

USE OF PRIOR INFORMATION AND
PROBABILISTIC IMAGE RECONSTRUCTION
FOR OPTICAL TOMOGRAPHIC IMAGING

by

HECTOR RICHARD ABRAHAM BASEVI

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
September 2014

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

Preclinical bioluminescence tomographic reconstruction is underdetermined. This work addresses the use of prior information in bioluminescence tomography to improve image acquisition, reconstruction, and analysis.

A structured light surface metrology method was developed to measure surface geometry and enable robust and automatic integration of mirrors into the measurement process. A mouse phantom was imaged and accuracy was measured at 0.2mm with excellent surface coverage.

A sparsity-regularised reconstruction algorithm was developed to use instrument noise statistics to automatically determine the stopping point of reconstruction. It was applied to *in silico* and *in simulacra* data and successfully reconstructed and resolved two separate luminescent sources within a plastic mouse phantom.

A Bayesian framework was constructed that incorporated bioluminescence properties and instrument properties. Distribution expectations and standard deviations were estimated, providing reconstructions and measures of reconstruction uncertainty. The reconstructions showed superior performance when applied to *in simulacra* data compared to the sparsity-based algorithm.

The information content of measurements using different sets of wavelengths was quantified using the Bayesian framework via mutual information and applied to an *in silico* problem. Significant differences in information content were observed and comparison against a condition number-based approach indicated subtly different results.

ACKNOWLEDGMENTS

This work would not have been possible without my supervisors Dr Iain B. Styles, Dr Hamid Dehghani, Prof Ela Claridge, and Prof Jon Frampton. I would particularly like to thank Iain, my primary supervisor, without whom this project would not have existed. I am extremely grateful to him for always finding the time to meet with me and providing a second opinion. His advice and assistance on all subjects has been invaluable. Hamid has been my second supervisor but has always been willing to participate, help, and come up with ideas. He has been a fount of knowledge on the diffuse optics community and literature. I would also have been much more lost with regards to Nirfast, and physically lost at conferences, without him. I would also like to thank them both for reading through this lengthy document more times than obligated to and at shorter notice than is reasonable.

I would like to thank Ela who has always been ready to listen and offer advice, and whose Graphics module, which I had unofficially attended, proved unexpectedly useful during my PhD. It has been helpful to have a supervisor who understands the technical aspects of the work but is also not too invested in it.

I would also like to thank Jon, for taking time out of his busy schedule to meet and offer a very different perspective on imaging. It is easy to lose sight of the big picture when working on computer simulations in the basement of the School of Computer Science, and Jon has provided great insight into the challenges and practicalities of actually performing preclinical biomedical science.

I have not mentioned James Guggenheim until now but his contributions have been just as important. James has been a fellow PSIBS student in the same area of bioluminescence

tomography. Working in the same research area and sharing an office for the second half of the PhD has meant that we have often worked closely together and have bounced many ideas off each other. It has been a great pleasure working with him. James had a strong involvement in the surface metrology work described later in this document, and indeed the surface metrology project arose from his dissatisfaction with my previous crude implementation of a Fourier-transform profilometry-based system. The design of the surface metrology imaging algorithms and techniques was largely mine, but James also contributed to the iterative design process and built and calibrated the system. If I had tried to do so it would likely have been a disaster. He designed and constructed a multi-modality imaging system (see his thesis for details) was used to acquire the data used for the majority of the work in this thesis.

I would like to thank my office mates and PSIBS colleagues for their friendship and assistance. To Andrew Palmer, Yuxuan Zhan, Richard Williams, Rory Steven, Eric Pitkeathly, Miriam Dowle, James Brown, Alan Race, Rachel Flight and everyone else, it has been an honour.

Robert Holt, Michael Jermyn, Kelly Michaelsen, Frederic Leblond, Kenneth Tichauer, Scott Davis, Brian Pogue, Venkataramanan Krishnaswamy and everyone else at Dartmouth College also deserve thanks for the interesting discussions, help with Nirfast, and for agreeing to image the XPM-2 phantom in your fluorescence imaging system.

Penultimately, I like to thank the past and current staff of the PSIBS Doctoral Training centre: Prof Michael Hannon, Dr Josephine Bunch, Helen Williams, Louise Oram, and Ann Smith, as well as all the academics involved in the PSIBS program. It has been a pleasure working with you and learning from you.

Finally, I would like to thank my parents for their constant support. I have not said so enough.

Funding for my PhD was provided by EPSRC grant EP/F50053X/1 through a studentship at the PSIBS Doctoral Training Centre at the University of Birmingham.

PUBLICATIONS ASSOCIATED WITH THIS WORK

A number of publications have resulted from the work presented in this thesis. Figures from these publications have been used in this thesis and references have been placed in the figure captions specifying the origin of the figures where that is the case.

A patent application with publication number *WO2014091214 A1* has been filed for the surface metrology technology described in chapter 2.

Journal papers

The surface metrology system presented in chapter 2 was published in the paper (Basevi et al. 2013b):

H. R. A. Basevi, J. A. Guggenheim, H. Dehghani, and I. B. Styles (2013). “Simultaneous multiple view high resolution surface geometry acquisition using structured light and mirrors”. In: *Optics Express* 21.6, pp. 7222–7239.

The surface metrology system was used in the process of generating a free space light propagation model (see Appendix A.5) via a method that was published in the paper (Guggenheim et al. 2013b):

J. A. Guggenheim, H. R. A. Basevi, I. B. Styles, J. Frampton, and H. Dehghani (2013). “Quantitative surface radiance mapping using multi-view images of light-emitting turbid media”. In: *Journal of the Optical Society of America A: Optics, Image Science, and Vision* 30.12, pp 2572–2584.

The reconstruction algorithm presented in chapter 3 was published in the paper (Basevi et al. 2012b):

H. R. A. Basevi, K. M. Tichauer, F. Leblond, H. Dehghani, J. A. Guggenheim, R. W. Holt, and I. B. Styles (2012). “Compressive sensing based reconstruction in bioluminescence tomography improves image resolution and robustness to noise”. In: *Biomedical Optics Express* 3.9, pp. 2131–2141.

The surface metrology system and reconstruction algorithm were utilised in the paper on the multi-modality imaging system described in Appendix A that was published in the paper (Guggenheim et al. 2013a):

J. A. Guggenheim, H. R. A. Basevi, J. Frampton, I. B. Styles, and H. Dehghani (2013). “Multi-modal molecular diffuse optical tomography system for small animal imaging”. In: *Measurement Science and Technology* 24.10, p 105405.

Conference papers

The reconstruction algorithm presented in chapter 3 was published in the abstract (Basevi et al. 2012a):

H. R. A. Basevi, J. A. Guggenheim, H. Dehghani, and I. B. Styles (2012). “Application of Compressive Sensing to Bioluminescence Tomography”. In: *Biomedical Optics*. Optical Society of America, BSu3A. 64.

The Markov Chain Monte Carlo-based reconstruction and uncertainty estimation technique presented in chapter 5 was published in the abstract (Basevi et al. 2014):

H. R. A. Basevi, J. A. Guggenheim, S. Taylor, H. Dehghani, and I. B. Styles (2014). “Towards Bayesian Reconstruction and Analysis in Bioluminescence Tomography via Markov Chain Monte Carlo Techniques”. In: *Biomedical Optics*. Optical Society of America, BM3A. 45.

The mutual information-based measurement wavelength optimisation method presented in chapter 6 was published in the paper (Basevi et al. 2013a):

H. R. A. Basevi, J. A. Guggenheim, H. Dehghani, and I. B. Styles (2013). “Information-theoretic method for wavelength selection in bioluminescence tomography”. In: *European Conference on Biomedical Optics*. Optical Society of America, 879909.

The reconstruction algorithm presented in chapter 3 was used with a method of dimensionality reduction for bioluminescence tomography reconstruction that was published in the paper (Styles et al. 2013):

I. B. Styles, H. R. A. Basevi, J. A. Guggenheim, H. Dehghani (2013). “Random matrix-based dimensionality reduction for bioluminescence tomography reconstruction”. In: *European Conference on Biomedical Optics*. Optical Society of America, p. 87990J.

CONTENTS

1	Introduction	1
1.1	Types of pre-clinical tomographic imaging	4
1.2	Bioluminescence tomographic imaging and reconstruction	7
2	Surface geometry measurement using structured light	15
2.1	Contributions to this work	15
2.2	Introduction	15
2.3	Survey of surface metrology systems used in optical tomographic imaging systems	27
2.4	Methods	33
2.4.1	Wrapped phase measurement	35
2.4.2	Phase unwrapping	36
2.4.3	Optical model to convert phase measurements into spatial coordinates	38
2.4.4	Mirror integration	42
2.4.5	Noise removal	48
2.4.6	Conversion from point clouds to meshes	50
2.4.7	System calibration	53
2.4.8	Experiments	54
2.5	Results and discussion	54
2.6	Conclusions	59
3	Application of Compressive Sensing-based techniques to Biolumines-	

cence Tomography	65
3.1 Contributions to this work	65
3.2 Introduction	65
3.3 Survey of reconstruction algorithms used in bioluminescence tomographic imaging	74
3.4 Methods	82
3.4.1 CSCG algorithm design	82
3.4.2 CSCG algorithm implementation	87
3.4.3 Experiments	88
3.5 Results and discussion	95
3.6 Conclusions	113
4 An example of Surface Capture and Compressive Sensing-inspired re- construction to Bioluminescence Tomography	119
4.1 Contributions to this work	119
4.2 Introduction	119
4.3 Methods	120
4.4 Results and discussion	126
4.5 Conclusions	139
5 Posterior sampling using Markov Chain Monte Carlo for Biolumines- cence Tomography	141
5.1 Introduction	141
5.2 Posterior distribution design	146
5.2.1 Knowledge of the imaging process	148
5.2.2 The likelihood	150
5.2.3 The prior distribution	151
5.2.4 Designing the likelihood function	152
5.2.5 Designing the prior distribution	157

5.3	Methods	158
5.3.1	Posterior distribution construction	158
5.3.2	Markov Chain Monte Carlo sampling	160
5.4	Results and discussion	162
5.5	Conclusion	165
6	Optimisation of bioluminescence tomography measurement wavelengths using mutual information as a metric of information content	168
6.1	Introduction	168
6.2	Survey of measurement optimisation methods in optical tomographic imaging	171
6.3	Methods	175
6.3.1	Design of the prior distribution	176
6.3.2	Design of the likelihood distribution	178
6.3.3	Mutual information calculation	179
6.3.4	Experiments	180
6.3.5	Validation	180
6.4	Results and discussion	181
6.5	Conclusions	185
7	Conclusions	189
	Appendices	202
A	The multi-modality imaging system used in this work	203
A.1	Contributions to this work	203
A.2	Introduction	203
A.3	Instrument hardware	204
A.4	Instrument software	208
A.5	The free space light propagation model	209
A.6	Image reconstruction hardware	210

A.7 Image reconstruction software	211
A.8 XPM-2 Phantom Mouse	211
A.9 Conclusions	213
List of References	214

LIST OF FIGURES

2.1	XPM-2 phantom imaged under structured illumination	34
2.2	Phase measurement dependence on number of phase shift images and signal strength	37
2.3	Diagram of a simplified surface metrology system	39
2.4	Example of dual photography	45
2.5	Correspondences between directly imaged and mirror pixels	47
2.6	XPM-2 phantom signal intensity, unwrapping error, and error mask for noise removal	51
2.7	Wrapped and unwrapped phase maps	55
2.8	Surface coverage and mesh of XPM-2 phantom	57
2.9	Error associated with XPM-2 mesh	58
2.10	Mouse cadaver photographs and phase maps	60
2.11	Mouse cadaver surface points and mesh	61
2.12	Height maps of a two pence coin	62
3.1	ℓ^1 -norm and ℓ^2 -norm solutions of a two dimensional problem with one measurement	71
3.2	Two dimensional BLT mesh	91
3.3	Examples of randomly generated bioluminescence sources	92
3.4	Dartmouth instrument detector positions on XPM-2 phantom	93
3.5	CSCG reconstruction error on binary CS problem	98
3.6	CSCG reconstruction error standard deviation on binary CS problem . . .	99

3.7	CSCG reconstruction sensitivity on binary CS problem	100
3.8	CSCG reconstruction specificity on binary CS problem	101
3.9	CSCG reconstruction error on two dimensional BLT problem, without column normalisation	102
3.10	CSCG reconstruction error on two dimensional BLT problem, with column normalisation	103
3.11	CSCG reconstruction error standard deviation on two dimensional BLT problem, without column normalisation	104
3.12	CSCG reconstruction error standard deviation on two dimensional BLT problem, with column normalisation	105
3.13	CSCG reconstruction sensitivity on two dimensional BLT problem, without column normalisation	106
3.14	CSCG reconstruction sensitivity on two dimensional BLT problem, with column normalisation	107
3.15	CSCG reconstruction total bioluminescence error on two dimensional BLT problem, without column normalisation	108
3.16	CSCG reconstruction total bioluminescence error on two dimensional BLT problem, with column normalisation	109
3.17	Reconstruction algorithm performance on three dimensional BLT problem, in the absence of noise	110
3.18	Reconstruction algorithm performance on three dimensional BLT problem, with 1% Gaussian noise	111
3.19	Reconstruction algorithm performance on three dimensional BLT problem, with 5% Gaussian noise	112
3.20	Reconstruction algorithm performance for XPM-2 phantom using data acquired at and by Dartmouth College	114
4.1	Images of XPM-2 phantom source A	127
4.2	Images of XPM-2 phantom source B	128

4.3	Images of XPM-2 phantom source C	129
4.4	Reconstructions of XPM-2 phantom source A	131
4.5	Reconstructions of XPM-2 phantom source A	132
4.6	Reconstructions of XPM-2 phantom source B	134
4.7	Reconstructions of XPM-2 phantom source B	135
4.8	Reconstructions of XPM-2 phantom source C	136
4.9	Reconstructions of XPM-2 phantom source C	137
5.1	Expectations and standard deviations estimated from posterior samples for source A	162
5.2	Expectations and standard deviations estimated from posterior samples for source B	163
5.3	Expectations and standard deviations estimated from posterior samples for source C	164
6.1	Test problem used for mutual information estimation	177
6.2	A sample from the prior distribution	178
6.3	Mutual information and validation for single wavelengths	182
6.4	Mutual information and posterior distribution standard deviation estimates for wavelength sets consisting of two wavelengths	184
6.5	Mutual information and validation for wavelength sets consisting of three wavelengths.	185
6.6	Condition numbers of Jacobians	186
A.1	Multimodal imaging instrument	206
A.2	C9100-14 ImageEM-1L EM-CCD camera quantum efficiency	207
A.3	10nm bandwidth interference filter transmittances	207
A.4	Mirror reflectivity	207
A.5	XPM-2 Phantom Mouse picture and optical properties	212

LIST OF TABLES

3.1	Reconstruction algorithm localisation error on three dimensional BLT problem	113
3.2	Reconstruction algorithm three dimensional localisation error and standard deviation on three dimensional BLT problem for source A	115
3.3	Reconstruction algorithm three dimensional localisation error and standard deviation on three dimensional BLT problem for source B	116
3.4	Reconstruction algorithm volume error on three dimensional BLT problem	116
4.1	Reconstructed XPM-2 source A total intensity	133
4.2	Reconstructed XPM-2 source B total intensity	138
4.3	Reconstructed XPM-2 source C total intensity	138
5.1	Reconstructed total intensity of XPM-2 sources A, B, and A and B together	165
A.1	XPM-2 phantom source intensity	213

ACRONYMS

ATP Adenosine Triphosphate.

BLI Bioluminescence Imaging.

BLT Bioluminescence Tomography.

CCD Charge Coupled Device.

CG conjugate gradient.

CS Compressive Sensing.

CSCG compressive sensing conjugate gradient.

CT Computed Tomography.

DA Diffusion Approximation.

DMD digital micromirror device.

DOT Diffuse Optical Tomography.

EM Electron Multiplying.

EM-CCD Electron Multiplying Charge Coupled Device.

FEM Finite Element Method.

FMT Fluorescence Molecular Tomography.

FWHM full-width-half-maximum.

GN Gauss-Newton.

GPU Graphics Processing Unit.

LASSO least absolute shrinkage and selection operator.

LCD liquid crystal display.

MAP maximum *a posteriori*.

MC Monte Carlo.

MCMC Markov Chain Monte Carlo.

MRI Magnetic Resonance Imaging.

NCG nonlinear conjugate gradient.

NNLS non-negative least squares.

PAT Photoacoustic Tomography.

PMT photomultiplier tube.

RAM Random Access Memory.

RTE Radiative Transfer Equation.

SNR signal-to-noise ratio.

TOF time of flight.

TV total variation.

VIS-NIR visible-near infra-red.

CHAPTER 1

INTRODUCTION

There are a number of types of experiment that are performed as part of preclinical biomedical research, including *in silico*, *in vitro*, *ex vivo*, and *in vivo* experiments. Each have their advantages and disadvantages.

In silico experiments involve computational models and the exploration of biomedical questions through simulation. Computational models do not of themselves require harm to a living organism and so are free of ethical issues. The scale of experimentation is also limited only by the computational resources available and so it is relatively straightforward to scale up experiments by purchasing more computer hardware. However, computational models are man-made and so are only as realistic as their design. Inevitably, these models are limited by their relative simplicity.

In vitro experiments involve tests conducted outside of the natural environment which is commonly replaced with a petri dish. Experiments are conducted on individual or several layers of cells, and so few ethical issues are involved. These experiments involve living organisms unlike computational modelling and so allow the examination of more complex biological systems which may not have been fully characterised. However, cells are grown in a controlled environment without physical structures and biological molecules originating from the cells' natural environment. This limitation affects the degree to which results obtained from *in vivo* experiments are applicable to whole organisms.

Ex vivo experiments involve tests conducted on larger and more representative bio-

logical systems than *in vitro* experiments but still involve experimentation outside of a living subject. These can involve tissues or whole organs which can in some cases be grown artificially but are often extracted from animal subjects. These biological systems are more realistic because they contain representative physical and biological structures, and vasculature. However, these systems still contain only a subset of the tissues of an organism and must be maintained artificially which can be difficult over long periods. Ethical issues arise here over the harvesting of these tissues.

In vivo experiments are the most biologically realistic experiments performed as part of preclinical research because they are performed using whole organisms. The major difference between *in vivo* preclinical research and clinical research is the choice of subject species. This type of research can produce the most powerful results and is often a necessary step before clinical testing of an experimental treatment. *In vivo* experiments can also be performed over long periods as the complexity of maintaining a suitable environment is reduced. These experiments involve the most serious ethic issues as they are conducted on whole organisms that can experience discomfort and pain. *In vivo* experiments may be necessary because some environments cannot be adequately replicated using *in silico*, *in vitro*, and *ex vivo* techniques.

Optical imaging offers a number of advantages that make it suitable to probing *in vitro*, *ex vivo*, and *in vivo* experiments. Optical manipulation is a mature field and it is possible to produce, manipulate, collect, and analyse light via a variety of methods. Light is also non-harmful in low intensities and so it can be used to probe biological systems without confounding experimental results. Certain optical techniques are appropriate for certain types of preclinical experiment. *In vitro* experiments involve single cells or a small number of cell layers and so scattering and absorption is limited. This enables direct microscopic imaging, and a number of techniques allow imaging of multiple features of cells at resolutions approaching and in some cases exceeding the diffraction limit of light. *Ex vivo* experiments may involve microscopy, but absorption and scattering of light are functions of tissue thickness and *ex vivo* tissue thickness may only allow surface imaging.

Background signals such as autofluorescence may also complicate analysis.

The issues of resolution and depth are more pronounced in *in vivo* experiments as it may be necessary to image through a whole organism. However, *In vivo* experiments are biologically valuable because of their realism and so it may be beneficial to perform preclinical *in vivo* optical imaging despite these limitations. Mouse models are commonly used in preclinical studies and tissue thickness in mice is on the order of several centimetres, which is within the limits of detection of optical imaging techniques (Ntziachristos et al. 2005). A number of optical probes have been developed which emit light and provide the ability to label particular cell types, organelles, and even gene expression. These labels may be exogenous and introduced via injection, temporarily endogenously manufactured via cellular machinery through the addition of plasmids, or permanently endogenously manufactured through the splicing of genetic material into the nuclear deoxyribose nucleic acid (DNA). Two primary types of these labels use the mechanisms of fluorescence and bioluminescence to produce light, and are available with a variety of emission spectra.

Ethical concerns limit the use of *in vivo* investigations due to the suffering of the subjects. Consequently it is necessary to maximise the efficiency of studies to obtain a maximum amount of information from a minimum number of subjects. Surface imaging is of limited use because the position of a light-emitting marker in tissue and its distance from the surface determines the proportion of light reaching the surface (Guggenheim et al. 2013b). Tomographic imaging techniques can address this need by accounting for the effects of tissue absorption and scattering to measure distributions of markers in three dimensions. This can localise the features of interest to particular structures and organs and quantitatively reconstruct marker concentration. However, tomographic reconstruction of fluorescent and bioluminescent markers is an active area of research because of the degree to which information provided by surface measurements is affected by tissue absorption and scattering.

This work explores potential solutions to the imaging reconstruction challenges of Bioluminescence Tomography (BLT) in the context of pre-clinical *in vivo* imaging. It

investigates the acquisition and utilisation of additional information and its use in enhancing the process of reconstructing quantitative three dimensional bioluminescence distributions. Measurements of subject surface geometry and prior knowledge of properties of bioluminescence distributions and instrumental characteristics are integrated into the reconstruction process, and are used to enable image reconstruction, improve image reconstruction quality, and produce metrics of reconstruction confidence to aid interpretation. The proposed improvements are a step towards reducing the number of experimental subjects used within biomedical preclinical studies without compromising study effectiveness.

1.1 Types of pre-clinical tomographic imaging

A variety of *in vivo* imaging techniques allow examination of a number of aspects of biology and physiology (Ntziachristos et al. 2005; Tremoleda et al. 2011; Wessels et al. 2007). The most widely used techniques that can image deep within tissue have limitations that affect their applicability to studies that require imaging of a number of subjects at multiple time points. Magnetic Resonance Imaging (MRI) provides high resolution imaging with good soft tissue contrast in terms of water content and is non-ionising, but is expensive. Computed Tomography (CT) also provides high resolution imaging, but has low soft tissue contrast and uses ionising radiation which can prevent imaging at multiple time points. Optical tomographic imaging techniques can potentially address this gap as light is non-ionising and optical instrumentation is relatively inexpensive.

Diffuse Optical Tomography (DOT) allows measurement of the absorption and scattering properties of tissue to depths of several centimetres via the use of exogenous light (Ntziachristos et al. 2005). Different tissues have varying optical absorption and scattering properties (Cheong et al. 1990) and so DOT can provide soft tissue contrast. Certain compounds such as oxygenated and deoxygenated haemoglobin, lipids, and water, dominate the absorption of many tissues and have spectrally dependent absorption properties (Alexandrakis et al. 2005; Srinivasan et al. 2003). The scattering of tissue can be approxi-

mated by Mie theory which assumes spherical scatterers and is dependent on scatter size, and also has a spectral dependence (Mourant et al. 1997). Consequently it is possible to reconstruct the concentrations of tissue chromophores and properties of tissue scatterers by using multiple wavelengths of probe light (Dehghani et al. 2008). This light may be introduced through optical fibres physically coupled to the surface of the subject or via a non-contact laser or digital projector. Light reaching the surface after passing through tissue may be collected using optical fibres physically coupled to the surface of the subject or via a non-contact camera sensor.

BLT allows quantification of the distribution of marked cells of interest through the use of a genetic modification that causes those cells to produce endogenous light. This requires a luciferase enzyme which is encoded by the genetic modification as well as oxygen, Adenosine Triphosphate (ATP), and a substrate chemical called luciferin (Seliger and McElroy 1960; White et al. 1971). There are a number of different luciferase enzymes with different chemical and spectral properties (Thorne et al. 2010). The emission spectrum of luciferase is also affected by environmental conditions (White et al. 1971) including body temperature (Liu et al. 2010). The substrate molecule, luciferin, is typically injected into the subject prior to imaging, but luciferin production can also be incorporated within the genetic modification (Contag and Bachmann 2002). The light produced by these cells propagates through tissue and is measured on the surface of the imaging subject. The surface light distribution can be used to infer the number of internal luciferin oxidation reactions, which is related to the concentration of luciferase and therefore to the number of cells of interest. This is accomplished through the inversion of a physical model of light propagation which maps internal bioluminescence to surface light.

Bioluminescence Imaging (BLI) also uses the light produced by these bioluminescent cells but does not involve reconstruction of the internal light distribution. Data analysis is conducted directly on surface bioluminescence measurements which are typically in the form of a camera image. A common application of BLI and BLT is to monitor the state of an artificially introduced tumour to examine the efficacy of a potential cancer treatment

(Edinger et al. 2003; Rehemtulla et al. 2000; Tremoleda et al. 2011). BLI does not require the use of sophisticated reconstruction techniques and is interpretable without specialist knowledge but does not account for tissue attenuation. BLT can provide superior information to BLI though the incorporation of knowledge of tissue optical properties but is used less frequently for pre-clinical studies due to the challenges of BLT image reconstruction. BLT and BLI have advantages in that several bioluminescent markers can be used simultaneously (Luker and Luker 2008), there is no need for excitation light, and there is very little background signal. However, it is difficult to manipulate bioluminescence emission in tissue whereas modalities such as DOT and Fluorescence Molecular Tomography (FMT) can modulate the probe light to gain additional information.

FMT is similar in concept to BLT in that the goal is to reconstruct the internal distribution of light emitting molecules (Leblond et al. 2010; Ntziachristos 2006). However, fluorescent molecules require excitation via an external light source, and the propagation of this excitation light through tissue is also subject to absorption and scattering. Fluorescent agents do not typically require metabolic activity to function and so cells of interest can be fluorescently labelled without genetic manipulation through the use of antibodies to deliver and bind the fluorescent agents to the targeted cells or cellular features of interest. There is often significant tissue autofluorescence (Ntziachristos 2006), unlike BLT.

Photoacoustic Tomography (PAT) utilises both light and acoustic waves to image features of interest (Cox et al. 2009; Laufer et al. 2009; Wang and Hu 2012). It exploits the localised heating that results from light absorption, which causes tissue expansion that in turn induces acoustic waves. Lasers can be used to introduce light into a subject, which is absorbed as it passes through tissue. Tissue absorption varies with the concentration of constituent chromophores and so different tissues absorb different amounts of light, are heated to varying degrees, and generate different intensities of acoustic waves. These waves can be detected using ultrasonic transducers placed on the surface of the subject. The localisation of sound waves can be performed to a higher resolution than pure optical

techniques (DOT, BLT, FMT) and so PAT can produce images with higher resolution for certain types of features. However, PAT is not sensitive to optical scattering and is relatively insensitive to fluorescence (Wang 2008; Wang and Hu 2012), and has high equipment costs. Haemoglobin is one of the primary tissue chromophores and so PAT is often applied to imaging vasculature.

Radioactive tracers can also be used to perform tomographic imaging (Tremoleda et al. 2011; Weber and Ivanovic 1999). Positron emission tomography (PET) imaging uses radioactive positron sources. These positrons annihilate with electrons present within tissue. The energy from these annihilation events produces gamma rays, which are detected by a detector with high temporal resolution. The created gamma rays observe conservation of momentum and so pairs of gamma rays travel in opposite directions, assuming that the momentum of the annihilating positron and electron cancel out, which is not true in practice. The presence and location of annihilation events can be calculated from the directions of gamma rays simultaneously incident on the detector. Radiative contrast agents can be used in conjunction with optical methods (Alexandrakis et al. 2005).

1.2 Bioluminescence tomographic imaging and reconstruction

BLT reconstruction is performed by modelling the propagation of light from a bioluminescent source to the locations at which surface measurements are made, and using this model to optimise the light source distribution so as to better fit the observed measurements. Modelling the propagation of light requires good knowledge of the optical properties of the medium through which the light travels, which can be obtained through DOT (Guggenheim et al. 2013a; Naser and Patterson 2010) or by inferring optical properties from structures present in other modalities (Feng et al. 2009). The modelling is typically performed by solving an equation for particle transport called the Radiative Transfer Equation (RTE), or an approximation such as the Diffusion Approximation (DA), or by

modelling the propagation of large numbers of individual photons using Monte Carlo (MC) techniques to produce statistical estimates of the surface light distribution (see Arridge and Hebden (1997); Arridge and Schweiger (1997) for a summary of forward models).

The first step in the reconstruction of a bioluminescence distribution is the measurement of surface light. This can be accomplished via a range of measurement device which can be broadly divided into contact and non-contact devices. Contact measurements require detectors to be placed on the surface of the imaging subject, and collect all the photons leaving the area of contact (assuming perfect detector efficiency and that the front surface of the detector is refractive index-matched to the subject tissue). For example, a photomultiplier tube (PMT) coupled to one end of an optical fibre constitutes a contact measurement device when the other end of the fibre is placed against the surface of the imaging subject, and is commonly used in DOT and FMT imaging (Hebden et al. 1997). Measurement at multiple wavelengths can improve the quality of reconstructed images (Alexandrakis et al. 2005; Chaudhari et al. 2005).

Non-contact measurement devices are placed away from the imaging subject, and so are less efficient at collecting light from a given area. However, a sensor chip such as a Charge Coupled Device (CCD) allows many measurements to be acquired simultaneously. This enables measurement of the surface light over a large portion of the surface of the imaging subject with high spatial resolution in conjunction with an appropriate system of lenses. The large area of measurement and high spatial resolution are advantageous because tissues are generally scattering dominated and so photons exit the imaging subject over a large fraction of its surface regardless of the compactness of the internal source. Knowledge of this surface light distribution can allow calculation of the internal source distribution.

There are a number of challenges associated with BLT imaging and image reconstruction.

- The number and type of measurements available for image reconstruction are limited by the surface coverage, spatial resolution and spectral resolution of the imaging

instrument.

- Measurements can only be acquired on the surface of the imaging subject rather than within tissue.
- Tissue is highly absorbing within the visible range which greatly attenuates signal.
- Tissue is scattering dominated, which results in destruction of directional and phase information and contributes to greater total absorption through increased photon path lengths.
- The duration for which a subject can be imaged is limited by restrictions on the length of time for which a subject can be kept sedated and so still for imaging, by substrate consumption (Burgos et al. 2003), and by the dynamics of other biological processes that affect metabolism.

These challenges mean that a set of surface measurement at a single wavelength is insufficient on its own to uniquely specify a bioluminescence distribution in general (Wang et al. 2004). Imaging at a number of wavelengths can improve uniqueness (Han and Wang 2007) and a number of these factors can be mitigated to a degree, but challenges inherent to the imaging process remain.

- The number of measurements can be increased through the use of a detector with greater spatial and/or spectral resolution. However, the imaging subject has a finite surface area and the emission spectrum is continuous and may be strongly peaked (Seliger and McElroy 1960). Increasing the number of measurements beyond a level dictated by the spatial point spread function of the imaging system results in increased measurement dependence, as does the use of smaller spectral bands in the same spectral region. Increasing the resolution of the imaging sensor beyond this point provides diminishing returns in terms of the total amount of information contained within the measurement set.

- Sub-surface measurements may be acquired via invasive means such as needle probes, but doing so may provoke a physiological response in the subject and so confound the conclusions of the biomedical study.
- High tissue absorption greatly reduces the fraction of bioluminescence that reaches the detector. Exposure time can be extended to increase signal-to-noise ratio (SNR), but cannot be extended indefinitely due to limitations in the length of time for which a subject can be imaged.
- High tissue scattering destroys directional information and causes light to exit the surface of the subject over a large region. The quantity of information collected can be increased by imaging over a larger region of the subject surface area, and the SNR can be increased by raising the exposure time, but the above limitations apply.

These challenges result in measurement sets that typically contain fewer measurements than the number of unknowns that parameterise the bioluminescence distribution. The measurements are also not pairwise independent due to scattering which reduces the information content of the measurement set. Surface signal strength is also typically low as it drops exponentially with the distance of the bioluminescent source from the surface, and bioluminescent signals may be further attenuated by particularly strongly absorbing organs. Measurements may also be contaminated by a statistical and instrument noise.

The second step in the reconstruction process involves modelling light propagation in tissue. This is necessary because the ability to calculate the fluence throughout the imaging subject that is associated with a given bioluminescence distribution allows calculation of the surface measurements which would arise from that distribution. A suitable reconstruction algorithm can then use the light propagation model to search for the bioluminescence distribution that produces measurements that best match those observed. There are a number of types of light propagation model that can be used in this situation of which the two most popular are the Finite Element Method (FEM) coupled with the DA, and the MC method coupled with statistical rules for the propagation of a single

photon (Arridge and Hebden 1997; Arridge and Schweiger 1997).

FEM enables approximate solution of the weak forms of differential equations by discretising the domain of interest into a number of disjoint region (Arridge et al. 1993). These regions are defined by nodes, which are vertices, and surface/volume elements, which are polygons or polyhedrons depending on the dimensionality of the problem, and the value of the function on each region is given by basis functions which interpolate between the values at the vertices on the edge of the region. The weak form of the differential equation can then be represented as a system of linear equations. This system of equations can then be inverted to find an approximation of the unknown function parametrised by the basis functions on the mesh by applying relevant boundary conditions. In the case of optical modalities the differential equation of interest is the RTE. This equation treats light as a particle and neglects wave effects such as interference and refraction. It describes the change in energy in a volume element resulting from a radiative source, divergence and convergence, extinction via scattering or absorption, and scattering into the volume element. Unfortunately this equation is computationally expensive to solve, and so it is common to solve an approximation. These approximations are typically expansions of the RTE in terms of spherical harmonics, with high-order terms discarded (Klose and Larsen 2006; Necedal and Wright 1999).

The DA is the most commonly used approximation, and consists of the zeroth and first order terms of the expansion (Dehghani et al. 2008; Schweiger and Arridge 2014). It is only valid when the effect of scattering dominates that of absorption. Fortunately, this is generally the case for light in biological tissues with the exception of low scattering tissues such as cerebrospinal fluid (Custo et al. 2006). Higher order expansions can provide improved reconstruction accuracy (Lu et al. 2011).

MC stochastically simulates the paths of an ensemble of individual photons as they propagate through tissue (Boas et al. 2002; Wang et al. 1995). The effect of scattering and absorption on an individual photon is simulated by drawing random samples from probability distributions to determine how far a photon travels before interacting with tissue,

whether it is absorbed or scattered, and in which direction it is scattered. The ensemble's statistical properties converge on the distribution of light throughout the imaging subject as the ensemble size grows.

The third step in the process involves finding the bioluminescence distribution that is most consistent with the surface light measurements subject to any other information available. It is possible to calculate objectively how well a given bioluminescence distribution fits a set of measurements by using a model of light propagation to calculate the surface light and measurements resulting from the particular bioluminescence distribution. The difference between the simulated and experimental measurement sets then provides an indication of the goodness of fit of the chosen bioluminescence distribution. The goodness of fit is a function of the measurement discrepancies and the characteristics of the imaging system.

The bioluminescence measurement process can be represented in the following form in the simplest FEM case:

$$\mathbf{J}\mathbf{x} = \mathbf{y} + \mathbf{e} \tag{1.1}$$

where \mathbf{J} is a matrix that describes how the measurements change as the bioluminescence distribution changes, \mathbf{x} is a vector representing the bioluminescence distribution, \mathbf{y} is a vector representing the set of measurements, and \mathbf{e} is a vector representing an additive noise contribution. This equation is linear because the RTE ignores wave properties of light such as interference. There are effectively an infinite number of possible solutions as bioluminescence distributions are represented as n -dimensional real numbers where n is the number of vertices in the FEM mesh. This means that it is impractical to exhaustively search for the best solution. In addition \mathbf{J} may have a non-empty null space and there are an infinite number of solutions for \mathbf{x} which satisfy eq. (1.1) when the rank of the measurement matrix \mathbf{J} is smaller than the length of \mathbf{x} . This is the case when there are fewer measurements than the length of x , or when several different subsets of measurements provide the same information. Unfortunately, both of these conditions are typically true in BLT. It is then necessary to make use of additional information to choose

one of the solutions that satisfy eq. (1.1). It is common to adopt a Tikhonov regularisation scheme (Dehghani et al. 2008), which is equivalent to the assumption that each of the elements of the bioluminescence distribution are drawn from a Normal distribution in the Bayesian interpretation. This modified problem can be solved in one step. However the inclusion of other types of information can require the use of an iterative algorithm.

Additional information can take multiple forms and can improve the accuracy of any of the previous steps. Prior knowledge can inform measurement protocol design in order to obtain the maximum amount of information possible. An example of this would be using knowledge of the location at which bioluminescent cells were injected to extrapolate to a region in which the cells are likely to be found and to acquire measurements for that region. The quality of light propagation models can also be improved via additional information, such as the optical properties of different tissues and their spatial distribution, and the geometry of the imaging subject surface. Finally, the process of image reconstruction can be made more accurate and robust through the incorporation of prior information such as the expected density and bioluminescence efficiency of cells (Lu et al. 2011), and knowledge of the distribution of cell colony shapes.

This thesis aims to improve the process of BLT through the use of prior information to enhance each of the three steps of BLT imaging. It has four aims:

- Perform accurate surface geometry measurement with a large field of view, to enable the use of physical models of light propagation in tissue and in free space. This will enable quantitative imaging.
- Use prior information to enable principled bioluminescence reconstruction without a need for manual tuning of algorithmic parameters. This will enable objective reconstruction.
- Use prior information to produce measures of reconstruction uncertainty. This will enhance interpretation and analysis of reconstructed images.
- Use prior information to optimise the imaging processes and even the biomedical

study protocols themselves. This will allow imaging studies to obtain a maximum amount of information using a minimum number of subjects.

A surface capture method is developed to measure the geometry of the imaging subject to high accuracy and with large field-of-view in chapter 2, and this is used to improve the accuracy of physical modelling. In chapter 3 knowledge of the expected compactness of bioluminescence distributions is used to justify an application of Compressive Sensing (CS) techniques to the BLT reconstruction problem. The developed surface metrology technique and CS algorithm are then applied to the reconstruction of bioluminescent sources within a mouse phantom in chapter 4. In chapter 5 prior information about bioluminescence distributions is introduced and represented in the form of a prior probability distribution. This is combined with knowledge of the noise characteristics of the imaging system via Bayes' theorem to form a posterior distribution. The posterior distribution is used to perform reconstruction of bioluminescence distributions and to quantify the reliability of reconstruction in a spatially resolved manner. In chapter 6 the measurement process is enhanced through the use of the posterior distribution and information theoretic techniques to determine the measurement protocol that will statistically provide the most information, taking into account the properties of the measurement instrument and the expected properties of the imaged bioluminescence distributions. Information on the primary imaging system and phantom used in this work can also be found in Appendix A.

CHAPTER 2

SURFACE GEOMETRY MEASUREMENT USING STRUCTURED LIGHT

2.1 Contributions to this work

The main contributions to this work were made by the author and James Guggenheim. The author was the primary contributor in the design of the imaging process, and implementation of the algorithms to process images and extract spatial coordinates. The author also contributed to the calibration of the hardware, and to calibration of the mirrors in particular. James Guggenheim also contributed to the design of the system, and was the sole contributor in the integration of the hardware and software of the surface metrology subsystem into the existing instrument (see Appendix A), and the major contributor in the calibration of the hardware. The author acquired the data used in this chapter with the exception of mouse cadaver data figs. 2.10 and 2.11 which were acquired by James Guggenheim.

2.2 Introduction

Modelling light propagation accurately for a biological subject requires information about the shape of the subject surface in order to calculate the surface light resulting from an interior bioluminescence distribution. Accurate knowledge of the surface is important

because an increase in path length of 1mm can result in a halving of the measured signal if the attenuation coefficient is 1mm^{-1} , which is lower than most biological tissues (Cheong et al. 1990).

A particular geometry can be enforced on a system, for example by the immersion of a subject in a cuboid tank filled with a lipid emulsion (Culver et al. 2003; Koenig et al. 2007). In that case the resulting geometry is a cuboid regardless of the original geometry of the subject, but it may not be feasible to immerse a subject in such a tank of liquid. It may be impractical to do so because of the necessity of sealed breathing apparatus, and because light must then travel through the additional liquid medium. The use of a scattering and absorbing lipid solution is common due to the requirement of a scattering-dominated regime for DA techniques. This may result in significant attenuation and scattering of the light signal beyond that induced by passage through the tissue of the subject, with an associated loss of signal strength and information.

A number of systems use “mouse holders” to fix mice in position (Deliolanis et al. 2007; Wang et al. 2006). These may enforce a certain geometry by physically moulding the mouse into a desired shape, but these may not be physiologically neutral and so may confound the study being conducted. It would also be possible to map surface measurements onto a mouse-shaped phantom or atlas if the subject were fixed in a known position. A number of methods could be employed to accomplish this but the natural variation in mouse size may render this unreliable. If the above systems are inadequate then it is necessary to measure the subject geometry.

A multi-modal DOT and BLT preclinical imaging instrument was developed by James Guggenheim (Guggenheim et al. 2013a) (see Appendix A) which did not possess the capacity to measure surface geometry. It was necessary to design and implement a geometry measurement subsystem without interfering with the existing system functionality. That subsystem is presented in this chapter.

A number of techniques can be used to measure surface geometry. These can be partitioned into the two categories of contact and non-contact techniques. Contact techniques

use the position of a physical probe in conjunction with a contact sensor to measure a number of points on a surface. Contact techniques are used in industrial manufacturing but are unsuitable for the instrument of interest because of the limited room for probe articulation within the confines of the system, and because tissue is elastic. Subject elasticity implies that the probe itself may alter the surface geometry of the subject. These techniques are also expensive due to the sensor, motor, and control hardware requirements and can also require significant measurement time, particularly if the probe is moved slowly so as to avoid deforming the subject.

Surface metrology systems that use non-contact techniques commonly use light as a probe. Surface information can also be extracted from data from other non-contact modalities such as MRI and CT if this data is being acquired for other reasons, such as to obtain internal physiological information such as bone and organ boundaries (Guo et al. 2010; Tichauer et al. 2011). Modalities such as MRI and CT may not be cost effective to use primarily for surface measurement with the exception of specific sample types which are difficult to image using light, such as translucent water-containing gels and polished metal surfaces respectively. The set of light-based non-contact methods can then be further partitioned into two sub-categories consisting of methods that employ triangulation, and methods that measure time of flight (TOF).

Triangulation requires data corresponding to a single location measured from multiple unique viewpoints. The location of interest can be determined provided that the origin and direction of each of the light rays is known by calculating the location of their intersection.

TOF techniques do not use multiple viewpoints but require a light source in a fixed location with respect to a sensor which itself has high temporal resolution. The light source or the sensor must also be directional. TOF techniques then measure the total time between the light leaving the source, reflecting off the point of interest and reaching the sensor. The directionality of the sensor or light source coupled with the time of flight enables the location of the point of reflection to be calculated. TOF devices include LIDAR and TOF range cameras. LIDAR systems are suited to imaging much larger

objects than the intended subjects of this imaging system due to the need for high temporal resolution as light travels $1m$ in $3.3ns$. LIDAR systems of sufficient accuracy do exist, but are not practical for this application due to hardware costs. TOF range cameras are becoming more popular and cost effective, but have limited spatial resolution and also require special hardware (Dorrington et al. 2007; Payne et al. 2009).

There are many types of non-contact triangulation techniques. The most minimalist is called stereophotogrammetry and utilises two cameras which are calibrated so that the directions of the light rays associated with each of the camera pixels in a pinhole camera model are known, as are the cameras' positions and orientations relative to each other (Kraus 2007). Each camera in the fixed configuration acquires an image. The two images of the subject from different viewpoints are then analysed to determine points common to the two images. Each of these points is associated with two lines in space and calculation of the intersection of these lines gives the three dimensional coordinates of these points. Mirrors can be used to acquire multiple viewpoints using a single camera, or to acquire a 360° view of the subject (Kuthirummal and Nayar 2006). The conceptual simplicity of this technique is attractive, but processing pairs of images to determine common points is non-trivial and dependent on the degree of surface detail and colour. Surface uniformity can limit the number of common points that can be found. This issue is relevant because the XPM-2 phantom used in this work does not possess significant surface detail (see fig. A.5a). These problems are not insurmountable but time constraints and reliability requirements necessitated a different approach.

Some techniques use projected light to simplify the process of finding pixel correspondences. The simplest of these consists of a camera and a laser that is coupled to a mirror galvanometer. The laser projects a spot onto the mirror galvanometer which is reflected onto the subject, and the pixel coordinates of the laser light spot in the resulting camera image are estimated. The location of the spot is then triangulated using the location and orientation of the laser, the location and orientation of the mirror galvanometer, and the location, orientation and pixel direction of the camera. The mirror galvanometer then

changes orientation which alters the position of the laser spot. The process is repeated until a sufficient or maximum number of points are acquired.

Laser point systems have an advantage in that the camera sensor can be placed at a set angle relative to the imaging lens called the Scheimpflug angle (Häusler and Heckel 1988), which is dependent on the relationship between the laser location and orientation and the camera location, and ensures that the imaged laser spot is in focus on the camera sensor regardless of its distance from the camera. However, maintaining this perfect focus for a moving laser beam requires additional articulation to rotate the camera along with the laser beam, and there is a further disadvantage in that acquiring a complete data set would be time consuming as a single point is acquired per image. Finally, two axis mirror galvanometers are expensive.

A natural extension of this type of system involves the replacement of the laser point with a laser line. This both greatly increases the number of points acquired in a single camera image and simplifies the necessary articulation of the mirror galvanometer to a single axis. However, the method can still be slow relative to systems that project two dimensional illumination.

Two dimensional structured light projection enables faster measurement. Lasers in conjunction with masks and fringe gratings were commonly used to generate structured light patterns but the introduction of digital projectors using liquid crystal display (LCD) or digital micromirror device (DMD) technology has allowed more diverse sets of patterns and sets containing multiple types of patterns to be used for three dimensional imaging (Geng 2011; Salvi et al. 2004).

Similarly to laser line scanning, it is possible to label each imaged camera pixel with the pixel line on the digital projector from which the imaged light originated. The labelling can be performed using binary encoding, whereby projector pixels are placed into a sequence of “on” or “off” states, which equate to maximum and minimum projected intensity respectively. Labels are encoded temporally using sequences of patterns (Geng 2011; Salvi et al. 2004). For example, the number 7 may be represented by 00111 in a 5-

bit encoding, which would equate to the pixel being “off” for the first two images and being “on” for the last three images, reading from left to right. Binary encoding requires $\lceil \log_2 n \rceil$ images to uniquely label the light from a projector with n pixel lines, and offers an advantage over laser line scanning which requires n images to label the same number of lines. Additionally, this technique is relatively insensitive to sources of noise such as shot noise and fluctuations in background illumination due to the need to distinguish only between “on” and “off” states.

However, it is not possible to arrange the optical components so that the entire pattern is in focus regardless of the surface geometry unlike laser line scanning. This limitation can affect the depth of field of the technique. For example, the effect of defocus on a pattern of alternating “on” and “off” pixels can be so great that the individual lines cannot be resolved. This can be mitigated by using stripes with widths of multiple pixels so that the centres of the stripes can be correctly decoded even if the edges of the stripes are blurred, but it is not possible to distinguish between projector lines within a given stripe as a consequence.

