue noise properties of triangular meshes. *Visualization and Computer Graphics, IEEE Transactions on*, 18(6):914–924, 2012.

[187] Jorge Ripoll and Vasilis Ntziachristos. Imaging scattering media from a distance: theory and applications of noncontact optical tomography. *Modern Physics Letters B*, 18(28n29):1403–1431, 2004.

[188] Xueli Chen, Xinbo Gao, Xiaochao Qu, Duofang Chen, Xiaopeng Ma, Jimin Liang, and Jie Tian. Generalized free-space diffuse photon transport model based on the influence analysis of a camera lens diaphragm. *Applied optics*, 49(29):5654–5664, 2010.

[189] Jorge Ripoll, Ralf B Schulz, and Vasilis Ntziachristos. Free-space propagation of diffuse light: theory and experiments. *Physical review letters*, 91(10):103901, 2003.

[190] Xueli Chen, Xinbo Gao, Xiaochao Qu, Jimin Liang, Lin Wang, DaâĂŹan Yang, Anikitos Garofalakis, Jorge Ripoll, Jie Tian, et al. A study of photon propagation in free-space based on hybrid radiosity-radiance theorem. *Optics Express*, 17(18):16266–16280, 2009.

[191] Xueli Chen, Xinbo Gao, Duofang Chen, Xiaopeng Ma, Xiaohui Zhao, Man Shen, Xiangsi Li, Xiaochao Qu, Jimin Liang, Jorge Ripoll, et al. 3d reconstruction of light flux distribution on arbitrary surfaces from 2d multi-photographic images. *Opt. Express*, 18(19):19876–19893, 2010.

[192] Xue-Li Chen, Heng Zhao, Xiao-Chao Qu, Duo-Fang Chen, Xiao-Rui Wang, and Ji-Min Liang. All-optical quantitative framework for bioluminescence tomography with non-contact measurement. *International Journal of Automation and Computing*, 9(1):72–80, 2012.

[193] Xueli Chen, Jimin Liang, Xiaochao Qu, Yanbin Hou, Shouping Zhu, Duofang Chen, Xinbo Gao, and Jie Tian. Mapping of bioluminescent images onto ct volume surface for dual-modality blt and ct imaging. *Journal of X-Ray Science and Technology*, 20(1):31–44, 2012.

[194] Xueli Chen, Xinbo Gao, Xiaochao Qu, Duofang Chen, Bin Ma, Lin Wang, Kuan Peng, Jimin Liang, and Jie Tian. Qualitative simulation of photon transport in free space based on monte carlo method and its parallel implementation. *Journal of Biomedical Imaging*, 2010:10, 2010.

[195] Kathryn E Luker, Laura Anne Mihalko, Bradley T Schmidt, Sarah A Lewin, Paramita Ray, Dmitry Shcherbo, Dmitriy M Chudakov, and Gary D Luker. In vivo imaging

of ligand receptor binding with gaussia luciferase complementation. *Nature Medicine*, 18(1):172–177, 2011.

[196] Per Christian Hansen and Dianne Prost O'Leary. The use of the L-curve in the regularization of discrete ill-posed problems. *SIAM Journal on Scientific Computing*, 14(6):1487–1503, 1993.

[197] Natalie Fortier, Guy Demoment, and Yves Goussard. GCV and ML methods of determining parameters in image restoration by regularization: Fast computation in the spatial domain and experimental comparison. *Journal of visual communication and image representation*, 4(2):157–170, 1993.

[198] Judit Chamorro-Servent, Juan Aguirre, Jorge Ripoll, Juan José Vaquero, and Manuel Desco. Feasibility of U-curve method to select the regularization parameter for fluorescence diffuse optical tomography in phantom and small animal studies. *Optics Express*, 19:11490, 2011.

[199] Vivek Venugopal and Xavier Intes. Adaptive wide-field optical tomography. *Journal of Biomedical Optics*, 18(3):036006–036006, 2013.

[200] Michael Jermyn, Brian Pogue, Subhadra Srinivasan, Scott Davis, and Hamid Dehghani. Robust algorithm for automated source placement in near-infrared diffuse imaging. In *Biomedical Optics*. Optical Society of America, 2010.

[201] Edward E Graves, Joseph P Culver, Jorge Ripoll, Ralph Weissleder, and Vasilis Ntziachristos. Singular-value analysis and optimization of experimental parameters in fluorescence molecular tomography. *JOSA A*, 21(2):231–241, 2004.

[202] Christopher H Contag and Michael H Bachmann. Advances in in vivo bioluminescence imaging of gene expression. *Annual Review of Biomedical Engineering*, 4(1):235–260, 2002.

[203] Alexander Cong, Wenxiang Cong, Yujie Lu, Peter Santago, Arion Chatziioannou, and Ge Wang. Differential evolution approach for regularized bioluminescence tomography. *Biomedical Engineering, IEEE Transactions on*, 57(9):2229–2238, 2010.

[204] Jinchao Feng, Chenghu Qin, Kebin Jia, Dong Han, Kai Liu, Shouping Zhu, Xin Yang, and Jie Tian. An adaptive regularization parameter choice strategy for multispectral bioluminescence tomography. *Medical Physics*, 38:5933, 2011.

[205] Ge Wang, Wenxiang Cong, Haiou Shen, Xin Qian, Michael Henry, Yue Wang, et al. Overview of bioluminescence tomography-a new molecular imaging modality. *Frontiers in Bioscience*, 13(1281-1293):190, 2008.

[206] Ge Wang, Haiou Shen, Wenxiang Cong, Shan Zhao, and Guo Wei Wei. Temperature-modulated bioluminescence tomography. *Optics Express*, 14(17):7852–7871, 2006.

[207] Irving H Blifford Jr et al. Factors affecting the performance of commercial interference filters. *Applied Optics*, 5(1):105–111, 1966.

[208] Andrea Bassi, Andrea Farina, Cosimo D'Andrea, Antonio Pifferi, Gianluca Valentini, and Rinaldo Cubeddu. Portable, large-bandwidth time-resolved system for diffuse optical spectroscopy. *Optics Express*, 15(22):14482–14487, 2007.

[209] Qun Zhu, Frederic Leblond, Fadi El-Ghussein, Brian W Pogue, and Hamid Dehghani. Development and evaluation of a time-resolved near-infrared fluorescence finite element model. In *SPIE BiOS*, pages 78960T–78960T. International Society for Optics and Photonics, 2011.

[210] Normand P Nantel, David G Nilson, Bradley W Rice, and Tamara L Troy. Tissue phantom calibration device for low level light imaging systems, 2009. US Patent 7,629,573.

[211] Miroslawa Poprawa-Smoluch, Jacob Baggerman, Hong Zhang, Huub PA Maas, Luisa De Cola, and Albert M Brouwer. Photoisomerization of disperse red 1 studied with transient absorption spectroscopy and quantum chemical calculations. *The Journal of Physical Chemistry A*, 110(43):11926–11937, 2006.

[212] Theodore Moffitt, Yin-Chu Chen, and Scott A Prahl. Preparation and characterization of polyurethane optical phantoms. *Journal of Biomedical Optics*, 11(4):041103–041103, 2006.