Conversely it is possible to label all pixels in an image simultaneously by projecting a continuous ramp pattern. The imaged pixel intensity then directly indicates the projector pixel line, and the continuous nature of the ramp limits the effect of defocus. It is necessary to acquire two more images in practice to calibrate for uneven subject reflectance. In this case it is possible to assign a projector line to each pixel but this value is affected by sources of noise within the camera and projector, and changes in background illumination.

Sinusoidal projection patterns use the phase associated with sinusoids to label pixels (Gorthi and Rastogi 2010; Salvi et al. 2004; Su and Chen 2001). These patterns address the issues associated with binary encodings and ramp patterns but add an additional level of complexity. Sinusoidal patterns are more robust to projector defocus than binary patterns, and in fact a Gaussian point spread function will convert a square wave into a sinusoid (Lei and Zhang 2009). Defocus also has the potential to improve measurement resolution by smoothing the discrete intensity levels of adjacent pixels to create a continuous function

between them, which allows the pattern to be treated as continuous rather than a series of discrete blocks (Taylor 2012). Additionally, the effect of the noise component on the resulting phase value may be restricted to the 2π interval within a period instead of the whole range. The additional level of complexity that accompanies these benefits is a consequence of the periodicity of sinusoidal patterns. This periodicity prevents a unique mapping between measured pattern intensities and pixel lines for patterns containing multiple periods, as the phase extracted from the intensities is wrapped onto the $(-\pi, \pi]$ interval. It is necessary to unwrap the image of phases to recover the full range of values.

Phase can be measured using a variety of methods which differ in the required number and type of patterns and in the processing operations applied to the resulting images. The first of these is Fourier transform profilometry (Su and Chen 2001; Takeda and Mutoh 1983), which can calculate wrapped phase from a single image of a pattern in the case of a parallel optical axes geometry. This is accomplished by treating the phase change due to a subject as a perturbation in a carrier signal with the spatial frequency of the undeformed pattern at a specific distance from the camera. The part of the spectrum centred around the fundamental frequency can be filtered out from the harmonics in Fourier space. It is then possible to separate the phase from the uneven surface reflectance by taking the imaginary part of the natural logarithm of the distribution. However, there is an unavoidable loss of high spatial frequency information within the filtering process which results in reduced accuracy for small features. Also, the filtered spectrum is still contaminated by background illumination unless the images were taken in complete darkness. Background illumination can be removed before the filtering step using a second image of a projected pattern with the same spatial frequency but at a phase shift of π from the first. Further, the use of an image of the pattern projected onto a reference plane allows the removal of the fundamental frequency component leaving only the modulation resulting from the subject. However, the loss of high spatial frequencies during the filtering operation is still unavoidable.

Phase can be extracted without a loss of high spatial frequencies by processing each

pixel separately. This can be accomplished using phase shift profilometry (Berryman et al. 2003; Srinivasan et al. 1984), which requires 3 or more patterns. The n patterns contain sinusoids with the same spatial frequency but with phase shifts of $\frac{2n\pi}{N}$, where $n \in \{0, 1, \dots, N-2, N-1\}$. The phase can then be extracted by using an approach based on a Phasor representation. This can require more images than Fourier transform profilometry but can reconstruct phase for spatial features with an imaged size of only one pixel because all pixels are treated independently.

Let g_n be a measured pixel intensity under projection of pattern number n . This intensity takes the form:

$$g_n = B + \frac{A}{2} \left(1 + \cos \left(\phi + \frac{2\pi n}{N} \right) \right) \quad (2.1)$$

where A is the maximum intensity of the projected pattern, B is a background intensity, ϕ is the phase being measured, and $\frac{2\pi n}{N}$ is the phase offset. The phase ϕ can be calculated from $\{g_n\}$ using the following formula (Bruning et al. 1974):

$$\phi = -\arctan \left(\frac{\sum_{n=0}^{N-1} g_n \sin \left(\frac{2\pi n}{N} \right)}{\sum_{n=0}^{N-1} g_n \cos \left(\frac{2\pi n}{N} \right)} \right) \quad (2.2)$$

Calculation of ϕ via a computer language must use a version of the `atan2` function which takes two arguments (the numerator and the denominator), in order to calculate the correct value of ϕ uniquely over a $[0, 2\pi)$ or $[-\pi, \pi)$ interval. This results in the following implementation:

$$\phi = -\text{atan2} \left(\sum_{n=0}^{N-1} g_n \sin \left(\frac{2\pi n}{N} \right), \sum_{n=0}^{N-1} g_n \cos \left(\frac{2\pi n}{N} \right) \right) \quad (2.3)$$

An increase in the number of phase offsets reduces the error in the calculated phase ψ (see fig. 2.2).

A variety of techniques can attempt to unwrap an image of wrapped phase without any additional images beyond those used to measure wrapped phase (Judge and Bryanston-

Cross 1994; Zappa and Busca 2008). These techniques require a reference point of known phase and examine local phase changes to determine where phase wrapping events occur. Intuitively, these are likely to be where the wrapped phases of two adjacent pixels are approximately π and $-\pi$ and phase changes are small elsewhere in the neighbourhood, assuming that wrapped phase belongs to the interval $(-\pi, \pi]$. This relies on the surface in that neighbourhood being continuous and the unwrapped phase change between adjacent pixels being significantly smaller than 2π . A reference value is needed because knowledge of phase wrapping events allows calculation of relative phase only. The multiple of 2π contained in the reference value is propagated outwards from the reference pixel and modified by $\pm 2\pi$ when a phase wrapping event is detected. The surface must be continuous in order to perform this operation correctly, which limits the applicability of this technique. In addition, a reference phase value must result from either a fixed point in the system or the use of an additional operation to label or otherwise measure the unwrapped phase at a point.

It is important to note that one cannot distinguish with certainty between a phase change from π to $-\pi$ in adjacent pixels due to a phase wrapping event and an actual reduction in phase of 2π without additional information. Similarly, a phase change from 0 to δ for small δ may result from a phase wrapping event where the phase changes by $2\pi + \delta$. In addition, a phase change of $2\pi + \delta$ cannot be distinguished from a phase change of $2n\pi + \delta$ for an integer n . These techniques cannot guarantee correctness but sophisticated algorithms use quality measures or optimisation techniques to better classify phase wrapping and non-phase wrapping events in order to minimise these types of mistakes (Bioucas-Dias and Valadao 2007; Su and Chen 2004; Zappa and Busca 2008). This is important because the process of propagating phase offsets from regions of unwrapped phase to regions of wrapped phase also propagates any previous unwrapping errors. Fourier transform profilometry-based phase measurements may be more straightforward to unwrap via these techniques because noise is spatially smoothed by the Fourier filtering step, unlike in phase shift profilometry.

Additional information can be used to provide a unique, independent unwrapping for each pixel. This can be accomplished by labelling each period so that the correct phase offset of $2n\pi$ can be assigned (Salvi et al. 2004; Sansoni et al. 1999). Colour can be used to label each period in a pattern without necessitating the use of additional images (Geng 2011). However, the number of colours that can be successfully used is dependent on the accuracy of the camera and projector, and on the spectral and spatial reflectance of the subject. In practice, successful use of colour may necessitate the use of multiple patterns or exploiting the topology of projected patterns. For example, a pattern with 32 periods over the whole projected pattern requires 32 colours to encode all the periods in a single pattern. It is possible that only 8 colours can be resolved reliably given the noise in the projection and imaging processes. Projecting two images using those 8 colours enables the encoding of 56 values, which is sufficient. Alternatively, it is possible to project a single pattern composed of 8 colours such that each triple of adjacent colours occurs only once. Surfaces with discontinuities and small features may not be suitable for this approach because the topology of the projected pattern may not be preserved in the resulting image.

The periods of a pattern can also be labelled using binary encoding. Binary encoding can allow individual pixels to be unwrapped independently of neighbouring pixels, but defocus can affect accuracy particularly at the boundaries between periods. The effect of defocus can be partially mitigated through the projection of a second set of inverted binary patterns, so that the difference between the pattern and its inverted version for each bit is used to determine whether it is “on” or “off” (Scharstein and Szeliski 2003).

Finally, it is possible to project multiple spatial frequencies of sinusoidal patterns and use previously unwrapped phase at lower frequencies to inform the unwrapping of phase at higher frequencies (Kinell 2004; Li et al. 2005; Saldner and Huntley 1997; Wang et al. 2009). Projection of a pattern with one or less period across the pattern provides a starting point as the phase for the entire pattern falls within the interval $(-\pi, \pi]$. The measurement error in the phase relative to the phase range is high for low frequencies, but

the measurement can be used to estimate the unwrapping offset for a higher frequency. This offset is $2n\pi$ where n is an integer. The error in the offset can be exactly zero due to the discrete set of possible offsets. Correct choice of the offset results in no propagation of noise from the lower spatial frequency pattern to the next higher spatial frequency pattern. This unwrapping process can be repeated with successively higher frequencies until a desired level of phase accuracy or the resolution limit of the projector and camera is reached. Unwrapping using multiple spatial frequencies can require more images than using a binary encoding but is less affected by camera or projector defocus.

Spatial coordinates for the pixels in an image may be calculated using a map of measured phase on the camera CCD and knowledge of the properties of the camera and projector. This is simplest in the case of orthographic projection, which necessitates positioning the projector at an infinite distance from the subject or using a columnating lens such as a Fresnel lens (Lanman et al. 2009). This may be impractical and so it may be necessary to treat the projector as an inverse pinhole camera at a finite distance from the imaging subject or to use a more complex optical model.

Simplified models allow calculation of subject height relative to a reference plane from phase measurements by taking advantage of a specific spatial arrangement of the camera and projector (Berryman et al. 2003). This simplifies the required calculations but does not provide exact phase, limits the configuration of the system and requires the measurement of a reference plane.

More general models have been proposed. Wang et al. (2006) presented a model that allowed arbitrary placement of a camera for a system using sinusoidal projection, phase shifting, and measurement of a background (carrier) plane. The intrinsic and extrinsic optical parameters of the camera and projector were not measured, and coefficients of a three parameter model were estimated using a least squares fitting process instead. Their method is not efficient in this situation because the optical parameters of the camera are needed to enable multiview imaging through the mirrors and for other calibration processes such as calculating the free space transfer matrix (see Appendix A.5). Wang

et al. (2009) presented an extension of this using a ten parameter model.

Mao et al. (2007) presented a Fourier-transform method with a model that allowed the camera and projector to be placed arbitrarily, but required full visibility of the field of view and required background plane measurements. Measurement error was found to vary in simulation depending on the relative positions of the camera and projector.

Several groups have integrated mirrors into surface metrology imaging. Epstein et al. (2004) integrated mirrors into binary pattern-based geometry measurement, and used coloured landmarks placed on the imaging platform and the mirrors to estimate their locations. Lanman et al. (2009) took a different approach and used a Fresnel lens to perform orthographic projection. This was an elegant approach that enabled imaging and projection through mirrors without concerns about interference but required the mirrors to be placed in fixed orientations relative to the camera, whereas the mirrors in the multi-modality system must not be fixed.

Several existing systems use subject silhouettes from different angles to estimate surface geometries (Deliolanis et al. 2007; Jiang et al. 2012; Lasser et al. 2008; Meyer et al. 2007; Zhang et al. 2013) but these require a rotation stage to move either the subject or the camera and detector. These systems can be inaccurate when measuring surfaces containing concave regions but this inaccuracy may be limited in the case of mouse imaging as few concave regions exist (Lasser et al. 2008).

The multi-modal imaging system used in this work (see Appendix A) was capable of DOT and BLT imaging, but was not capable of measuring subject surface geometry. It was possible to circumvent this limitation for some objects of known shape. For example, a low resolution representation of the surface geometry was provided by the manufacturer in the case of the XPM-2 Phantom Mouse (see Appendix A.8) and it would be possible to map camera measurements onto this model using ray tracing. However this mapping would be subject to uncertainty in the knowledge of camera and projector parameters, mirror locations, and subject location, and would not be possible when imaging an arbitrary subject. This necessitated the design, construction and integration of a surface metrology

subsystem into the instrument, which is described in this chapter.

2.3 Survey of surface metrology systems used in optical tomographic imaging systems

A variety of surface metrology systems are either in use in existing optical tomographic imaging systems, or are used in conjunction with an existing optical tomographic imaging system. These systems may be dedicated to surface metrology or use data from other imaging modalities such as CT and MRI. A representative set of imaging systems are discussed in this section.

Chaudhari et al. (2005) presented a multi-view bioluminescence imaging system composed of an IVIS 200 two-dimensional imaging system, plexiglass stage, and four fixed mirrors which provided three additional views of the subject. These views were of the left and right sides of the subject viewed through a single mirror, and of the bottom side of the subject viewed through the plexiglass stage and two mirrors. This design is sensitive to bioluminescent sources on the underside of the animal due to the visibility of that side, and the use of mirrors potentially allows DOT illumination to be projected onto the underside of the subject as in the system described in Appendix A. However, a projector for DOT illumination would provide poor resolution for this application due to the field of view of the projector, the distance of the projector from the subject, and the size of the projection region.

Subject surface geometry was not measured via this system and was calculated from micro-CT imaging data acquired using a separate instrument. Silhouettes calculated from the geometry data were compared with optical images acquired using the bioluminescence imaging system to register the two data sets using affine transformations. These transformations were then used to map surface bioluminescence measurements onto the subject surface geometry. Reconstruction of the internal bioluminescence distributions was carried out using a hyperspectral forward model and demonstrated good localisation of an

implanted brain tumour.

The use of mirrors to perform four-view bioluminescence imaging provides sensitivity to bioluminescence sources in a variety of locations. However, the lack of an integrated optical surface metrology system necessitated the use of an additional micro-CT instrument and registration of the two data sets. This process potentially contains a number of sources of error.

- The use of two separate imaging systems necessitates the physical transport of the imaging subject and stage between the two systems. It is possible for the subject resting position to change during this process, which may introduce an uncorrectable discrepancy between the two data sets.
- The registration of micro-CT data to the optical data is subject to errors due to the focal plane of the optical imaging system. The four views are acquired in the same photograph but the optical path length is different for each view due to the presence of the mirrors. In particular the bottom view was visibly out of focus relative to the other views in a presented phantom image, which affects the reliability of the silhouette extraction needed to perform registration.
- The distance and angle of the subject with respect to the camera, and the position of the focal plane affect the imaging sensitivity (Guggenheim et al. 2013b). This is difficult to correct without knowledge of the positioning of the subject with respect to the system optical components.

The authors stated that the views visible through mirrors were scaled relative to the top view to account for the differences in path length, and that the focus of the bottom view can be improved by moving the stage vertically. This is appropriate for the presented block phantoms due to the planar nature of the phantom surfaces but is not possible for an arbitrarily shaped object. However, the authors showed good localisation of reconstructed bioluminescence despite these limitations demonstrating that good qualitative results may be achieved using two separate imaging systems.

Cao et al. (2010) presented a combined BLT and CT system composed of a CCD camera, three mirrors, an x-ray tube, and a flat panel x-ray detector. All components were mounted onto a gantry which rotated around the imaging subject. Two of the mirrors were used to provide additional views of the imaging subject. The integration of both modalities into a single system allowed image acquisition in a single coordinate system. This avoided the necessity of co-registering images taken from different modalities based on feature similarities.

Kumar et al. (2008) presented a combined surface metrology and FMT imaging system consisting of a camera, a stage, a laser coupled to an optical fibre and a two-axis translation stage, a mirror, and a commercial 3-D Facecam 100 3D camera. This camera provided measurement accuracy of 100 μ m. The mirror was used to redirect light towards the 3D camera, and had to be placed or removed when changing from FMT imaging to surface metrology imaging or vice versa, respectively. FMT measurements were mapped onto surface geometry via a registration process that involved finding an affine transformation for the surface geometry using fiducial markers and surface metrology measurements of the stage itself.

An advantage of this approach is that the commercial 3D camera provides good accuracy and can be integrated into the system without requiring expertise in surface metrology, and the measurement accuracy has been previously quantified by the manufacturer.

A disadvantage of this approach is the necessity of manipulating the mirror that redirects the light to the 3D camera, either manually or through the use of a translation stage which adds complexity and cost to the system. The optical measurements also have to be registered to the surface geometry, which is a source of error. This can be avoided if the optical properties, position and orientation of the camera used for optical measurements can be measured in the coordinate system of the 3D camera. Performing this calibration would add additional complexity to the system, particularly in the case where the position of the redirection mirror is not fixed.

Li et al. (2009) presented a combined surface metrology and FMT imaging system.

This instrument consisted of a camera, conical mirror, three line lasers on a linear stage for surface metrology measurement, and a scanning laser for fluorescence excitation. The mirror provided 360° surface coverage and the three line lasers enabled surface metrology of three sides of the mouse. The subject was placed on a transparent stage and the bottom side of the subject was assumed to be flat. The surface metrology subsystem illuminated the subject directly, whilst the fluorescence excitation subsystem illuminated the subject via reflection off the mirror.

The surface metrology subsystem provided width measurement to an accuracy 1mm when imaging an object of width 30mm. Spurious points resulting from stage reflections were manually removed. Fluorescence distributions were reconstructed for a block phantom using single and multi-wavelength data sets and demonstrated good localisation ability particularly in the multi-wavelength case. It was also demonstrated that a conical mirror increased the light collection efficiency compared to a planar mirror.

This system provided excellent surface coverage for fluorescence imaging, and good coverage for surface metrology imaging. The line lasers allowed a simple calculation of spatial location to reasonable accuracy. The error in subject width measurement may have resulted from errors in mirror calibration rather than a limitation of the technique. The use of the same camera for FMT imaging and surface metrology measurement avoided the need to register surface geometry onto the fluorescence images.

Surface geometry was calculated assuming that laser light is incident on the subject surface, incident on the mirror and then recorded by the camera. However it is possible for there to be multiple reflections potentially generating incorrect surface points. This is a particular risk due to the use of a conical mirror rather than a set of planar mirrors.

The surface metrology resolution and imaging time is also a function of the number of steps of the linear stage that are acquired. It is possible to increase resolution by increasing the number of steps but this is linearly related to the imaging time. The requirement of a linear stage could be removed through the use of a digital projector but the maximum surface metrology resolution would then be determined by the native resolution of the

projector.

Deliolanis et al. (2007); Lasser et al. (2008) presented a system that used subject silhouettes and back projection to reconstruction surface geometry. The system consisted of a camera, a rotation stage and a transparent plate onto which the subject was fixed, a photoluminescent acrylic plate for silhouette illumination, and a laser for fluorescence excitation. The photoluminescent plate was used as background illumination for silhouette imaging. A number of silhouettes were acquired at different rotation angles and a recursive volume carving algorithm was used to create an octree of voxels that were marked either as belonging to the subject or empty. Voxels were marked as belonging to the subject if they were contained within the silhouette for all of the rotation angles, and marked as empty if they were not contained within the silhouette for any of the rotation angles. Voxels were subdivided if neither condition was satisfied and the voxel size was greater than a resolution threshold.

This method is conceptually simple but requires a rotatable stage or imaging system. In the former case the subject must be firmly fixed in place to prevent subject movement and in the latter case a complex hardware configuration is required. An image is acquired at each measurement angle and this may require a large total imaging time. Also, this type of method cannot reconstruct certain types of concave surface and so consequently is vulnerable to geometry measurement errors.

Jiang et al. (2014); Jiang et al. (2012) utilised imaging techniques from light field imaging (Ng et al. 2005) to acquire surface geometry and luminescence measurements using a camera sensor fitted with an array of microlenses. Each of the microlenses split the incident light onto a small grid of pixels in a manner that provided angular resolution. Each of these grids contained an image of a region of the subject and the grids for adjacent microlenses contained similar information as their fields of view overlap. Photo-consistency measurements were used to identify common points in pixel grids which were used to reconstruct depth as in photogrammetry. Data collected from this sensor was used to reconstruct silhouettes using back projection, and using photo-consistency. The

results were positive and the photo-consistency-based technique produced volume error rates of 2%. An advantage of techniques that use photo-consistency over silhouettes is that they can reconstruct concave surfaces. However, photo-consistency measures require visible detail in images which requires non-uniformity in surface smoothness or colour.

Baum et al. (2012) presented a combined DOT and surface metrology system composed of a laser, projector and polariser, camera and polariser, and a rotating stage for the subject. Binary patterns that formed a Gray code labelling were projected onto the subject to enable surface geometry reconstruction. Polariser were placed in front of the projector and camera and used to remove photons incident on the camera that were not polarised in the same manner as those exiting the projector polariser. This reduced the contribution of photons that scattered multiple times through tissue because these scattering processes do not preserve polarisation.

A number of commercial optical imaging instruments either include an integrated surface metrology subsystem or allow that functionality to be added to a basic imaging instrument. Biospace Lab (Paris, France) produce an optical imaging system (Biospace Lab 2014c) which can be upgraded with a module that adds mirrors to enable multiple view optical imaging (Biospace Lab 2014b), and a pico-projector to enable surface metrology (Biospace Lab 2014a). PerkinElmer (Massachusetts, USA) produce an optical imaging system called IVIS 200 (PerkinElmer 2014) which includes a scanning laser galvanometer that projects structured light patterns for surface metrology. The system uses Fourier transform profilometry and quality-guided phase unwrapping. This system has been used for BLT imaging (Kuo et al. 2007). A number of patents have been filed on imaging instruments that use structured light for surface geometry measurement and mirrors and/or rotating components, such as Geng (2007); Nilson et al. (2007); Ntziachristos and Ripoll (2011); Rice et al. (2010); Stearns et al. (2010).

2.4 Methods

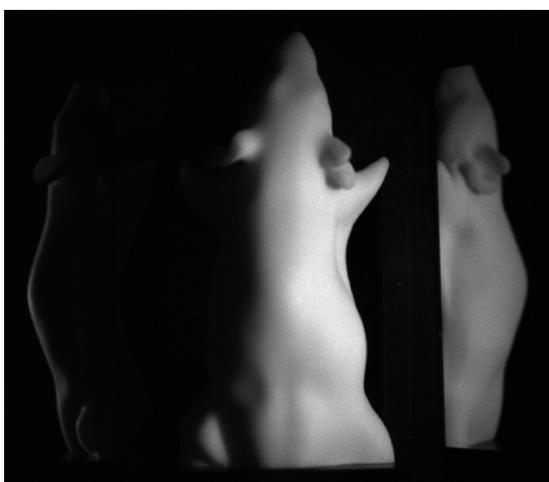
The imaging instrument required a surface metrology system that could be integrated into the frame of the instrument without affecting the existing components and imaging protocols. The initial functionality required of the surface metrology system was the ability to image a subject from a single viewpoint. This did not include the mirrors. Surface light measurements from the mirrors would be integrated later using ray tracing techniques to map measurements from the CCD to the mirrors and then to the surface previously created using the surface metrology data. The field of view of a single-view surface metrology system would not necessarily cover the entire field of view visible through the mirrors and so might prevent the use of some measurements. It became desirable to integrate the mirrors into the system if possible when this limitation became evident. It was also necessary to implement the subsystem with a limited budget. These conditions prevented the use of a contact measurement system or the purchase of an existing commercial system.

The initial design choice was between a stereophotogrammetry system or a structured light system. The use of a stereophotogrammetry system would necessitate the addition of at least one camera and light source. The task of finding pixel correspondences would be complex, in particular because of the homogeneity in the surface texture of the XPM-2 phantom. Consequently stereophotogrammetry was not explored further.

The surface metrology system design added two consumer-grade projectors to provide structured illumination and utilised the existing Hamamatsu CCD camera as the imaging device (see Appendix A). Two projectors were used to provide increased surface coverage. The use of the existing camera avoided the addition of a second camera and enables a direct mapping between camera pixels and points on the surface of the imaging subject. This map allows direct assignment of surface light measurements to surface locations for DOT and BLT, although the use of a free space light transfer model provides a superior mapping (Guggenheim et al. 2013b).

The process of projecting patterns and imaging was implemented using LabVIEW and MATLAB. Imaging is performed using the full CCD which produces images of 1024×1024

pixels, although images can be binned by the camera hardware or by post-processing. A short time delay was implemented between the command to project an image and opening the camera shutter to begin acquiring an image to compensate for latency in the display system. The fields of view of the projectors overlap and so the projectors are used sequentially resulting in two sets of images. A $10nm$ bandpass filter is used to attenuate the received signal to avoid image saturation. Figure 2.1 shows images of projections from projector 1 onto the XPM-2 phantom. The projector does not illuminate the entirety of the subject and so there is almost no signal through one of the mirrors. This necessitates the use of two projectors.



(a) XPM-2 phantom signal intensity imaged with projector 1



(b) XPM-2 phantom imaged with projector 1 projecting a sinusoidal pattern

Figure 2.1: The XPM-2 phantom imaged with the system camera under illumination from projector 1. The use of a single projector can be seen to produce limited coverage. Two mirrors are present in the images, although the second is almost invisible due to the limited projector coverage. Images from Basevi et al. (2013b).

The first iteration of the system used a Fourier transform profilometry technique to measure wrapped phase. The system used two pattern images with a π phase difference to acquire phase maps of the subject and of a reference plane. A local phase unwrapping algorithm was used, on the assumption that a part of the subject (specifically the paws) would be close enough to the reference plane that it would be possible to propagate a reference phase value from the platform onto the subject. The simplified cross axis model was used to extract height from the phase data.

Initial testing indicated that this design was unlikely to be sufficient. Surface discontinuities between the subject and the platform prevented the use of a reference phase point on the platform, and the phase unwrapping process was prone to error. This assessment prompted an iterative process to refine the system which resulted in change to every aspect of the design.

2.4.1 Wrapped phase measurement

The method used to measure wrapped phase was changed from using Fourier transform-based measurement to using phase shifting to preserve high spatial frequencies and so enable the resolution of small features. Phase shifting also provided two advantages relative to Fourier transform-based measurement that enabled the integration of mirror data into the surface metrology system. The addition of mirrors is equivalent to the creation of a virtual camera for each mirror which are reflections of the real cameras about the mirror planes. Firstly, a standard use of Fourier transform-based measurement would require the measurement of a reference plane for each of these virtual cameras that is orthogonal to the axis of the virtual camera rather than the real camera. This would require the creation and placement of additional physical planes to calibrate the method, and so would either require the mirrors to be fixed in place or require measurements of reference planes for each imaging subject. This is not required for phase shifting measurements. Secondly, the spatial frequency of the reference plane of each virtual camera may be different from that of the real camera. This would necessitate image segmentation of the parts of the image corresponding to the real camera and to the two virtual cameras and then applying separate Fourier filtering operations for each view. This is not required for phase shifting measurements.

A total number of 6 phase shifting patterns were used to measure phase at each spatial frequency. The global phase shift between adjacent patterns was $\frac{2\pi}{6}$. The minimum number of patterns necessary to measure phase in this manner is 3 but an increase in the number of patterns used reduces the error in the phase measurement. The estimated

dependence of phase measurement accuracy on the number of phase shift patterns and signal strength can be found in fig. 2.2. Measured wrapped phase lies in the interval $(-\pi, \pi]$. Figure 2.2 indicates that the number of phase shift images is important to measurement accuracy for very low signal strengths but the standard deviation of phase error is uniformly low for signal strengths of greater than 16 photons. In this case, low standard deviations were defined to be less than 5% of the phase range, which is equivalent to a standard deviation of below 0.3. The choice of 6 patterns is likely to be more than necessary given the 16-bit detector and high projector signal strength, but it is more important that the measurement process be robust than time-efficient. This is primarily because it is not possible to repeat a measurement due to the limited time during which a subject can be kept under sedation and the time requirements of the other imaging modalities.

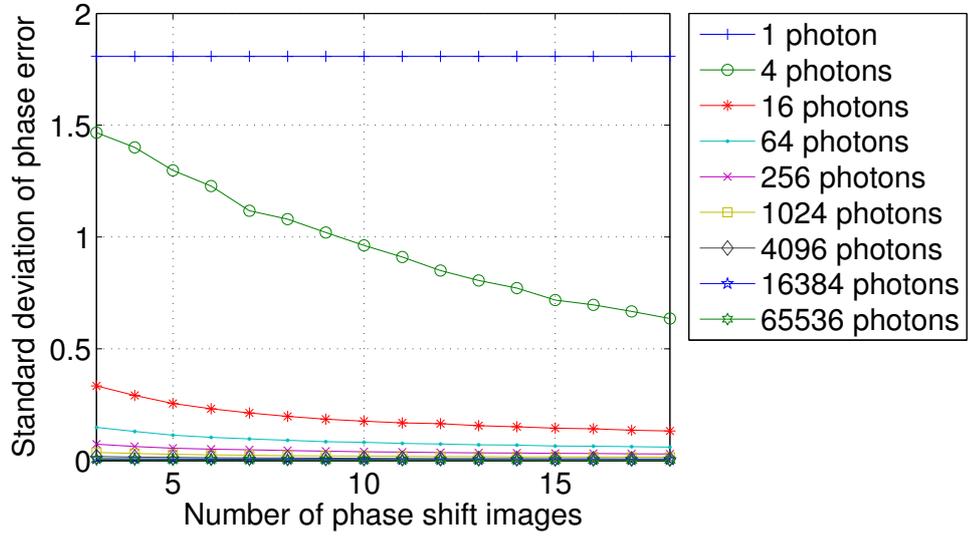
2.4.2 Phase unwrapping

The phase unwrapping method was changed to use multiple spatial frequencies instead of using local phase unwrapping because of the presence of discontinuities, particularly when integrating mirrors into the process. The chosen sequence of spatial frequencies has the form:

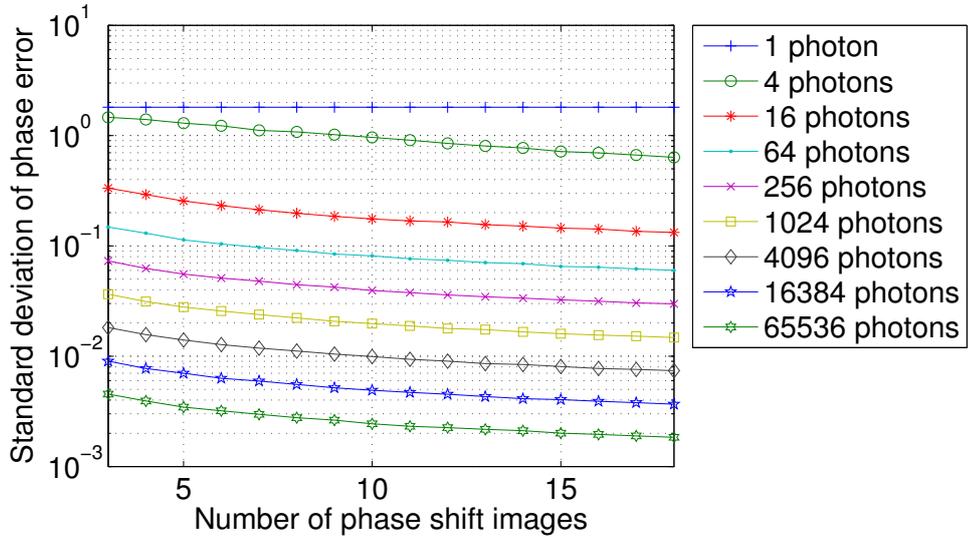
$$f_i = 2^{i/2} \text{ for } i \in \{0, 1, \dots, n-1, n\} \quad (2.4)$$

where 2^n is the spatial frequency in units of periods per pattern. A total of 14 frequencies are used resulting in 84 patterns and a maximum frequency of 90 periods per pattern. This results in a wavelength of 7 pixels when considering the horizontal resolution of the projector (640 pixels). The number of unwrapping images scales efficiently as projector resolution or maximum frequency increases because of the geometric progression of frequencies.

The phase unwrapping method begins with the frequency $f_0 < 1$ which does not require unwrapping, and extrapolates upwards. This extrapolation provides an estimate of the



(a) Phase measurement standard deviation dependence on signal strength and number of phase shifting images



(b) Phase measurement standard deviation dependence on signal strength and number of phase shifting images, using a log scale

Figure 2.2: The phase measurement dependence on signal strength and the number of phase shifting images in terms of measurement standard deviation, in radians. These dependences were estimated from measurement sets containing 10 000 samples assuming a perfectly sensitive detector and shot noise. Samples were obtained using MC sampling.

$2n\pi$ offset that best fits the measured phase. The method also tests for the possibility of sufficient noise in the phase measurement process to induce a phase wrapping event, which is especially probable at the extremes of each period. For example, noise in the measurements may be sufficient to cause a measurement of -0.9π to become -1.1π which is calculated as 0.9π as a result of the periodic nature of the cosine function. It is more appropriate to consider the measurement to be -1.1π than 0.9π if the extrapolated phase from lower frequency phase maps predicts a phase measurement of -0.8π . This results in a difference of 2π in the offset that is chosen. The measurement becomes more accurate at each frequency if it is assumed that the $2n\pi$ offsets have been chosen correctly at the previous frequencies because the noise in the measurement is bounded between the constant range $(-\pi, \pi]$ while the measurement range is increasing, as $\phi \in (-f\pi, f\pi]$.

2.4.3 Optical model to convert phase measurements into spatial coordinates

The crossed axes approximation was removed to improve the surface metrology subsystem and enable the incorporation of the mirrors into the surface metrology system. This action was taken because the crossed axes approximation was inaccurate and restricted the placement of the optical components in the system, including the camera and projectors. Moving to an inverse pinhole camera model for the projector improves performance due to increased model accuracy. Secondly, the general pinhole model allows arbitrarily placed cameras and projectors, and enables arbitrary virtual camera and thus mirror placement by extension. Thirdly, the pinhole model does not require reference plane measurement. The projector illumination is represented as a continuous phase field, which allows the use of projector light rays that do not originate from the centre of a projector pixel.

A schematic of a simplified version of the system can be found in fig. 2.3, which has been reduced to a single camera and projector. The camera and projector are drawn on-axis to simplify visualisation, but the following derivation is valid for arbitrarily placed cameras and projectors as all vectors and calculations are carried out in three dimensional

space.

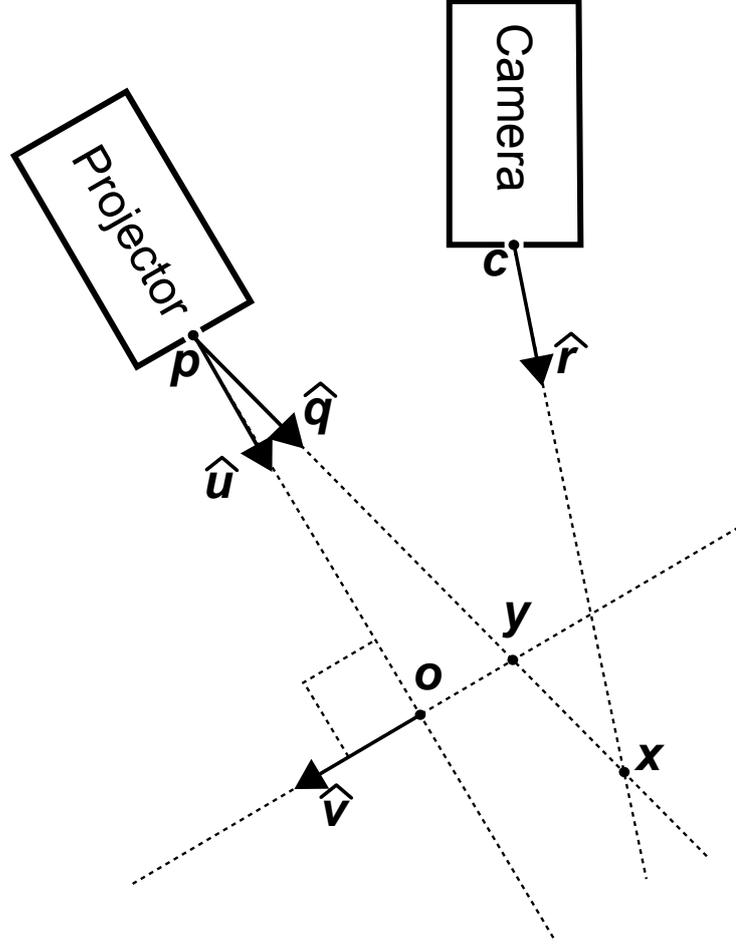


Figure 2.3: Diagram of a simplified surface metrology system. \mathbf{c} and \mathbf{p} are the locations of the camera and projector pupils respectively. Point \mathbf{o} is a point on a plane orthogonal to the axis of the projector, which is in direction $\hat{\mathbf{u}}$ on which the spatial frequency, and direction as encoded by $\hat{\mathbf{v}}$, of the projected pattern is known. The direction of the pattern is the direction in the plane of maximal increasing pattern phase. $\hat{\mathbf{r}}$ is a ray originating at the camera pupil, corresponding to a camera pixel. The unknown point \mathbf{x} lies somewhere along the line containing \mathbf{c} and spanned by $\hat{\mathbf{r}}$. \mathbf{y} is the intersection point of the plane containing \mathbf{o} and the line between \mathbf{x} and \mathbf{p} (spanned by $\hat{\mathbf{q}}$). For simplicity, this schematic is given in two dimensions, but this representation is also valid in three dimensions. Image and caption adapted from Basevi et al. (2013b).

The camera and projector pupils are at positions \mathbf{c} and \mathbf{p} respectively. The phase distribution is increasing in direction $\hat{\mathbf{v}}$ and constant in the direction orthogonal to $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$. Point \mathbf{o} lies on the axis of the projector and also on a plane orthogonal to the axis of the projector and parallel to the direction of increasing phase. The axis of the projector is in the direction $\hat{\mathbf{u}}$. The pattern has a given spatial frequency f within the plane defined

by point \mathbf{o} and normal direction $\hat{\mathbf{u}}$, such that the phase at a point \mathbf{y} within the plane is given by:

$$\phi(\mathbf{y}) = 2\pi f \hat{\mathbf{v}} \cdot (\mathbf{y} - \mathbf{o}) \quad (2.5)$$

Note that point \mathbf{y} does not necessarily need to lie on the plane defined by points \mathbf{p} , \mathbf{o} , and $\mathbf{o} + \hat{\mathbf{v}}$ for eq. (2.5) to be valid.

Point \mathbf{x} is a point in space with unknown coordinates which is imaged by the camera. The direction of the ray from the camera pupil \mathbf{c} to point \mathbf{x} is $\hat{\mathbf{r}}$. Direction $\hat{\mathbf{r}}$ can be calculated using the intrinsic and extrinsic parameters of the camera and the location of the pixel in pixel coordinates associated with that direction. Point \mathbf{y} lies on the line segment between \mathbf{p} and \mathbf{x} . The direction $\hat{\mathbf{q}}$ is:

$$\hat{\mathbf{q}} = \frac{1}{|\mathbf{x} - \mathbf{p}|}(\mathbf{x} - \mathbf{p}) \quad (2.6)$$

Point \mathbf{y} can then be represented as the intersection of the line defined by \mathbf{p} and $\hat{\mathbf{q}}$ and the plane defined by \mathbf{o} and $\mathbf{o} - \mathbf{p}$:

$$\mathbf{y} = \mathbf{p} + \frac{(\mathbf{o} - \mathbf{p}) \cdot (\mathbf{o} - \mathbf{p})}{\hat{\mathbf{q}} \cdot (\mathbf{o} - \mathbf{p})} \hat{\mathbf{q}} \quad (2.7)$$

If this is represented using $\hat{\mathbf{u}}$ then eq. (2.7) becomes:

$$\mathbf{y} = \mathbf{p} + \frac{|\mathbf{o} - \mathbf{p}|}{\hat{\mathbf{q}} \cdot \hat{\mathbf{u}}} \hat{\mathbf{q}} \quad (2.8)$$

Equation (2.8) can then be substituted into eq. (2.5):

$$\phi = 2\pi f \hat{\mathbf{v}} \cdot \left(\mathbf{p} + \frac{|\mathbf{o} - \mathbf{p}|}{\hat{\mathbf{q}} \cdot \hat{\mathbf{u}}} \hat{\mathbf{q}} - \mathbf{o} \right) \quad (2.9)$$

But $\hat{\mathbf{v}}$ is orthogonal to $\mathbf{p} - \mathbf{o}$, so:

$$\phi = \frac{2\pi f |\mathbf{o} - \mathbf{p}|}{\hat{\mathbf{q}} \cdot \hat{\mathbf{u}}} \hat{\mathbf{v}} \cdot \hat{\mathbf{q}} \quad (2.10)$$

Further substituting eq. (2.6) into eq. (2.9) produces:

$$\phi = 2\pi f \hat{\mathbf{v}} \cdot \left(\mathbf{p} + \frac{|\mathbf{o} - \mathbf{p}|}{(\mathbf{x} - \mathbf{p}) \cdot \hat{\mathbf{u}}} (\mathbf{x} - \mathbf{p}) \right) \quad (2.11)$$

Rearranging for \mathbf{x} results in:

$$\left(\frac{\phi}{2\pi f} \hat{\mathbf{u}} - |\mathbf{o} - \mathbf{p}| \hat{\mathbf{v}} \right) \cdot \mathbf{x} = \left(\frac{\phi}{2\pi f} \hat{\mathbf{u}} - |\mathbf{o} - \mathbf{p}| \hat{\mathbf{v}} \right) \cdot \mathbf{p} \quad (2.12)$$

If \mathbf{k} is defined:

$$\mathbf{k} = \frac{\phi}{2\pi f} \hat{\mathbf{u}} - |\mathbf{o} - \mathbf{p}| \hat{\mathbf{v}} \quad (2.13)$$

and \mathbf{x} is parametrised in terms of \mathbf{c} , α and $\hat{\mathbf{r}}$:

$$\mathbf{x} = \mathbf{c} + \alpha \hat{\mathbf{r}}$$

then eq. (2.12) can be rewritten as:

$$\mathbf{k} \cdot (\mathbf{c} + \alpha \hat{\mathbf{r}}) = \mathbf{k} \cdot \mathbf{p}$$

Solving for α produces:

$$\alpha = \frac{\mathbf{k} \cdot (\mathbf{p} - \mathbf{c})}{\mathbf{k} \cdot \hat{\mathbf{r}}}$$

Which implies that the equation for the unknown point x is:

$$\mathbf{x} = \mathbf{c} + \frac{\mathbf{k} \cdot (\mathbf{p} - \mathbf{c})}{\mathbf{k} \cdot \hat{\mathbf{r}}} \hat{\mathbf{r}} \quad (2.14)$$

Equations (2.13) and (2.14) allow calculation of the spatial coordinates corresponding to a phase value given:

- The position of the camera pupil, \mathbf{c} .
- The direction of the camera pixel vector, $\hat{\mathbf{r}}$.

- The position of the projector pupil, \mathbf{p} .
- The orientation of the projector, $\hat{\mathbf{u}}$.
- The direction of increasing pattern phase, $\hat{\mathbf{v}}$.
- The spatial frequency of the pattern within a defined plane, f .
- The distance of the defined plane from the projector, $|\mathbf{o} - \mathbf{p}|$.

This formula is valid for any configuration of camera and projector, and is efficient to calculate because the calculation can be vectorised and only uses the absolute value, addition, subtraction, multiplication and division operations.

A cloud of points in space is created by applying eq. (2.14) to every point which is both directly visible to the camera and is reflecting sufficient light from the projector. This formula can also be used to calculate the positions of surface points observed through the mirrors by reflecting the \mathbf{c} and $\hat{\mathbf{r}}$ vectors appropriately about the plane of the mirror.

2.4.4 Mirror integration

Two 45° mirrors were already used to extend the surface coverage of the surface metrology system for DOT and BLT imaging to improve imaging of deep sources. The system was extended to utilise the mirrors for surface geometry measurement, which expanded the coverage of the surface metrology system to match that of the DOT and BLT components and provide a direct mapping between CCD pixels associated with the mirrors and the surface of the subject, to support DOT and BLT imaging.

Shape and size may vary between subjects, as may resting position upon the platform. The mirrors are not fixed to the platform so that they can be moved to adapt for subject variation in order to maximise the quality and coverage of imaging for each subject. Use of the mirrors for surface geometry measurement involves reflecting the camera about the plane of the mirror to create a virtual camera, but the position of the mirrors is not

known *a priori*. It was necessary to incorporate a mirror calibration method to acquire this information.

The subject is placed in the system at the beginning of the imaging procedure and then the mirrors are placed on opposite sides of the subject torso, as close to the subject as possible in order to maximise light collection. After this point neither the subject nor the mirrors can be moved as the component arrangement must remain identical for surface metrology, DOT, and BLT imaging. The platform and the walls of the instrument are black in order to minimise reflected light within the system. This greatly reduces the visibility of the mirrors to the camera above. Consequently, manual identification of the mirrors through selection of image landmarks is difficult and unreliable. An extension of the surface metrology system was developed to allow semi-automatic calibration of the location and orientation of the mirrors.

Dual photography is a technique that enables the roles of camera and projector to be computationally reversed. Specifically, given a scene illuminated by a projector and imaged by a camera, it is possible to calculate the image of that scene if the projector was a camera, and the camera was a projector. This allows both a change of imaging perspective and illumination. This is made possible through a mapping between camera and projector pixels. Camera pixels map to projector pixels, unlike the surface metrology system, where camera pixels map to lines of projector pixels.

It is possible to create such a mapping by using the surface metrology system to measure two phase maps resulting from different $\hat{\mathbf{v}}$ vectors. In the simplest case, this may consist of a vertical and a horizontal phase map. These phase maps act as a two dimensional coordinate system mapping camera pixels to projector pixels. Figure 2.4 shows such a map of a scene composed of a set of lego blocks marked with a pen and two mirrors, acquired using the imaging system. Figures 2.4a and 2.4b show the projected phase in orthogonal directions. The phase range is greater than 2π and is transformed into a set of sinusoidal projection images for phase measurement. Figures 2.4c and 2.4d shows measured phase in the orthogonal directions, which provides a many-to-one mapping

between the camera pixels and the projector pixels. Figure 2.4e shows a camera image of the scene under uniform projection. This image can be transformed onto the projector pixels using the previously acquired mapping, which results in an image of the scene from the projector's perspective, as can be seen in fig. 2.4f. It can be seen that the information from the direct view and from the view through the mirror combines to form a single view of the object being imaged.

Each camera pixel maps to at most one projector pixel, but this is not true in reverse. The addition of mirrors to the imaging system allows the same region to be imaged from two different perspectives and as a consequence one projector pixel may map to two camera pixels. However, both camera pixels correspond to the same point in space and this fact can be used to estimate the mirror positions. Projecting orthogonal phase maps doubles the imaging time and results in a phase coordinate system that maps camera pixels to projector pixels. The phase maps can be searched to find pairs of pixels that share a common phase coordinate. The image is divided through manual selection of two lines into three disjoint regions consisting of:

- Pixels that image the subject directly.
- Pixels that image the subject through mirror 1.
- Pixels that image the subject through mirror 2.

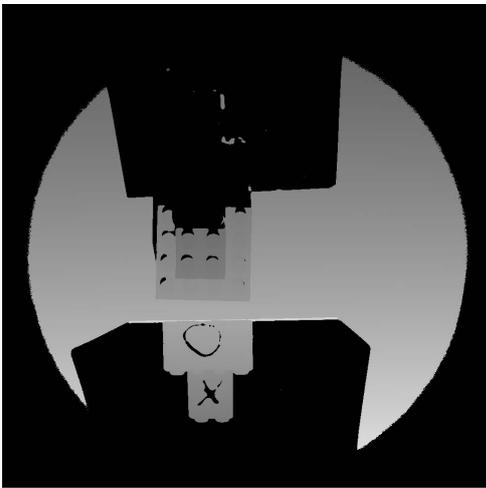
Each of these regions may contain pixels that do not image the subject but this does not affect the analysis as they can be excluded using a signal strength criterion. A subset of the pairs of pixels selected earlier will contain one pixel from the direct region and one pixel from the mirror 1 region. The spatial coordinates for the pixels from the direct region may be calculated immediately and these provide coordinates that must be intersected by the rays from the camera to the pixels in the mirror 1 region. The mirror location and orientation can then be optimised automatically to minimise the minimum distance between these rays and the calculated coordinates.



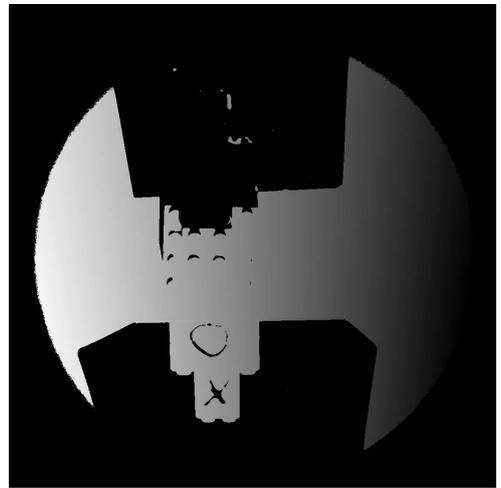
(a) Projector phase in first dimension



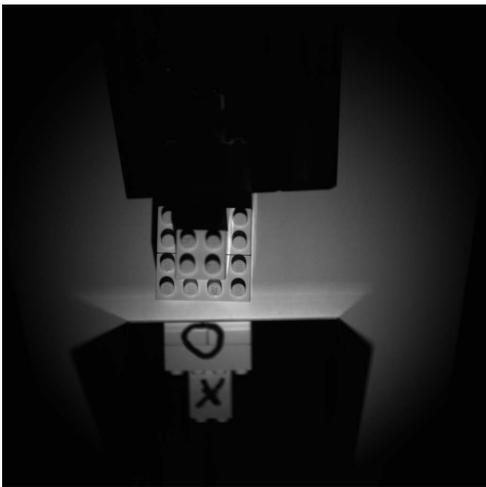
(b) Projector phase in second dimension



(c) Measured phase in first dimension



(d) Measured phase in second dimension



(e) Measured signal



(f) Measured signal transformed as if the roles of the camera and projector were reversed

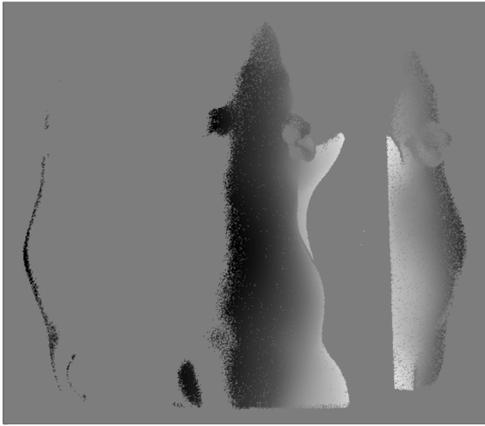
Figure 2.4: Use of two phase maps to map between projector pixels and camera pixels. Figure 2.4f is the result of this mapping being applied to fig. 2.4e.

The mirror is defined in terms of an arbitrary location on the mirror and the normal vector to the mirror. This definition specifies an infinite plane. The true mirror has finite extent, but the edges of the mirror may not be visible which prevents the optimisation of a plane with finite extent. This necessitates the manual segmentation of pixels into those belonging to the direct image and the mirror image as the mirror is modelled as an infinite plane during the optimisation process. The optimisation task is accomplished by randomly choosing an initial mirror location and orientation and using the `fminunc` function in MATLAB to minimise an objective function which calculates the sum of the minimum distances of the camera rays from the calculated coordinates. The objective function is evaluated by simulating the imaging process and reflecting the camera rays using the input mirror configuration. The `fminunc` function performs an unconstrained optimisation process using a quasi-Newton algorithm. This process is repeated multiple times using different random initial mirror configurations and the optimal result is chosen by using the objective function as a metric.

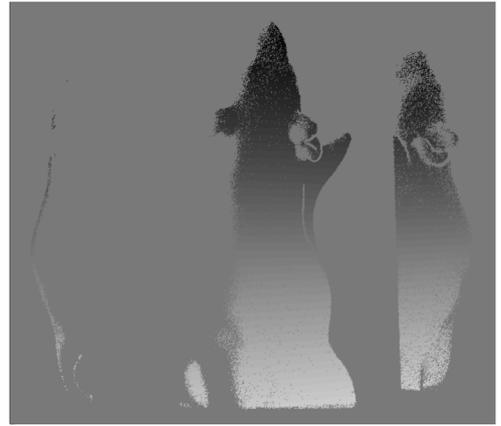
This method performs best when there are a large number of correspondences between the direct and mirror images. This was the case for all of the subjects imaged during this work due to the placement of the camera and projectors, the size of the mirrors, and the subject sizes and shapes.

Figure 2.5 shows an example of imaging using orthogonal phase maps, and of correspondences between pixels that image the surface directly and pixels that image the surface indirectly via a mirror. The phase map in fig. 2.5a is the phase map that is normally used in surface geometry measurement. The phase map in fig. 2.5b is orthogonal to the phase map in fig. 2.5a and the two maps together provide a unique labelling for a set of points on the subject surface. The second phase map (fig. 2.5b) can also be used to measure surface geometry if desired. A selection of the resulting correspondences is present in fig. 2.5c. The selection displayed has been kept small to aid visual interpretation but 500 correspondences are used to calculate mirror locations and orientations.

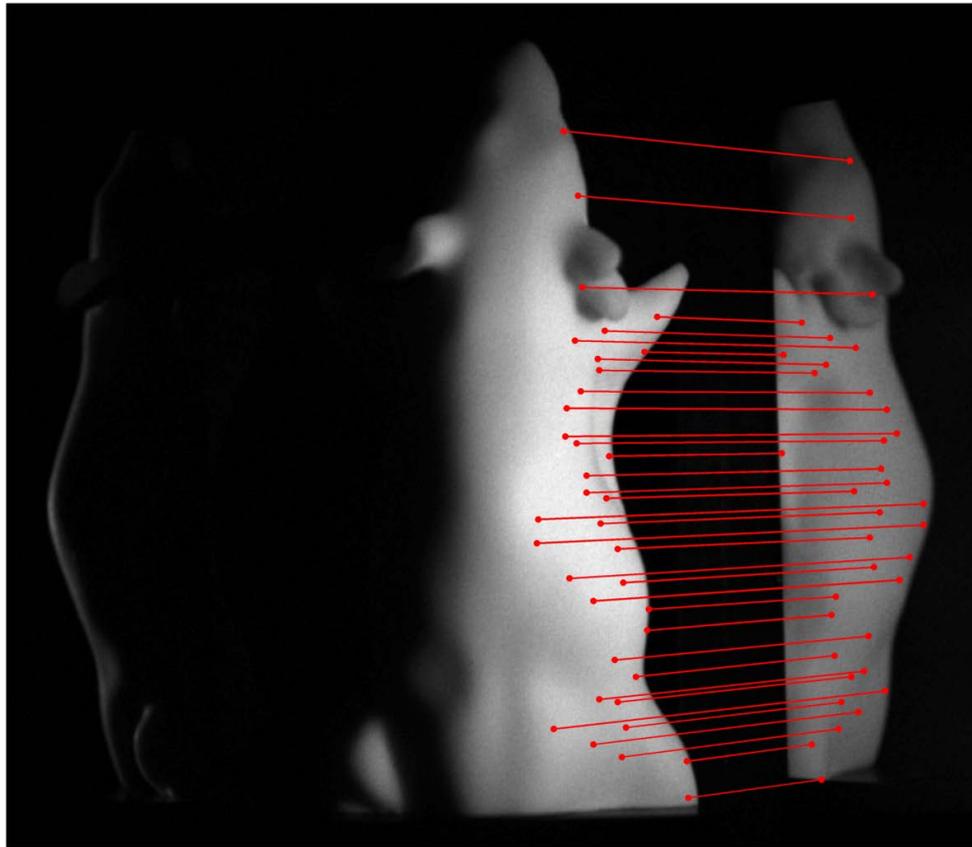
The parameters of virtual cameras are calculated using the optimised mirror configu-



(a) Measured phase map using horizontal phase



(b) Measured phase map using vertical phase



(c) Correspondences between direct and mirror pixels calculated from phase maps

Figure 2.5: A selection of correspondences between pixels imaged directly and pixels imaged through a mirror for the XPM-2 phantom. Two orthogonal phase maps were used to calculate these correspondences. These correspondences can then be used to calculate the location and orientation of the mirror.

rations. These virtual cameras are then used to convert phase for mirror pixels into spatial coordinates. The locations and orientations of all components are defined in a common coordinate system and so the calculated spatial coordinates for pixels in all three views are calculated in the same coordinate system and do not need post-processing such as registration.

The positions of the mirrors in conjunction with the calibration data of the projectors can be used to determine which pixels of the projectors may be projected onto the mirrors. Light from a projector reflecting off a mirror may be incident on the imaging subject. This would create a spurious signal as the optical model assumes that projector light travels in a straight line until it is incident on the imaging subject. This is avoided by using the calibration data to calculate which projector pixels may be projected onto the mirrors and setting the intensity of those pixels to zero in all of the projected patterns.

2.4.5 Noise removal

The previous sections have addressed the tasks involved in converting projector signals into phase measurements and then spatial coordinates, but all camera measurements contain a degree of noise, which may dominate for some pixels. A proportion of pixels also do not measure any direct projector illumination as the projection does not cover the entire field of view of the camera. Signal across the field of view of the camera is dependent on:

- Whether the projector illumination directly reaches the surface that is imaged by the pixel.
- The distance of the surface from the projector and camera.
- The angle of the surface to the projector and camera.
- The spectral absorption properties of the surface.

- The spectral reflectance properties of the surface, and the ratio of specular to diffuse reflection.

In addition, the signal may be contaminated by:

- Light from sources other than the projector, which may fluctuate in intensity over time.
- Movement of the surface during the exposure period.
- Light from the projector that has been reflected multiple times due to reflection off the instrument itself or due to a concave surface.
- Light from the projector that has travelled into the subject tissue, reflected internally, and then emerged from the surface.
- Defocus in the camera and/or the projector with respect to the surface.
- Saturation of the camera CCD due to excessive exposure.

Signal contamination can affect the accuracy of phase measurement, and light that has undergone multiple reflections either internally or externally to the surface can create spurious signals. It is necessary to quantify noise and exclude pixels with unsatisfactory measurements from the final point cloud.

There are two primary quantities that can be used to distinguish pixels with acceptable and unacceptable levels of noise, which are the signal strength and the signal fidelity. The simplest measure of signal strength is the increase in the pixel intensity when all of the pixels in the projector are switched from an “off” state to an “on” state. SNR can also be calculated based on these values if desired. However, alternative measures of signal quality are useful because the projections consist of structured light. The image of the projected light is affected by camera and projector defocus, and high spatial frequency structures are more strongly influenced by defocus than are low spatial frequency structures. Surface orientation relative to the camera and projector may also result in an aliasing effect.

The phase shifted pattern images can also be used to calculate the intensity of the cosine wave projected for each pixel in the same manner as the phase of the cosine wave can be calculated. The equation to calculate the signal is:

$$A = \frac{4}{N} \sqrt{\left(\sum_{n=0}^{N-1} g_n \sin \left(\frac{2\pi n}{N} \right) \right)^2 + \left(\sum_{n=0}^{N-1} g_n \cos \left(\frac{2\pi n}{N} \right) \right)^2} \quad (2.15)$$

A measure of unwrapping error can also be generated using the differences at each frequency between the estimated phase extrapolated from the next lower frequency phase map and the calculated phase. The absolute values of all the phase unwrapping errors can be summed to generate a metric of unwrapping accuracy.

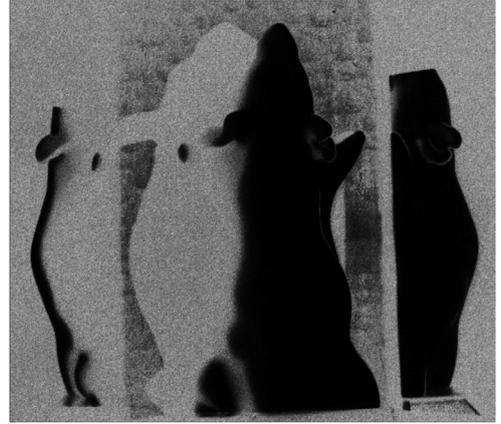
Figure 2.6 shows the structured light signal and unwrapping error from imaging of the XPM-2 phantom. Figure 2.6a shows the structured light signal for the phantom. The supporting platform cannot be seen and some features at which there are discontinuities, such as the front leg as viewed through the mirror, do not show a great dip in signal. Figure 2.6b shows the metric of unwrapping error for the XPM-2 phantom. Different features are visible here such as regions of the background platform and the spatial variations in signal intensity are not as pronounced. The application of a threshold of $0.075(N_f - 1)$ to the unwrapping error produces the mask in fig. 2.6c, where N_f is the number of frequencies. The threshold used is conservative and is equivalent to 2.4% of the maximum phase measurement error at each frequency but data quality is prioritised over data quantity. The resulting mask excludes some pixels that are likely of acceptable quality but importantly excludes regions around surface discontinuities such as the right ear, right ribcage, and the region above right leg. Segmentation using the structured light signal would not isolate these features.

2.4.6 Conversion from point clouds to meshes

The surface metrology system produces one or more clouds of points in three dimensional space. Each point represents the coordinates of a single pixel, or a group of pixels if the



(a) XPM-2 phantom structured light signal intensity imaged with projector 1



(b) XPM-2 phantom unwrapping error imaged with projector 1



(c) XPM-2 phantom unwrapping error mask imaged with projector 1

Figure 2.6: Noise removal for an XPM-2 phantom imaged with the system camera under illumination from projector 1. Dark pixels in fig. 2.6b represent pixels with low phase unwrapping error. Pixels in white in fig. 2.6c represent pixels that are considered of sufficiently high quality to use in mesh generation.

camera sensor is binned to reduce camera integration times at the expense of resolution. The point clouds must be converted into a surface mesh in order to visualise the data set as a single object. The point clouds must be converted into a volume mesh in order to use FEM techniques to model light propagation.

Surface meshes consist of vertices and faces. The vertices are points in three dimensional space that may or may not be the points from a point cloud data set. The faces are polygons defined by a number of vertices that lie on the edges of the polygons, and define regions of surface. The simplest face type is defined by three vertices and takes the shape of a triangle. The next simplest is defined by four vertices and takes the shape of a rectangle. Faces typically possess straight edges but edges can be curved in the case where an edge is defined in terms of three or more vertices. Triangular faces were used through this work.

Volume meshes consist of vertices, faces, and elements. Vertices and faces take the same form as those in surface meshes. Elements consist of a collection of faces that enclose a region of volume. The simplest of these are tetrahedral elements, which are defined by four triangular faces and four vertices. Tetrahedral elements were used throughout this work.

The coin data sets in this chapter (fig. 2.12) consist of a single view and were converted into surface meshes by taking advantage of the topology of the CCD sensor using the `delaunay` method in MATLAB. This function creates a set of triangular faces that form a convex hull over a set of points that lie in a two dimensional space. The points in the single-view point set are defined in three dimensional space, but the imaging process projects them onto the two-dimensional coordinate system of the camera sensor.

Meshlab was used to create surface meshes for the remaining data sets which consist of multiple views. Here `delaunay` cannot be used to create a surface mesh as point clouds may overlap and represent different regions of the subject surface. Meshlab contains a `Surface Reconstruction: Poisson` filter which was used to create the surface meshes (Kazhdan et al. 2006). The Poisson surface reconstruction algorithm requires surface nor-

mal vectors, which can be estimated by creating surface meshes for the point clouds associated with individual views and calculating the associated normal vectors. Alternatively, the Meshlab filter `Compute normals for point sets` can be used. The Poisson surface reconstruction algorithm computes a function which takes the value of 1 inside of the mesh and 0 outside, for which an isosurface of the algorithm forms the mesh surface. The normal vectors associated with the points are treated as samples of the gradient of this function, from which the function itself can be estimated.

2.4.7 System calibration

The process of calibration of the camera and projector is semi-automatic and takes advantage of the presence of a motorised lab jack to image a fixed pattern placed orthogonally to the camera axis at a number of distances from it. In the case of the camera calibration this pattern consists of a grid of black dots on a white background. The dot spacing and heights of the lab jack are known, and the dot spacing as measured in camera pixels increases as the distance between the camera and the pattern decreases. It is possible to calculate the height of the camera pupil from the platform by extrapolating the decrease in the physical size of the region imaged by each pixel to find the point at which it becomes zero.

It is then possible to calibrate for the projector positions and orientations in a similar manner given the previous camera calibration. This process also involves patterns of dots on a platform at a number of heights, but in this case the dots are projected by the projector. Dots are projected individually at a number of heights to allow automatic segmentation and labelling of dot position. The camera calibration also allows a three dimensional spatial position to be assigned to each dot. It is possible to fit for projector pupil location and orientation given these positions and a knowledge of the locations on the projector on which these originated.

2.4.8 Experiments

The system was validated quantitatively *in simulacra* via the use of an XPM-2 Phantom Mouse (Appendix A.8). This phantom is an object with known surface geometry, and optical properties similar to tissue. The phantom was imaged using the surface metrology system and all components were used. Consequently, two mirrors were placed to enable acquisition of three views simultaneously, and two projectors were used sequentially resulting in a total of six views. These six views were integrated into a mesh using Meshlab's `Poisson surface reconstruction` filter (Kazhdan et al. 2006). The set of manufacturer-provided XPM-2 surface points was manually reduced to 878 points which were expected to lie on the surface of the recovered mesh. The manufacturer-provided surface points were then rigidly registered to the recovered mesh and the minimum distance between each of the points and the recovered mesh was calculated.

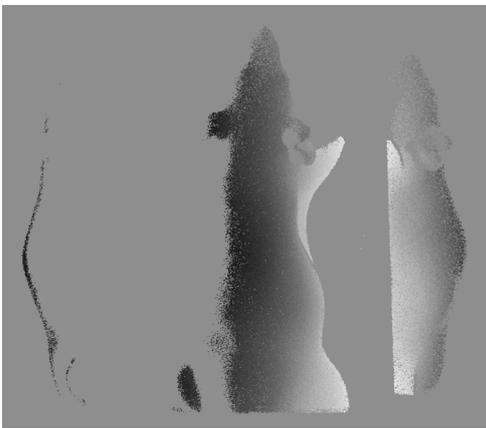
Secondly, a mouse cadaver was imaged using the system to demonstrate that it can successfully be used to image tissue. The imaging protocol was identical to that used to image the XPM-2 phantom.

Finally, the system was used to image the front and reverse sides of a two pence coin. The edge of the coin were shiny and reflected sufficient projector light to saturate the camera. A black whiteboard marker was applied to the edge to reduce reflection to an acceptable level.

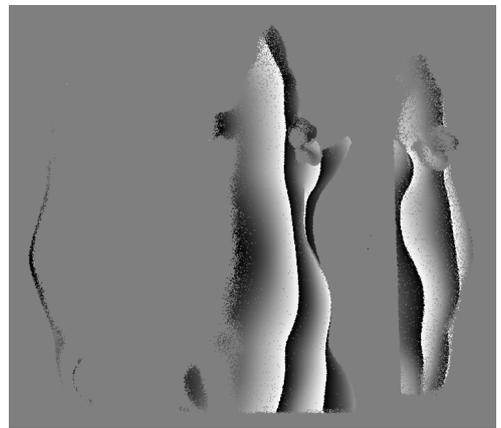
2.5 Results and discussion

The XPM-2 phantom was imaged using the surface metrology system. Images of the phantom under illumination from projector 1 can be found in fig. 2.1, and phase maps for projector 1 can be found in fig. 2.7. Pixels within the phase maps with excessive noise have been set to zero. It can be seen that this also removes pixels at surface discontinuities from the perspective of the camera, as these pixels image signal from multiple surfaces. Figures 2.1 and 2.7 demonstrate that a single projector does not provide sufficient coverage

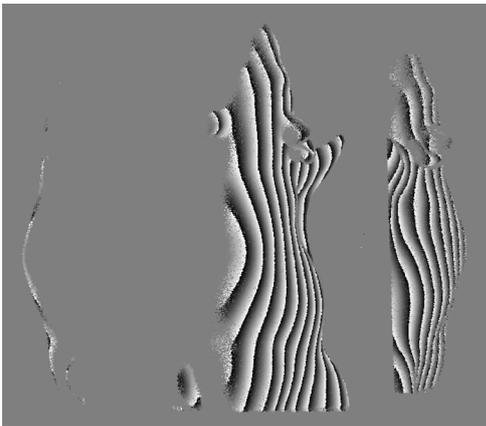
to measure the entire exposed surface of the subject, and the high resolution projections (fig. 2.1b) and wrapped phase maps (fig. 2.7c) indicate that spatial phase unwrapping would be difficult for this example as separate fringes merge in the images due to camera resolution and the angle of view. The low frequency phase map (fig. 2.7a) has not been unwrapped but does not contain any discontinuities as the unwrapped phase range is within the $[-\pi, \pi)$ interval. The high frequency phase map unwrapped using information from lower frequencies (fig. 2.7d) is smooth and appears similar to the low frequency phase map, but possess a higher SNR.



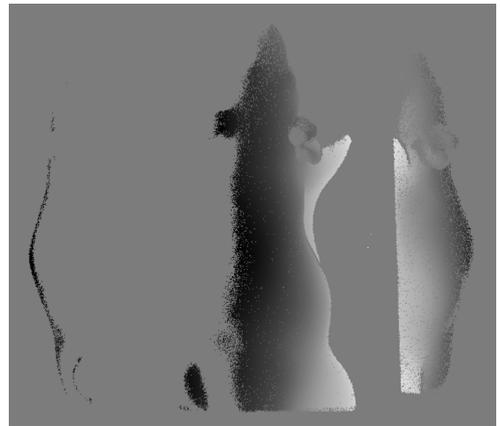
(a) Low frequency phase map, wrapped



(b) Medium frequency phase map, wrapped



(c) High frequency phase map, wrapped



(d) High frequency phase map, unwrapped

Figure 2.7: Wrapped phase maps of the XPM-2 phantom, and an unwrapped high frequency phase map. Images from Basevi et al. (2013b).

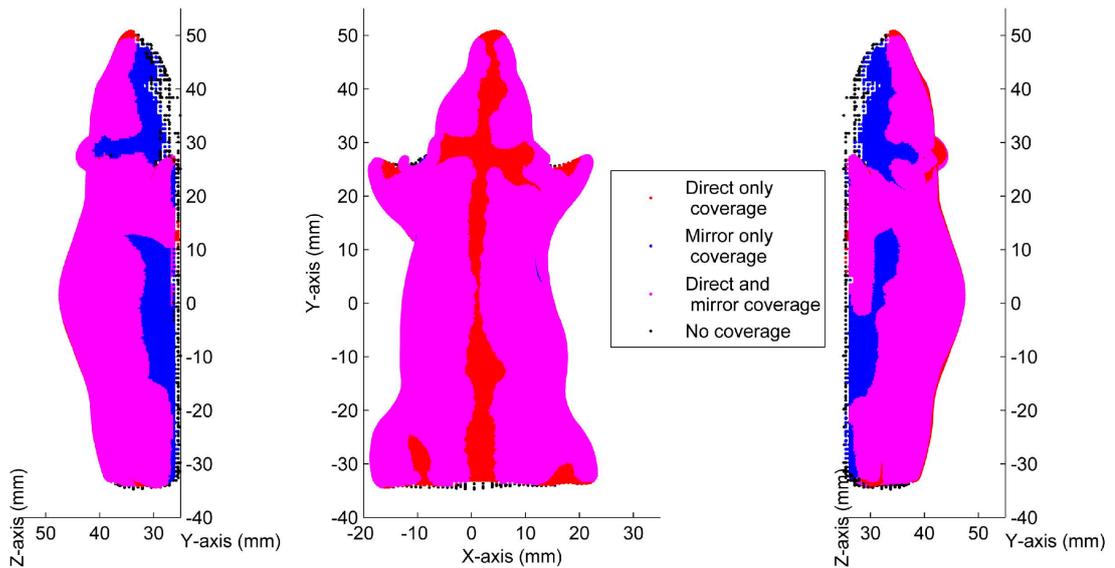
Figure 2.8 shows the coverage of the instrument in terms of direct views and mirrors, and the resulting surface mesh. The sections in red in fig. 2.8a are visible only to the

direct views, the sections in blue are visible only to the mirror views, the sections in pink are visible to both views, and the sections in black are not visible to any views. Regions low on the subject are only visible to the mirror views. These may be significant to bioluminescence reconstruction in the case where bioluminescent sources are deep within the subject, as these regions would produce strong signal due to short path lengths between source and surface.

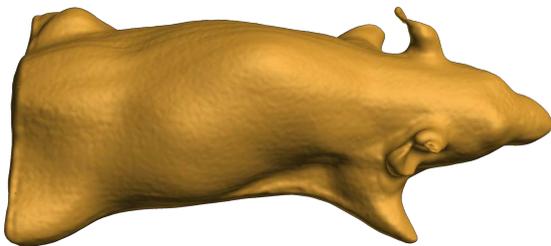
The resulting surface mesh can be seen in fig. 2.8b. This mesh is detailed and contains the major features present in the XPM-2 phantom. The mesh is qualitatively good when compared to a photograph of the phantom in fig. 2.8c. The mesh is missing part of the subject's left ear and lower jaw. The former is due to the complex ear shape. This prevents complete coverage of the ear and the meshing algorithm experienced difficulties processing the point clouds in that location as a result. The lower jaw is missing because it was not visible to the camera directly but more importantly was in shadow with respect to the illumination from both projectors.

A quantitative examination of the mesh quality was conducted by comparing points from a manufacturer provided mesh to the mesh generated by the surface metrology system. Points deemed to be outside of the field of coverage of the surface metrology system were manually removed leaving 878 points to compare with the mesh, which can be seen in fig. 2.9a. The set of points was rigidly registered to the phantom mesh to preserve the shape of the point set and then the distance between each of the manufacturer-provided points and the phantom was calculated. The results can be seen in fig. 2.9c, which indicates that 50% of points lie within 0.1mm of the mesh surface, 82% of points lie within 0.2mm of the mesh surface, and 96% of points lie within 0.4mm of the mesh surface. A small number of outlier points are found up to 1.3mm away, and these correspond to complex features that the system was unable to image with sufficient fidelity, such as the left ear.

A mouse cadaver was imaged as part of a BLT experiment, and the surface metrology results qualitatively show that the system is capable of imaging animal subjects in addition



(a) Surface coverage of the system when imaging the XPM-2 phantom

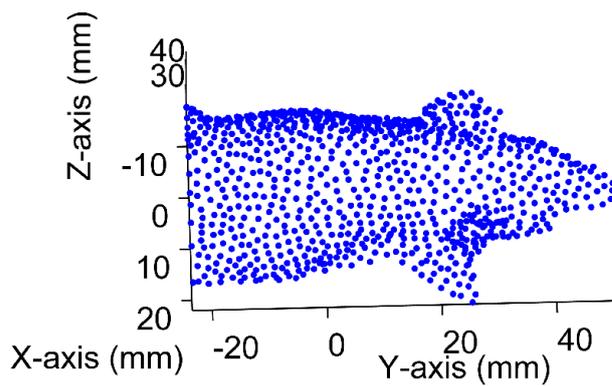


(b) XPM-2 phantom mesh calculated from all point clouds

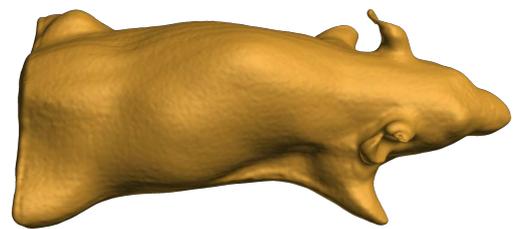


(c) Photograph of XPM-2 phantom

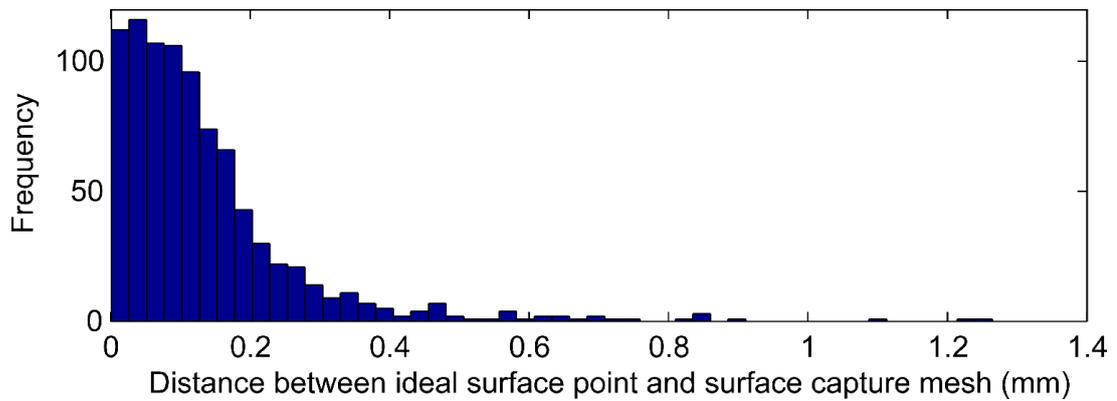
Figure 2.8: Surface coverage of the system when imaging the XPM-2 phantom, a mesh calculated from all the point clouds, and a photograph of the XPM-2 for comparison. Images from Basevi et al. (2013b).



(a) Ideal points from manufacturer-provided mesh used to calculate error



(b) XPM-2 phantom mesh used in error calculation



(c) Error between ideal surface points and measured surface mesh

Figure 2.9: Ideal points taken from a manufacturer-provided mesh, and the distance between those points and the surface of the surface metrology mesh. Images from Basevi et al. (2013b).

to phantoms. Figure 2.10 shows the cadaver under illumination from each projector, and the measured phase maps. Sufficient signal is obtained to acquire the phase maps for the vast majority of the surface using the pair of projectors and pair of mirrors. Figure 2.11 shows the point clouds for all the views superimposed, and the resulting mesh. The point clouds in fig. 2.11a indicate that all of the views provide some useful information in this case. The mesh calculated from the point clouds has been rendering as a wire frame in fig. 2.11b. A wire frame rendering consists of the triangle edges but does not include the triangle faces. Regions with a high density of edges result from a high density of surface points, and regions with a low density of edges were calculated from fewer surface points and therefore the surface approximation is inferior. Regions with a low density of edges exist in this example between the hind legs and the tail, and the underside of the animal as a result of poor visibility. However, the subject surface is in general qualitatively excellent.

Finally, the ability of the surface metrology system to resolve fine spatial features was tested by imaging a two pence coin. The coin is made of bronze and was minted in 1979. It has a total height of 1.85mm. Images and height maps of the coin can be seen in fig. 2.12. The surface detailing is clearly visible in the height maps, which supports the previous result indicating that surface error is on the order of 100s of micrometres. A set of regular spatial artefacts can also be observed from the height maps in figs. 2.12b and 2.12d. These form a grid that aligns with the projected pixel orientations in the projector. The cause of these artefacts is unknown, but the size of these artefacts is small enough that they can be neglected for the subjects for which this system was designed.

2.6 Conclusions

The surface metrology system was developed to address the need for a method of measuring the surface geometry of subjects of BLT imaging that:

- does not affect the existing hardware within the system.



(a) Mouse cadaver illuminated by projector 1



(b) Mouse cadaver illuminated by projector 2

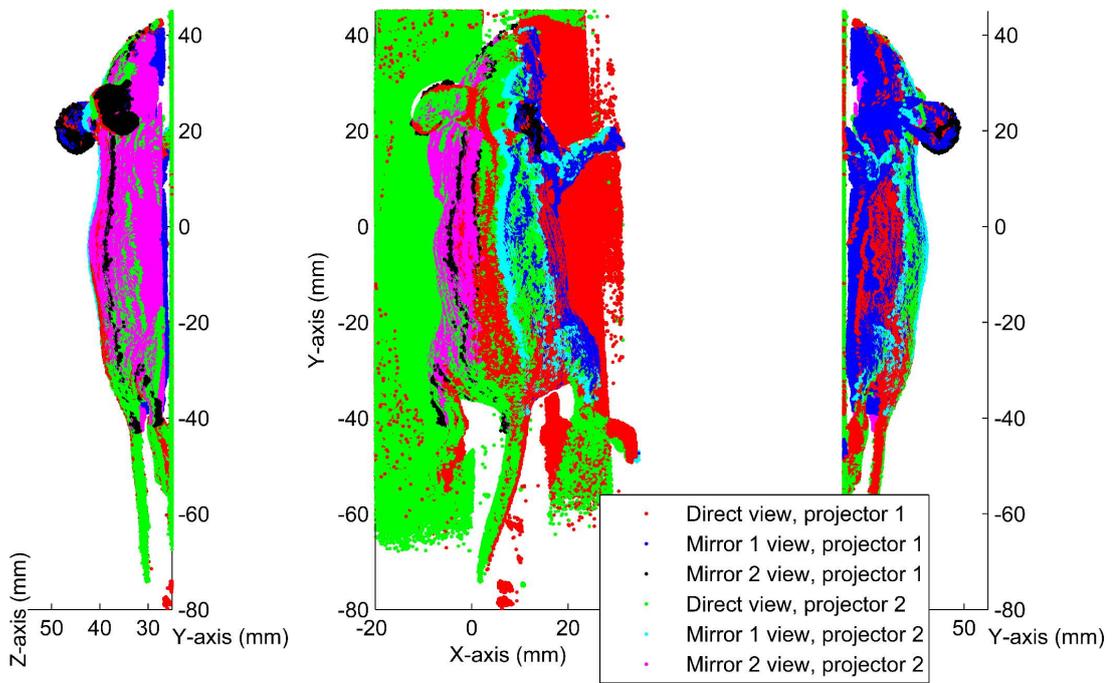


(c) Mouse cadaver phase map for projector 1

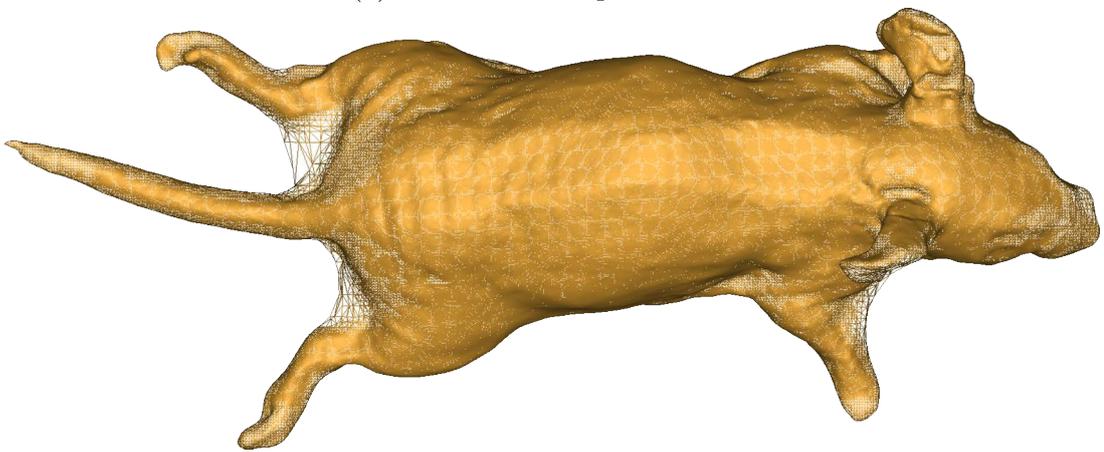


(d) Mouse cadaver phase map for projector 2

Figure 2.10: Mouse cadaver illuminated by projectors and associated phase maps. Images from Basevi et al. (2013b).



(a) Mouse cadaver point clouds

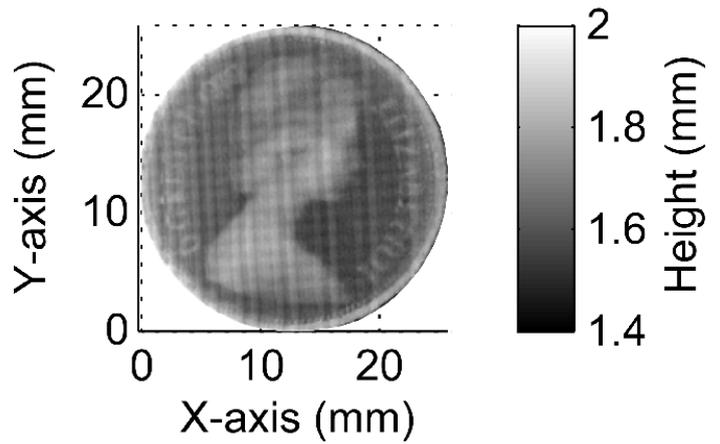


(b) Mouse cadaver mesh

Figure 2.11: Mouse cadaver point clouds and final mesh displayed as a wire frame. Images from Basevi et al. (2013b).



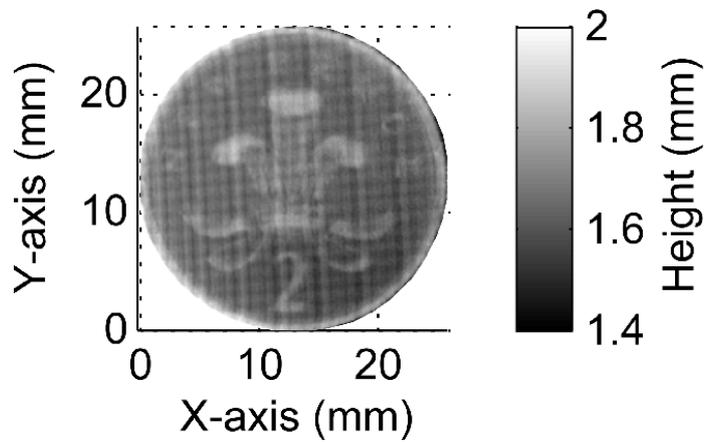
(a) Two pence coin illuminated image, heads up



(b) Two pence coin height map, heads up



(c) Two pence coin illuminated image, tails up



(d) Two pence coin height map, tails up

Figure 2.12: Surface metrology of a 1979 two pence coin cast in bronze. Bronze two pence coins have a height of 1.85mm. Images from Basevi et al. (2013b).

- can require additional hardware to be integrated into the system provided that existing hardware is not affected.
- can provide a mapping from pixels of the existing system camera to imaging subject surface coordinates.
- can utilise the existing system mirrors to measure geometry visible through the mirrors, and provide a mapping between system camera pixels and surface points visible through the mirrors.

The surface metrology system presented in this chapter meets these requirements and additionally allows the position of the system mirrors, which are also used in BLT imaging, to be calibrated from expanded surface metrology data. This method of mirror calibration enables flexible mirror placement at the time of surface metrology and BLT imaging so that the mirror positioning can be optimised for the individual size and orientation of each imaging subject, and the position and orientation calculated after imaging is completed.

The surface metrology system was tested *in simulacra* using the XPM-2 phantom, which is used in the majority of this work. Imaging results can be found in figs. 2.1 and 2.7 to 2.9, and show that integration of the mirrors provides increased surface coverage and error in the resulting surface mesh is less than 0.1mm in 50% of points tested, less than 0.2mm in 82% of points tested, and less than 0.4mm in 96% of points tested.

The system was further applied to image a mouse cadaver as can be seen in figs. 2.10 and 2.11. This demonstrates that the system can image biological tissue in addition to the polymer resin material that constitutes the XPM-2 phantom.

Finally, a two pence coin was imaged to demonstrate the system's ability to measure fine features and coin surface detail was successfully recovered. Regular grid artefacts were observed, but the magnitude of these artefacts is insignificant in the context of the size of the intended imaging subjects.

This system has two limitations. Firstly, it is necessary to manually segment acquired images to separate the direct view from the mirror views. This could be accomplished

automatically through the use of markers placed on the mirrors. The mirror calculation is not calculated from landmarks selected in the segmentation process so the effect of user subjectivity is minimal.

Secondly, the projectors must not project through the mirrors or two conflicting signals may be observed at points on the subject surface. This is linked to the mirror segmentation problem as knowledge of the position of the mirror is necessary to modify projected images so as not to project onto the mirror. A marker-based system would probably be necessary to address this robustly.

The system performs well enough to acquire high quality data despite these limitations and is used in *in simulacra* imaging for the remainder of this work.

CHAPTER 3

APPLICATION OF COMPRESSIVE SENSING-BASED TECHNIQUES TO BIOLUMINESCENCE TOMOGRAPHY

3.1 Contributions to this work

The *in simulacra* XPM-2 phantom data used in this chapter was collected by Kenneth Tichauer, Frederic Leblond, and Robert Holt using the Dartmouth College instrument. The non-negative least squares (NNLS) algorithm was written by James Guggenheim.

3.2 Introduction

Surface light resulting from a bioluminescent source varies linearly in intensity with respect to the strength of the source. A number of techniques can be used to model the processes involved in light propagation through tissue but FEM models are typically used to represent the problem (Arridge et al. 1993; Arridge 1999; Arridge and Hebden 1997; Arridge and Schweiger 1997; Dehghani et al. 2008; Schweiger and Arridge 2014) and will be used throughout this work. FEM models consist of a set of nodes and a set of elements (see fig. 3.2a for an example). The nodes are defined by a spatial location, absorption and scattering values, and a bioluminescence value. Absorption and scattering values may be expressed in terms of bulk absorption and scattering, or parametrised in terms of the

concentrations of tissue chromophores and scatterer properties. The elements are defined by a set of nodes. The bioluminescence distribution at a location within an element is calculated by interpolating the bioluminescence values at each of the nodes that form the element. This results in a continuous, piecewise-differentiable bioluminescence distribution over the region of space defined by the sets of elements. This representation can then be used to solve the RTE or the DA. The relationship between bioluminescence values and modelled surface measurements takes the form:

$$\mathbf{J}\mathbf{x} = \tilde{\mathbf{y}} \quad (3.1)$$

Where \mathbf{x} is the vector of bioluminescence values at n nodes, $\tilde{\mathbf{y}}$ is the vector of m measurements predicted by the physical model of light propagation and associated with bioluminescence distribution \mathbf{x} , and \mathbf{J} is a matrix called the Jacobian which is made up of the derivatives $\partial\tilde{y}_i/\partial x_j$. These derivatives are calculated from the physical model of light propagation and provide a map from bioluminescence distributions to predicted surface measurements. If measurements are acquired at multiple wavelengths then the Jacobians and measurements for each wavelength can be concatenated to form one system of linear equations that encompasses all the wavelengths (Chaudhari et al. 2005).

The BLT reconstruction problem is ill-posed (Wang et al. 2004), as are other optical tomographic modalities (Arridge and Lionheart 1998). A well-posed inverse problem:

1. Has an objective function with a zero.
2. Has an objective function with a single zero.
3. Has an objective function for which the position of the zero changes continuously with the measurements.

In other words, there is a continuous bijective function between the set of all possible bioluminescence distributions and the set of all possible surface measurements. A typical BLT reconstruction problem does not meet this criterion because the problem is underdeter-

mined and so the objective function has multiple zeros. The problem is underdetermined for a number of reasons:

- The unknowns that parametrise the bioluminescence distribution typically outnumber the known surface measurements, and so the \mathbf{J} matrix is rectangular ($m < n$). This is a consequence of the number of nodes and elements necessary to parametrise the target volume with sufficient resolution, and of instrumental limitations.
- The set of vectors that map changes in the bioluminescence distribution to changes in the measurements is not a linearly independent set, which means that the rows and columns of \mathbf{J} are not linearly independent ($\text{rank } \mathbf{J} < m$). This is a consequence of the diffusive nature of tissue and the limitation of measurement locations to those on the surface of the volume of interest. Measurements can be acquired at multiple wavelengths to improve reconstruction quality (Alexandrakis et al. 2005; Lu et al. 2011), but the Jacobians for different wavelengths are not independent.
- The matrix \mathbf{J} and measurements \mathbf{y} are contaminated with sources of error. Errors in \mathbf{J} are a consequence of inaccuracies in the physical model of light propagation and the parametrisation of the volume of interest. Sources of error in \mathbf{y} result from statistical fluctuations of signals and sources of noise within the measurement instrument.

The most common objective function used to represent the BLT reconstruction problem is:

$$f(\mathbf{x}; \mathbf{J}, \mathbf{y}) = \|\mathbf{y} - \mathbf{J}\mathbf{x}\|_2^2 \quad (3.2)$$

A solution \mathbf{x} that minimises eq. (3.2) produces predicted measurements that best match those observed when used as an input to the forward model under certain assumptions on the noise distribution, and in the ideal case where the physical model of light propagation is perfectly accurate and the measurements are noise-free then $f(\mathbf{x}; \mathbf{J}, \mathbf{y}) = 0$. In a statistical interpretation this assumption is that the measurement noise is normally

distributed. However, \mathbf{J} has a null space with dimension: $\dim \text{null} = n - \text{rank } J \leq n - m$. This implies that there are an infinite number of solutions.

Tikhonov regularisation is commonly employed to make problems tractable. In its general form (Kilmer et al. 2007), Tikhonov regularisation modifies the objective function to take the following form:

$$f(\mathbf{x}; \mathbf{J}, \mathbf{y}) = \|\mathbf{y} - \mathbf{J}\mathbf{x}\|_{\Sigma_y^{-1}}^2 + \|\mathbf{x} - \mathbf{x}_0\|_{\Sigma_x^{-1}}^2 \quad (3.3)$$

where \mathbf{x}_0 is a vector of the same size as \mathbf{x} , Σ_x and Σ_y are covariance matrices for the bioluminescence distribution and measurements respectively, and $\|\mathbf{x}\|_Q = \sqrt{\mathbf{x}^\top \mathbf{Q} \mathbf{x}}$. The regularisation term penalises bioluminescence distributions that deviate from \mathbf{x}_0 , weighted by the matrix Σ_x .

In a Bayesian interpretation, this is equivalent to the assumption that x is drawn from a multivariate normal distribution with a mean of \mathbf{x}_0 and a covariance of Σ_x . The choice of \mathbf{x}_0 and Σ_x allows the incorporation of prior knowledge into the reconstruction problem. Knowledge of spatial structure can be incorporated into the vector \mathbf{x}_0 through spatial variation of the mean, and into the matrix Σ_x via the diagonal and off-diagonal elements. These correspond to the spread of values and correlations between different spatial locations respectively. This method implicitly assumes that \mathbf{x} is drawn from a multivariate Normal distribution in a Bayesian interpretation but this may or may not be justified depending on the characteristics of the particular imaging study of interest. Equation (3.3) is often simplified to:

$$f(\mathbf{x}; \mathbf{J}, \mathbf{y}) = \|\mathbf{y} - \mathbf{J}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x} - \mathbf{x}_0\|_2^2 \quad (3.4)$$

where λ is a regularisation weighting. This form is used when there is a lack of prior information about the bioluminescence distribution and noise characteristics of the measurement instrument. The ideal value of λ is generally not known *a priori* and is often chosen empirically. A number of other regularisation methods exist.

Total variation (TV) also incorporates knowledge of groups of pixels that are likely to be more correlated than in general, but this does not arise from specific biological knowledge (Feng et al. 2012; Gao and Zhao 2010). Tikhonov regularisation favours solutions similar to \mathbf{x}_0 , whereas TV biases the solution so that nearby nodes take similar values without directly biasing what that value should be. Consequently, TV enforces a spatially smooth solution.

The least absolute shrinkage and selection operator (LASSO) is a regularisation method that restricts the solution space to the region $\|\mathbf{x}\|_1 \leq k$ (Tibshirani 1996). The use of the ℓ^1 -norm restricts the solution to an n -diamond (see fig. 3.1 for a two dimensional example) rather than an n -sphere in the case of the ℓ^2 -norm, where n is the dimensionality of the space. This has the effect of biasing the solution of an underdetermined problem towards sparsity and in a Bayesian interpretation is equivalent to the assumption that \mathbf{x} follows a Laplace distribution (Park and Casella 2008). This property is exploited in the technique of CS.

CS can in some cases provide an exact solution to an underdetermined linear problem (Baraniuk 2007) by taking advantage of prior knowledge of the spatial distribution of \mathbf{x} and the properties of \mathbf{J} . Specifically, \mathbf{x} must be sparse in a known basis, and the rows and columns of \mathbf{J} must be distributed and independent (Candes and Romberg 2007; Rauhut 2010). Mathematical results within the field of CS have shown that a sparse problem can be solved exactly from a smaller number of knowns than the dimensionality of the problem. If a problem cannot be solved exactly, then it may still be possible to place a bound on the deviation of the solution from the true value (Candes 2008). These results are powerful and have allowed an improvement in reconstruction quality in existing imaging instruments such as MRI (Lustig et al. 2007), as well as allowing novel imaging protocols and instruments such as single pixel cameras (Baraniuk 2008).

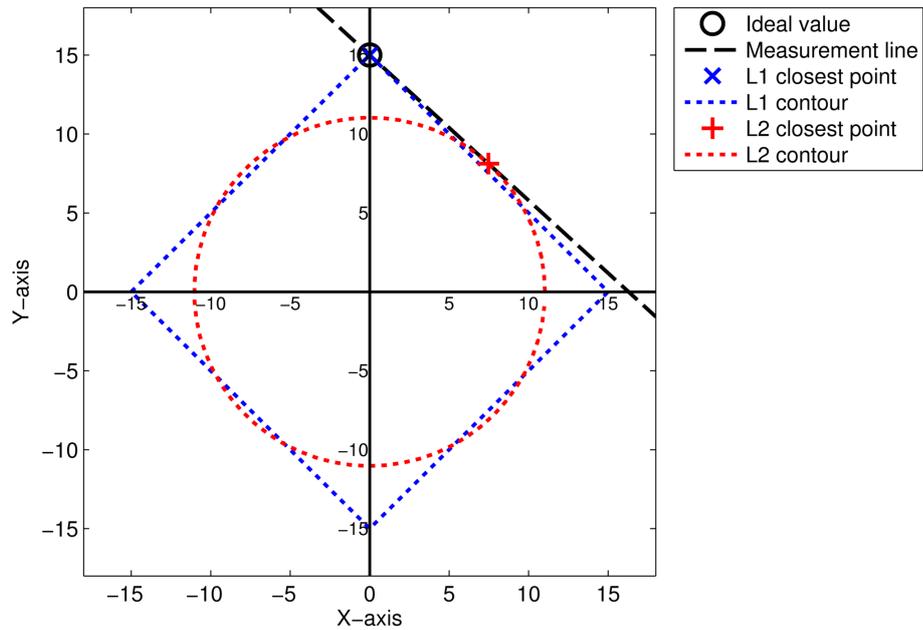
A two dimensional example of ℓ^1 -norm and ℓ^2 -norm reconstruction for a sparse solution can be found in fig. 3.1. A linear measurement/equation in two dimensions restricts the space of solutions to lie on a line. The intersection of two different measurement lines is

then a single point, and so two measurements are sufficient to specify a unique solution in two dimensions. If only one measurement is available then all points lying on the measurement line are possible solutions. The ℓ^1 -norm and ℓ^2 -norm can be used to select a single solution by choosing the point with the lowest norm.

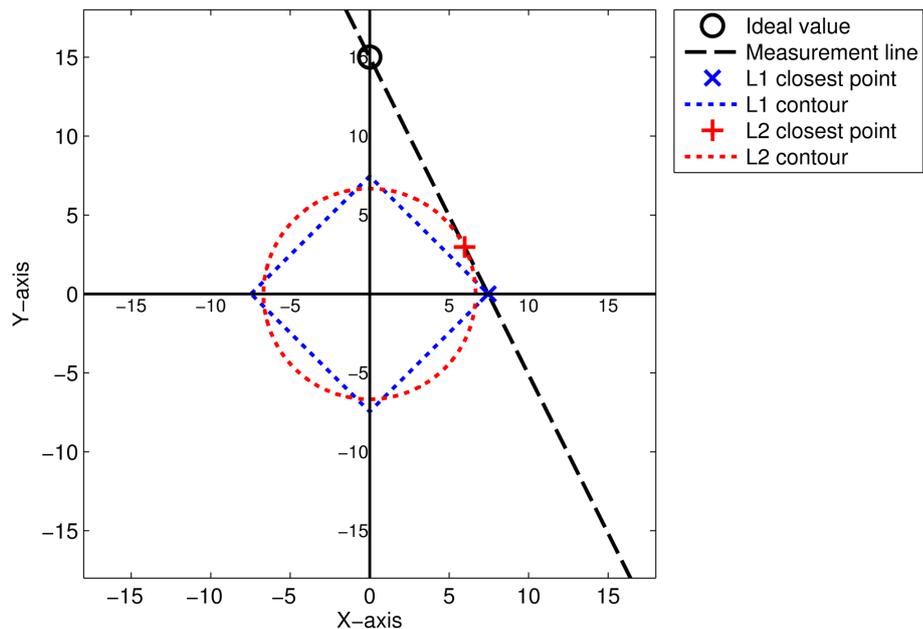
Figure 3.1a shows an example where the ℓ^1 -norm provides an exact solution and the ℓ^2 -norm does not. The dotted blue and red lines represent contours of equal norm value for the ℓ^1 -norm and ℓ^2 -norm respectively. This is a diamond for the ℓ^1 -norm and a circle for the ℓ^2 -norm. The ideal value is sparse as its X-axis component is zero. The diamond-shaped ℓ^1 -norm contour intersects with the measurement line at the position of the ideal value. The circular ℓ^2 -norm contour intersects with the measurement line at a position that contains non-zero X-axis and Y-axis components.

Figure 3.1b shows an example where neither the ℓ^1 -norm nor the ℓ^2 -norm provides an exact solution. The ℓ^1 -norm contour intersects with the measurement line on the Y-axis rather than the X-axis, and the ℓ^2 -norm contour again intersects with the measurement line at a position that contains non-zero X-axis and Y-axis components. The ℓ^1 -norm has a probability of 0.5 of providing an exact solution and a probability of 0.5 of providing an incorrect solution containing a zero X-axis component. This results from the shape of the ℓ^1 -norm contour which intersects the measurement line at $x = 0$ or $y = 0$ except in the case where the measurement line is at an angle of $\pi/4$ rad or $3\pi/4$ rad to the X-axis. Half of the remaining angles result in a ℓ^1 -norm contour which intersects the measurement line at $x = 0$, and the other half result in a ℓ^1 -norm contour which intersects the measurement line at $y = 0$. The ℓ^2 -norm contour only intersects the measurement line at $x = 0$ or $y = 0$ in the case where the measurement line is at an angle of 0 rad or $\pi/2$ rad to the X-axis.

A number of conditions have been derived that are sufficient to enable CS-based reconstruction such as the Restricted Isometry Property (Candes 2008), but to the best of the author's knowledge a set of necessary and sufficient conditions has not been formalised. Intuitively the set of measurements should be distributed in the sparse basis and



(a) ℓ^1 -norm and ℓ^2 -norm solutions of a two dimensional problem with one measurement where the ℓ^1 -norm solution is correct



(b) ℓ^1 -norm and ℓ^2 -norm solutions of a two dimensional problem with one measurement where the ℓ^1 -norm solution is incorrect

Figure 3.1: ℓ^1 -norm and ℓ^2 -norm solutions of a two dimensional problem with one measurement. The space of possible solutions permitted by the measurement are represented by the black dashed line, and the ℓ^1 -norm and ℓ^2 -norm solutions and contours of minimal distance that touch the measurement line are represented in blue and red respectively. Figure 3.1a shows a case where the random measurement is such that the ℓ^1 -norm solution is exact, and fig. 3.1b shows a case where the random measurement is such that neither distance metric produces the correct solution.

orthogonal. This implies that the measurement set provides information on all the dimensions of the signal and each measurement within the set provides unique information. These conditions are satisfied by a number of measurement types including Fourier space measurement, such as MRI (Lustig et al. 2008), and random binary measurement, such as single pixel cameras (Baraniuk 2008). BLT matrices do not satisfy these conditions. The sensitivity to bioluminescence in tissue drops exponentially with tissue depth and so the measurements are not well distributed. The diffusive nature of tissue also leads to dependent measurements. However, in certain cases BLT distributions possess a property that implies that techniques used in CS can potentially improve the quality of BLT image reconstruction. This property is that of sparsity.

Bioluminescence used in preclinical studies is typically produced by cells that are naturally or are genetically altered to express genes that code for the luciferase enzyme and potentially also the luciferin substrate (Contag and Bachmann 2002; Dothager et al. 2009; Luker and Luker 2008; Prescher and Contag 2010; Thorne and Contag 2005). The cells may not originate from the subject but if they do they must be removed for the genetic alteration process. Consequently, the cells must be reintroduced prior to imaging and this is typically performed by injection, or incision and implantation into the resulting cavity. The cells are then initially confined to a small volume within the subject, and thus the distribution of bioluminescence is sparse. These cells may divide and grow over time, or die, depending on the environmental conditions. The distribution becomes less sparse as the cells reproduce, but this is a gradual process. The distribution also becomes less sparse if cells spread through the blood stream and move to other sites within the subject, but this is also a gradual process and the signal limitations of BLT resulting from the attenuation of tissue prevent the detection of small colonies of cells. Intuitively, this implies that techniques that exploit sparsity may provide superior reconstruction quality.

A number of studies have examined the potential of sparse reconstruction techniques on BLT reconstruction quality (see section 3.3). These studies show good reconstruction quality but in general do not address the weighting of the sparsity component of the

objective function in a robust manner. This weighting is usually determined empirically, and this is only practical when the ideal solution is known such as in the case of *in silico* and *in simulacra* experiments. Image reconstruction consistency would be improved if the sparsity weighting could be determined in an objective manner.

To that end, a reconstruction algorithm was designed to employ sparsity in an objective manner. A conjugate gradient (CG) algorithm (Shewchuk 1994) was modified to include a sparsity term. The sparsity term is initially weighted to dominate over the measurement mismatch term so that the initial minimum of the objective function is the zero vector. This sparsity is then iteratively decreased until the sparsity weighting becomes negligible and the algorithm has converged, or the measurement mismatch is below a criterion defined by the noise level of the imaging system. The algorithm converges on a solution to the standard non-sparse objective function (eq. (3.4)) in the former case, and produces a sparse solution that satisfies the measurements to within acceptable measurement discrepancy in the latter case. The algorithm was tested on *in silico* and *in simulacra* examples. This algorithm will be referred to as compressive sensing conjugate gradient (CSCG). To the author's knowledge this algorithm was the first sparsity-based reconstruction to be applied to the problem of BLT that incorporates knowledge of instrument noise into the reconstruction process and uses that knowledge to automatically choose an appropriate sparsity weighting within a single execution of the algorithm. The CSCG algorithm uses the result of reconstructions with high sparsity weightings to inform the reconstructions at lower sparsity weightings to increase reconstruction efficiency and does not need to reconstruct using lower sparsity weightings than the sparsity weight finally chosen.

3.3 Survey of reconstruction algorithms used in bioluminescence tomographic imaging

A variety of reconstruction algorithms are used in optical tomographic imaging. BLT is linear in terms of the relationship between bioluminescence intensity and measurement magnitude and so algorithms that solve systems of linear equations can potentially be used to reconstruct bioluminescence. However bioluminescence problems are typically underdetermined and regularisation terms are added to the objective function to improve the quality of solutions. Some regularisation terms are non-linear and in general non-linear optimisation algorithms are required to minimise the augmented objective function. A representative selection of reconstruction algorithms are presented in this section that are used to reconstruct BLT distributions or similar types of optical tomographic images.

Lu et al. (2011) presented reconstruction of BLT multispectral data of a tumour xenograft mouse model. The physics of light propagation were modelled using the DA and also third order spherical harmonics (SP_3), and reconstructed using a quasi-Newton method with a least squares objective function constrained by an upper bound on bioluminescence intensity. The reconstruction was constrained to a 22mm slice of the subject torso. CT and PET images were also acquired and used for validation.

The source centre was reconstructed with accuracy of 1.5mm in each axis using an SP_3 physical model and 4mm in each axis. The PET imaging showed the presence of a number of less intense sources, but this was explained as non-specific probe uptake leading to differences between the true BLT and PET distributions. The SP_3 physical model provided superior reconstruction to the DA model.

Their results demonstrate that it is possible to reconstruct bioluminescent sources without the use of complex regularisation. The use of a permissible source region reduced the size of the problem by approximately 40%, and the upper bound on source intensity also provided a constraint, although the value of this constraint was not specified. It was possible to reconstruct BLT using the DA, although the SP_3 physical model provided superior accuracy. In general however, it may not be possible to use such a small permissible

region.

Alexandrakis et al. (2005) presented a combined PET and BLT system which utilised an FEM-based DA physical model to model light propagation. An expectation maximisation algorithm was used to reconstruct bioluminescence sources. The expectation maximisation algorithm required empirical selection of a stopping criterion which the authors indicated as a weakness of the technique.

Chaudhari et al. (2005) presented a multispectral reconstruction technique that used an singular value decomposition (SVD) to perform measurement compression, and minimised an objective function consisting of a least squares measurement discrepancy term, a basic Tikhonov regularisation term equivalent to the term in eq. (3.4) with $\mathbf{x}_0 = \mathbf{0}$, and a penalty term to penalise unphysical negative solutions. The regularisation term weighting was selected using the L-curve technique (Hansen and O’Leary 1993) and the penalty term weighting was selected empirically. It was noted that large penalty weightings negatively affected reconstruction convergence. A heuristic preconditioner was used to increase convergence speed. Their results indicated that the use of a multispectral model improved reconstruction quality.

The use of a penalty function was not ideal because this allows negative solutions, which are unphysical. The penalty weight μ was chosen to vary during the reconstructions process as $\mu_i = 0.001 \times \max x_i$ for iteration i . The generalisability of this formula to other BLT reconstruction problems was not explored.

Cong et al. (2010) presented a reconstruction method for BLT that used an evolutionary computation algorithm called differential evolution (DE) (Storn and Price 1997). Their implementation with DE differs from conventional BLT reconstruction algorithms in a number of ways:

1. Bioluminescence distributions were parametrised as a finite number of point sources defined by a spatial location and source power. These point sources were not necessarily located at the nodes of the FEM mesh and might be fewer in number than the number of nodes in the FEM mesh.

2. The algorithm state at each iteration consisted of a number of separate bioluminescence distributions instead of just one.
3. The algorithm did not make use of the derivatives of the objective function to generate new bioluminescence distributions. Instead, new bioluminescence distributions were generated by combining three existing bioluminescence distributions.
4. Existing bioluminescence distributions were replaced by new bioluminescence distributions if this would produce a reduction in objective function value. Each new bioluminescence distribution was compared to one existing bioluminescence distribution.
5. Distributions were generated under the assumption of a specified number of point sources, and two point sources were combined if the distance between the points was smaller than a specified value.

The authors' representation greatly reduces the dimensionality of the reconstruction problem. This offsets the limited number of and dependence of measurements. The use of a population preserves variation in the bioluminescence distributions for a number of iterations until the algorithm converges. This allows the algorithm to explore suboptimal regions in the objective function and so allows the search to be more global. The results showed good reconstruction quality in *in silico* and *in simulacra* experiments. The authors also state that reconstruction bioluminescence for an *in vivo* experiment showed good agreement with tumours observed via dissection.

This algorithm regularised the reconstruction process by reducing the solution space using a parametrisation. This both improves the conditioning of the problem and reduces the computational complexity, but is only suitable when it can be assumed that the bioluminescence distribution consists of a relatively small number of point sources and the number of these points sources is known *a priori* to a degree.

Feng et al. (2009) presented a maximum *a posteriori* (MAP) algorithm to solve the BLT reconstruction problem using a Gaussian Markov random field prior and permissible

source region. This prior penalised differences in adjacent mesh elements, and has also been applied to DOT and FMT reconstruction (Milstein et al. 2004; Ye et al. 1999). Measurement noise was modelled assuming that shot noise is dominant. A gradient-based optimisation algorithm then optimised for each mesh element independently and sequentially within each operation.

In silico experiments were performed using a heterogeneous mouse model and MC software to generate measurements. The permissible source region was assigned to a slice of the lung. Reconstruction results were compared with and without a Bayesian approach although the non-Bayesian reconstruction algorithm was not specified. The results indicated that reconstructed source location and density were superior using the Bayesian approach, that the algorithm could reconstruct over a number of different SNRs, and that the algorithm could reconstruct two sources simultaneously. *In simulacra* experiments were performed using a homogeneous block phantom, and successfully reconstructed a single source although with poor localisation.

The MAP reconstruction algorithm provided improved performance and the Bayesian approach allows flexible and principled regularisation, but the weight matrix for the Gaussian Markov random prior used in the experiments was not specified. The matrix can be designed using prior structural information, but the weight matrix choice was not fully justified. In addition, the restriction of the permission source region to a slice of the lung greatly reduces the solution space and computational complexity, but such a strong restriction may not be practical in the absence of prior information.

Feng et al. (2012) presented reconstruction algorithm for BLT using TV regularisation and the split Bregman method, and a non-negativity constraint. The regularisation is conceptually similar to the Gaussian Markov random prior used in the Feng et al. (2009) in that it penalises differences between elements. The Gaussian Markov random prior is however more flexible in the types of correlations between regions that it can represent. For example, the Gaussian Markov random prior could represent a correlation between spatially separated regions resulting from the presence of a blood vessel

that passes through both regions, whereas TV regularisation is purely local. The split Bregman method was used to efficiently solve reconstruction problems involving TV regularisation through the use of an auxiliary variable and splitting the optimisation problem into a number of sub-problems that separate the measurement discrepancy and TV terms of the objective function into different sub-problems. The TV regularisation weighting was chosen empirically and fixed and it was noted that this choice was usually successful. *In silico* experiments were performed on a two dimensional circular example in multiple cases and compared to algorithms with regularisation using the ℓ^1 -norm and ℓ^2 -norm. It was demonstrated that the presented algorithm produced superior performance.

Lu et al. (2009) presented a quasi-Newton reconstruction algorithm for BLT that used sparsity-based regularisation via the ℓ^1 -norm, with non-negativity and maximum bound constraints. It should be noted that the paper authors referred to the technique as a CS method but CS is as much dependent on the properties of the measurement matrix as it is on the type of regularisation, and BLT measurement matrices have not been shown to meet the requirements of CS. The ℓ^1 -norm regularisation weights were chosen empirically to maximise observed imaging quality, and larger regularisation weights were used for measurements with lower SNR.

In silico experiments were carried out on a block phantom and compared to algorithms without regularisation, and with ℓ^2 -norm regularisation. Reconstruction performance of the ℓ^1 -norm-regularised algorithm was superior to the other algorithms, particularly in the case of low measurement SNR. *In simulacra* experiments were carried out using a mouse-shaped phantom with a single point light source, and all methods were able to reconstruct the source. The unregularised algorithm reconstructed the most compact source but contained strong artefacts. The ℓ^2 -norm-regularised algorithm reconstructed a very distributed source, and the ℓ^1 -norm-regularised algorithm reconstructed a compact source with slightly less intense artefacts.

Naser and Patterson (2011b) presented reconstruction algorithms for BLT and FMT that used the ℓ^1 -norm for the measurement discrepancy term rather than for regularisa-

tion, and used an iteratively shrinking permissible source region. These algorithms were an extension of Naser and Patterson (2010).

The algorithms normalise the contribution of each measurement wavelength to the reconstruction process by dividing the Jacobians and measurements for each wavelength by the maximum measurement at that wavelength. This is done to prevent the reconstruction process being dominated by measurements at the most sensitive wavelength. However, this is not fully addressing the issue of measurement weighting because it may not be appropriate for each measurement wavelength to contribute equally to the reconstruction. For example, if the SNR of one wavelength is twice that of another wavelength, normalising the contribution of both wavelengths would cause the reconstruction to under-fit for the high SNR wavelength and over-fit for the low SNR wavelength. Instead, measurements should ideally be weighted by some measure that is a function of measurement SNR.

The reconstruction algorithm minimised an ℓ^1 -norm measurement discrepancy term subject to a non-negativity constraint, upper bound on bioluminescence at each location, and a permissible source region. This permissible source region initially included all nodes within the subject excluding nodes less than one path length from the surface. The permissible source region was shrunk at each iteration, and the fraction of the nodes with lowest intensities was discarded. This fraction was a fixed fraction of the current size of the permissible source region. The algorithm was iterated for a fixed number of iterations until the size of the permissible source region was reduced to a chosen value. This required prior knowledge of the size of the bioluminescent source.

An advancement on these algorithms was presented in Naser and Patterson (2011a) which multiplied eq. (3.1) on the left by \mathbf{J}^\top and used an eigenvector decomposition to reduce the reconstruction problem to one which was not underdetermined. This was accomplished by discarding eigenvectors which had negligible eigenvalues. The criterion for discarding eigenvalues was empirically determined but was dependent on measurement SNR. The algorithmic requirement of prior selection of the final size of the permissible

source region was retained.

Yu et al. (2010) presented a BLT reconstruction algorithm that incorporated a ℓ^1 -norm regularisation into the objective function. A variant of the LASSO algorithm was used to perform the reconstruction process which iteratively selected an increasing subset of the columns of the Jacobian to use in the next iteration of the reconstruction. No justification was given for the values chosen for regularisation weighting.

Zhang et al. (2011) presented a primal-dual interior-point BLT reconstruction algorithm that made use of ℓ^1 -norm-based regularisation. The reconstruction formulation optimised the sparsity of the reconstruction within the null space of the Jacobian matrix as an equality between experimental and simulated measurements was enforced. A logarithmic barrier term was used to keep reconstruction values positive. This method does not require an explicit choice of regularisation weighting, but does not account for measurement SNR. The implementation of the algorithm terminates reconstruction when the ℓ^2 -norm of the measurement discrepancy vector falls below a threshold, so the ℓ^1 -norm of the bioluminescence distribution may not be fully minimised.

The presented algorithm was compared against a least squares conjugate gradient algorithm in *in silico* and *in vivo* experiments. A permissible source region was used to confine the reconstruction to the lung in the *in silico* case and to a small cuboid region in the *in vivo* experiments. The presented algorithm displayed superior performance in both cases, and particularly in an *in silico* experiment involving two separate bioluminescence sources. However the use of such a permissible region may not be possible in practice without prior information from additional imaging modalities.

Zhang et al. (2012) presented a reconstruction algorithm for BLT that made use of sparsity-based regularisation and a Poisson measurement likelihood distribution. Sparsity-based regularisation was implemented via the ℓ^1 -norm and ℓ^0 -norm separately. Reconstruction using the ℓ^1 -norm was implemented via an adaptation of an expectation maximisation algorithm which optimised the objective function, and reconstruction using the ℓ^0 -norm was implemented via an orthogonal matching pursuit algorithm which optimised

the measurement discrepancy term whilst growing a support set for the bioluminescence distribution. The regularisation weighting was chosen empirically for the ℓ^1 -norm-based reconstruction in order to maximise observed performance. The authors referred to the process of maximising the likelihood as MAP estimation. However, the regularisation was not formulated as a prior distribution and so a posterior distribution is not explicit.

The algorithms were tested against an existing ℓ^1 -norm-based CS reconstruction algorithm on an *in silico* CS problem involving Bernoulli matrices. The three algorithms produced comparable accuracy although the orthogonal matching pursuit algorithm was significantly more computationally expensive than the other algorithms. The algorithms were then tested on *in silico* BLT problems involving one and two bioluminescent sources. The orthogonal matching pursuit algorithm produced the best results and the other two algorithms failed on both reconstruction problems.

He et al. (2011) presented a BLT reconstruction algorithm that used ℓ^1 -norm-based sparsity regularisation and a iteratively refined permissible source region. The resolution of the FEM mesh was increased in this region which enabled compact bioluminescence sources to be represented with higher resolution and increased the accuracy of the physical model of light propagation within this region. The approach to permissible source region refinement is similar to Naser and Patterson (2011b) in that reconstruction is initially performed on a large region, and regions with low bioluminescence are removed. The approach differs from that of Naser and Patterson (2011b) in that the criterion for region removal is a threshold of 0.2 on the ratio of the local bioluminescence to the maximum bioluminescence, while a fixed fraction of nodes are removed at each iteration in the case of Naser and Patterson (2011b). This enables the use of this approach when prior knowledge of bioluminescence distribution size is unavailable. However, this approach may conflict with the use of ℓ^1 -norm-based regularisation because the iterative mesh refinement serves to make the bioluminescence distribution less sparse on the chosen basis, because more nodes are added to the region where the bioluminescence distribution is expected to be found. A truncated Newton interior-point algorithm was used to reconstruct the

bioluminescence distribution at each iteration.

The algorithm was tested in *in silico* experiments involving cylindrical and mouse-shaped phantoms. Regularisation weightings were empirically chosen to maximise observed performance, and initial permissible source regions significantly smaller than the problem size were chosen in each case. ℓ^1 -norm and ℓ^2 -norm-based regularisation was compared in the case of the cylindrical phantom, and the ℓ^1 -norm-based algorithm provided superior results. The experiment involving the mouse-shaped phantom demonstrated that the presented algorithm could reconstruct a compact bioluminescence source in this situation.

Chen et al. (2014) presented a reconstruction algorithm for BLT that makes use of the $\ell^{1/2}$ -norm to enforce solution sparsity more strongly than if using the ℓ^1 -norm. A non-negativity constraint and maximum bound on bioluminescence value was also applied, and the objective function was minimised using an interior point algorithm. The regularisation weighting was empirically chosen.

The performance of the proposed algorithm was compared against that of a number of algorithms on an *in silico* mouse-shaped phantom. Regularisation types in these comparison algorithms included ℓ^1 -norm-based and ℓ^2 -norm-based regularisation. The proposed algorithm was found to provide superior reconstruction performance.

3.4 Methods

3.4.1 CSCG algorithm design

The objective function to be minimised is similar to eq. (3.3) with the ℓ^2 -norm regularisation component replaced with an ℓ^1 -norm component. The ℓ^0 -norm provides the most direct measure of sparsity, but is discontinuous. The ℓ^1 -norm is continuous and differentiable except at zero and also maximises sparsity, so is commonly used instead of the

ℓ^0 -norm when sparsity is desired. The resulting function takes the form:

$$f(\mathbf{x}; \mathbf{J}, \mathbf{y}) = \|\mathbf{y} - \mathbf{J}\mathbf{x}\|_{\Sigma_y^{-1}}^2 + \lambda \|\mathbf{T}\mathbf{x}\|_1 \quad (3.5)$$

where λ is a regularisation weight and matrix \mathbf{T} is the basis within which \mathbf{x} is expected to be sparse.

Equation (3.5) reduces to the following form in the absence of knowledge of the instrument noise characteristics and under the assumption that \mathbf{x} is sparse in a Cartesian coordinate system:

$$f(\mathbf{x}; \mathbf{J}, \mathbf{y}) = \|\mathbf{y} - \mathbf{J}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1 \quad (3.6)$$

In this work x was assumed to be spatially sparse, and the regularisation term of the objective function was modified in order to make it also differentiable at zero. The final version of the objective function took the form:

$$f(\mathbf{x}; \mathbf{J}, \mathbf{y}) = \|\mathbf{y} - \mathbf{J}\mathbf{x}\|_{\Sigma_y^{-1}}^2 + \lambda \sum_{m=1}^M \sqrt{x_i^2 + \delta} \quad (3.7)$$

A CG algorithm formed the base of the CSCG algorithm. The CG algorithm can directly solve linear problems with symmetric positive-definite matrices. The objective functions for these problems have a global minimum. The algorithm is initialised at an arbitrary position and moves in directions of decreasing objective value until it converges on the solution. The solution is reached in N iterations for a problem of dimension N . The direction of movement at an iteration is conjugate to the direction of movement at all other iterations with respect to \mathbf{J} , so $\mathbf{d}_i^\top \mathbf{J} \mathbf{d}_j = \mathbf{0}$, where \mathbf{d}_i and \mathbf{d}_j are direction vectors for iterations i and j , and $i \neq j$. The distance to move in that iteration can be calculated directly so as to minimise the objective function in that direction.

The objective function in this problem is non-linear due to the sparsity regularisation term, and the \mathbf{J} matrix is neither symmetric nor positive-definite. The CG algorithm can be adapted to solve non-linear problems and these algorithms are called nonlinear conjugate gradient (NCG) methods. NCG methods are iterative, and despite the word

“conjugate” in the name do not necessarily move in conjugate directions. These methods behave like the CG algorithm on quadratic problems, but behaviour is less well defined on other problems.

Firstly, the distance to move in a particular direction to minimise the objective function cannot be calculated directly. Line search methods are typically used instead. These evaluate the objective function at a number of distances along the search direction and approximate an optimal movement distance from those values. The search directions may also lose conjugacy which can affect convergence. The search direction can be replaced with the gradient if it is determined that the iteration-to-iteration reduction in objective function value is insufficient to correct for this. Finally, the CG algorithm is a local greedy reconstruction algorithm in the sense that it considers only the local properties of the objective function at each iteration and always moves in a direction of decreasing objective function value. Consequently, it may converge to a local minimum in the objective function if one exists rather than a global minimum.

The choice of λ in eqs. (3.5) to (3.7) is crucial as this determines the sparsity weighting and thus the sparsity of the final solution. The final solution is the zero vector in the extreme where $\|\mathbf{y} - \mathbf{J}\mathbf{x}\|_{\Sigma_y^{-1}}^2 \ll \lambda\|\mathbf{T}\mathbf{x}\|_1$. The final solution is a local minimum of the measurement error in the other extreme where $\|\mathbf{y} - \mathbf{J}\mathbf{x}\|_{\Sigma_y^{-1}}^2 \gg \lambda\|\mathbf{T}\mathbf{x}\|_1$. The CSCG algorithm is deterministic and so the solution onto which the algorithm will converge depends only on the initialisation parameters of the algorithm including the initial position. It is dependent on the regularisation weighting λ because it affects the shape of the objective function and may change the position of the minima. Further, if λ changes during the execution of the algorithm then the intermediate values of λ may affect the choice of minimum to which the algorithm converges.

The use of a variety of λ values may be necessary because the ideal λ value is not known *a priori*. It is possible to execute the algorithm to convergence using a number of λ values separately and then choose an optimal reconstruction from the resulting set of reconstructions, but this approach is problematic. Specifically, the task of choosing an

optimal reconstruction from a set is non-obvious.

Selection can be conducted in the simplest case by taking the reconstruction that is associated with the smallest discrepancy in measurements. Alternative approaches include comparing the reconstructed image to some prior knowledge of expected image features, and choosing the image that possesses the least apparent noise in the reconstructed image.

Selection of the reconstruction associated with the smallest measurement discrepancy is sub-optimal in the unavoidable presence of measurement noise as this reconstruction is the most likely to have over-fitted to the noise. Selection of the image most consistent with expected image features is not possible in general as the range of expected image features will vary depending on the specific experiment and unexpected but real image features are not only possible but can be more informative scientifically than expected features. Selection of the image that possesses the least noise in the reconstructed image is not currently practical because noise features resulting from the imaging modality and reconstruction technique have not been well characterised and so any such selection is likely to be subjective and unreliable as a result, excluding extreme examples.

The proposed solution involves a stopping criterion that balances the simultaneous objectives of maximising reconstruction image sparsity and minimising the measurement discrepancy with knowledge of the noise characteristics of the measurement process. This process always contains a degree of noise, such as the statistical variation in the probe signal. A criterion can be chosen based on the noise sources in the measurement process to separate the set of measurement discrepancies arising from the set of possible reconstruction images into discrepancies that are consistent with noise, and discrepancies that are not. This separates the set of possible reconstruction images in the same manner. The reconstructions that produce measurement discrepancies inconsistent with noise can then be discarded and the desired reconstruction chosen from the remaining reconstructions. Measurement noise can no longer be used as a discrimination criterion between the remaining reconstructions as this would likely overfit for the noise in the measurements. The criterion of sparsity can now be used to find the sparsest of the remaining

reconstructions which can be selected as the final reconstruction to present to the user.

It is now necessary to consider the process of finding this reconstruction given the underdetermined nature of the problem and the lack of prior knowledge of the optimal value of λ . It is possible to perform multiple reconstructions with different values of λ and choose the reconstruction with the largest value of λ that meets the measurement discrepancy criterion. However, this is inefficient. Instead, a sequence of reconstructions can be performed for decreasing values of λ , and the result of a reconstruction then used as the initial position of the next reconstruction to prevent repetition of effort.

Choosing the first value of λ such that the sparsity component is dominant over the measurement discrepancy component ensures that the initial reconstruction is the zero vector $\mathbf{0}$, and is contained in the set of reconstructions that are associated with measurement discrepancies that are inconsistent with measurement noise (otherwise the measurements are consistent with no bioluminescence present in the subject and the reconstruction process is complete). The algorithm then reduces λ and starts the next reconstruction at the solution of the previous reconstruction, and so searching for a less sparse solution in an efficient manner because the old solution is likely to be near to the new solution. After a number of reconstructions the position of the algorithm may move from the set of reconstructions with measurement discrepancies that are inconsistent with measurement noise to the set of reconstructions that are consistent with measurement noise. At this point the current position of the algorithm is a solution with high sparsity that also meets the measurement error criterion, as the algorithm generally moves from a position of higher sparsity to lower sparsity due to the relaxation of the sparsity term. It is possible that the algorithm never reaches a satisfactory solution if λ is reduced until the sparsity contribution is negligible and the algorithm converges on a local minimum that does not satisfy the measurement discrepancy criterion. If so, the current position of the algorithm is chosen as the final reconstruction result.

3.4.2 CSCG algorithm implementation

The CSCG algorithm was loosely based on Lustig et al. (2007), in that it is a NCG algorithm that optimises an objective function containing an ℓ^1 -norm regularisation term which has been approximated as in eq. (3.7). However, a different line search method is utilised and the algorithm contains a number of additions detailed below. The algorithm optimises the objective function eq. (3.5). If the Σ or \mathbf{T} matrices are not provided, identity matrices are used instead. The algorithm can be set to terminate when each measurement associated with the current reconstruction position is on average within $k\sigma$ of the experimental measurement, where σ is the standard deviation for each measurement assuming that the noise component of each measurement is independent of the other measurements, and k is a non-negative real number. The algorithm is initialised so that the ratio of the sparsity term and the measurement discrepancy terms is by default 1×10^5 , using $\mathbf{J}^\top \mathbf{y}$ and the zero vector $\mathbf{0}$ as inputs respectively. This equates to the formula:

$$\lambda = \frac{10^5 \|\mathbf{y} - \mathbf{J}\mathbf{0}\|_{\Sigma_y^{-1}}^2}{\sum_{m=1}^M \sqrt{(\mathbf{J}^\top \mathbf{y})_m^2 + \delta}} \quad (3.8)$$

The algorithm is executed until the objective function change is below an arbitrary threshold, chosen here to be 1×10^{-5} of the objective function value at the previous iteration, at which point the value of λ is reduced. By default it is reduced to $1/\sqrt{2}$ of its previous value. The reconstruction process continues until the measurement discrepancy term is satisfied or the sparsity weighting is reduced to 1×10^{-15} of its original value.

During each iteration the algorithm performs a quadratic line search. It tests values in the current search direction until the points are obtained such that the middle point takes a value lower than the first and last points. A quadratic function is fit to these three points and the minimum of this function is used as the new position of the algorithm. The Fletcher-Reeves formula is used to generate a new conjugate search direction. The value of λ is decreased and the conjugate search direction is reset to the gradient when the algorithm converges.

The algorithm allows the rows of the \mathbf{J} matrix to be normalised to equalise the contribution of each measurement to the reconstruction process. This is not meaningful in the case where measurement noise is known, as the use of a measurement noise covariance matrix Σ provides a principled scaling to balance the contribution of each measurement using the measurement noise.

The algorithm also allows the columns of the \mathbf{J} matrix to be normalised during the reconstruction process. This is a numerical means of compensating for the exponential decrease in measurement sensitivity with tissue depth but has no physical meaning. Let \mathbf{w} be a vector containing the ℓ^2 -norms of the columns of the \mathbf{J} matrix. Let \mathbf{W} be the diagonal matrix containing the elements of \mathbf{w} . Let $\tilde{\mathbf{x}} = \mathbf{W}\mathbf{x}$, and $\tilde{\mathbf{J}} = \mathbf{J}\mathbf{W}^{-1}$. The CSCG algorithm is used to solve for $\tilde{\mathbf{x}}$ using $\tilde{\mathbf{J}}$ and \mathbf{y} . The desired vector \mathbf{x} is then calculated using the formula $\mathbf{x} = \mathbf{W}^{-1}\tilde{\mathbf{x}}$.

The pseudocode for the CSCG algorithm can be found in algorithm 1.

Algorithm 1 CSCG algorithm

```

Normalise the columns of the measurement matrix if applicable
Calculate initial and final sparsity weightings
Set initial algorithm position to the zero vector
Set initial search direction to initial gradient
while sparsity weighting > final sparsity weighting & measurement discrepancy unacceptable do
    while not converged & measurement discrepancy unacceptable do
        Perform quadratic line search to find new position
        Calculate new search direction from old search direction and current gradient
    end while
    Reduce sparsity weighting
    Set new search direction to current gradient
end while
Correct the reconstruction for the column normalisation if applicable

```

3.4.3 Experiments

The CSCG algorithm was first tested using a CS problem in the absence of noise to verify that it is capable of reconstructing underdetermined linear problems involving measurement matrices which satisfy the requirements of CS.

To that end, random binary measurement matrices were utilised. The elements within a random binary matrix have a probability p to take the value 1, independent of the values of all other matrix elements, and a probability of $1 - p$ to take the value 0. The specific random binary matrices used were such that exactly half of the elements of each row of the matrix took the value 1. These were applied to random two dimensional images of 256 pixels of varying sparsity ranging from 4 non-zero pixels to 256 non-zero pixels. All non-zero pixels were assigned the value 4096. A set of measurement matrices with sets of measurement ranging from 2 measurements to 254 measurements were applied to these images, and the resulting set of measurements was reconstructed using CSCG to measure algorithm performance. Neither row nor column normalisation were utilised. Twenty random images were tested for each pair of degree of image sparsity and number of measurements. The measurement-related stopping criterion was set such that the measurement discrepancy should be reduced to zero, due to the absence of noise.

The resulting images were analysed to extract four quantities:

- The image pixel error relative to the true magnitude of the non-zero pixels, averaged over all the pixels in an image and over all the images for a given degree of image sparsity and number of measurements.
- The image pixel standard deviation relative to the true magnitude of the non-zero pixels, averaged over all the pixels in an image and over all the images for a given degree of image sparsity and number of measurements.
- The image pixel sensitivity, defined as the number of true non-zero pixels identified by the reconstruction divided by the total number of true non-zero pixels, averaged over all the pixels in an image and over all the images for a given degree of image sparsity and number of measurements.
- The image pixel sensitivity, defined as the number of true zero pixels identified by the reconstruction divided by the total number of true zero pixels, averaged over all

the pixels in an image and over all the images for a given degree of image sparsity and number of measurements.

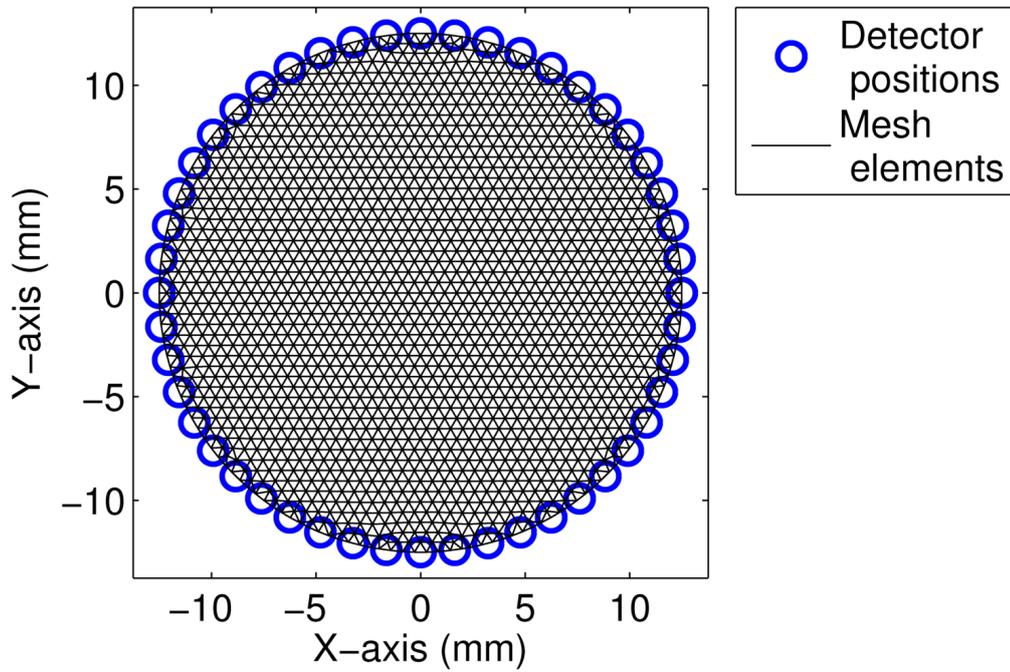
Pixels are classified as non-zero if their value was greater than or equal to half the true non-zero pixel value of 4096.

The CSCG algorithm was next utilised on the same problem in the presence of Gaussian and shot noise. This experiment examined the effect of noise on CSCG reconstruction and of the measurement-related stopping criteria on reconstruction performance. To simulate an 16-bit imaging device with perfect quantum efficiency, measurements for all test cases were normalised so that the maximum measurement took the value of 45,875, 70% of the maximum 16-bit number (65,535). Gaussian noise was added to the normalised measurements with zero mean and variance of $\mathbf{y} + 10$, to simulate shot noise and an small additive background Gaussian noise source. The stopping criteria tested were:

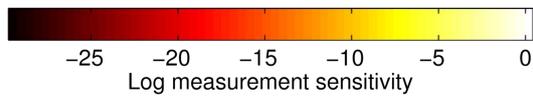
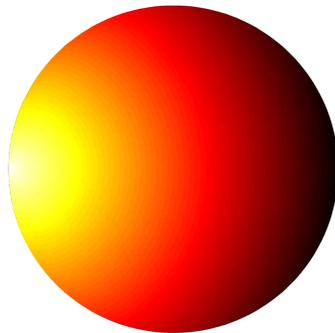
- All measurements associated with the reconstruction match the experimental measurements exactly.
- The measurements associated with the reconstruction deviate from the experimental measurements by one standard deviation on average.
- The measurements associated with the reconstruction deviate from the experimental measurements by two standard deviations on average.

These experiments were conducted under the same conditions as the noise-free case, and the results were analysed in the same way.

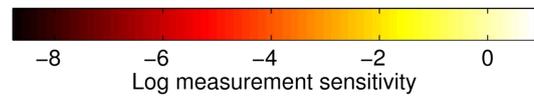
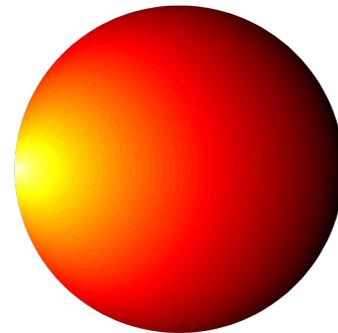
Following testing on the binary CS problem, the CSCG algorithm was applied to a two dimensional BLT reconstruction problem. Sixteen detectors were applied to a circular mesh of diameter 25cm consisting of 1,785 nodes and 3,418 elements, which can be seen in fig. 3.2. Measurements were simulated using Jacobian matrices calculated using Nirfast for wavelengths of 560, 580, 600, 620, and 640. These Jacobian matrices were used for both forward and inverse models. Example measurement sensitivities can be found in figs. 3.2b and 3.2c.



(a) Two dimensional BLT mesh



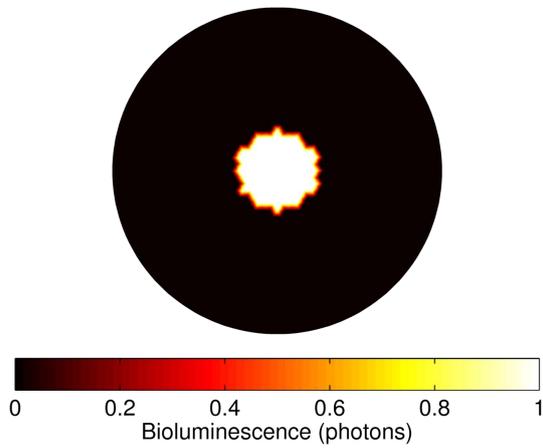
(b) Log sensitivity for a measurement at a wavelength of 560nm



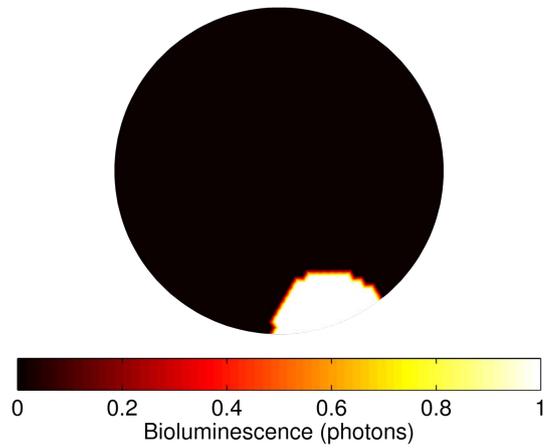
(c) Log sensitivity for a measurement at a wavelength of 640nm

Figure 3.2: Two dimensional BLT mesh used for CSCG testing. The optical properties of this mesh are homogenous, and arise from tissue that contains 40% water and 0.01 μ mol of oxygenated and deoxygenated haemoglobin.

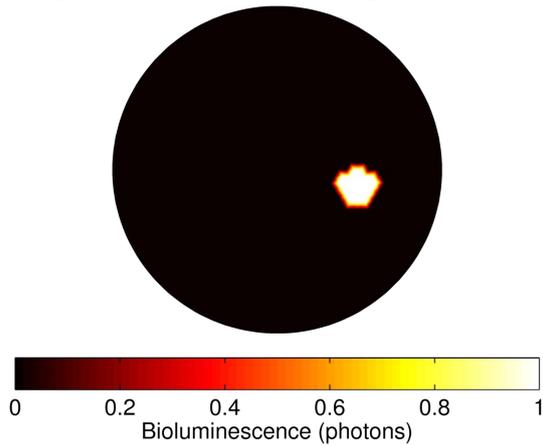
Random bioluminescence sources were generated for this mesh (see fig. 3.3) and CSCG reconstruction performance for sources of different sizes and depths was evaluated using 20 random sources for each source size and depth pair. Algorithm performance was also evaluated in the presence of shot noise and for different stopping criteria in the same way as the binary CS problem. In addition to the examination of mean error, standard deviation, sensitivity and specificity, the algorithm was also evaluated on the error in the total amount of bioluminescence recovered relative to the ideal amount of bioluminescence. The total amount of bioluminescence was calculated by taking the sum of the bioluminescence at each of the mesh nodes.



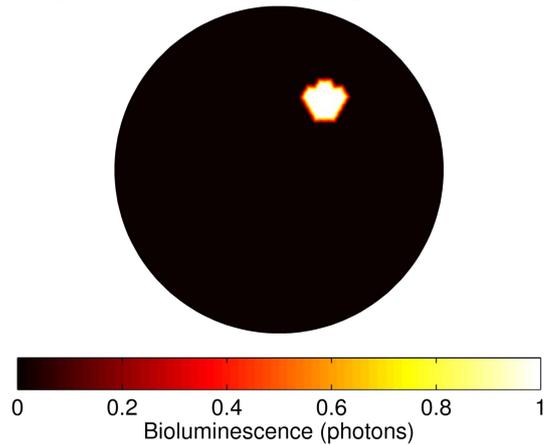
(a) Sample bioluminescence distribution at a depth of 12.5mm consisting of 100 nodes



(b) Sample bioluminescence distribution at a depth of 0.0mm consisting of 100 nodes



(c) Sample bioluminescence distribution at a depth of 6.25mm consisting of 25 nodes



(d) Sample bioluminescence distribution at a depth of 6.25mm consisting of 25 nodes

Figure 3.3: Examples of random bioluminescence sources generated for a two dimensional BLT mesh used for CSCG testing.

The algorithm was next tested on a BLT reconstruction problem *in silico* and *in simulacra*. The experiment was conducted using a multimodal imaging instrument designed and based at Dartmouth College (Leblond et al. 2011). The system is a dual modality contact imaging system. It uses CT to acquire information on surface geometry and detector positions, and uses optical fibres coupled to PMTs for optical measurement. The system was treated as a black box for the purposes of this test to examine the algorithm performance in the absence of knowledge about the imaging instrument. A case study using the University of Birmingham imaging instrument in the presence of instrument knowledge can be found in chapter 4.

Instrument measurement calibration, mesh generation, and measurement placement was carried out by the collaborators at Dartmouth College. A visualisation of the experimental configuration can be found in fig. 3.4 The resulting mesh was used to calculate a Jacobian matrix. Measurements with unphysical values of below 0 were removed. The resulting measurement matrix and measurements were processed by the CSCG algorithm without knowledge of the instrument noise. The algorithm was executed without a measurement discrepancy termination criterion *in silico* and *in simulacra* and column normalisation was not utilised, although row normalisation was utilised.

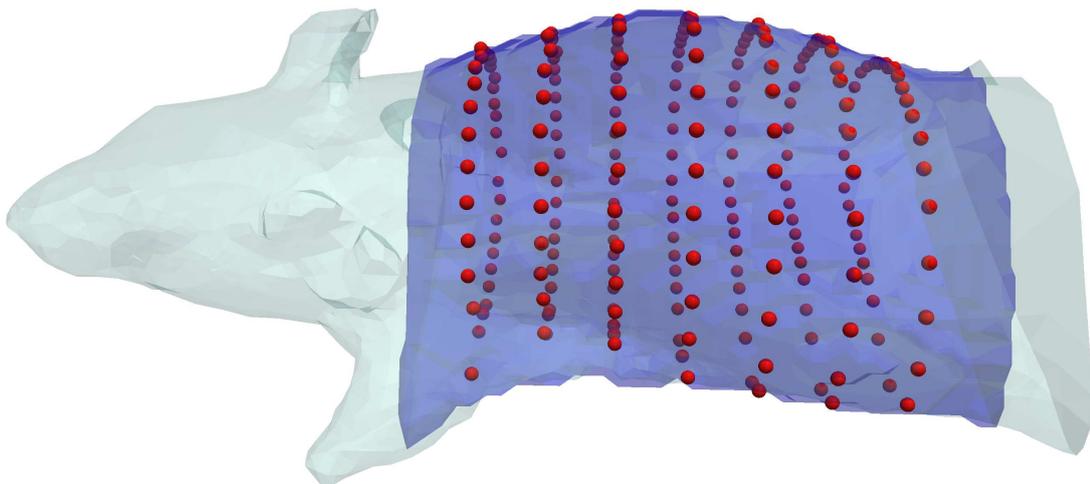


Figure 3.4: Dartmouth imaging configuration and mesh. The light blue mesh is a mesh of the full mouse phantom provided by the manufacturer. The dark blue mesh is of the torso of the phantom and was generated using the Dartmouth data. The red spheres indicate detector positions. Image from Basevi et al. (2012b).

In the case of the *in silico* experiments, measurements were generated using manufacturer provided values of source intensity and estimated source sizes. Reconstruction was tested with ideal measurements and with measurements contaminated with artificially introduced measurement noise. This noise was additive Gaussian noise with a mean of 0 and a standard deviation of 1% or 5% of the measurement magnitude. Thirty unique noise samples were tested for each noise level. Separate meshes were used for measurement generation and for image reconstruction to avoid performing an inverse crime. A high resolution mesh consisting of 93,018 nodes was utilised for measurement generation, while a low resolution mesh consisting of 10,171 nodes was utilised to calculate the measurement matrix from which the bioluminescence distribution was reconstructed. It is not possible to exactly represent the bioluminescence distribution in the high resolution basis using the low resolution basis as a consequence of the difference in mesh resolution.

Two benchmark algorithms were utilised for comparison. These algorithms are the Gauss-Newton (GN) (Nocedal and Wright 1999) and NNLS (Lawson and Hanson 1995) algorithms. Both algorithms incorporate basic Tikhonov regularisation and minimise the following objective function:

$$f(\mathbf{x}; \mathbf{J}, \mathbf{y}) = \|\mathbf{y} - \mathbf{J}\mathbf{x}\|_2^2 + \lambda\|\mathbf{x}\|_2^2 \quad (3.9)$$

The NNLS algorithm adds the additional constraint that $x_i \geq 0$, and can in some cases be used to solve CS problems (Foucart and Koslicki 2014). The NNLS algorithm was based on the NNLS algorithm found in (Lawson and Hanson 1995) with the addition of regularisation.

Quantitative analysis was not possible in the *in simulacra* experiment due to lack of instrument information. Two measures of reconstruction performance were examined in the *in silico* experiments where all experimental parameters are known. These measures were:

- The centre of the bioluminescence distribution for each source. These centres were

chosen to be the positions of the nodes of maximum bioluminescence in the regions around the expected source positions.

- The volume enclosed by the distributions. These were estimated by calculating the volume of the regions enclosing the node with maximum luminescence such that that the bioluminescence values within the regions is equal or greater than half of the maximum luminescence.

3.5 Results and discussion

Figures 3.5 to 3.8 show the performance of the CSCG algorithm on a typical CS reconstruction problem in the absence of noise, and in the presence of noise dominated by shot noise. The algorithm performance was measured in terms of the error between the true and reconstructed values (fig. 3.5), the standard deviation of the error between the true and reconstructed values (fig. 3.6), the sensitivity with which the algorithm selects non-zero pixels (fig. 3.7), and the specificity with which the algorithm selects non-zero pixels (fig. 3.8). The pixel error (fig. 3.5a) indicates that the pixel value is underestimated for measurement matrices with few measurements, and that this underestimation is more severe for images with many non-zero pixels. The lower-right half of the plot displays a uniformly low reconstruction error, indicating that CSCG reconstruction is near-exact in this region, and that there are two reconstruction regimes. The first regime is found in the upper left half of the plot where the number of measurements is not significantly larger than the number of unknowns, and within this region the algorithm provides poor reconstruction performance that degrades as the number of measurements is reduced. The second regime is found in the lower right of the plot where the number of measurements is significantly larger than the number of unknowns, and within this region the algorithm provides excellent reconstruction performance independently of the number of measurements. This pattern is supported by the standard deviation (fig. 3.6a), which shows the same two regions and a gradually increasing performance with the number of

measurements in the region of poor performance.

The algorithm sensitivity (fig. 3.7a) displays a sharp transition between the two regimes of poor and exact reconstruction, as seen in CS applications (Donoho and Tanner 2009). This transition indicates that more measurements than non-zero pixels are needed for reliable reconstruction and that the ratio of measurements to non-zero pixels decreases as the number of non-zero pixels increases. For example approximately 100 measurements are necessary for reliable reconstruction of an image with 50 non-zero pixels, resulting in a ratio of 2. This ratio decreases as the number of non-zero pixels increases, and approximately 200 measurements are needed for reliable reconstruction of an image with 150 non-zero pixels, resulting in a ratio of 1.33.

The algorithm specificity (fig. 3.8a) is less informative because at high levels of sparsity the vast majority of the pixels are zero and a sparse reconstruction which selects a small number of non-zero pixels will not greatly affect the ratio of the number of correctly identified zero pixels to the true number of zero pixels whether the reconstruction is correct or entirely incorrect. This experiment demonstrates that CSCG behaviour is consistent with that of a CS algorithm and is optimising for sparsity.

CSCG performance was also tested in the presence of noise. The types of noise were chosen to simulate shot noise in an 16-bit imaging system with a perfect quantum efficiency in addition to an insignificant Gaussian background with a standard deviation of 10 photons. Many scientific imaging devices including the instrument primarily used in this work provide 16-bit measurement or greater so this test is a representative example of the effect of shot noise. The primary CCD camera used in this work (see Appendix A.3) provides 16-bit measurement. It can be seen that performance here is inferior to that of the noise-free case. The standard deviation is not greatly affected, but the average pixel error increases (becomes more negative), and the sensitivity and specificity are degraded. The region in which near-exact reconstruction is possible is reduced. This indicates that it is still possible to reconstruct near-perfectly in the presence of noise, but more measurements are required than in the noise-free case.

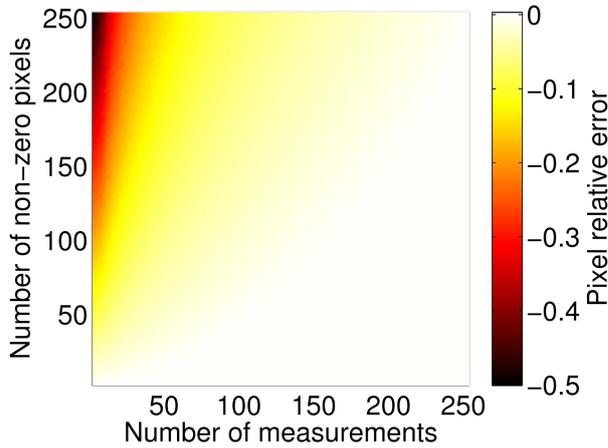
CSCG performance was also tested on the same problem with the addition of a non-zero measurement discrepancy termination criterion. This criterion is that the measurement discrepancies are on average within one standard deviation. The results show a decrease in sensitivity, a decrease in specificity, and an increase in reconstruction error. Pixel standard deviation is largely unchanged.

This pattern is repeated when the measurement discrepancy criterion is increased to two standard deviations for the same problem. Sensitivity and specificity are further reduced. These results suggest that the acquisition of a large set of measurements for reconstruction is particularly important when applying a measurement noise criterion to noisy data.

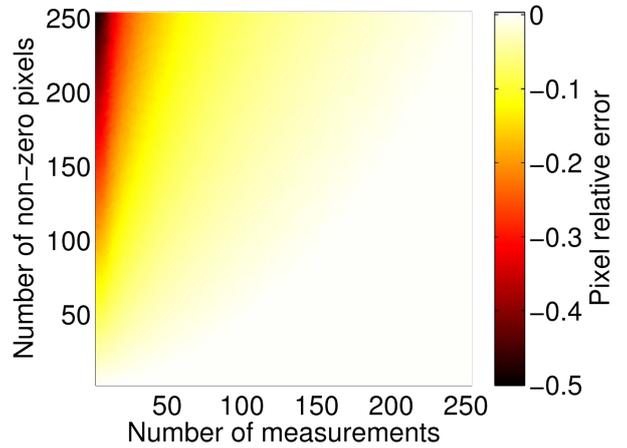
The performance of CSCG was also validated on a two dimensional BLT problem *in silico*. Performance was also examined in the case when the columns of the Jacobian measurement matrix are normalised prior to reconstruction. Column normalisation attempts to remove the bias towards reconstructing shallow sources due to the exponentially decreasing measurement sensitivity with source depth. It was assumed that the bioluminescent source is a single roughly spherical source. The exact shape of the source was determined by the internal structure of the circular BLT mesh. Random bioluminescence sources of strength I with n mesh nodes and a depth d for a mesh with radius R were generated by sampling random angles θ and setting the bioluminescence of the n nodes nearest the point with the polar coordinates (θ, R) to the value I .

Figure 3.9 shows the average node reconstruction error as a function of source size and depth. Algorithm performance is good for depths less than 10mm, but performance is significantly poorer at greater depths. There also appears to be a sharp boundary between regions of good and regions of poor reconstruction performance unlike in the CS case. The position of this boundary shifts when noise is added, but the choice of stopping criterion does not significantly affect average reconstruction error.

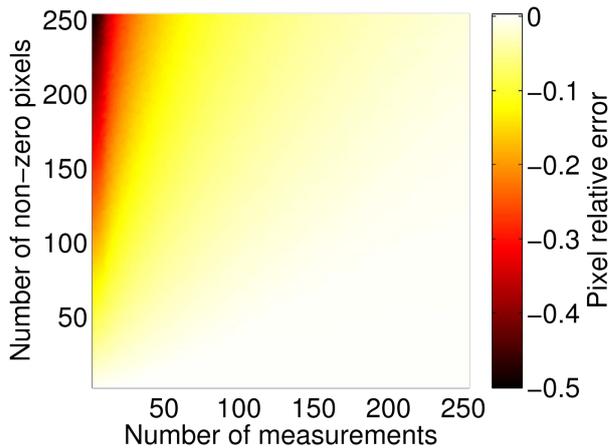
Figure 3.10 shows the average node reconstruction error in the case where columns of the Jacobian matrix have been normalised prior to reconstruction. The reconstruction



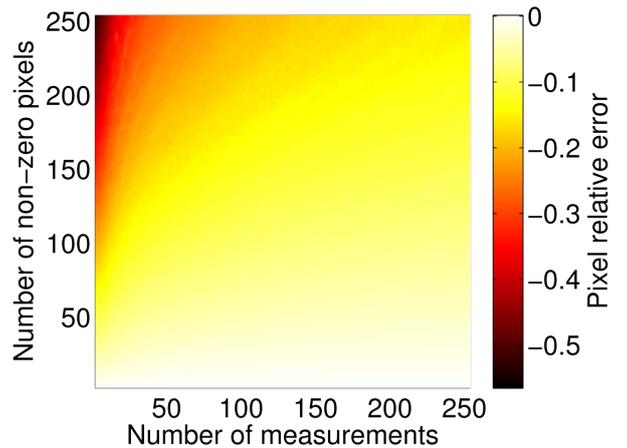
(a) CSCG average pixel error



(b) CSCG average pixel error, with shot noise and termination criterion of 0σ

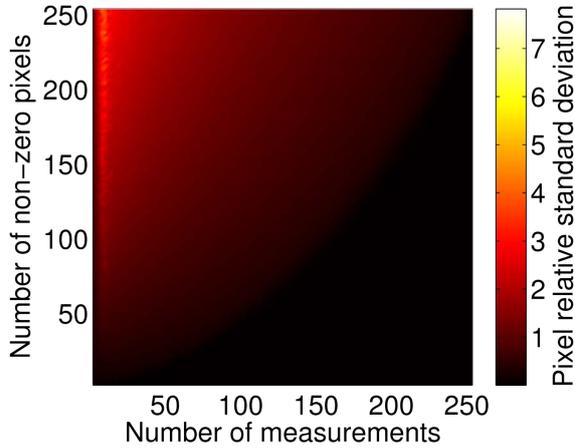


(c) CSCG average pixel error, with shot noise and termination criterion of 1σ

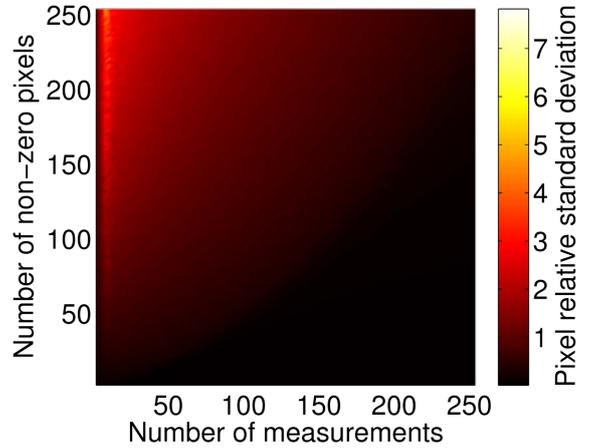


(d) CSCG average pixel error, with shot noise and termination criterion of 2σ

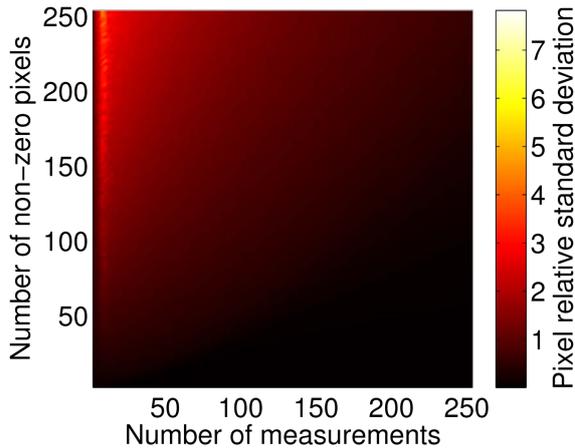
Figure 3.5: CSCG performance on a CS reconstruction problem involving random binary measurement for various degrees of sparsity and measurement set sizes. Twenty random images were reconstructed for every pair of image sparsity and measurement set size. The pixel error averaged over all image pixels and all images for a given image sparsity and measurement set size is shown as a fraction of the true non-zero pixel value.



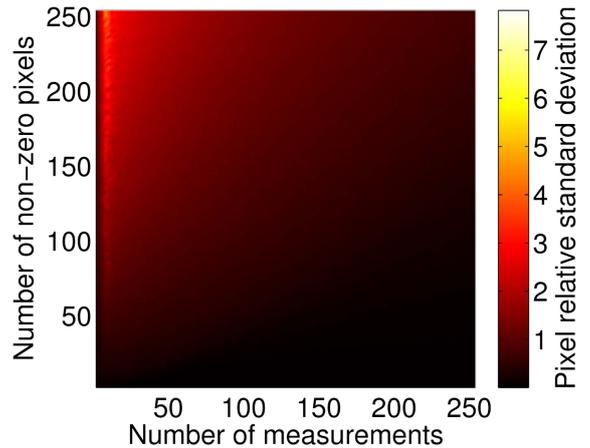
(a) CSCG average pixel standard deviation



(b) CSCG average pixel standard deviation, with shot noise and termination criterion of 0σ

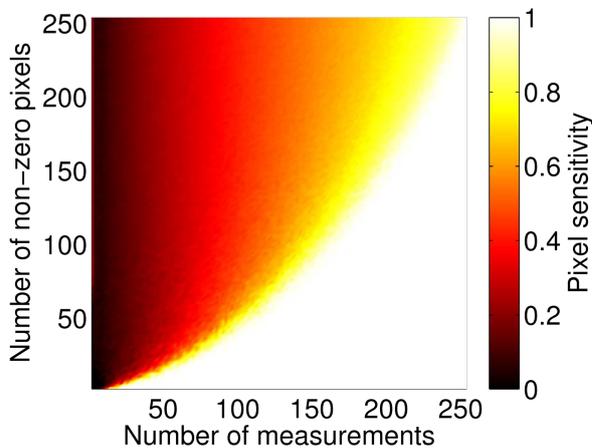


(c) CSCG average pixel standard deviation, with shot noise and termination criterion of 1σ

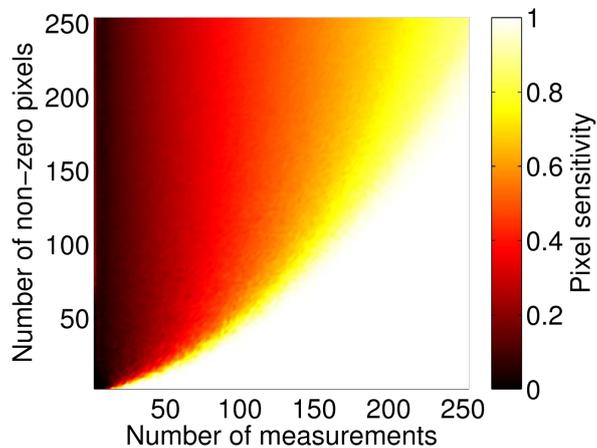


(d) CSCG average pixel standard deviation, with shot noise and termination criterion of 2σ

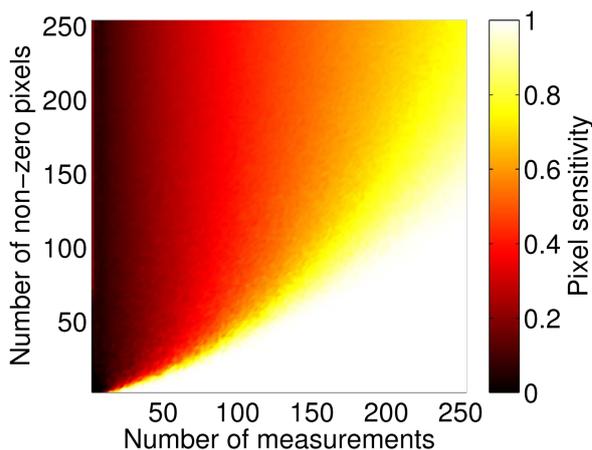
Figure 3.6: CSCG performance on a CS reconstruction problem involving random binary measurement for various degrees of sparsity and measurement set sizes. Twenty random images were reconstructed for every pair of image sparsity and measurement set size. The pixel standard deviation averaged over all images for a given image sparsity and measurement set size is shown as a fraction of the true non-zero pixel value.



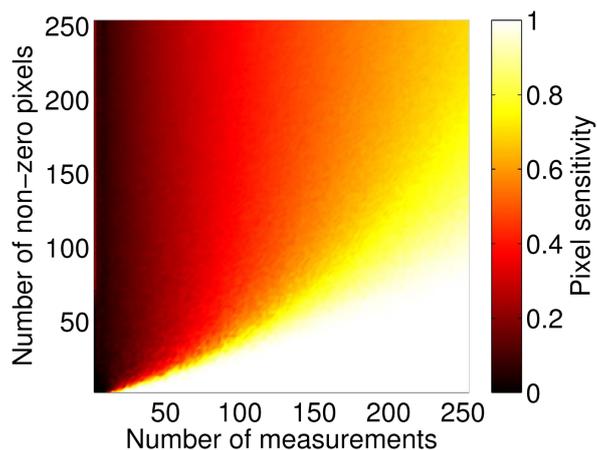
(a) CSCG average pixel sensitivity



(b) CSCG average pixel sensitivity, with shot noise and termination criterion of 0σ

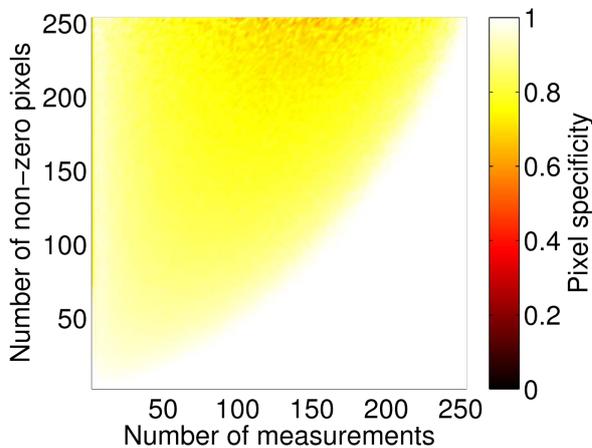


(c) CSCG average pixel sensitivity, with shot noise and termination criterion of 1σ

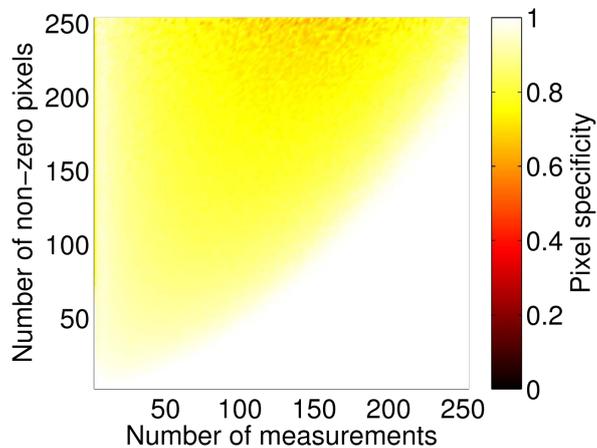


(d) CSCG average pixel sensitivity, with shot noise and termination criterion of 2σ

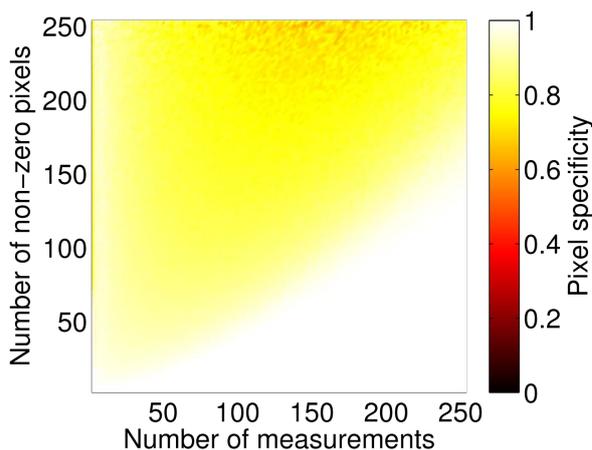
Figure 3.7: CSCG performance on a CS reconstruction problem involving random binary measurement for various degrees of sparsity and measurement set sizes. Twenty random images were reconstructed for every pair of image sparsity and measurement set size. The pixel sensitivity averaged over all images for a given image sparsity and measurement set size is shown as a fraction of the true non-zero pixel value. Pixels were classified into zero pixels and non-zero pixels for the purpose of the sensitivity measure. Non-zero pixels were defined to be pixels that took values greater than or equal to half of the true non-zero pixel value.



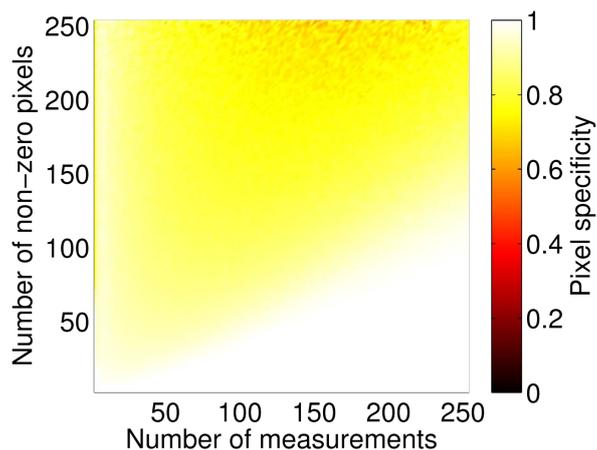
(a) CSCG average pixel specificity



(b) CSCG average pixel specificity, with shot noise and termination criterion of 0σ



(c) CSCG average pixel specificity, with shot noise and termination criterion of 1σ



(d) CSCG average pixel specificity, with shot noise and termination criterion of 2σ

Figure 3.8: CSCG performance on a CS reconstruction problem involving random binary measurement for various degrees of sparsity and measurement set sizes. Twenty random images were reconstructed for every pair of image sparsity and measurement set size. The pixel specificity averaged over all images for a given image sparsity and measurement set size is shown as a fraction of the true non-zero pixel value. Pixels were classified into zero pixels and non-zero pixels for the purpose of the sensitivity measure. Non-zero pixels were defined to be pixels that took values greater than or equal to half of the true non-zero pixel value.

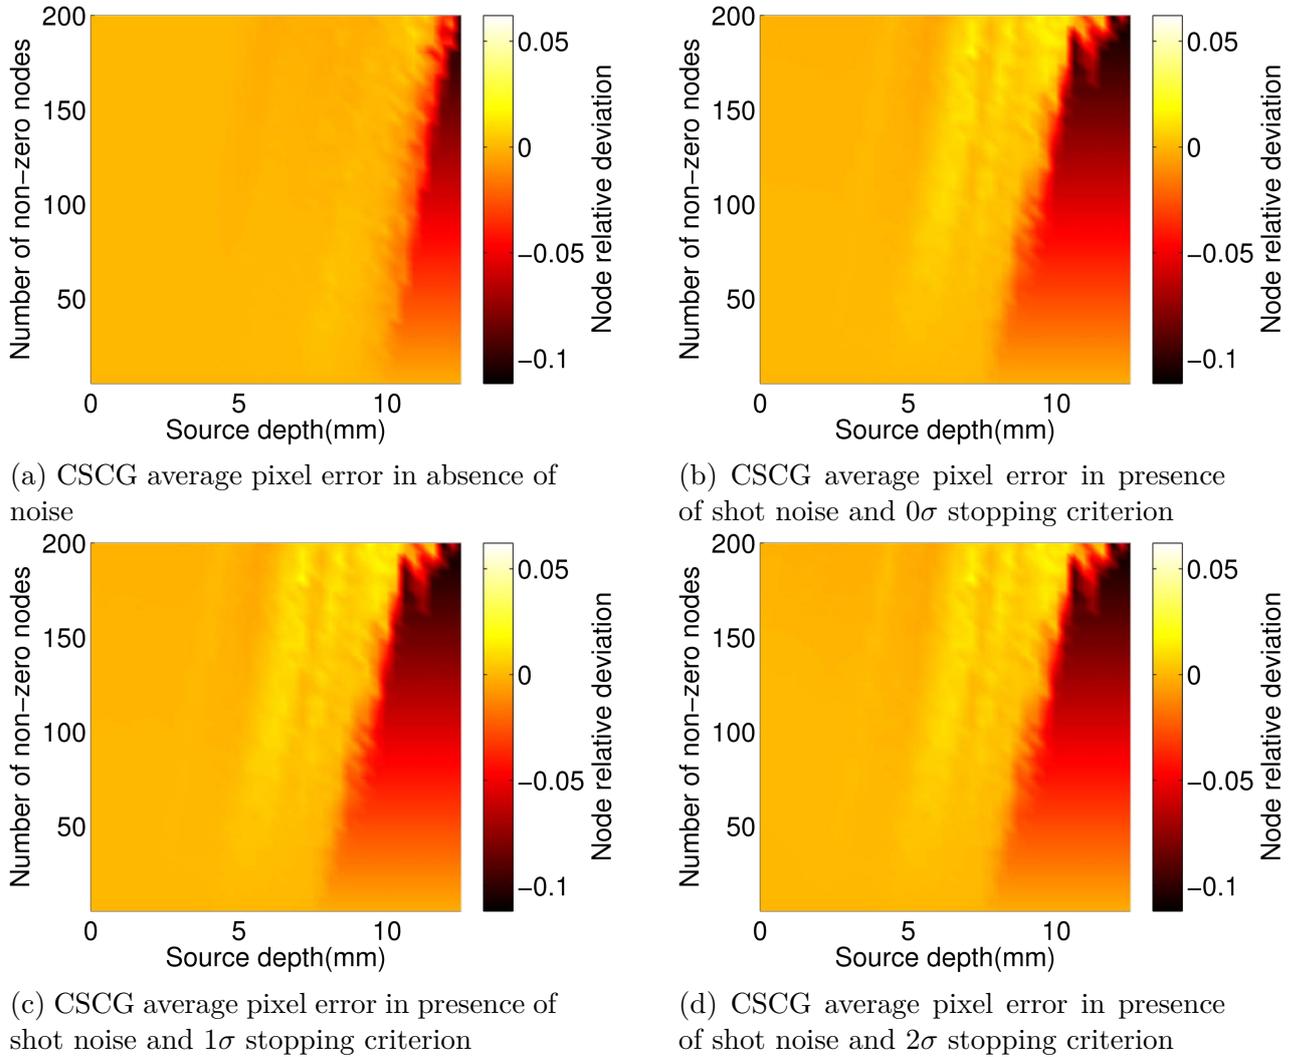


Figure 3.9: CSCG average nodal error on a two-dimensional BLT reconstruction problem for various source sizes and depths. Column normalisation was not applied to the measurement matrix. Twenty images containing randomly placed sources were reconstructed for every pair of source size and source depth.

error is uniformly smaller in this case, and is greatest for shallow sources. This is consistent with the expected effect of column normalisation. The addition of noise increases the error for shallow sources but does not greatly affect deep sources. However, the use of a non-zero stopping criterion has a larger, negative effect on the reconstruction performance in this case.

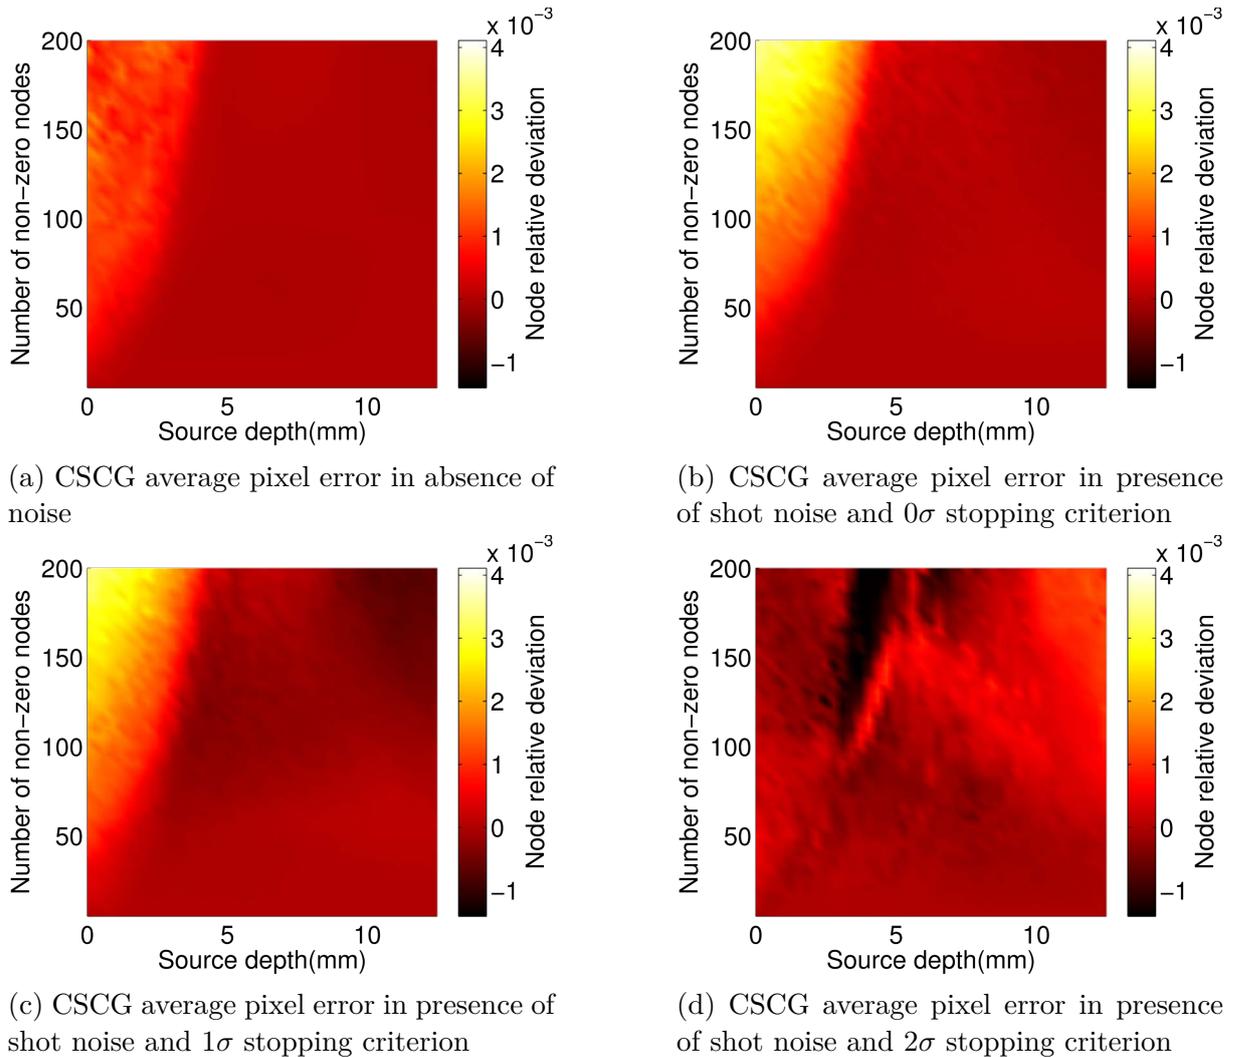


Figure 3.10: CSCG average nodal error on a two-dimensional BLT reconstruction problem for various source sizes and depths. Column normalisation was applied to the measurement matrix. Twenty images containing randomly placed sources were reconstructed for every pair of source size and source depth.

The average standard deviation of reconstructions is shown in fig. 3.11. There are structures visible in the visualisation which may arise from the internal structure of the mesh. Algorithm performance is dissimilar to the CS case. Again, reconstruction per-

formance is superior for shallow, large sources based on this metric. The addition of noise counter-intuitively appears to reduce the standard deviation which is a metric of the consistency of reconstruction. This metric is not significantly affected by the stopping criterion.

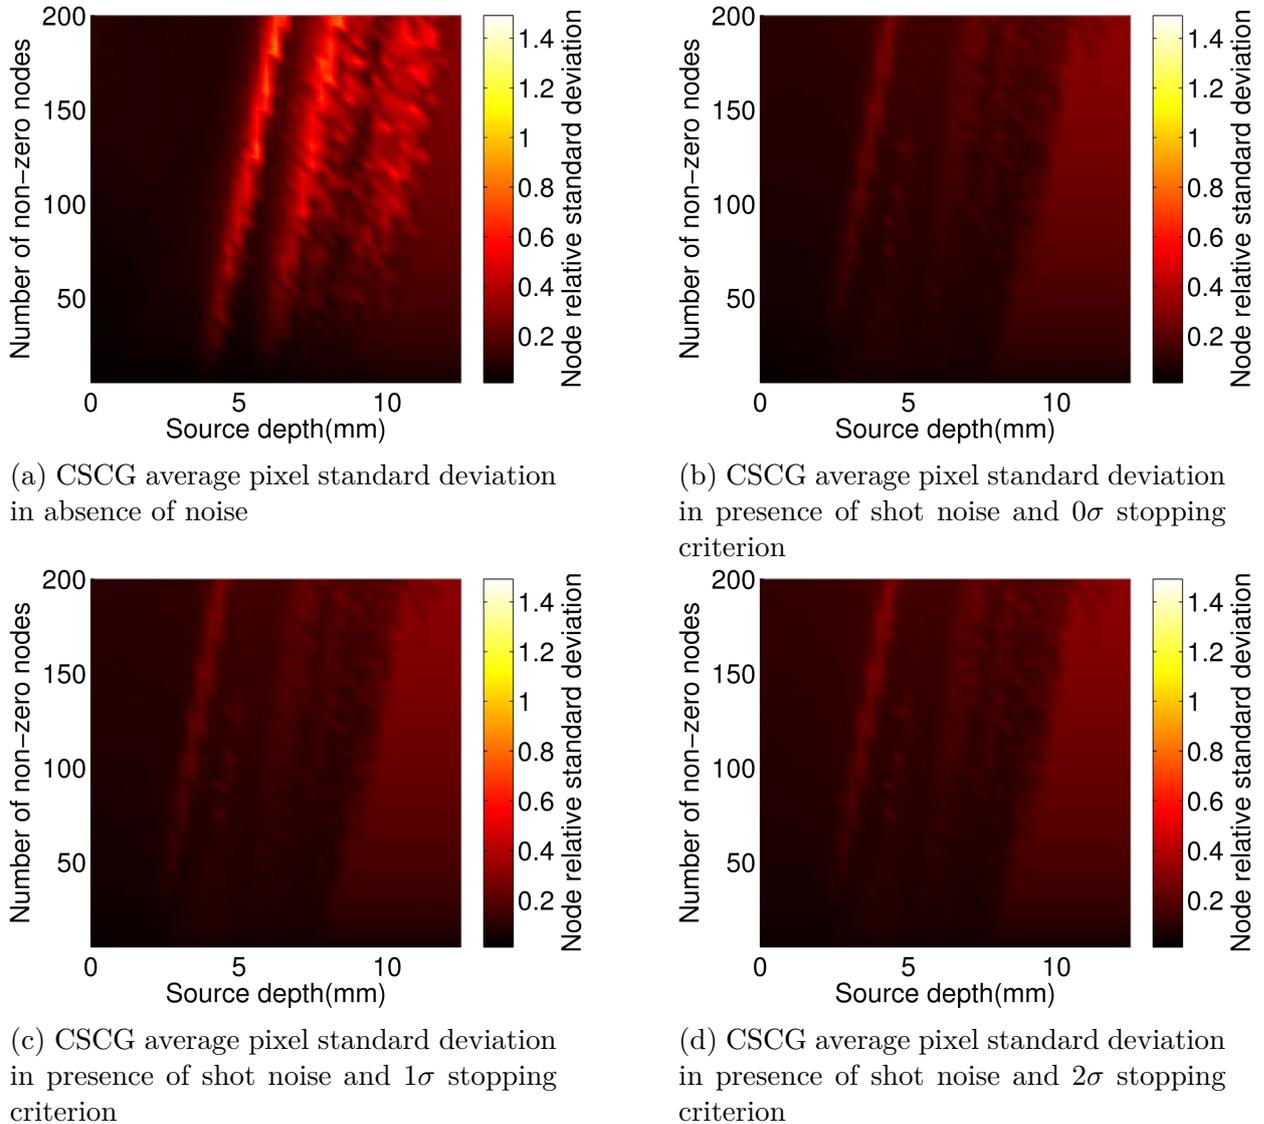


Figure 3.11: CSCG average nodal standard deviation on a two-dimensional BLT reconstruction problem for various source sizes and depths. Column normalisation was not applied to the measurement matrix. Twenty images containing randomly placed sources were reconstructed for every pair of source size and source depth.

The average standard deviation of reconstructions with the addition of column normalisation is shown in fig. 3.12. The visible trends differ from the case without column normalisation (fig. 3.11) and again the results show superior performance for deep sources

and inferior performance for shallow sources. The addition of noise affects reconstruction accuracy but does not change the general trends. The performance with the 2σ stopping criterion however is appreciably different to the other cases.

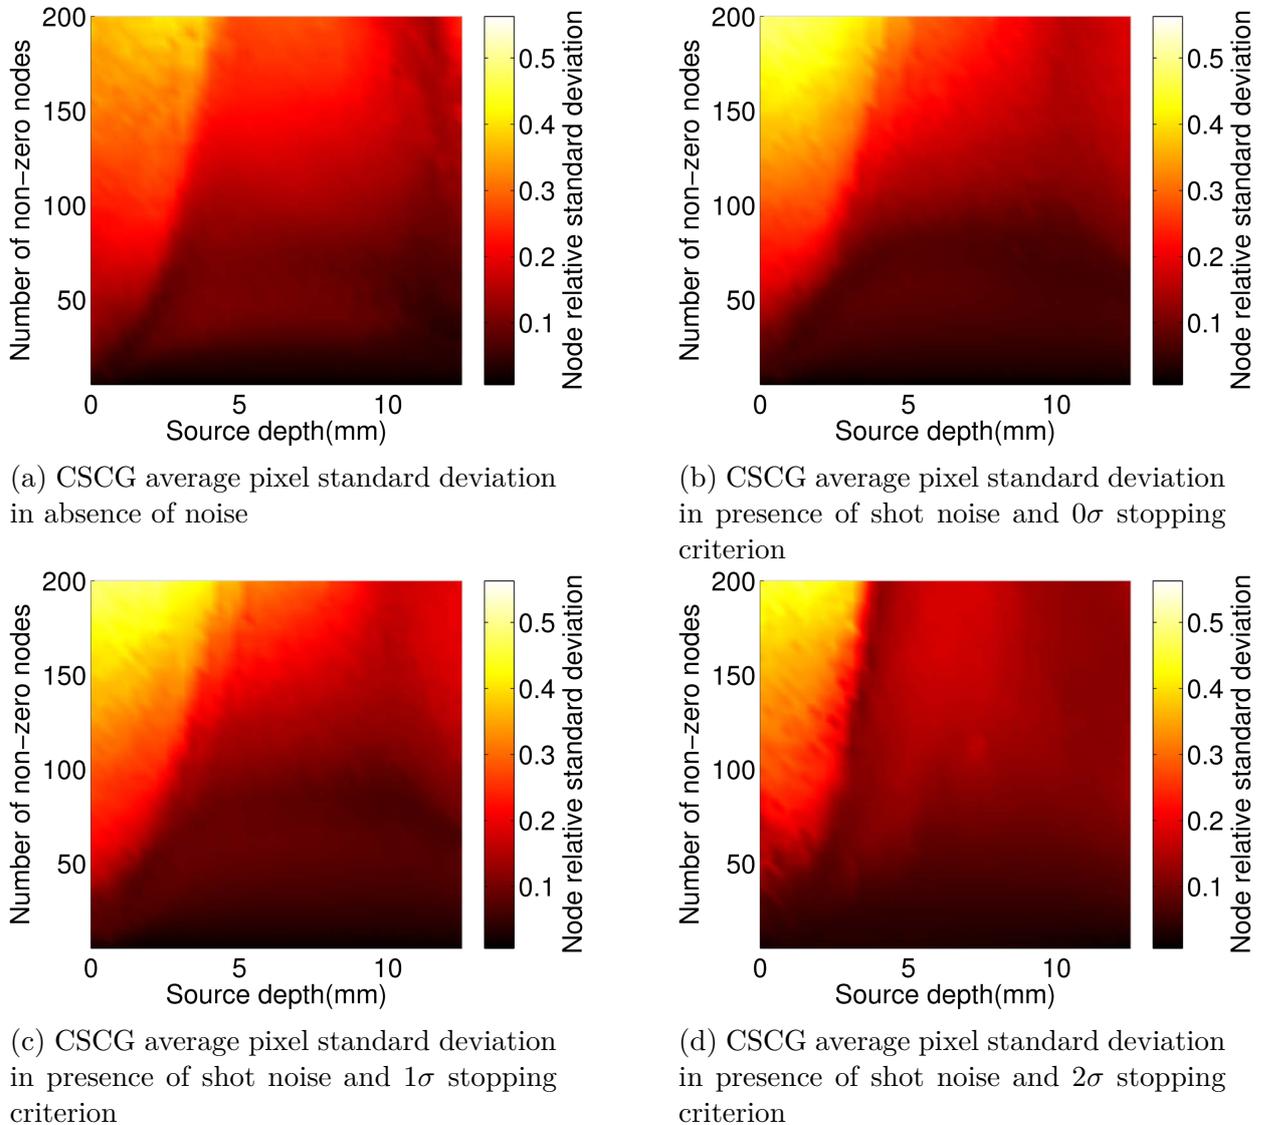


Figure 3.12: CSCG average nodal standard deviation on a two-dimensional BLT reconstruction problem for various source sizes and depths. Column normalisation was applied to the measurement matrix. Twenty images containing randomly placed sources were reconstructed for every pair of source size and source depth.

Figure 3.13 shows the algorithm sensitivity in the absence of column normalisation. Here again structures are visible, and algorithm performance drops dramatically to zero as the source depth, and size to an extent, go beyond a threshold. The algorithm reconstructs with sensitivity of approximately 0.8 when source depth is sufficiently low and

source size is sufficiently high. The addition of noise leads to a decrease in the variability of reconstruction performance on source depth and size. However the position of the threshold at which performance drops dramatically does not change.

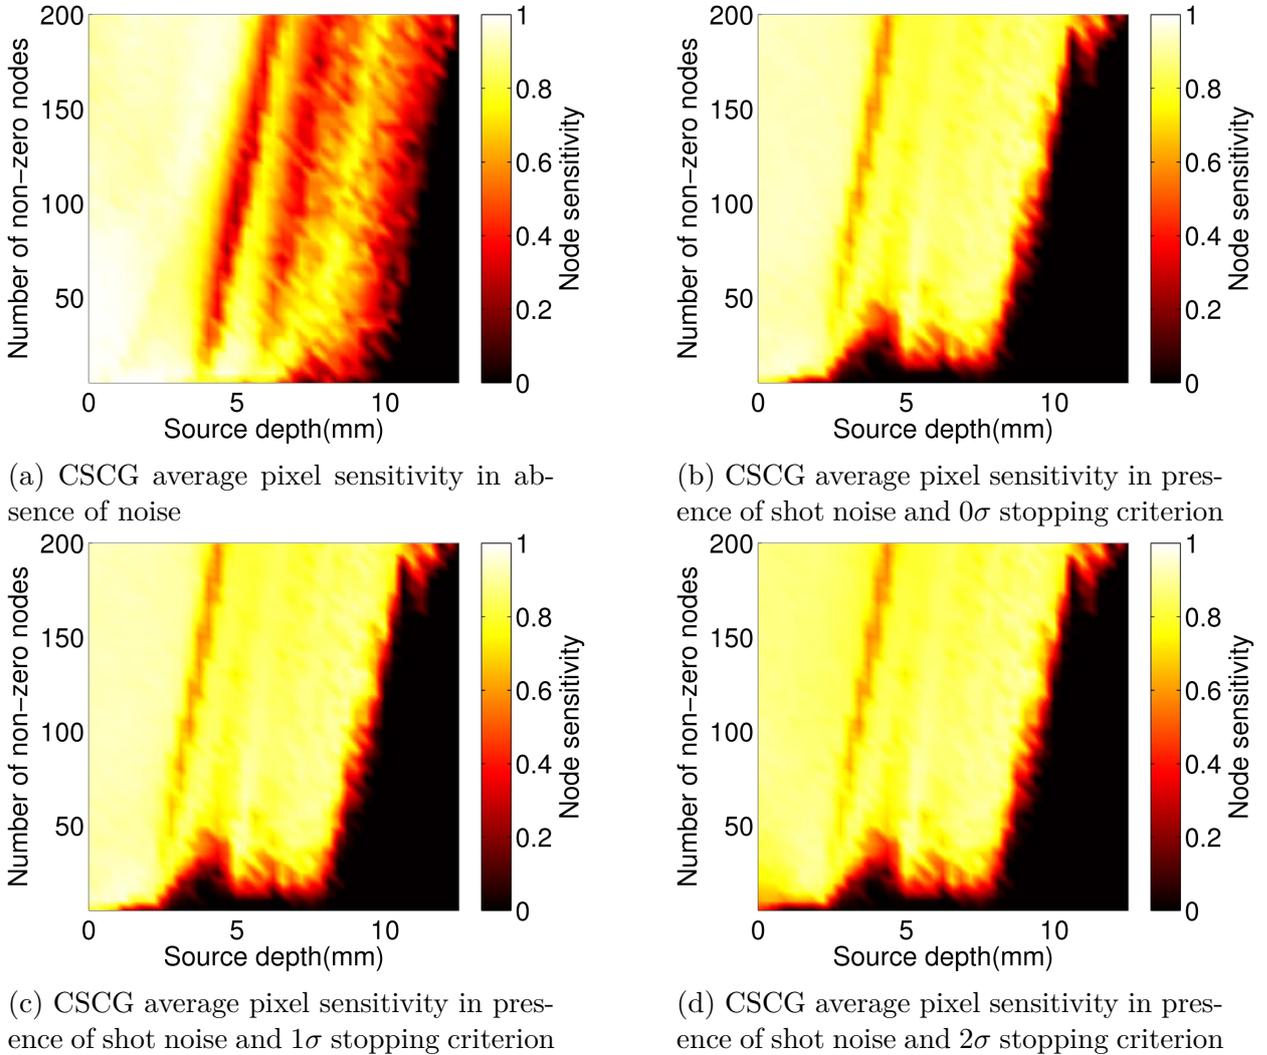


Figure 3.13: CSCG average nodal sensitivity on a two-dimensional BLT reconstruction problem for various source sizes and depths. Column normalisation was not applied to the measurement matrix. Twenty images containing randomly placed sources were reconstructed for every pair of source size and source depth.

The addition of column normalisation also dramatically changes the algorithm sensitivity, as can be seen in fig. 3.14. Here the reconstruction sensitivity is above or in the region of 0.4 everywhere, for both shallow and deep sources. Performance is however superior for deeper sources. The noise free case also shows a region of poor reconstruction performance for sources with sizes of between 20 and 50 nodes and depths of between 3

and 10mm which is not present in algorithm performance in the presence of noise. The use of a non-zero stopping criterion does not significantly affect reconstruction performance when the criterion is 1σ , whereas the use of 2σ does produce significant changes. The noise free, and noisy 0σ and 1σ cases also possess small regions of near-perfect sensitivity for very small source sizes.

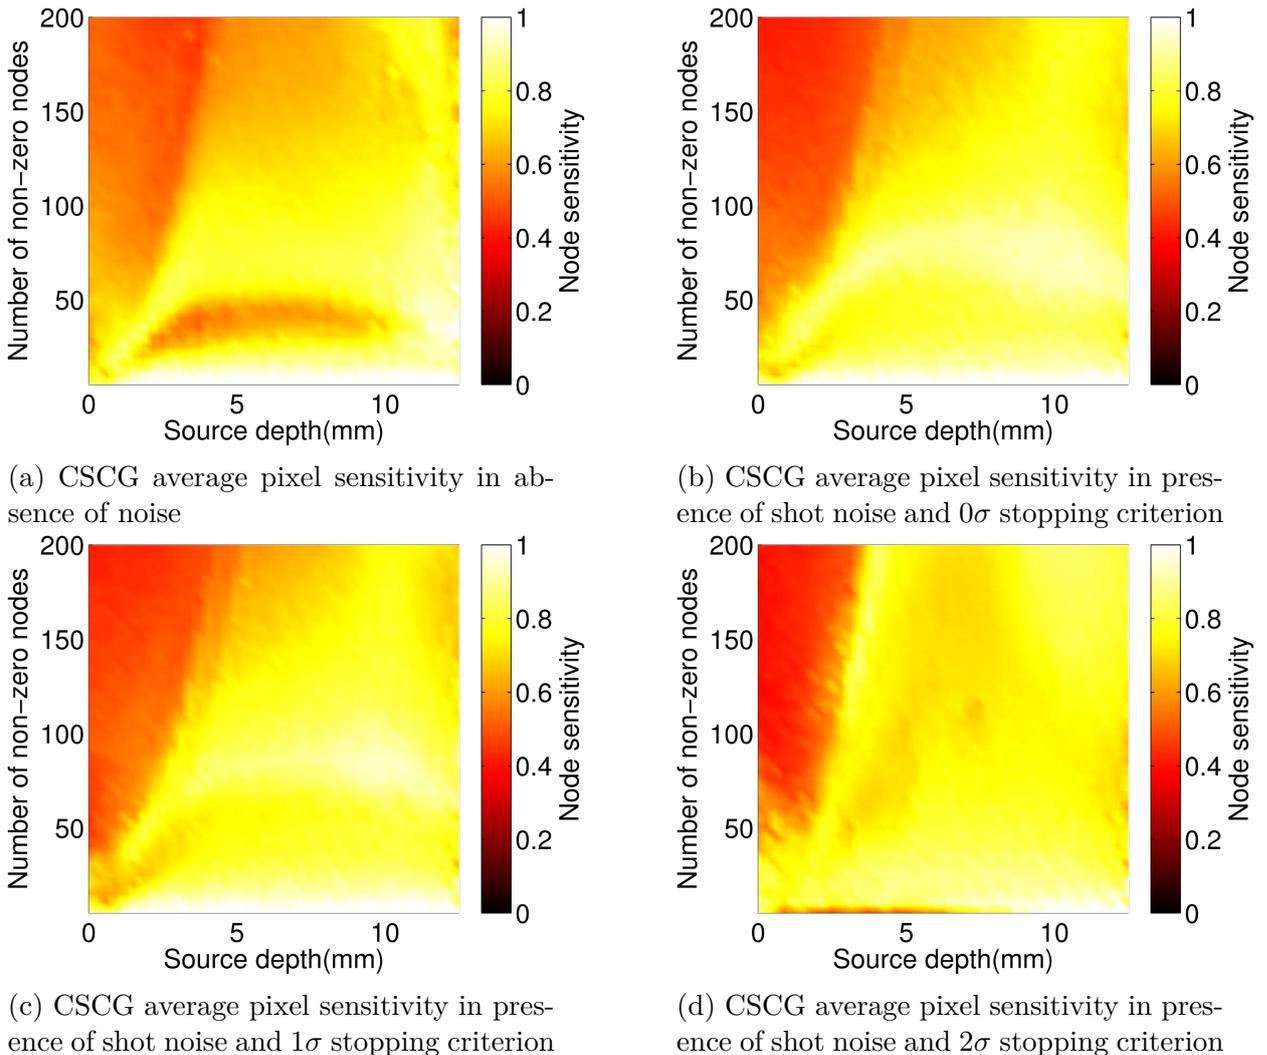


Figure 3.14: CSCG average nodal sensitivity on a two-dimensional BLT reconstruction problem for various source sizes and depths. Column normalisation was applied to the measurement matrix. Twenty images containing randomly placed sources were reconstructed for every pair of source size and source depth.

The algorithm specificity for reconstruction with and without column normalisation is uniformly near-perfect. Consequently these results have not been shown in figure form because no contrast would be observable.

Figure 3.15 shows the mean total bioluminescence error as a fraction of the true value, without using column normalisation. The results show a similar distribution to that of the sensitivity. Shallow sources are reconstructed well, but deep sources are reconstructed very poorly and there is a sharp transition between the regimes of good and poor reconstruction.

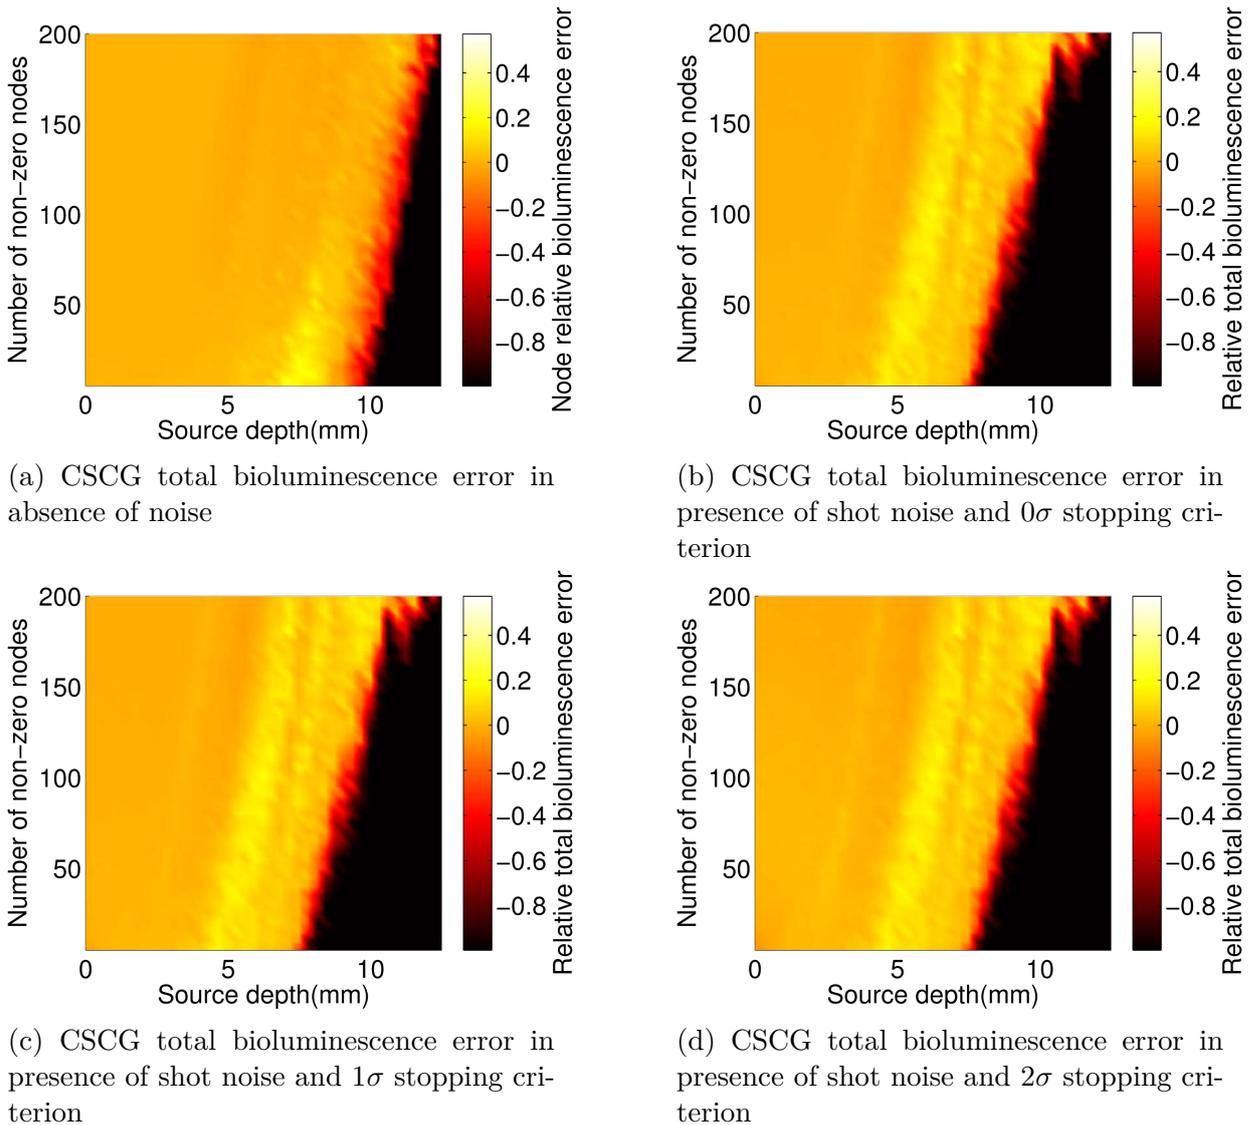


Figure 3.15: CSCG total bioluminescence error on a two-dimensional BLT reconstruction problem for various source sizes and depths. Column normalisation was not applied to the measurement matrix. Twenty images containing randomly placed sources were reconstructed for every pair of source size and source depth.

Figure 3.16 also shows the mean total bioluminescence error in the case where column normalisation is used. The mean reconstructed total bioluminescence is within 10% of

the true total bioluminescence for sources of all size and depth, despite the sensitivity of approximately 0.4 for shallow sources (fig. 3.14).

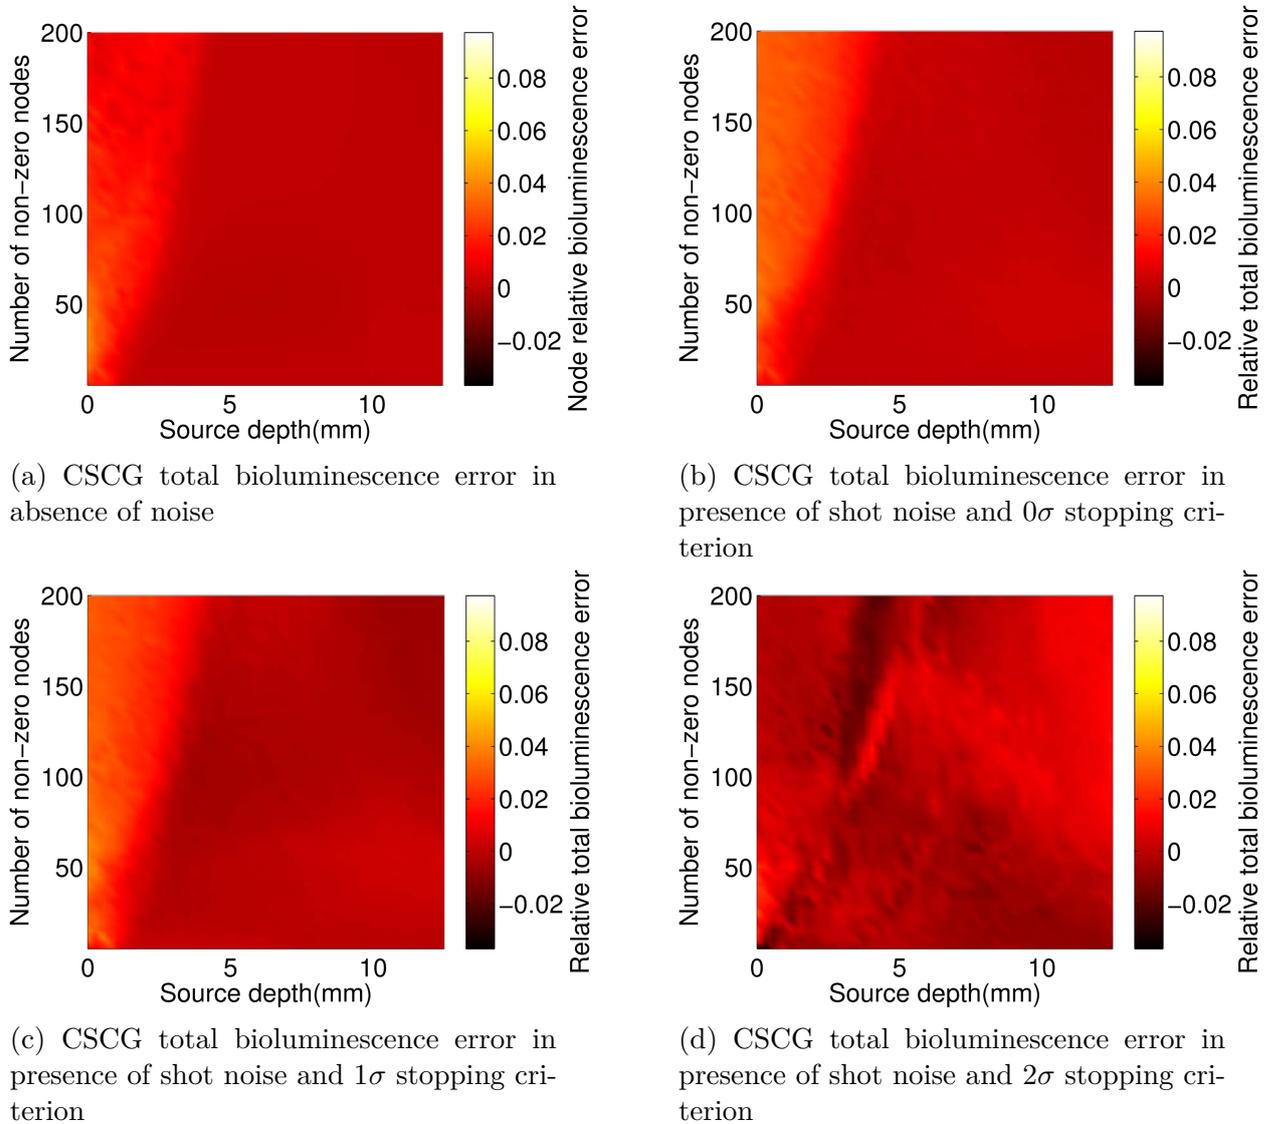


Figure 3.16: CSCG total bioluminescence error on a two-dimensional BLT reconstruction problem for various source sizes and depths. Column normalisation was applied to the measurement matrix. Twenty images containing randomly placed sources were reconstructed for every pair of source size and source depth.

Figure 3.17 shows an example of GN, NNLS, and CSCG reconstruction in the absence of measurement noise. The images are a projection of the three dimensional bioluminescence distributions. All reconstruction algorithms successfully separate the two sources. The shape of the GN reconstruction (fig. 3.17b) does not agree well with the target (fig. 3.17a) and each source is split into a number of smaller sources with poor locali-

sation. The shape of the NNLS reconstruction (fig. 3.17c) is in better agreement with that of the target, but the sources are less spatially compact than the target and so the intensities are reduced. The CSCG reconstruction (fig. 3.17d) displays superior source compactness and intensity.

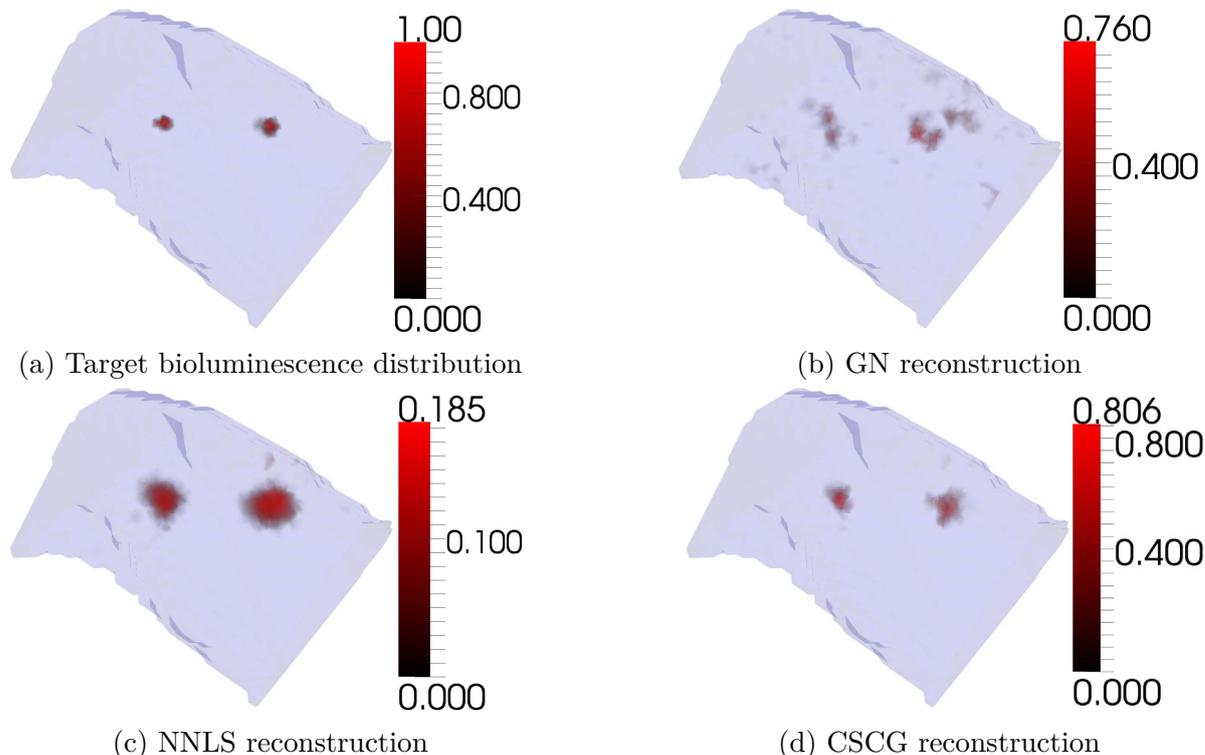


Figure 3.17: Reconstruction performance *in silico* in the absence of measurement noise. Images from Basevi et al. (2012b).

Figure 3.18 shows an example of GN, NNLS, and CSCG reconstruction in the presence of 1% Gaussian measurement noise. Here NNLS and CSCG successfully separate the two sources. GN performance is poor (fig. 3.18b) and it is not possible to visually distinguish the sources from reconstruction artefacts. The shape of the NNLS reconstruction (fig. 3.18c) is again in good agreement with that of the target, but the sources are similarly less spatially compact and the intensities are reduced. The addition of 1% noise has not significantly affected the reconstruction. The CSCG reconstruction (fig. 3.18d) also behaves similarly to the noise-free case, but the source compactness is slightly reduced.

Figure 3.19 shows an example of GN, NNLS, and CSCG reconstruction in the presence of 5% Gaussian measurement noise. The performance of all algorithms is visibly affected

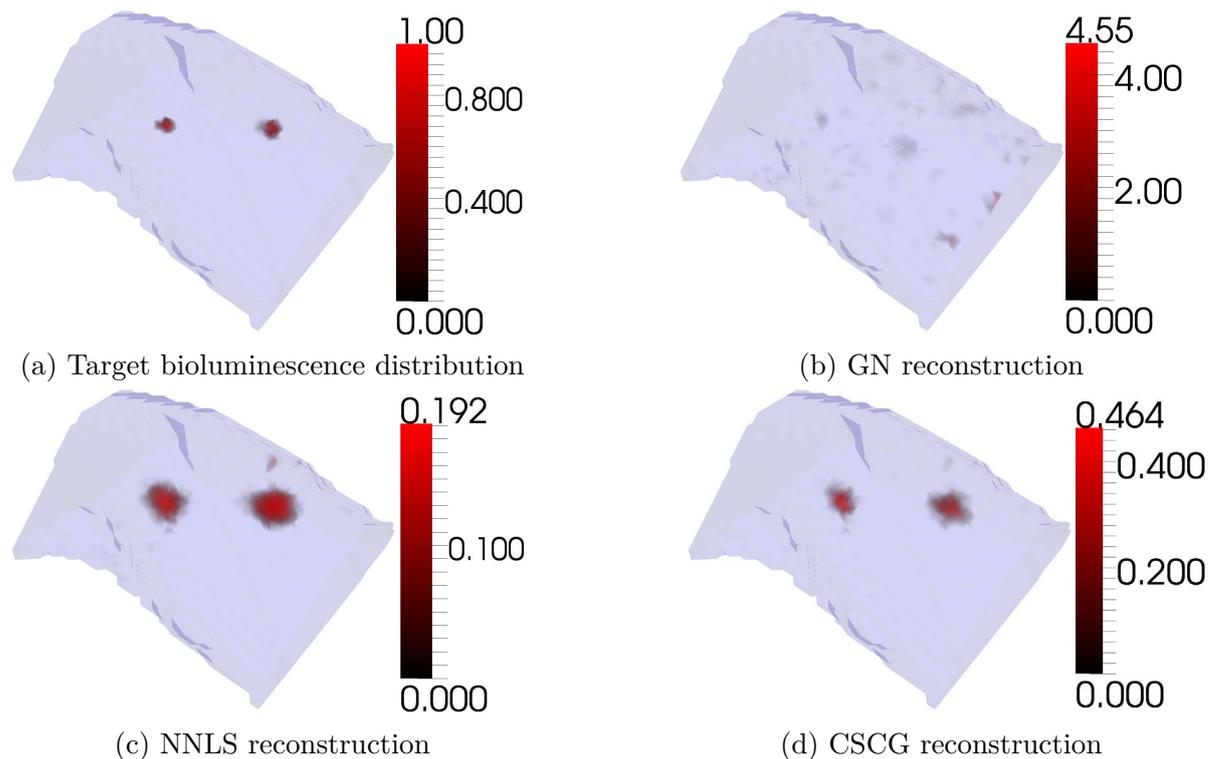


Figure 3.18: Reconstruction performance *in silico* in the presence of 1% Gaussian measurement noise. Images from Basevi et al. (2012b).

by the degree of noise present. GN performance is poor (fig. 3.19b) and again it is not possible to visually distinguish the sources from reconstruction artefacts. The NNLS reconstruction (fig. 3.19c) is in reasonable agreement with that of the target, but the source shapes have become significantly deformed. The same applies to the CSCG reconstruction (fig. 3.19d), which is also more compact than the NNLS reconstruction.

Tables 3.1 to 3.3 show the localisation error of the two sources. Source A is the source shown on the left in figs. 3.17 to 3.19, and source B is the source on the right. Source A is nearer to the surface and source B is deeper.

Table 3.1 shows the distance between the location of the true source and the average of the reconstructed location over the 30 reconstructions using different noise samples. The GN algorithm performance is universally inferior to that of the NNLS and CSCG algorithms. CSCG performance is superior to that of NNLS, although the difference is small. The error in three dimensions along with the standard deviation can be found in tables 3.2 and 3.3.

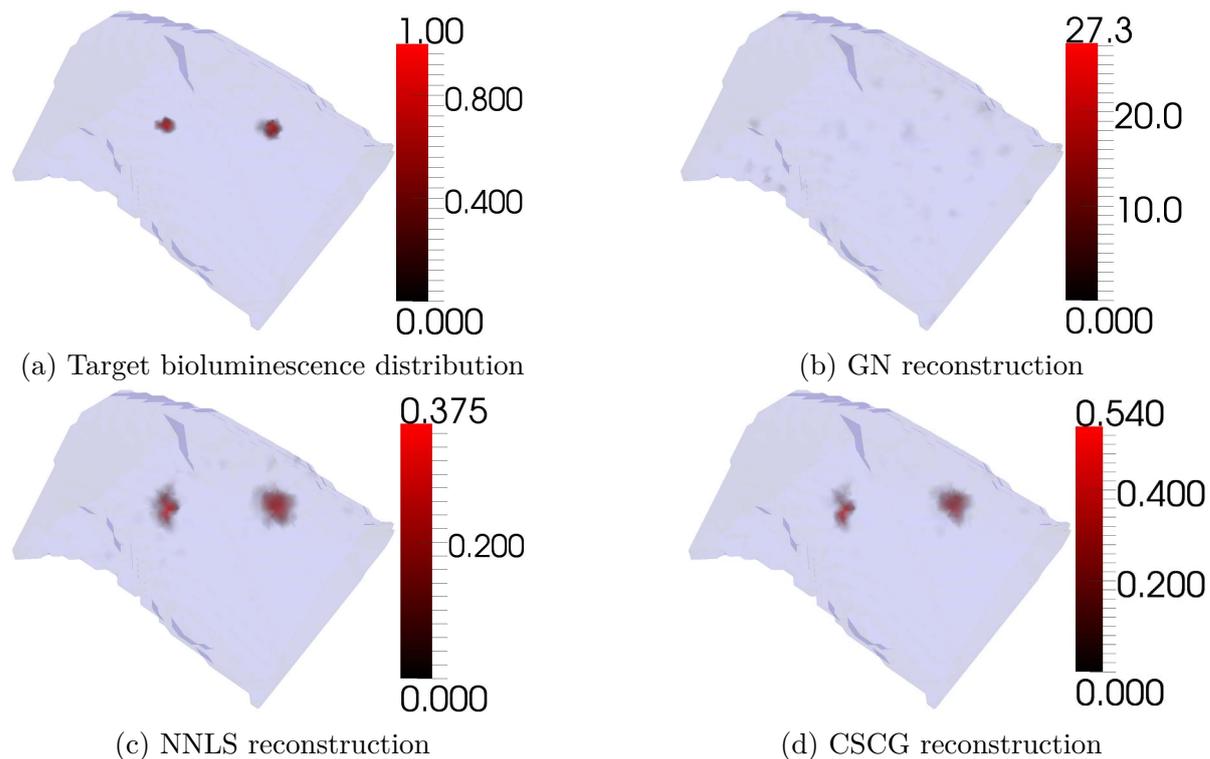


Figure 3.19: Reconstruction performance *in silico* in the presence of 5% Gaussian measurement noise. Images from Basevi et al. (2012b).

Table 3.2 shows the error in three dimensions for source A. Here again the GN algorithm performance is uniformly inferior to that of the other algorithms. The results for NNLS and CSCG are similar, but suggest that error in the z-axis is greater than that for the x-axis and y-axis. This may be a result of the asymmetry of the volume of interest.

The results for source B (table 3.3) are similar. The GN algorithm performance is inferior to that of NNLS and CSCG. NNLS and CSCG performance is similar, but z-axis localisation is inferior to x-axis and y-axis localisation, particularly in the case of high noise.

Table 3.4 contains the average volume of the reconstructed sources, as a percentage of the true value. The volume is a measure of compactness, and here CSCG outperforms NNLS. The GN results are unreliable in the context of its poor localisation ability, and CSCG reconstruction produces volumes that are closer to the true value than NNLS except for the case of source A with 5% noise.

Figure 3.20 shows the result of a reconstruction of *in simulacra* data acquired using

Noise	Source A			Source B		
	GN	NNLS	CSCG	GN	NNLS	CSCG
0%	1.8	0.6	0.6	3.6	1.6	0.9
1%	4.0	1.0	0.5	2.9	0.9	0.7
5%	5.3	1.7	1.7	2.8	1.7	1.1

Table 3.1: Simulation localisation error (mm) of mean reconstructed source centre. Source locations were chosen to be the locations of the nodes of maximum bioluminescence within a region around the expected source location. Source A is the left-most source, and source B is the right-most source, as displayed in figs. 3.17 to 3.19. At each noise level, 30 samples of noisy measurements were generated. The location of each source was taken to be the location of the maximum value in a region (a sphere of radius 8mm) around the true source centre. Table from Basevi et al. (2012b).

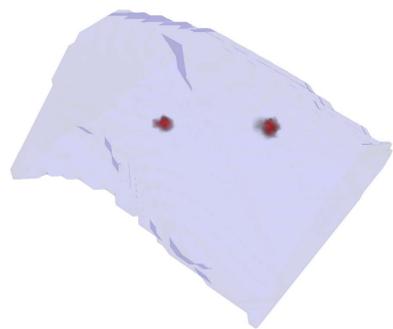
the Dartmouth College instrument. Figure 3.20a shows approximate expected bioluminescence distribution. The source locations are provided by the manufacturer but the source sizes are unknown. The source locations are also potentially subject to a degree of error due the unknown optical fibre placement and orientation. Reconstruction intensities are non quantitative due to lack of information about the measurement process.

Reconstruction performance is uniformly inferior to the *in silico* experiments. The GN algorithm does not reconstruct sources near the expected locations and the bioluminescence distribution is dominated by artefacts. Both NNLS and CSCG reconstruct two separate sources in approximately the right locations, but both contain noise artefacts. The CSCG reconstruction is visibly more compact than the NNLS reconstruction.

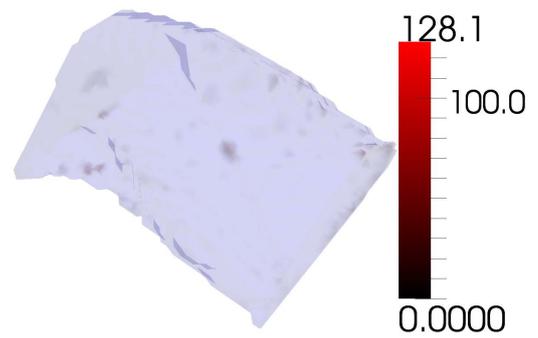
3.6 Conclusions

The CSCG algorithm was developed to allow reconstruction of BLT problems where bioluminescent sources are expected to be spatially sparse. This is performed through a reconstruction technique that balances reconstruction sparsity against measurement discrepancy in a way that incorporates information on measurement noise and automatically determines the correct sparsity weighting for the reconstruction process.

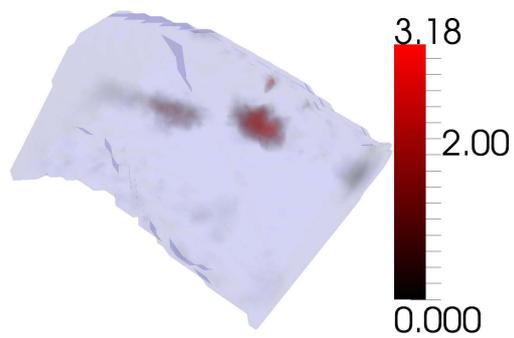
The algorithm was tested on a CS measurement problem in the presence and absence of



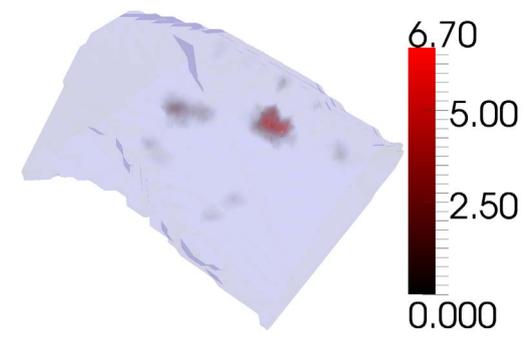
(a) Approximate target bioluminescence distribution



(b) GN reconstruction



(c) NNLS reconstruction



(d) CSCG reconstruction

Figure 3.20: Reconstruction performance *in simulacra* using data acquired via an instrument at Dartmouth College. Images from Basevi et al. (2012b).

		Source A								
Noise		GN			NNLS			CSCG		
0%		$\begin{bmatrix} 0.5 \\ 1.3 \\ -1.1 \end{bmatrix}$			$\begin{bmatrix} -0.5 \\ -0.2 \\ -0.1 \end{bmatrix}$			$\begin{bmatrix} -0.5 \\ -0.2 \\ -0.1 \end{bmatrix}$		
1%		$\begin{bmatrix} 0.2 \\ -0.5 \\ -4.0 \end{bmatrix}$	\pm	$\begin{bmatrix} 1.6 \\ 2.3 \\ 1.4 \end{bmatrix}$	$\begin{bmatrix} 0.0 \\ 0.1 \\ -1.0 \end{bmatrix}$	\pm	$\begin{bmatrix} 0.4 \\ 0.9 \\ 0.4 \end{bmatrix}$	$\begin{bmatrix} -0.3 \\ 0.2 \\ -0.4 \end{bmatrix}$	\pm	$\begin{bmatrix} 0.4 \\ 0.6 \\ 0.5 \end{bmatrix}$
5%		$\begin{bmatrix} -0.6 \\ -0.8 \\ -5.2 \end{bmatrix}$	\pm	$\begin{bmatrix} 2.6 \\ 3.4 \\ 1.0 \end{bmatrix}$	$\begin{bmatrix} 0.1 \\ -0.1 \\ -1.7 \end{bmatrix}$	\pm	$\begin{bmatrix} 0.9 \\ 0.5 \\ 0.4 \end{bmatrix}$	$\begin{bmatrix} 0.5 \\ -0.0 \\ -1.7 \end{bmatrix}$	\pm	$\begin{bmatrix} 0.8 \\ 0.5 \\ 0.3 \end{bmatrix}$

Table 3.2: Simulation localisation error and standard deviation (mm) of reconstructed source A centre. Source locations were chosen to be the locations of the nodes of maximum bioluminescence within a region around the expected source location. Source A is the left-most source as displayed in figs. 3.17 to 3.19. At each noise level, 30 samples of noisy measurements were generated. The location of each source was taken to be the location of the maximum value in a region (a sphere of radius 8mm) around the true source centre.

shot noise-dominated measurement noise. The results indicated that the CSCG algorithm behaves as a CS algorithm and that the minimum measurement set size required for near-exact imaging increases significantly in the presence of noise and in particular when using measurement discrepancy to determine a final sparsity weighting. This demonstrated the capabilities of the algorithm but it should be noted that BLT measurement matrices do not satisfy the criteria of CS due to the attenuating properties of tissue resulting from absorption and scattering.

Absorption causes measurement sensitivity to reduce exponentially with source depth in tissue with respect to the detector. Scattering causes measurements to become dependent, which limits the quantity of information that can realistically be obtained in an experiment. Increasing the measurement density on the subject is of limited use because increased measurement density necessitates decreasing detector-detector distances, which decreases measurement independence due to scattering.

The algorithm was then tested on a BLT problem *in silico*. Reconstruction quality was good but not perfect in the noise-free case. It was expected that the reconstruction quality would be imperfect because measurements were generated on a higher resolution

		Source B					
Noise	GN		NNLS		CSCG		
0%	$\begin{bmatrix} -3.5 \\ 1.0 \\ -0.1 \end{bmatrix}$		$\begin{bmatrix} 1.5 \\ 0.3 \\ -0.2 \end{bmatrix}$		$\begin{bmatrix} -0.5 \\ -0.2 \\ -0.7 \end{bmatrix}$		
1%	$\begin{bmatrix} -2.9 \\ -0.1 \\ 0.5 \end{bmatrix}$	$\pm \begin{bmatrix} 2.9 \\ 1.8 \\ 1.2 \end{bmatrix}$	$\begin{bmatrix} 0.9 \\ 0.1 \\ -0.1 \end{bmatrix}$	$\pm \begin{bmatrix} 0.5 \\ 0.2 \\ 1.6 \end{bmatrix}$	$\begin{bmatrix} 0.5 \\ -0.1 \\ -0.6 \end{bmatrix}$	$\pm \begin{bmatrix} 0.7 \\ 0.2 \\ 0.4 \end{bmatrix}$	
5%	$\begin{bmatrix} -1.1 \\ -0.5 \\ 2.5 \end{bmatrix}$	$\pm \begin{bmatrix} 5.1 \\ 3.0 \\ 2.3 \end{bmatrix}$	$\begin{bmatrix} 0.3 \\ -0.0 \\ 1.7 \end{bmatrix}$	$\pm \begin{bmatrix} 1.1 \\ 0.7 \\ 2.6 \end{bmatrix}$	$\begin{bmatrix} 0.6 \\ -0.3 \\ 0.9 \end{bmatrix}$	$\pm \begin{bmatrix} 1.2 \\ 0.8 \\ 1.4 \end{bmatrix}$	

Table 3.3: Simulation localisation error and standard deviation (mm) of reconstructed source B centre. Source locations were chosen to be the locations of the nodes of maximum bioluminescence within a region around the expected source location. Source B is the right-most source as displayed in figs. 3.17 to 3.19. At each noise level, 30 samples of noisy measurements were generated. The location of each source was taken to be the location of the maximum value in a region (a sphere of radius 8mm) around the true source centre.

Noise	Source A			Source B		
	GN	NNLS	CSCG	GN	NNLS	CSCG
0%	43%	340%	74%	86%	740%	96%
1%	36%	290%	110%	72%	660%	150%
5%	26%	82%	57%	58%	330%	230%

Table 3.4: Simulation mean volume (mm^3) (see Table 3.1) as a percentage of the true volume. The volume of each source was taken to be the connected volume enclosed by half the maximum value in the region, containing the location of the maximum value. Table from Basevi et al. (2012b).

mesh than that used for reconstruction. This does not allow the initial distribution to be represented exactly in the lower resolution basis, and potentially adds error to the physical model of light propagation. The localisation ability of both NNLS and CSCG was good, but the NNLS algorithm reconstructed significantly larger and less intense sources than those reconstructed by CSCG. This result was positive and indicates that it is possible to accurately reconstruct two bioluminescent sources simultaneously.

Noise was then added to the *in silico* measurements to simulate a more realistic scenario. In the presence of 1% Gaussian noise the reconstruction quality of CSCG was degraded but still acceptable. Source size was larger than in the noise-free case but the localisation was still good. NNLS reconstruction was similarly affected, and GN reconstruction became very unreliable. In the presence of 5% Gaussian noise the shape of the reconstructed sources for NNLS and CSCG were distorted and source size was also affected, but localisation was still good and both sources could be clearly distinguished. 5% noise is higher than would be expected for the system discussed in earlier chapters, but the noise level of the Dartmouth system is unknown.

Finally, the algorithms were used to reconstruct the *in simulacra* measurements from the Dartmouth system. These measurements are subject to a great amount of uncertainty, including the measurement placement, coupling with the surface, and measurement noise, and the physical model/mesh generation is subject to error in the segmentation of the CT data, so it was expected that the results would be inferior to the *in silico* reconstructions. In practice, both NNLS and CSCG were able to resolve the two sources, which was a pleasing result. The CSCG reconstruction also reconstructed the sources more compactly and with reduced noise. As expected, it was not possible to distinguish the sources from the GN reconstruction.

CSCG is a gradient-based reconstruction algorithm and so does not require calculation of other matrices such as inverse matrices or matrices of higher derivatives. Consequently the Jacobian matrix is the largest object that must be stored in memory. This matrix has *number of measurements* \times *number of mesh nodes* elements, each of which requires

8B of storage using a standard double precision floating-point data format. The CSCG algorithm required approximately 2 minutes to reconstruct the data acquired by the Dartmouth system using an four core desktop computer containing 8GB of Random Access Memory (RAM) (see Appendix A.6). The NNLS algorithm required approximately 40 seconds, and the GN algorithm required approximately 14 minutes for the same reconstruction problem. The reconstruction time in general for all algorithms will vary with the number of measurements and unknowns, and may also vary depending on the accuracy of the physical model of light propagation and the level of noise present in the measurements.

CSCG uses matrix-vector multiplications and vector inner product calculations which can be parallelised. These operations were not explicitly parallelised within the algorithm code. However, the algorithms were written in the MATLAB programming language which automatically parallelises those operations where suitable computing hardware is available. The line search operation can also be parallelised to an extent by evaluating the objective function simultaneously at multiple locations. These locations correspond to multiple possible execution paths as only one of them is ultimately chosen. However, the algorithm is fundamentally iterative and so there is limited scope for algorithm parallelisation. The GN and NNLS algorithms are also iterative and have the same limitations.

These results show that CSCG and to an extent NNLS are capable of simultaneously reconstructing two bioluminescence sources *in silico* and *in simulacra* which suggests that this may hold true when applied to the imaging system discussed in Appendix A and used elsewhere in this work. See chapter 4 for a case study of the application of CSCG and the surface metrology system detailed in chapter 2 to a set of *in simulacra* data acquired with the non-contact imaging system in Appendix A. In this case knowledge of the imaging system allows the measurement discrepancy criterion to be applied.

CHAPTER 4

AN EXAMPLE OF SURFACE CAPTURE AND COMPRESSIVE SENSING-INSPIRED RECONSTRUCTION TO BIOLUMINESCENCE TOMOGRAPHY

4.1 Contributions to this work

The BLT and surface metrology measurements were collected by Shelley Taylor. The free space transfer matrix was calculated by James Guggenheim.

4.2 Introduction

A surface metrology system for measurement of imaging subject surface geometry has been described in chapter 2. Data from this system enables the creation of a tetrahedral volume mesh that can act as a basis for parametrisation of FEM-based modelling of light propagation and bioluminescence reconstruction. An algorithm for reconstruction of sparse bioluminescence distributions has been introduced in chapter 3. This algorithm enables reconstruction of bioluminescence when represented as a linear system of equations, such as is the case when using an FEM model. It uses an expectation of bioluminescence sparsity to regularise the reconstruction problem to overcome the under-determined nature of BLT problems, and further can use a principled stopping criterion to choose a final

solution if the noise characteristics of the imaging instrument are known. These two works have previously been examined in isolation; this chapter presents an example of their combined use to reconstruct bioluminescence distributions *in simulacra* using a multi-modal imaging instrument detailed in Appendix A. The instrument noise characteristics can be estimated through knowledge of its constituent components.

4.3 Methods

The surface metrology system presented in chapter 2 was used to acquire surface geometry information. Imaging was performed using 10nm bandwidth filters with central wavelengths at 560nm, 580nm, 590nm, 600nm, 620nm, and 640nm. Surface metrology data was acquired without using camera binning, but the images of bioluminescence were acquired using 4×4 camera binning to reduce exposure time. The camera exposure time was estimated by imaging for a short period of 10s and then using the resulting image to calculate the expected exposure time at which the maximum image pixel value would reach 40 000 counts. The camera digitisation is 16-bit and so the maximum readout value is 65 535, but a significantly lower ideal value was chosen to prevent accidental CCD saturation due to a poor exposure estimate. Measurements were performed on the 29th of July, 2013, and repeated on the 1st of August. The XPM-2 phantom was left in the same position between measurements and their repeats, and so a single mesh and free space transfer matrix was used in reconstruction of measurements from both days.

The free space transfer matrix (see Appendix A.5) contains the sensitivity of camera pixels to surface points on the volume mesh, taking into account the placement of the optical components of the instrument. It accounts for the effect of distance, thin lenses, and apertures on the camera pixel sensitivity. The matrix is estimated using a discretisation and ray tracing approach.

The Jacobian matrices corresponding to the propagation of light through tissue were calculated once the subject volume mesh and free space transfer matrix had been obtained

from surface metrology data. The Nirfast software package was used to calculate these matrices, and previously measured homogeneous tissue properties (fig. A.5) were used in the calculations at each wavelength. The filters used for imaging had bandwidths of 10nm and it was assumed that a Jacobian matrix calculated for the filter central wavelength was a reasonable approximation for the sum of the Jacobian matrices at all wavelengths, weighted by the source emission (fig. A.5), filter transmittance (fig. A.3), camera quantum efficiency (fig. A.2), and mirror reflectivity if appropriate which was set to 0.98 at all wavelengths. The Jacobian matrix for the central wavelength of each filter was multiplied by the integral of the product of the source emission spectrum, filter transmittance, camera quantum efficiency, and mirror reflectivity to account for the effect on the number of measured photons of the spectral source strength and spectral efficiency of the optical components. The Jacobian matrix was also multiplied by the number of counts per electron, which was 1.67.

There are several possible methods of reconstructing a bioluminescence distribution \mathbf{x} given a Jacobian matrix \mathbf{J} , a free space transfer matrix \mathbf{T} , and a set of camera pixel counts \mathbf{c} . The relationship between the Jacobian matrix and bioluminescence distribution is $\mathbf{J}\mathbf{x} = \mathbf{y}$ where \mathbf{y} is the light leaving the surface of the subject. The relationship between the camera pixel counts and the light leaving the surface of the subject is $\mathbf{T}\mathbf{y} = \mathbf{c}$. The first possible method involves calculating the surface light using the free space transfer matrix, and then calculating the bioluminescence using the surface light and the Jacobian matrix. This approach is problematic however in the case where the free space transfer matrix is itself underdetermined, and that was the case here. Reconstruction of the surface light given an underdetermined free space transfer matrix which requires regularisation, but the appropriate choice of regularisation type is not obvious. In this case the additional underdetermined surface light reconstruction problem can potentially negatively affect the process of solving the subsequent bioluminescence reconstruction problem.

The second method involves combining the two equations to form a composite equa-

tion, which can then be solved in a single step. This equation is:

$$k\mathbf{T}\mathbf{J}\mathbf{x} = \mathbf{c} \quad (4.1)$$

where k is a multiplicative calibration factor that accounts for the spectral source emission, optical component transmittance and reflectivity, and for the camera quantum efficiency. This factor enables quantitative reconstruction in bioluminescence units of photon/s. This method offers the advantage of using the camera measurements directly. Doing so simplifies the treatment of noise, as the noise processes involved in CCD measurements are relatively well defined. Conversely, it is unclear how noise would propagate from the camera measurement vector \mathbf{c} to the surface light vector \mathbf{y} in the case where the free space transfer matrix \mathbf{T} is underdetermined and regularisation must be used in the process of reconstructing \mathbf{y} .

Six imaging wavelengths were used, but only measurements with an estimated SNR of 20dB or more were used for each wavelength. However, there were still too many measurements remaining to process using the available computing hardware. Also, the free space transfer matrix \mathbf{T} was underdetermined resulting in an increase in the underdeterminancy of the calibrated product $k\mathbf{T}\mathbf{J}$. The measurement data for each wavelength was compressed using two different methods and reconstructed using CSCG.

The first method iteratively combines pairs of measurements for which the cosine of the angle between associated Jacobian rows is greater than 0.9, which equates to an angle of less than 0.451rad or 25.8°. The angle between Jacobian rows is an indication of measurement similarity (Freiberger and Scharfetter 2009). A cosine of 0 indicates that two measurements are orthogonal, and a cosine of 1 or -1 indicates that two measurements provide the same information. In this case all elements of the Jacobian are positive and so the cosine is always non-negative. Summing together two Jacobian rows that are most similar reduces the size of the Jacobian whilst preserving more information than discarding one of the rows.

The threshold used to decide whether or not to sum two rows was chosen empirically. The cosines of angles between all pairs of measurements were calculated at each iteration and pairs of measurements not meeting the threshold were discarded. All remaining pairs were ordered by the cosine of the angle in descending order. The most similar pair of measurements were combined via addition and all remaining pairs containing one of the combined pair were also discarded. This process continued until no pairs remained. Finally, the cosines of angles between all pairs of measurements were recalculated and the process repeated until no pairs of measurements met the criterion.

The noise properties of the combined measurements were calculated in the same manner, because the sum of a set of independent normally-distributed random variables is itself normally distributed. If an original set of random variables $\{r_i\}$ are distributed with properties $r_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$, where i is an index, then a weighted sum of these random variables is distributed: $\sum_{ii} k_i r_i \sim \mathcal{N}(\sum_{ii} k_i \mu_i, \sum_{ii} k_i \sigma_i^2)$. The variance of the combined measurements is then the sum of the variances of the individual measurements.

The second method makes use of a random projection technique (Bingham and Man- nila 2001; Halko et al. 2011), which reduces the number of rows of the free space transfer matrix-Jacobian products. This is accomplished through the multiplication on the left by a random orthogonal basis matrix with properties designed to preserve the distance between vectors outside of the null space of $k\mathbf{TJ}$. Noise estimates for the compressed data were calculated by making use of the approximation that the noise is dominated by normally-distributed noise, and so the weighted sum is also normally-distributed. Consequently the properties of the distributions for the compressed data can be calculated using the random projection matrix. If \mathbf{R} is a random projection matrix, and $\tilde{\mathbf{R}}$ is such that $\tilde{R}_{ij} = |R_{ij}|$, then the projected measurements are $\mathbf{R}\mathbf{c}$ and the projected covariance is $\tilde{\mathbf{R}}\sigma$ where σ is the vector of measurement variances. The resulting measurements are no longer independent but the covariance between pairs of measurements is low relative to the measurement variances. This is a consequence of the high dimensionality of the measurement space and the properties of the random projection matrix which is distributed

so each row contains significant contributions from the majority of measurements. The resulting covariance matrix was assumed to be diagonal for the purpose of reconstruction using the assumption of low measurement covariance in order to reduce processing and storage costs.

Finally, the compressed calibrated free space transfer matrix and Jacobian product, and the compressed camera measurements at each wavelength were concatenated to form a single set of linear equations. This complete set was itself then compressed again to reduce inter-wavelength measurement dependencies. The measurements projected here are not independent because of the first projection step but the covariance matrix is again assumed to be diagonal allowing the same formula to be used. This approximation is not ideal but was justifiable because of the high dimensionality of the problem and beneficial because of the reduction in computation expense.

The final set of linear equations is calculated:

$$\mathbf{R}_{\lambda_0} \begin{bmatrix} \mathbf{R}_{\lambda_1} \begin{bmatrix} k_{\lambda_1,d} \mathbf{T}_d \\ k_{\lambda_1,m_1} \mathbf{T}_{m_1} \\ k_{\lambda_1,m_2} \mathbf{T}_{m_2} \\ \vdots \end{bmatrix} \mathbf{J}_{\lambda_1} \\ \mathbf{R}_{\lambda_6} \begin{bmatrix} k_{\lambda_6,d} \mathbf{T}_d \\ k_{\lambda_6,m_1} \mathbf{T}_{m_1} \\ k_{\lambda_6,m_2} \mathbf{T}_{m_2} \end{bmatrix} \mathbf{J}_{\lambda_6} \end{bmatrix} \mathbf{x} = \mathbf{R}_{\lambda_0} \begin{bmatrix} \mathbf{R}_{\lambda_1} \mathbf{c}_{\lambda_1} \\ \vdots \\ \mathbf{R}_{\lambda_6} \mathbf{c}_{\lambda_6} \end{bmatrix} \quad (4.2)$$

where \mathbf{R}_{λ_i} are the random projection matrices for wavelength i , and \mathbf{R}_{λ_0} is the random projection matrix for the concatenation of wavelength data. The matrix \mathbf{J}_{λ_0} can be

defined:

$$\mathbf{J}_{\lambda_0} = \mathbf{R}_{\lambda_0} \begin{bmatrix} \mathbf{R}_{\lambda_1} \begin{bmatrix} k_{\lambda_1,d} \mathbf{T}_d \\ k_{\lambda_1,m_1} \mathbf{T}_{m_1} \\ k_{\lambda_1,m_2} \mathbf{T}_{m_2} \\ \vdots \end{bmatrix} \mathbf{J}_{\lambda_1} \\ \mathbf{R}_{\lambda_6} \begin{bmatrix} k_{\lambda_6,d} \mathbf{T}_d \\ k_{\lambda_6,m_1} \mathbf{T}_{m_1} \\ k_{\lambda_6,m_2} \mathbf{T}_{m_2} \end{bmatrix} \mathbf{J}_{\lambda_6} \end{bmatrix}$$

and the vector \mathbf{c}_{λ_0} can be defined:

$$\mathbf{c}_{\lambda_0} = \mathbf{R}_{\lambda_0} \begin{bmatrix} \mathbf{R}_{\lambda_1} \mathbf{c}_{\lambda_1} \\ \vdots \\ \mathbf{R}_{\lambda_6} \mathbf{c}_{\lambda_6} \end{bmatrix}$$

The vector of variances \mathbf{v}_{λ_0} for the compressed measurements \mathbf{c}_{λ_0} is calculated:

$$\mathbf{v}_{\lambda_0} = \text{abs}(\mathbf{R}_{\lambda_0}) \begin{bmatrix} \text{abs}(\mathbf{R}_{\lambda_1}) \mathbf{v}_{\lambda_1} \\ \vdots \\ \text{abs}(\mathbf{R}_{\lambda_6}) \mathbf{v}_{\lambda_6} \end{bmatrix} \quad (4.3)$$

Where \mathbf{v}_{λ_i} is the vector of variances for measurement vector \mathbf{c}_{λ_i} , and $\text{abs}(\mathbf{M})_{ij} = |M_{ij}|$.

The CSCG algorithm was used to reconstruct the bioluminescence distribution in all cases. Reconstruction was performed for each of the two data compression methods, in the presence and absence of column normalisation. The noise for each compressed measurement was provided to the algorithm, and the stopping criterion was chosen so that the measurement discrepancy was on average one standard deviation away from the measured value.

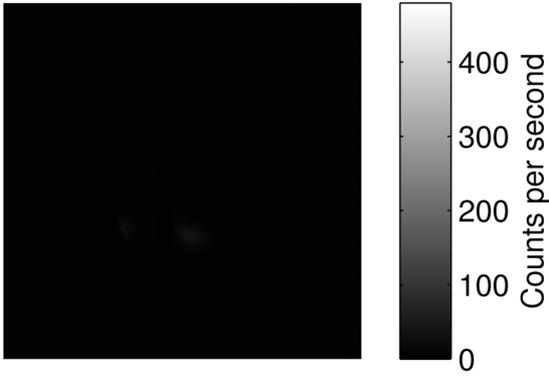
4.4 Results and discussion

Figures 4.1 to 4.3 show the camera measurements for the XPM-2 phantom with source A only switched on, source B only switched on, and both sources switched respectively. The state where both sources are switched on will be referred to as source C. The measurements were acquired using 10nm bandwidth filters with central wavelengths at 560nm, 580nm, 590nm, 600nm, 620nm, 640nm. The imaging wavelength greatly affects the intensity of the measured signal: the maximum signal intensity was measured at 600nm, and the signal intensity measured at 560nm was least. This reflects the complex interplay of source, tissue, and optical component properties. 600nm is not the wavelength of peak source emission (567nm), nor is it the wavelength of peak camera quantum efficiency (578nm), nor is it a minimum of XPM-2 scattering or absorption.

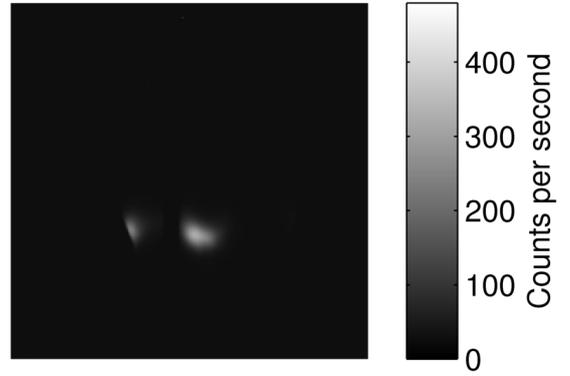
The signal intensity from source B is also approximately half that of source A, despite source B emitting photons at a greater rate (8.09×10^{10} photon/s compared to 8.40×10^{10} photon/s). This illustrates the importance of taking tissue optical properties and source location into account when interpreting BLI data.

Figure 4.4 shows a set of four reconstructions of the data for bioluminescent source A acquired on the 29th of July. Each reconstruction is displayed on a row of the figure. Reconstructions were performed with and without Jacobian column normalisation on data compressed using measurement summation and on data compressed using random projection. The left column of the figure contains visualisations of reconstructions with the colour scale range calculated using all of the nodes. The right column of the figure contains visualisations which exclude the nodes lying on the surface of the subject and where the colour scale range similarly excludes surface nodes. The locations of the tunnels containing the optical fibres are shown translucently. The upper tunnel corresponds to source A and the lower tunnel corresponds to source B (see fig. A.5b). The ends of the optical fibres should be positioned at the ends of the tunnels.

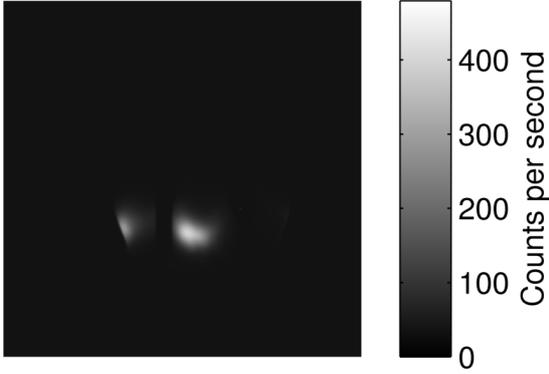
Several trends are visible:



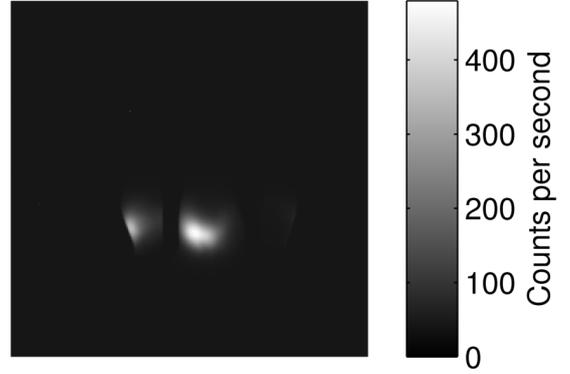
(a) XPM-2 phantom bioluminescence images in counts per second at 560nm



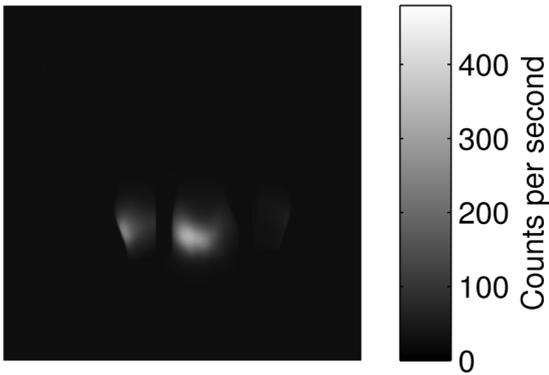
(b) XPM-2 phantom bioluminescence images in counts per second at 580nm



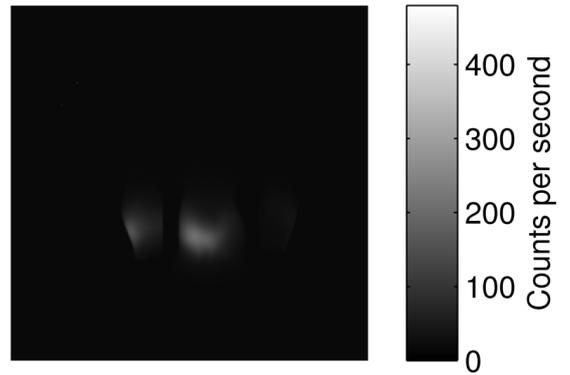
(c) XPM-2 phantom bioluminescence images in counts per second at 590nm



(d) XPM-2 phantom bioluminescence images in counts per second at 600nm

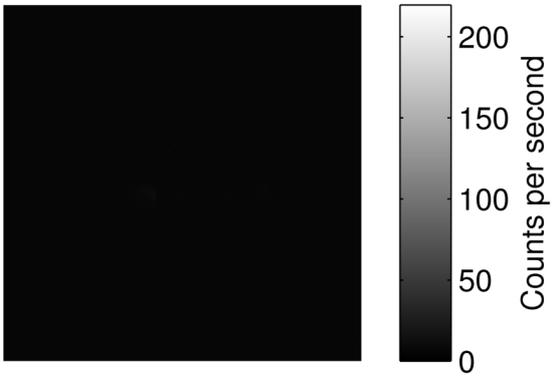


(e) XPM-2 phantom bioluminescence images in counts per second at 620nm

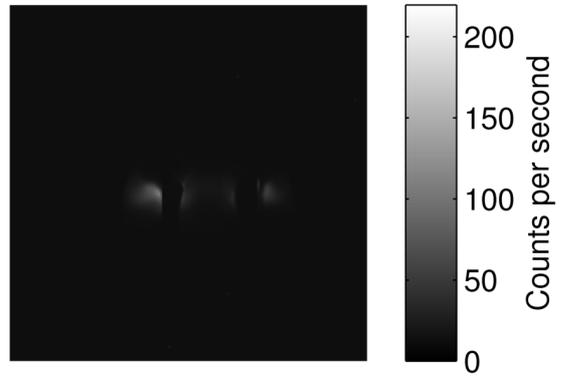


(f) XPM-2 phantom bioluminescence images in counts per second at 640nm

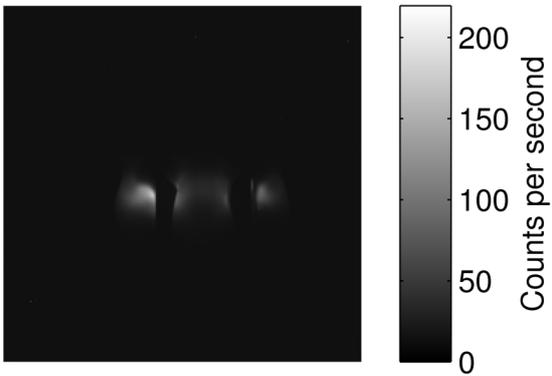
Figure 4.1: Images of XPM-2 phantom with source A switched on. The images were acquired on the 29th of July with camera binning set to 4×4 .



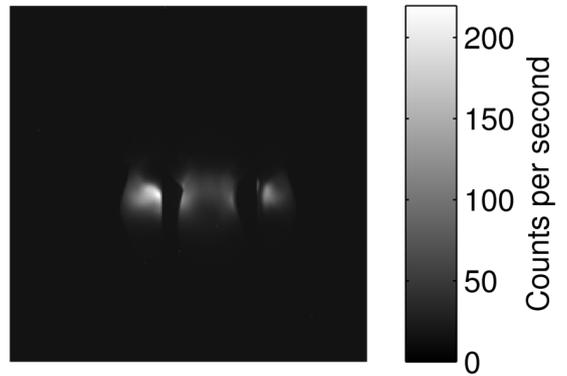
(a) XPM-2 phantom bioluminescence images in counts per second at 560nm



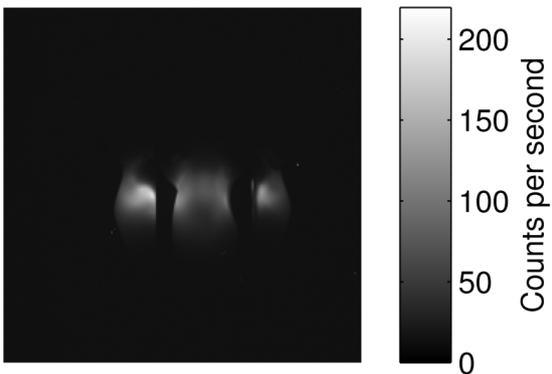
(b) XPM-2 phantom bioluminescence images in counts per second at 580nm



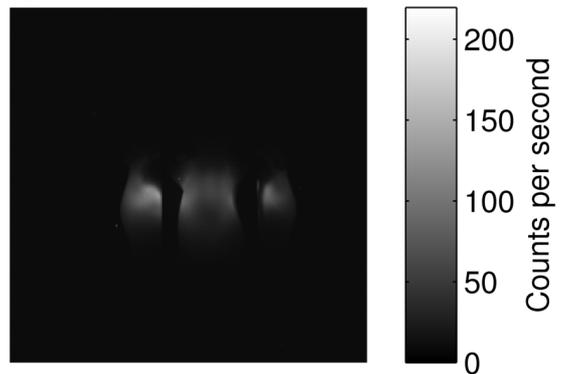
(c) XPM-2 phantom bioluminescence images in counts per second at 590nm



(d) XPM-2 phantom bioluminescence images in counts per second at 600nm

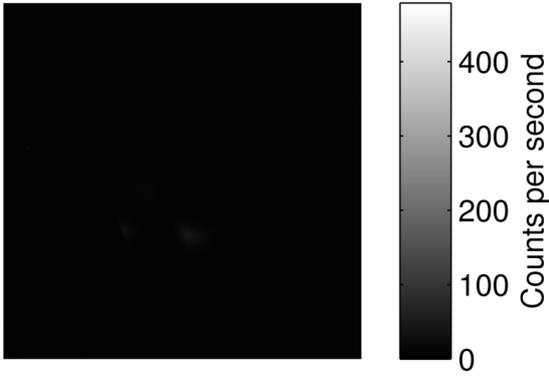


(e) XPM-2 phantom bioluminescence images in counts per second at 620nm

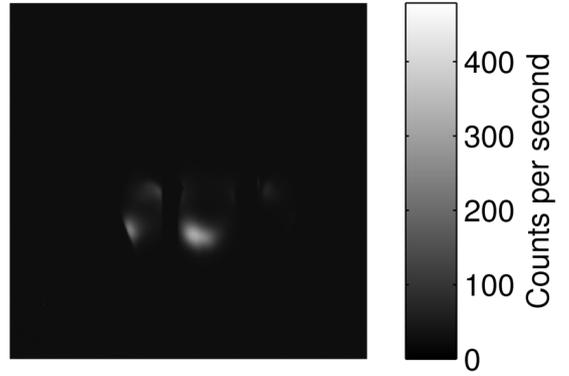


(f) XPM-2 phantom bioluminescence images in counts per second at 640nm

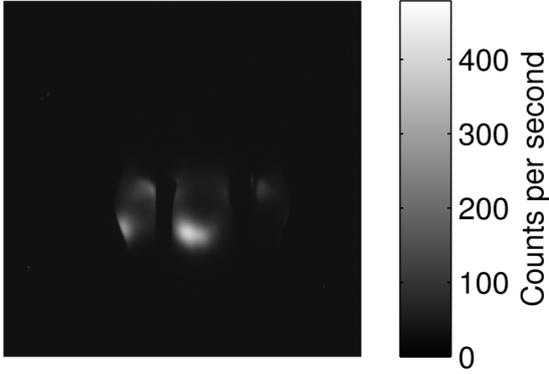
Figure 4.2: Images of XPM-2 phantom with source B switched on. The images were acquired on the 29th of July with camera binning set to 4×4 .



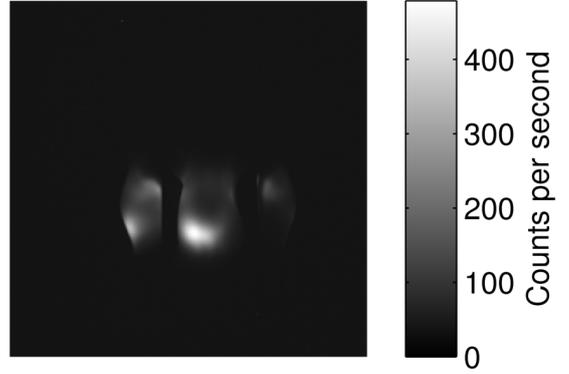
(a) XPM-2 phantom bioluminescence images in counts per second at 560nm



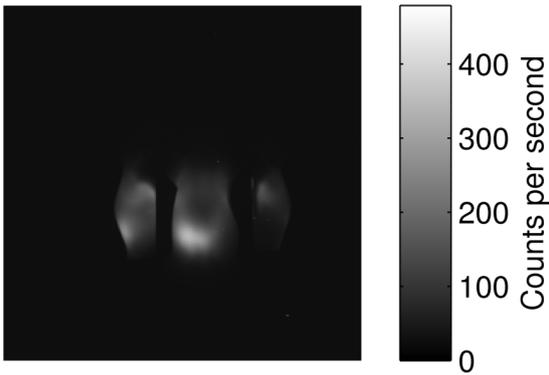
(b) XPM-2 phantom bioluminescence images in counts per second at 580nm



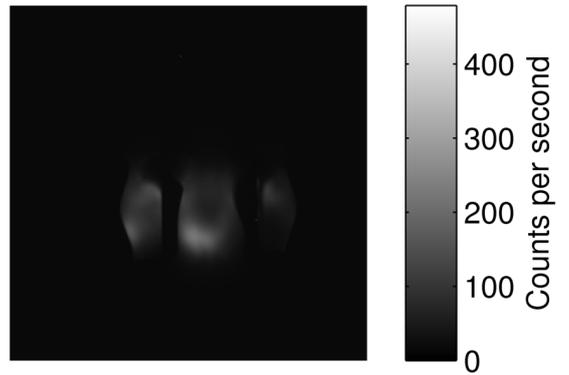
(c) XPM-2 phantom bioluminescence images in counts per second at 590nm



(d) XPM-2 phantom bioluminescence images in counts per second at 600nm



(e) XPM-2 phantom bioluminescence images in counts per second at 620nm



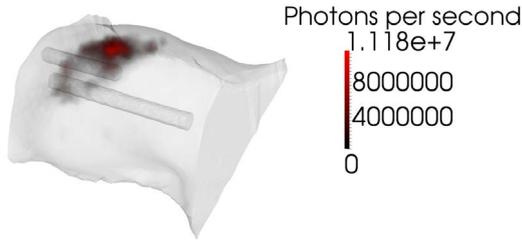
(f) XPM-2 phantom bioluminescence images in counts per second at 640nm

Figure 4.3: Images of XPM-2 phantom with source C switched on. The images were acquired on the 29th of July with camera binning set to 4×4 .

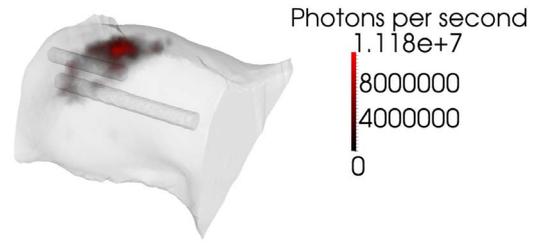
- Measurement compression using measurement summation is less effective than random projection. This may result from random projection preserving more of the information contained in the original matrix, but may also result from the changes to the structure of the matrix introduced by the random projection matrix.
- Column normalisation produces improved internal reconstruction quality, but also produces high magnitude surface artefacts that can affect the interpretation of results through their effects on visualisations. These artefacts are found on the surface of the subject in regions of low sensitivity and so can be removed by excluding the surface layer.
- It was possible to reconstruct the location and size of the bioluminescent source well, but only in the case where random projection and column normalisation were applied.

Figure 4.5 shows the result of reconstruction of a second set of data acquired for source A, on the 1st of August. These reconstructions show very similar trends and demonstrate the consistency of the surface metrology and CSCG techniques applied to the multimodality imaging system. Here again the use of column normalisation results in high magnitude surface artefacts in regions of low sensitivity. Their removal enables the true source to be visualised. The use of random projection measurement compression is again superior in performance to measurement summation, and the combination of random projection and column normalisation is optimal if surface nodes are excluded.

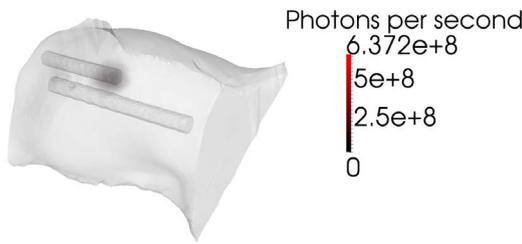
Figure 4.6 shows reconstructions of source B. This source is lower within the subject and thus produces lower surface signal, as can be seen in figs. 4.1 and 4.2. This reduction in signal is reflected in the reconstruction quality, which is inferior to that of source A reconstruction. The source is reconstructed closer to the surface than the position of the optical fibre. However, the same trends are observed: random projection is superior to measurement summation, and column normalisation improves performance if the reconstruction values for the surface nodes are discarded. The techniques produce a good



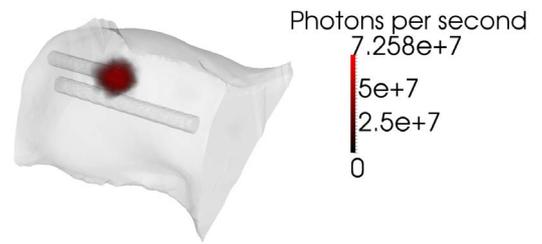
(a) Bioluminescence reconstruction without column normalisation using measurement summation



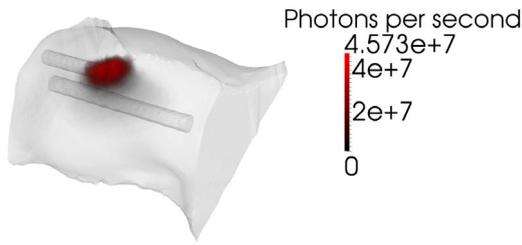
(b) Bioluminescence reconstruction without column normalisation using measurement summation with surface nodes removed



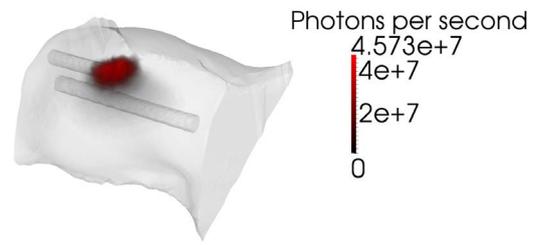
(c) Bioluminescence reconstruction with column normalisation using measurement summation



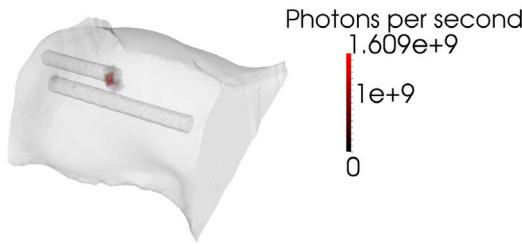
(d) Bioluminescence reconstruction with column normalisation using measurement summation with surface nodes removed



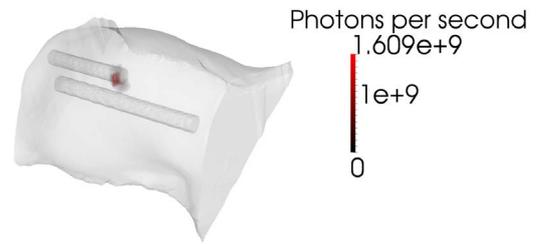
(e) Bioluminescence reconstruction without column normalisation using random projection



(f) Bioluminescence reconstruction without column normalisation using random projection with surface nodes removed

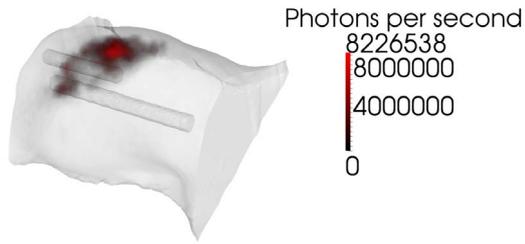


(g) Bioluminescence reconstruction with column normalisation using random projection

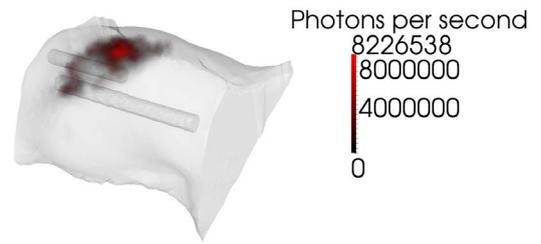


(h) Bioluminescence reconstruction with column normalisation using random projection with surface nodes removed

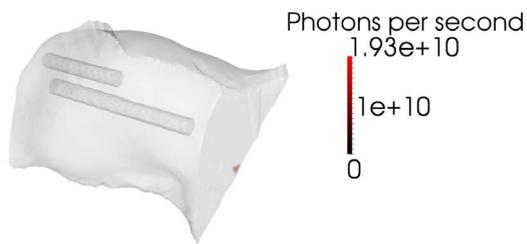
Figure 4.4: XPM-2 phantom reconstructions from measurements acquired on the 29th of July via CSCG with source A switched on.



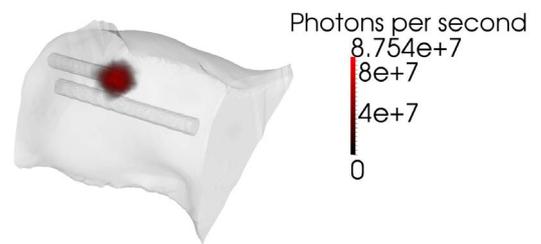
(a) Bioluminescence reconstruction without column normalisation using measurement summation



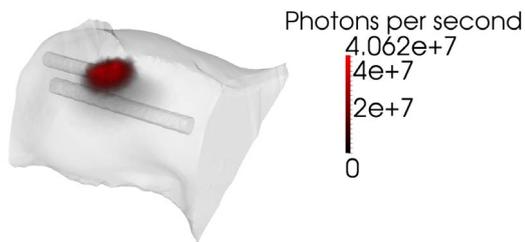
(b) Bioluminescence reconstruction without column normalisation using measurement summation with surface nodes removed



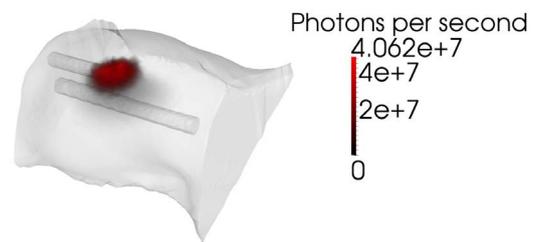
(c) Bioluminescence reconstruction with column normalisation using measurement summation



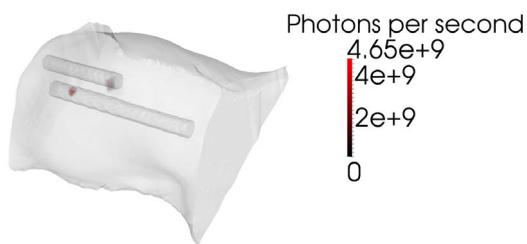
(d) Bioluminescence reconstruction with column normalisation using measurement summation with surface nodes removed



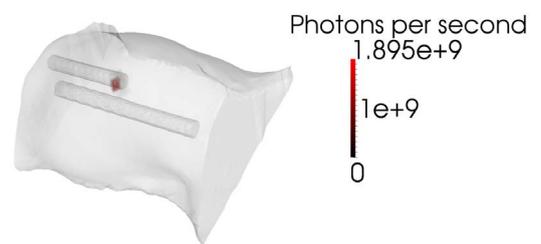
(e) Bioluminescence reconstruction without column normalisation using random projection



(f) Bioluminescence reconstruction without column normalisation using random projection with surface nodes removed



(g) Bioluminescence reconstruction with column normalisation using random projection



(h) Bioluminescence reconstruction with column normalisation using random projection with surface nodes removed

Figure 4.5: XPM-2 phantom reconstructions from measurements acquired on the 1st of August via CSCG with source A switched on.

reconstruction for source B in fig. 4.6h if the surface nodes are removed, but source B is reconstructed as being significantly less intense than source A, whereas source B is actually more intense.

The same trends are visible in the reconstructions of the repeated measurements acquired on the 1st of August, shown in fig. 4.7.

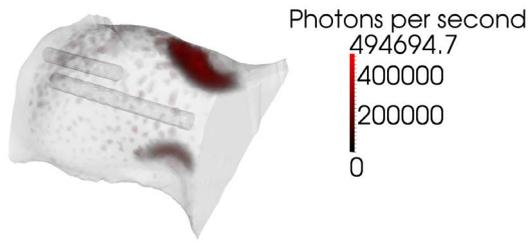
Figures 4.8 and 4.9 show the performance of the techniques on reconstruction of source C (sources A and B switched on simultaneously). Imaging was performed with both sources switched on, and the results show that it is possible to differentiate the two sources. The same trends as before are observed regarding random projection and measurement summation, and regarding column normalisation and surface artefacts. Source A is again reconstructed to have a significantly higher intensity than source B. Source B is also reconstructed closer to the surface than the position of the optical fibre.

Table 4.1 shows the total reconstructed intensity for source A. The true total intensity is 8.09×10^{10} photon/s, and the reconstructed intensities are consistently lower by one to two orders of magnitude. The reconstructed intensities all fall within two orders of magnitude, despite the large differences in the shapes of the reconstructed distributions in figs. 4.4 and 4.5. The variation in intensities is within one order of magnitude when excluding the contribution from surface nodes.

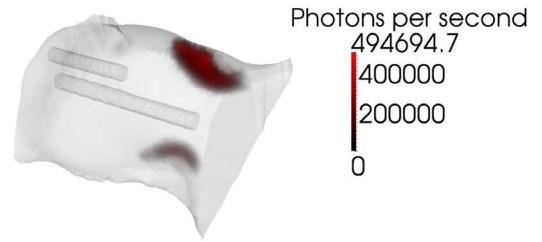
Measurement compression	Column normalisation	Measurement set	
		29 th of July, 2013	1 st of August, 2013
Summation	No	3.3×10^8 (2.9×10^8)	2.8×10^8 (2.4×10^8)
Summation	Yes	3.5×10^8 (1.3×10^9)	2.9×10^{10} (1.3×10^9)
Random projection	No	1.7×10^9 (1.7×10^9)	1.6×10^9 (1.5×10^9)
Random projection	Yes	4.4×10^9 (2.3×10^9)	1.1×10^{10} (2.3×10^9)

Table 4.1: Total reconstructed intensity of XPM-2 source A in units of photon/s. The ideal value is 8.09×10^{10} photon/s. The values in brackets represent the total intensity of the bioluminescence distribution after the removal of the contribution from the surface nodes.

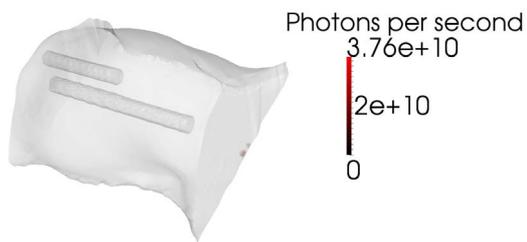
Table 4.2 shows the total reconstructed intensity for source B. This source is deeper within tissue (with respect to the measurement locations) and reconstruction variability



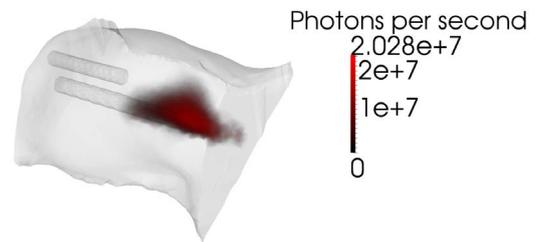
(a) Bioluminescence reconstruction without column normalisation using measurement summation



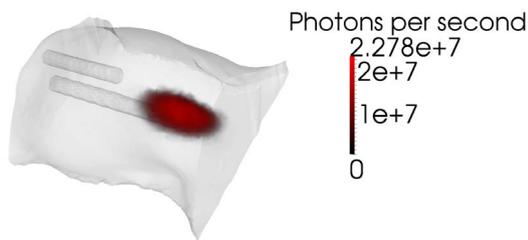
(b) Bioluminescence reconstruction without column normalisation using measurement summation with surface nodes removed



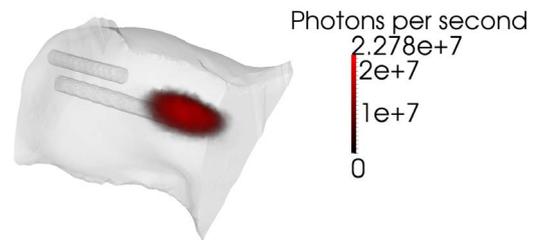
(c) Bioluminescence reconstruction with column normalisation using measurement summation



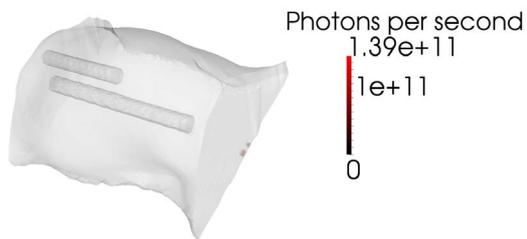
(d) Bioluminescence reconstruction with column normalisation using measurement summation with surface nodes removed



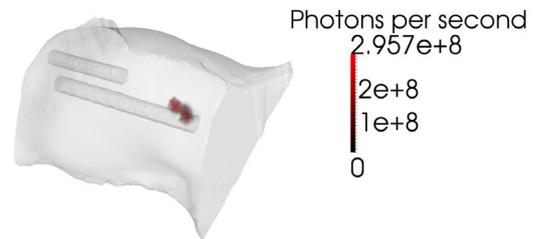
(e) Bioluminescence reconstruction without column normalisation using random projection



(f) Bioluminescence reconstruction without column normalisation using random projection with surface nodes removed

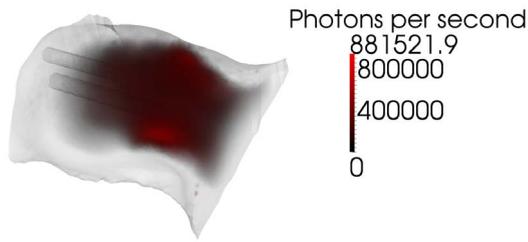


(g) Bioluminescence reconstruction with column normalisation using random projection

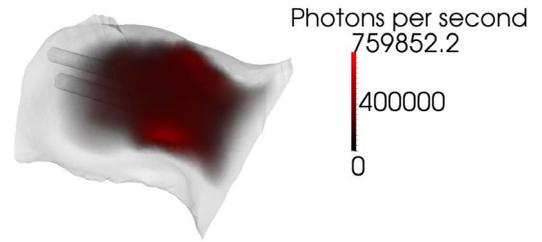


(h) Bioluminescence reconstruction with column normalisation using random projection with surface nodes removed

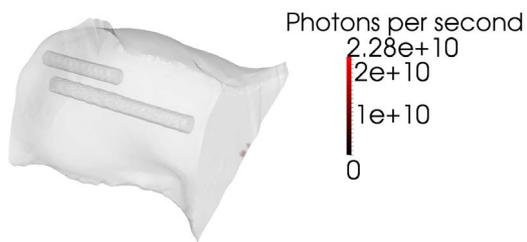
Figure 4.6: XPM-2 phantom reconstructions from measurements acquired on the 29th of July via CSCG with source B switched on.



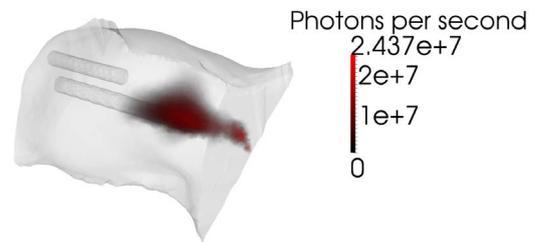
(a) Bioluminescence reconstruction without column normalisation using measurement summation



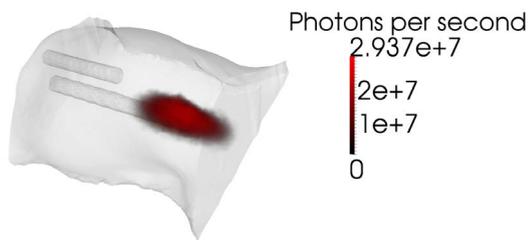
(b) Bioluminescence reconstruction without column normalisation using measurement summation with surface nodes removed



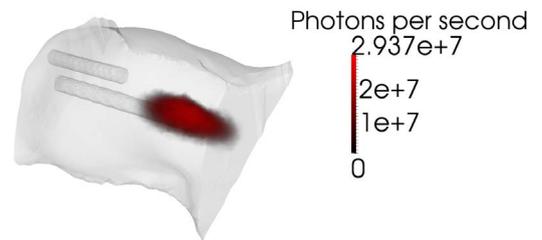
(c) Bioluminescence reconstruction with column normalisation using measurement summation



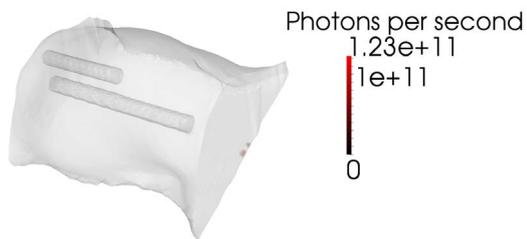
(d) Bioluminescence reconstruction with column normalisation using measurement summation with surface nodes removed



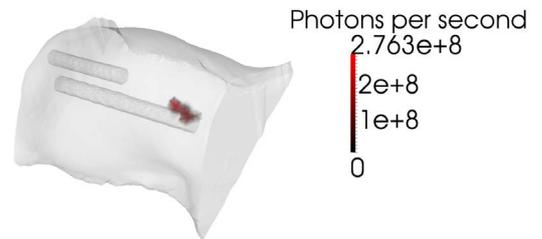
(e) Bioluminescence reconstruction without column normalisation using random projection



(f) Bioluminescence reconstruction without column normalisation using random projection with surface nodes removed

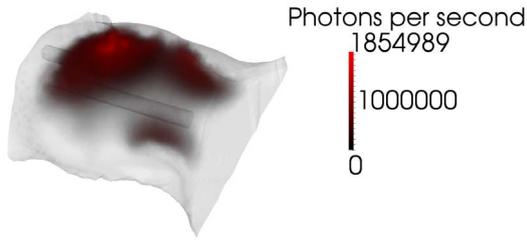


(g) Bioluminescence reconstruction with column normalisation using random projection

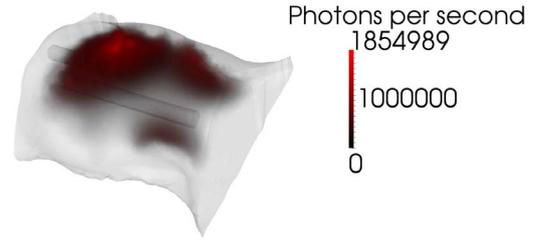


(h) Bioluminescence reconstruction with column normalisation using random projection with surface nodes removed

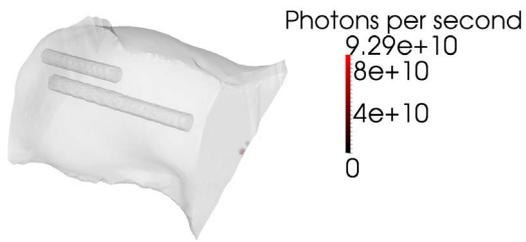
Figure 4.7: XPM-2 phantom reconstructions from measurements acquired on the 1st of August via CSCG with source B switched on.



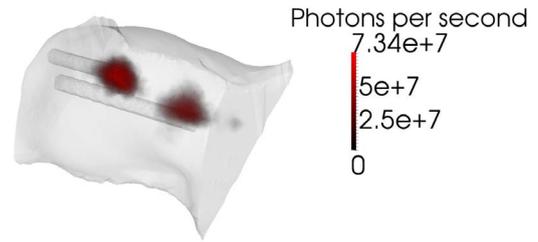
(a) Bioluminescence reconstruction without column normalisation using measurement summation



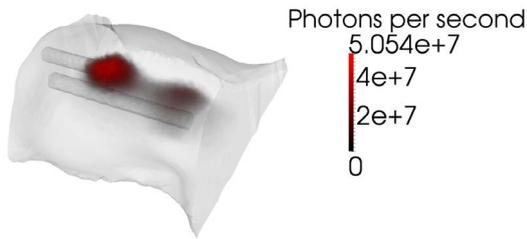
(b) Bioluminescence reconstruction without column normalisation using measurement summation with surface nodes removed



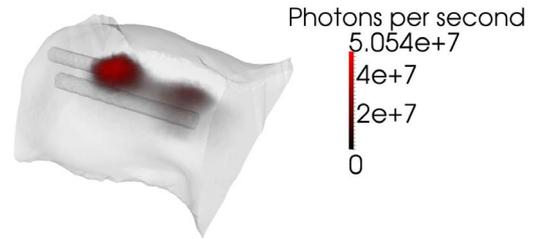
(c) Bioluminescence reconstruction with column normalisation using measurement summation



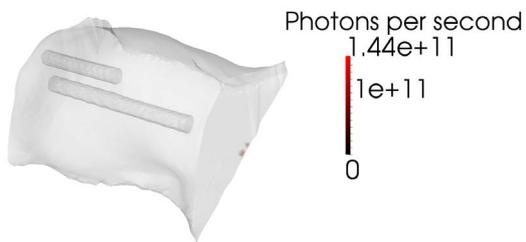
(d) Bioluminescence reconstruction with column normalisation using measurement summation with surface nodes removed



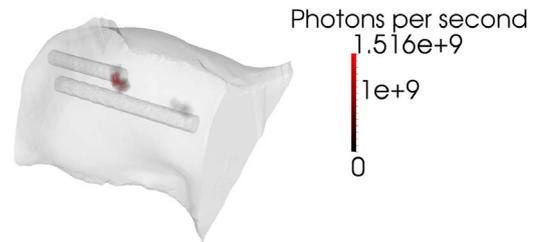
(e) Bioluminescence reconstruction without column normalisation using random projection



(f) Bioluminescence reconstruction without column normalisation using random projection with surface nodes removed

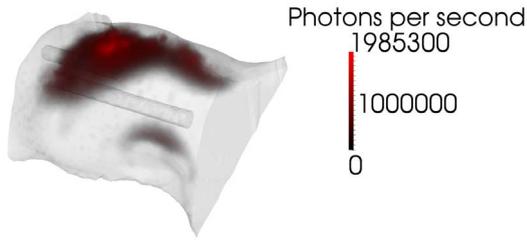


(g) Bioluminescence reconstruction with column normalisation using random projection

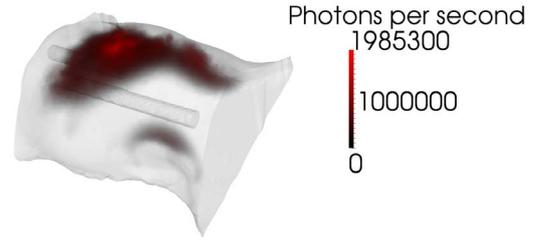


(h) Bioluminescence reconstruction with column normalisation using random projection with surface nodes removed

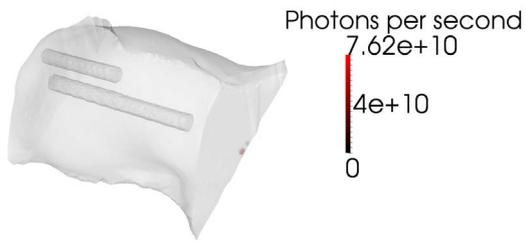
Figure 4.8: XPM-2 phantom reconstructions from measurements acquired on the 29th of July via CSCG with source C switched on.



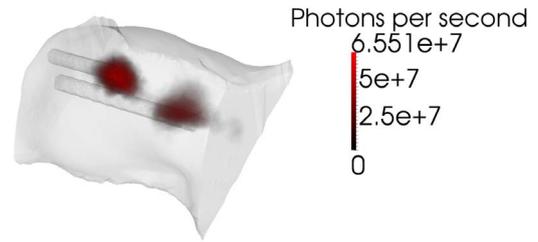
(a) Bioluminescence reconstruction without column normalisation using measurement summation



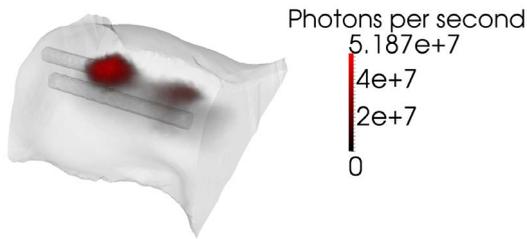
(b) Bioluminescence reconstruction without column normalisation using measurement summation with surface nodes removed



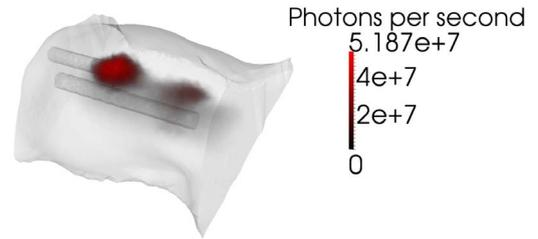
(c) Bioluminescence reconstruction with column normalisation using measurement summation



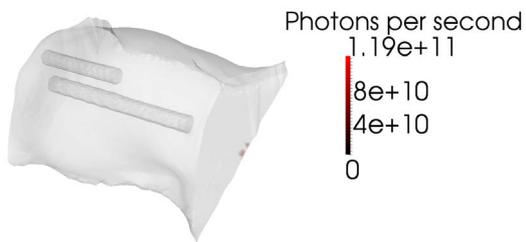
(d) Bioluminescence reconstruction with column normalisation using measurement summation with surface nodes removed



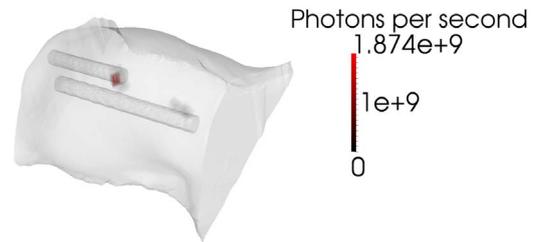
(e) Bioluminescence reconstruction without column normalisation using random projection



(f) Bioluminescence reconstruction without column normalisation using random projection with surface nodes removed



(g) Bioluminescence reconstruction with column normalisation using random projection



(h) Bioluminescence reconstruction with column normalisation using random projection with surface nodes removed

Figure 4.9: XPM-2 phantom reconstructions from measurements acquired on the 1st of August via CSCG with source C switched on.

is greater here. Intensities fall within four orders of magnitude, and range from three orders of magnitude below the true value to an order of magnitude above. The variation in intensities is within two orders of magnitude when excluding the contribution from surface nodes.

Measurement compression	Column normalisation	Measurement set	
		29 th of July, 2013	1 st of August, 2013
Summation	No	1.1×10^8 (7.8×10^7)	5.6×10^8 (5.3×10^8)
Summation	Yes	6.9×10^{10} (1.9×10^9)	4.7×10^{10} (1.7×10^9)
Random projection	No	2.1×10^9 (2.1×10^9)	2.7×10^9 (2.7×10^9)
Random projection	Yes	2.8×10^{11} (9.7×10^8)	2.5×10^{11} (1.0×10^9)

Table 4.2: Total reconstructed intensity of XPM-2 source B in units of photon/s. The ideal value is 8.4×10^{10} photon/s. The values in brackets represent the total intensity of the bioluminescence distribution after the removal of the contribution from the surface nodes.

Table 4.3 shows the total reconstructed intensity for source C. Intensities fall within four orders of magnitude and range from three orders of magnitude below the true value to an order of magnitude above. The variation in intensities is within one order of magnitude when excluding the contribution from surface nodes.

Measurement compression	Column normalisation	Measurement set	
		29 th of July, 2013	1 st of August, 2013
Summation	No	7.7×10^8 (7.4×10^8)	4.7×10^8 (4.4×10^8)
Summation	Yes	1.5×10^{11} (2.7×10^9)	1.3×10^{11} (2.8×10^9)
Random projection	No	2.9×10^9 (2.9×10^9)	2.9×10^9 (2.9×10^9)
Random projection	Yes	3.0×10^{11} (3.4×10^9)	2.5×10^{11} (3.2×10^9)

Table 4.3: Total reconstructed intensity of XPM-2 source C in units of photon/s. The ideal value is 1.649×10^{11} photon/s. The values in brackets represent the total intensity of the bioluminescence distribution after the removal of the contribution from the surface nodes.

4.5 Conclusions

The results show that it is possible to use the surface metrology and CSCG techniques to reconstruct bioluminescence distributions from non-contact measurements at multiple wavelengths. The two sources contained within the XPM-2 phantom were successfully reconstructed separately (figs. 4.4 to 4.7), and were distinguishable when reconstructed simultaneously (figs. 4.8 and 4.9). However, source B was consistently and incorrectly reconstructed with lower intensity than source A. The use of column normalisation additionally resulted in large reconstruction artefacts lying on the surface of the mesh in regions with lowest measurement sensitivity.

These artefacts may arise from the effect of noise in the measurements being amplified by the increase in the sensitivity of the least sensitive surface nodes. The increase in sensitivity results in the CSCG algorithm applying larger changes to these nodes, but the measurements contain little information about the bioluminescence at these locations because of their distance from the measurement locations. Consequently the changes to these nodes are primarily associated with the noise component of the measurements. The resulting bioluminescence is then further amplified by the correction for column normalisation that is applied after the CSCG reconstruction (see section 3.4).

The results indicate that the optimal choice of reconstruction parameters is dependent on whether removing the contribution of surface nodes from the reconstruction can be justified. In the case where it can be justified, the optimal reconstruction strategy (from those explored here) is to use random projection measurement compression and column normalisation. The use of random projection without column normalisation is optimal in the case where removal of contributions from surface nodes cannot be justified. A possible justification for the removal of the contribution of surface nodes is that the DA is not accurate for surface sources, and so these nodes would be unreliable in any case.

The choice of random projection without column normalisation or of random projection with column normalisation in conjunction with the removal of the contribution of surface nodes, also leads to very similar reconstructed distributions and total biolumi-

nescence for the two data sets. This suggests that the reconstruction process may be consistent to the level of noise found within the imaging instrument.

Despite large differences in reconstructed distributions (for example compare fig. 4.4a with fig. 4.4g) the reconstructed total bioluminescence falls within a narrow range (particularly when excluding reconstructed surface bioluminescence). This suggests that total bioluminescence may be reconstructed even if the shape of the distribution cannot be.

This reconstruction problem is higher-dimensional than that described in chapter 3. There are approximately 10 times the number of measurements and 5 times the number of unknowns. As a consequence the CSCG algorithm required approximately 30 minutes to perform reconstruction for this problem.

In conclusion, this example demonstrates that the surface metrology system in conjunction with the CSCG algorithm can enable reconstruction of a single or dual bioluminescence sources within a mouse phantom. The system noise was incorporated into the reconstruction process and manual selection of a stopping criterion was unnecessary. However, practical issues involving data quality and the use of measurement compression and column normalisation numerical techniques affected the reconstruction quality. A recommended reconstruction protocol would be to use random projection-based measurement compression and column normalisation and to discard the reconstruction at surface nodes, if this can be justified.

CHAPTER 5

POSTERIOR SAMPLING USING MARKOV CHAIN MONTE CARLO FOR BIOLUMINESCENCE TOMOGRAPHY

5.1 Introduction

The CSCG algorithm successfully reconstructed solutions to *in silico* and *in simulacra* problems without requiring manual selection of a termination criterion in chapters 3 and 4, but factors such as column normalisation and measurement compression influenced reconstruction quality. It would be advantageous to use a reconstruction method that is not affected by a need to choose algorithmic parameters without a physical justification, such as whether or not to use column normalisation, or alternatively where those choices are made based on prior information about the experiment rather than on subjective criteria. The production of measures of reconstruction reliability such as confidence intervals would also aid interpretation and potentially provide the ability to isolate noise artefacts such as those present in CSCG reconstruction using column normalisation (see chapter 4) in a principled manner.

A number of methods can enable automatic selection of algorithmic parameters, such as generalised cross validation, the L-curve method (Chaudhari et al. 2005; Hansen and O’Leary 1993; He et al. 2010), and the U-curve method (Chamorro-Servent et al. 2011). However, it would be desirable to use a reconstruction method for which all regularisation

is formed and weighted using prior knowledge specific to the reconstruction problem of interest rather than algorithm performance, and which is compatible with a range of possible types of regularisation. A Bayesian formulation (Robert 2001) for the reconstruction problem can meet these criteria. In such a formulation the problem can be regularised by incorporating any available prior knowledge about uncertainty in the process of measurement and in the potential solutions. This method is compatible with a range of potential types of information through the consistent representation of all information in the form of probability distributions.

The probability distribution which must be analysed to perform reconstruction is the posterior distribution $P(\mathbf{x}|\mathbf{c})$, and gives the probability of the bioluminescence distribution taking the value \mathbf{x} given the set of camera measurements \mathbf{c} . This distribution is usually not known directly. Instead, this distribution can be calculated from information about the measurement process which is contained in the likelihood distribution $P(\mathbf{c}|\mathbf{x})$, and *a priori* information about the space of bioluminescence distributions which is contained in the prior distribution $P(\mathbf{x})$. The posterior distribution can be used for bioluminescence distribution reconstruction, which in the conceptually simplest case consists of finding the MAP solution (Ahn et al. 2008; Milstein et al. 2004). This solution is the solution with the largest posterior probability. However, the posterior distribution contains more information than the MAP solution and this additional potentially useful information can also be extracted.

The MAP solution is useful but does not convey the possible range of solutions. The posterior distribution may be multi-modal and as such there may be a different solution which has a similar posterior probability to the MAP solution. The expectation of the posterior distribution is a single bioluminescence distribution which simultaneously conveys the range of likely solutions. In the case above the expectation would contain a superposition of these two solutions weighted by their relative probabilities.

Other moments of the posterior distribution can provide additional information. The spread of the posterior distribution in terms of the standard deviation or covariance pro-

vides a measure of the uncertainty in the reconstructed value of \mathbf{x} . This measure is spatially resolved, so can show locations where the value of the bioluminescence distribution is particularly uncertain.

It is however non-trivial to calculate these quantities from a posterior distribution in general. In this work Markov Chain Monte Carlo (MCMC) sampling was applied to the posterior distribution for the data reconstructed in chapter 4. This produced a number of samples drawn from the posterior distribution which were then used to estimate the distribution expectation and standard deviation.

A Markov Chain (Robert 2001) is a system that stochastically moves between a number of states and has no memory. The probability of the system moving to a particular state is dependent only on the current state and not the history of previous states. A Markov Chain can potentially become a stationary chain over time. The probability of a stationary chain being found in a particular state is independent of time, so the probability distribution describing the state of a stationary Markov chain is fixed.

MC sampling is used to approximate numerical quantities via random sampling. A famous example of this involves approximating the value of π by repeatedly dropping ball bearings randomly onto a drawing of a circle touching the inside of a square. The ratio of ball bearings inside the circle to the ratio of ball bearings inside the square allows an estimation of π , because the area of the circle and square are πr^2 and $4r^2$ respectively. The quality of the estimate improves with the number of ball bearings used, assuming infinitesimally small ball bearings. MC sampling can be used to approximate the integral of a function by uniformly sampling the value of the function within the integration domain. The average of these samples is an estimate of the value of the integral.

MCMC sampling attempts to acquire samples of a random variable distributed according to a particular probability distribution (Robert 2001). It does so by constructing a Markov chain for which the stationary probability distribution is the target probability distribution. The states of this Markov Chain are then samples from the distribution. One of the most famous MCMC algorithms is the Metropolis-Hastings (MH) algorithm

(Hastings 1970).

MH sampling avoids the calculation of the value of the distribution of interest at a given location so long as it is possible to calculate the value of a function proportional to the distribution. In this case, MH sampling avoids the need to calculate the value of $P(\mathbf{c})$. This is accomplished by using ratios of probabilities.

Assume that a Markov Chain is at position \mathbf{x}_i corresponding to a bioluminescence distribution at a given time point/iteration i . A proposal distribution $P(\mathbf{x}_{i+1}|\mathbf{x}_i)$ can be used to generate a new position \mathbf{x}_{i+1} . Given the two bioluminescence distributions \mathbf{x}_i and \mathbf{x}_{i+1} , and a set of measurements \mathbf{c} , the ratio of the posterior distributions is:

$$\begin{aligned} q(\mathbf{x}_{i+1}|\mathbf{x}_i) &= \frac{P(\mathbf{x}_{i+1}|\mathbf{c})}{P(\mathbf{x}_i|\mathbf{c})} = \frac{P(\mathbf{c}|\mathbf{x}_{i+1})P(\mathbf{x}_{i+1})/P(\mathbf{c})}{P(\mathbf{c}|\mathbf{x}_i)P(\mathbf{x}_i)/P(\mathbf{c})} \\ &= \frac{P(\mathbf{c}|\mathbf{x}_{i+1})P(\mathbf{x}_{i+1})}{P(\mathbf{c}|\mathbf{x}_i)P(\mathbf{x}_i)} \end{aligned} \quad (5.1)$$

The MH algorithm uses this ratio $q(\mathbf{x}_{i+1}|\mathbf{x}_i)$ to determine the probability with which the Markov Chain moves from state \mathbf{x}_i to state \mathbf{x}_{i+1} . In the case where the proposal distribution is symmetric (i.e. $P(\mathbf{x}_{i+1}|\mathbf{x}_i) = P(\mathbf{x}_i|\mathbf{x}_{i+1})$), the Markov Chain moves to state \mathbf{x}_{i+1} with probability 1 if $q(\mathbf{x}_{i+1}|\mathbf{x}_i) \geq 1$, and with probability $q(\mathbf{x}_{i+1}|\mathbf{x}_i)$ if $q(\mathbf{x}_{i+1}|\mathbf{x}_i) < 1$. The proposal distribution will ideally be the posterior distribution itself, but other functions can be used. The degree to which the proposal distribution approximates the posterior distribution will determine the speed at which the Markov Chain reaches the stationary distribution, which is the distribution that must be sampled.

The proposal distribution is often chosen to be a local distribution such as a normal distribution in practice, as opposed to a global distribution such as the uniform distribution. The use of a poor global proposal distribution would lead to extremely poor convergence, but the use of a poor local proposal distribution is not so detrimental to convergence as the use of a local distribution imposes a memory on the Markov Chain. MH algorithm is commonly run for a burn-in period during which no samples are retained. Ideally the samples acquired after the burn-in period will be independent of the initial

position of the Markov Chain. Samples can also be discarded during the sampling process to reduce the correlation in the samples. The parameters of the proposal distribution such as the standard deviation for the normal distribution can affect the rate at which samples are accepted. These parameters can be tuned during the burn-in period to optimise the sample acceptance rate (Roberts et al. 1997).

Reconstruction algorithms for BLT have already been surveyed in section 3.3 and so will not be addressed again here. However, Reinbacher-Köstinger et al. (2011) is of note as it addresses the issues of reconstruction and reconstruction reliability for the closely related imaging modality of FMT, and does so using similar methods to those presented here.

Reinbacher-Köstinger et al. (2011) applied MCMC sampling to an *in silico* FMT problem which consisted of a two dimensional circular subject containing two fluorescence sources. The imaging system was assumed to include a laser source and camera, although it was assumed that excitation illumination and imaging could be applied to any and all points on the subject surface. The measurement noise distribution was assumed to be Gaussian with a 1.5% noise level. The authors used Gaussian distributions to represent fluorescence sources, although the prior distributions of the parameters of the Gaussians were not specified. MCMC sampling via the MH algorithm (Hastings 1970) was used to acquire samples of the source centres, widths, concentrations, and life-times from the posterior distribution. Reconstruction means and standard deviations were calculated, and source centres and life-times were estimated well. The uncertainty in the source concentrations and widths was greater than in the source locations and life-times, but the two quantities were correlated in such a way as to estimate well the total quantity of fluorophore.

This work differs from that presented in Reinbacher-Köstinger et al. (2011) in terms of:

- its application to BLT.
- the use of a more complex physical model of the measurement process.

- the use of a more complex prior distribution.
- modifications to the MCMC MH algorithm.
- its application to an *in simulacra* problem with real measurement data.

5.2 Posterior distribution design

BLT image reconstruction quality is affected by a number of factors including biases and errors in the measurement and physical modelling processes, and ideally as many of these as is practical should be taken into account in the reconstruction process. A probability distribution is a natural representation for each of these factors and these probability distributions can be integrated to form a posterior distribution quantifying the probability of every possible bioluminescence distribution given a set of measurements, the distributions of biases and errors, and prior information that does not relate to the measurement process. If a bioluminescence distribution is represented by \mathbf{x} and belongs to some space of possible bioluminescence images \mathbf{X} , and a measurement set is represented by \mathbf{c} and belongs to some space of possible bioluminescence measurements \mathbf{C} , then the desired posterior distribution is $P(\mathbf{x}|\mathbf{c})$ and provides the probability of a specific bioluminescence distribution given a specific measurement set. It is then possible to find the value of \mathbf{x} which maximises the posterior distribution and present that as the reconstructed image. This is the MAP solution. However, the posterior distribution allows for a more powerful analysis that can provide additional information to aid clinical interpretation of the MAP and other solutions. This information may include:

- The expectation $E [P(\mathbf{x}|\mathbf{c})]$, which if differing from the MAP solution would suggest asymmetry or the existence of multiple modes in the posterior distribution.
- The distribution covariance, which would quantify the uncertainty in the reconstructed \mathbf{x} under the assumption that the posterior distribution is unimodal.

- The values and relative probabilities of the multiple modes, should the posterior distribution be multimodal.
- The sizes and shapes of the basins of attraction for the multiple modes, should the posterior distribution be multimodal.

The posterior distribution can be decomposed into the normalised product of the likelihood $P(\mathbf{c}|\mathbf{x})$, which is the probability of a specific measurement set given a specific bioluminescence distribution and includes biases and errors in the measurements and physical modelling, and the prior distribution $P(\mathbf{x})$, which is the probability of a specific bioluminescence distribution within the subject population in the absence of information provided by a specific measurement set. The relationship between the distributions is given by Bayes' Theorem (Robert 2001):

$$P(\mathbf{x}|\mathbf{c}) = \frac{P(\mathbf{c}|\mathbf{x})P(\mathbf{x})}{P(\mathbf{c})} \quad (5.2)$$

This equation allows the posterior distribution to be constructed from information about the imaging processes contained in the likelihood, and information about the biological processes contained in the prior distribution. The denominator $P(\mathbf{c})$ normalises the integral of the posterior distribution over all \mathbf{x} to unity, and can be calculated by marginalising the joint distribution $P(\mathbf{x}, \mathbf{c})$ over the space of bioluminescence distributions using the likelihood and prior distribution:

$$\begin{aligned} P(\mathbf{c}) &= \int_{\mathbf{x} \in X} P(\mathbf{x}, \mathbf{c}) d\mathbf{x} \\ &= \int_{\mathbf{x} \in X} P(\mathbf{c}|\mathbf{x})P(\mathbf{x}) d\mathbf{x} \end{aligned} \quad (5.3)$$

The measurements, \mathbf{c} , are fixed after an experiment has been conducted and so $P(\mathbf{c})$ is a constant. Some reconstruction algorithms are able to avoid direct calculation of $P(\mathbf{c})$ as a consequence of this by using ratios of probabilities (in the case of MCMC sampling, for

example) or by finding the bioluminescence distribution \mathbf{x} which satisfies:

$$\mathbf{x} = \arg \max_{\tilde{\mathbf{x}} \in X} P(\mathbf{c}|\tilde{\mathbf{x}})P(\tilde{\mathbf{x}}) \quad (5.4)$$

This solution is also the MAP solution as $P(\mathbf{c}|\mathbf{x})P(\mathbf{x})$ is proportional to $P(\mathbf{x}|\mathbf{c})$. The decomposition of the posterior distribution into the likelihood function and prior distribution allows the design of the posterior distribution to be separated into design of the likelihood function, and design of the prior distribution.

The likelihood function gives the probability of a particular set of measurements given a particular bioluminescence distribution. As a result, it represents the knowledge and uncertainty about the physical processes of imaging. In the case of BLT, this includes the physical model of photon propagation through tissue, the physical model of photon propagation from the surface of the imaging subject to the detector, and the physical model of measurement at the detector. The likelihood can be described as the probability that the differences between the observed measurements and a set of modelled measurements assuming \mathbf{x} arise as a result of the uncertainties in the modelling or detection processes.

The prior distribution gives the probability of observation of a particular bioluminescence distribution based upon prior information about the characteristics of the bioluminescent cells and the history of the imaging subject, and does not include knowledge provided by a set of subject measurements. Information used in the prior distribution may include the initial location and length of time since implantation of the bioluminescent cells, the growth rate and bioluminescence efficiency of the bioluminescent cells, and the genotype and phenotype of the imaging subject. This distribution has the effect of regularising the mathematical inversion problem.

5.2.1 Knowledge of the imaging process

It shall be assumed that all imaging is carried out using the system detailed in Appendix A. The first stage in the imaging process involves the creation of photons with a range of

wavelengths through the oxidisation of luciferin in the presence of the enzyme luciferase, oxygen, and ATP (Seliger and McElroy 1960). These photons may be detected if they propagate to the surface of the subject. Tagged cells will express a variable amount of luciferase enzyme, and will contain a quantity of luciferin that depends on the quantity injected, the injection site and the local and global vasculature, and the time since the luciferin injection (Burgos et al. 2003). The rate of photon production may also be limited by the availability of oxygen and ATP if the tagged cells are under metabolic stress.

Photons propagate in random directions once they are produced and interact with tissue primarily via elastic scattering, which causes them to change propagation direction, and via absorption, by which they are primarily converted into thermal energy. Scattering interactions dominate in most tissues for visible and infra-red radiation, but few photons reach the surface of the imaging subject due to the sheer number of interactions within tissue. The angular distribution of photon propagation more than a few scattering lengths away from bioluminescent sources is effectively uniform within reasonably homogeneous tissue except near the subject surface where it becomes biased towards the surface (Dehghani et al. 2008).

The photons that leave the surface propagate through free space and the majority of these do not enter the camera and are absorbed by the environment. A fraction of photons are reflected off the two 45° mirrors placed on the stage. Some photons pass through the interference filter that forms the first element of the optical measurement system. Interference filters transmit photons with efficiency dependent on the photon wavelength and angle of incidence (Lissberger and Wilcock 1959). The filters used by the system have a bandwidth of 10nm. The transmission efficiency is greatly reduced but is not necessarily negligible for photons with wavelengths that differ by more than 5nm from the central wavelength. The dependence of transmission efficiency and central wavelength on angle of incidence is not easily modelled. Ideally the imaging instrument should be optimised to minimise the spread of angles of incidence, but the use of mirrors to maximise subject surface coverage greatly increases this spread. It was also impractical

to add optical components to collimate the incident light and decollimate the transmitted light. The interference filters are circular and of finite size and so additionally act as an aperture.

The photons that are transmitted through the interference filter are then refracted through a system of lenses designed to bring a plane passing through or near to the imaging subject into focus on the instrument CCD. This process is subject to various types of lens aberration which should be minimal, and vignetting, which can be significant. Photons which leave the lens system then enter the camera body.

Finally, some of the photons that enter the camera body reach the CCD, and a fraction of these create electrons (electron-hole pairs) in discrete capacitor regions, each one of which constitutes an image pixel. Electrons can also be created spontaneously by the thermal energy present in the device. After a pre-chosen integration time, the number of electrons present in each capacitor region is read off subject to read error and converted into an integer value between 0 and 65 535 in the case of the 16-bit detectors used in the instrument, with an associated loss of information due to the discretisation and truncation processes.

5.2.2 The likelihood

The likelihood function should ideally incorporate all of the uncertainty present in the imaging process. This may include:

- Statistical uncertainty in the number of photons produced per second per unit of substrate and enzyme.
- Uncertainty in the optical properties of the subject tissue, which affects the number and distribution of photons that reach the surface according to the model of photon propagation.
- Uncertainty in the geometry of the subject surface.

- Uncertainty in the CCD electron generation.
- Uncertainty in the CCD readout process.

It should also incorporate models for the photon propagation and measurement processes so that the measurement process can be simulated. These include:

- Photon propagation through strongly absorbing but scattering-dominated tissue.
- Photon propagation through free space and the optical components in the system.
- Electron generation within the CCD.
- CCD read-out.

For a given bioluminescence distribution the physical models can be used to model the distribution of surface light resulting from the propagation of bioluminescent photons through tissue, the distribution of light incident on the CCD resulting from the free-space propagation of photons from the surface, the generation of electrons by the photons, and the conversion of electrons into recorded counts. The likelihood is then the probability that the discrepancy between the simulated and actual measurements results from the sources of uncertainty.

5.2.3 The prior distribution

The prior distribution should ideally incorporate all available information about bioluminescence distributions, both in general and in terms of the specific biomedical study that is being performed. In the case of general bioluminescence distributions, this may include the efficiency of the bioluminescence oxidation reaction in various conditions and the structure of the subject's vasculature, which may affect the local concentrations of the substrate luciferase over time as well as the concentrations of other components such as ATP and oxygen. Information specific to the particular biomedical study being conducted may include the initial location and number of injected cells, the type of injected

cells and their expected proliferation patterns, and the distribution of number of copies of the bioluminescence genes/operons within the cells. Together, this information can be used to infer the directions and patterns in which the cells are most likely to grow, their bioluminescence efficiency, and the temporal and spatial availability of substrates to fuel the bioluminescence reaction.

In practice little of this information will be known although many of these quantities may be experimentally measured or imaged given sufficient time, expertise, and instrumentation. In summary, the prior is the probability that a specific bioluminescence distribution will arise given all available knowledge (and uncertainty) on bioluminescence, cell proliferation, subject biology and physiology, and study methodology.

5.2.4 Designing the likelihood function

The likelihood function is dependent on the imaging instrument and the surface geometry of the imaging subject. The geometry of the imaging subject may change during the imaging session due to respiration or movement, and this adds an unavoidable element of uncertainty as the surface metrology subsystem cannot be utilised simultaneously to the BLT or DOT subsystems. A passive system could theoretically be implemented to analyse a separate video recording of the subject, extract the chest expansion and model the change in geometry, but this has not been addressed in this work.

It has also been assumed in this work that the optical properties of the subject are known which then enables exact calculation of surface light to the accuracy of the light propagation model, but in practice these properties are subject to an unknown level of error. Consequently there is an unknown level of error in the modelling of the surface light distribution resulting from internal bioluminescence.

The properties of the imaging instrument are better characterised and so these have been incorporated into the likelihood function as far as is practical. A model of free space light propagation is used to map the surface bioluminescence distribution onto the CCD pixels (Guggenheim et al. 2013b). This incorporates the major optical components of the

system including the mirrors, the filter wheel, the lens, and the camera.

- The mirrors have been incorporated into the free space model through the use of virtual cameras. However, the finite extent of the mirrors, which causes them to act as large apertures, has not been modelled for simplicity. This will have a limited effect on the reconstruction performance because the surface light distribution is limited in spatial extent due to the attenuation of tissue, and the positioning of the mirrors includes the majority of this in the final image.
- The filters and filter wheel differentially attenuate the signal as a function of wavelength via the installed filters and acts as an aperture due to the limited spatial extent of the filters. These have been incorporated into the free space model.
- The lens has been modelled as a thin lens in the free space model, and also acts as an aperture.
- The camera is the final optical component in the free space model, and has been incorporated as an imaging plane.

The result of the free space model is an optical transfer matrix that maps the discretised surface light distribution onto the pixels of the CCD but does not include sources of absorption such as the reflectance efficiency of the mirrors and the spectrally dependent absorption of the filters. This matrix can be extremely underdetermined which implies that it may not be possible to directly and uniquely reconstruct the surface light from the CCD measurements without additional information. Regularisation methods can be used to alleviate this problem to an extent (Guggenheim et al. 2013b), but an unprincipled choice of regularisation method can introduce additional error in the final result. Consequently in this work surface light measurements have not been directly reconstructed using the optical transfer matrix. Instead, the optical transfer matrix has been incorporated into a model that maps from an internal bioluminescence distribution to the CCD measurements directly.

If the imaging wavelength is λ , the direct view is represented by a subscript d , the view through the first mirror is represented by a subscript m_1 , the view through the second mirror is represented by a subscript m_2 , the optical transfer matrix is \mathbf{T} , the Jacobian matrix is \mathbf{J} , the internal bioluminescence distribution is \mathbf{x} , the CCD measurements are \mathbf{c} , \mathbf{e} is some noise component, and \mathbf{b} represents the camera background (i.e. the number of counts recorded on the camera using identical imaging settings but in the absence of light), then:

$$\begin{bmatrix} k_{\lambda,d} \mathbf{T}_d \\ k_{\lambda,m_1} \mathbf{T}_{m_1} \\ k_{\lambda,m_2} \mathbf{T}_{m_2} \end{bmatrix} \mathbf{J}_\lambda \mathbf{x} + \mathbf{b} + \mathbf{e} = \begin{bmatrix} \mathbf{c}_{\lambda,d} \\ \mathbf{c}_{\lambda,m_1} \\ \mathbf{c}_{\lambda,m_2} \end{bmatrix} \quad (5.5)$$

where $k_{\lambda,\{d,m_1,m_2\}}$ is a multiplicative factor that incorporates the attenuation due to optical component absorption and the multiplication factor to convert from electrons to camera counts. Explicitly, $k_{\lambda,\{d,m_1,m_2\}}$ is the following:

$$k_{\lambda,d} = t_\lambda q_\lambda r \quad (5.6)$$

$$k_{\lambda,m_1} = m_1 k_{\lambda,d} \quad (5.7)$$

$$k_{\lambda,m_2} = m_2 k_{\lambda,d} \quad (5.8)$$

where t_λ is the filter transmittance at the wavelength λ , q_λ is the quantum efficiency of the camera at the wavelength λ , and r is the conversion factor between electrons and counts. Equation (5.5) can be rewritten as:

$$\mathbf{T}_\lambda \mathbf{J}_\lambda \mathbf{x} + \mathbf{b} + \mathbf{e} = \mathbf{c}_\lambda \quad (5.9)$$

where:

$$\mathbf{T}_\lambda = \begin{bmatrix} k_{\lambda,d} \mathbf{T}_d \\ k_{\lambda,m_1} \mathbf{T}_{m_1} \\ k_{\lambda,m_2} \mathbf{T}_{m_2} \end{bmatrix}$$

$$\mathbf{c}_\lambda = \begin{bmatrix} \mathbf{c}_{\lambda,d} \\ \mathbf{c}_{\lambda,m_1} \\ \mathbf{c}_{\lambda,m_2} \end{bmatrix}$$

By combining the two problems of calculating the surface light distribution from CCD measurements and of calculating the internal bioluminescence distribution from surface light intensity, it is possible to:

- use the noise characteristics of the CCD directly rather than estimating the noise characteristics of the surface light distribution through propagation of the CCD noise characteristics via the optical transfer matrix, which as discussed above is flawed.
- use the priors on the bioluminescence distribution to indirectly regularise the optical transfer matrix in the reconstruction process.

The full problem can be expressed as:

$$\mathbf{TJx} + \mathbf{b} + \mathbf{e} = \mathbf{c} \tag{5.10}$$

where

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_{\lambda_1} \\ \vdots \\ \mathbf{T}_{\lambda_N} \end{bmatrix}$$

$$\mathbf{c} = \begin{bmatrix} \mathbf{c}_{\lambda_1} \\ \vdots \\ \mathbf{c}_{\lambda_N} \end{bmatrix}$$

and there are N wavelengths.

The propagation of light through tissue is modelled using the FEM model and the propagation of light from the tissue-air interface to the surface of the CCD is modelled using the free space model. The final piece is the model of the CCD, which maps the

distribution of incident photons to a discrete number of counts for each pixel supplied to the connected computer, and models the distribution of noise on those counts. This model has been largely constructed from manufacturer-provided information. The camera has a spectrally dependent quantum efficiency and the resulting electrons are then converted to counts with an efficiency that is less than unity.

There are a number of sources of noise within this process, including thermal and readout noise, but due to the quality of the camera hardware the dominant source of noise is shot noise. Shot noise is the statistical variation in the number of photons entering the camera per second, and follows a Poisson distribution. For sufficiently many input photons per second the Poisson distribution can be approximated by the Normal distribution with a mean of n and standard deviation of \sqrt{n} , where n is the average number of photons per second. This carries computational advantages as the Poisson distribution is discrete and asymmetric, while the normal distribution is continuous and symmetric. The normal distribution is also attractive because it takes the form of an exponential, which simplifies calculation of products of probabilities which can be performed through summation of exponents. This has a second advantage as it allows the calculation of products of probabilities so small that they cannot be accurately stored using a primitive type within a finite precision computer. It will be seen in later chapters that this is a practical necessity when calculating joint probabilities for thousands of measurements.

The likelihood for the full model, which incorporates both the Jacobian and the optical transfer matrix, is then the probability that the discrepancy between the predicted measurements and the observed measurements comes from the noise distribution. As the primary component of the noise distribution is shot noise, the noise distribution itself is strongly dependent on the predicted measurements.

Given eq. (5.9) and the Normal distribution approximation, the likelihood is then:

$$P(\mathbf{c} = \mathbf{y} - \mathbf{b}) \text{ for } \mathbf{c} \sim \mathcal{N}(\mathbf{TJ}\mathbf{x}, \text{diag}(\mathbf{TJ}\mathbf{x})) \quad (5.11)$$

5.2.5 Designing the prior distribution

The choice of prior is important as it contains additional knowledge that should regularise the inverse problem and so make it tractable. Ideally the prior should eliminate the null space of the problem so that there is a unique most probable bioluminescence distribution. It is important to remember that the region of the posterior distribution around that solution is also of interest, and it may be that the total probability within the basin of attraction of another local optimum is greater than that within the basin of attraction of the global optimum, which complicates the concept of the “solution” of the problem.

The prior can take the form of an explicit probability distribution, but choosing a basis with which to represent the solution that spans a strict subspace of the full solution space would also be imposing a prior. For example, the restriction of possible bioluminescence distributions to those consisting of a sum of a finite number of Gaussian distributions is equivalent to a prior distribution that assigns zero probability to bioluminescence distributions that do not satisfy this criterion. However, this can also be accomplished by representing the bioluminescence distributions using a Gaussian basis. The resulting parametrisation then allows the efficient calculation of integrals of the prior because it is not necessary to integrate over regions of the prior that are zero due to that restriction. This also applies to other methods of probing probability distributions such as MCMC methods.

The prior should contain knowledge regarding the distribution of possible bioluminescence distributions. This knowledge will shape the posterior distribution to an extent. Knowledge may be experiment-specific, and it is also possible for the introduction of false knowledge to bias the posterior distribution away from the true answer whilst simultaneously making the problem more well-posed. The prior must then be problem-specific, and if little is known then it is better that the prior be non-informative rather than incorrectly-informative.

It is possible to use arbitrary distributions when performing *in silico* experiments, but it is known that the plastic phantom used to acquire *in simulacra* measurements contains

two compact luminescent sources. One natural way to incorporate this knowledge into a prior is to use a sum of one or two Gaussians to represent the space of possible solutions, and solve for the parameters of these Gaussians. This representation has been used for the *in simulacra* experiments and has also been used for the *in silico* experiments in chapter 6. This prior naturally would not transfer to all possible pre-clinical experiments of potential interest, but may have sufficient validity to be useful for a subset of them. It may be possible to find a more general prior which can be tuned through hyperparameters by clinicians for the particular *in vivo* experiments to which it is being applied. However, that is beyond the scope of the current work.

5.3 Methods

The data used in chapter 4 was also used in this work and so the noise treatment was identical. Measurement sets compressed by summation were used in this work. Noise was again assumed to be normally distributed, and instrument calibration and Jacobian calculation was carried out as in chapter 4.

5.3.1 Posterior distribution construction

The posterior distribution $P(\mathbf{x}|\mathbf{c})$ is calculated from the prior distribution $P(\mathbf{x})$, the likelihood $P(\mathbf{c}|\mathbf{x})$, and the measurement probability $P(\mathbf{c})$. It was necessary to design the prior and likelihood distributions for this work. It was not necessary to design the measurement probability because MCMC sampling uses ratios of probabilities, and the probability $P(\mathbf{c})$ is a constant for any given measurement set \mathbf{c} . This avoids the need to calculate the measurement probability because contributions from $P(\mathbf{c})$ cancel out.

The prior probability distribution reflects prior knowledge about the space of possible bioluminescence sources. In this case, the XPM-2 phantom contains two luminescent sources that arise from light-emitting diode (LED) light propagating through optical fibres to the ends of the fibres within the XPM-2 phantom. The source shape was approximated

by a Gaussian distribution with finite support specified by a radial threshold. The distributions for sources A and B individually were assumed to consist of one Gaussian source, and the distribution for sources A and B together, which will be referred to as source C, was assumed to consist of two Gaussian sources. The resulting parametrisation for a luminescent source was as follows:

$$g(\mathbf{p}; \mu, \sigma, I, r) = \begin{cases} I \exp\left(-\frac{\|\mathbf{p}-\mu\|_2^2}{2\sigma^2}\right) & \text{for } \|\mathbf{p}-\mu\|_2^2 \leq r \\ 0 & \text{for } \|\mathbf{p}-\mu\|_2^2 > r \end{cases} \quad (5.12)$$

where \mathbf{p} is a point in three dimensional space, μ is the position of the source centre in three dimensions, σ is the width of the distribution and is a scalar, I is the maximum value of the distribution, and r is the size threshold of the distribution.

The source centre, width, and intensity were assumed to be independent and so the prior distribution was formed by taking the product of the prior distributions for each of the parameters μ , σ , I , and r .

- No knowledge was assumed about the position of the source centre μ .
- The intensity I was assumed to follow a folded normal distribution with mean of 8.09×10^{10} photon/s and standard deviation of 8.09×10^{10} photon/s.
- The width σ was assumed to follow a folded normal distribution with mean of 5mm and standard deviation of 5mm.
- The size threshold r was assumed to follow a folded normal distribution with mean of 5mm and standard deviation of 5mm.

A folded normal distribution takes the following form:

$$P(x) = \begin{cases} \frac{1}{\sigma\sqrt{2\pi}} \left(\exp\left(-\frac{(-x-\mu)^2}{2\sigma^2}\right) + \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \right) & \text{for } x \geq 0 \\ 0 & \text{for } x < 0 \end{cases} \quad (5.13)$$

The prior knowledge of the intensity distribution is based on the source strength for source A. However, the standard deviation of the distribution was chosen to be the same as the mean in order to provide limited information on the intensity and allow the reconstruction process to choose intensities significantly higher and much lower than the true value.

The prior knowledge of the width and size threshold restrict the maximum size of the source, but do not influence the minimum size of the source to the same degree as the standard deviation is again the same as the mean. The size of the tunnel containing the optical fibre is approximately 2mm in radius.

The likelihood distribution was constructed under the assumption that the measurement noise was Gaussian as in chapter 4. Measurement noise was further assumed to be independent. Consequently, the likelihood took the following form:

$$\begin{aligned}
 P(\mathbf{c}|\mathbf{x}) &= \prod_{i=1}^N P(c_i|\mathbf{x}) \\
 &= \prod_{i=1}^N \frac{1}{\sigma_i \sqrt{2\pi}} \exp \left(-\frac{(c_i - (\mathbf{J}_i \mathbf{x})^\top \mathbf{J}_i \mathbf{x})^2}{2\sigma_i^2} \right)
 \end{aligned} \tag{5.14}$$

where the measurement set consists of N measurements and σ_i is the standard deviation for measurement i . In the case of shot noise $\sigma_i = \sqrt{(\mathbf{J}_i \mathbf{x})^\top \mathbf{J}_i \mathbf{x}}$.

5.3.2 Markov Chain Monte Carlo sampling

Calculation of the value of the distribution $P(\mathbf{c})$ is non-trivial and can be computationally expensive (see chapter 6). Calculating moments of the posterior distribution such as the expectation by integrating functions of the posterior is also potentially computationally expensive and non-trivial. The expectation and standard deviation of the posterior distribution were estimated from samples obtained using the MCMC sampling technique to avoid these two issues.

The MH algorithm was used to perform sampling in this work. The proposal distri-

bution was designed to be a mixture of the normal distribution which is local and the prior distribution which is global, thus allowing the proposal distribution to produce large jumps with low probability ($p = 0.1$). A burn-in length of 10 000 samples was used, and blocks of 250 samples were used to optimise the properties of the proposal normal distribution to produce an ideal sample acceptance rate of 0.25. A total of 250 samples were acquired for each measurement set and Markov Chain after randomly discarding half of the samples. A total of four separate Markov Chains were executed for each measurement set, resulting in a total of 1000 samples per measurement set. Each Markov Chain was initialised using a sample from the prior distribution. The posterior expectation and standard deviation were estimated from the samples by calculating the sample mean and standard deviation respectively.

The size of the measurement sets caused practical computational difficulties in the calculation of the likelihood distribution $P(\mathbf{c}|\mathbf{x})$. The likelihood distribution is dependent on the chosen bioluminescence distribution \mathbf{x} , which is generated by the proposal distribution and the prior distribution when the algorithm is initialised. There is then no guarantee that the particular instance of \mathbf{x} is similar to the bioluminescence distribution that resulted in \mathbf{c} and so each of the measurements may be unlikely. The measurement sets were also composed of more than a thousand measurements. Consequently the product of all of the measurement probabilities (see eq. (5.14)) commonly evaluated to zero when evaluated using a double precision floating point data type. The natural logarithms of all probabilities were calculated directly, stored, and used in all subsequent calculations to mitigate this. This was possible because the likelihood distribution is a product of normal distributions and the prior distribution is a product of folded normal distributions, and both distributions are expressed in terms of exponentials.

5.4 Results and discussion

Figure 5.1 shows posterior expectations and standard deviations for source A calculated for two data sets. In both cases the source is reconstructed near to the expected location which is the tunnel exit, and smaller than or equal in width to the size of the tunnel. However, the reconstructed source widths for the two data sets are different. The reconstruction for the 1st of August data set is wider and the standard deviation is larger. The standard deviation is on the order of 0.01% and 1.0% of the expectation value for the 29th of July and 1st of August data sets respectively. The estimated standard deviation is too low to account for the differences in the expectations for the two data sets which ideally should be identical.

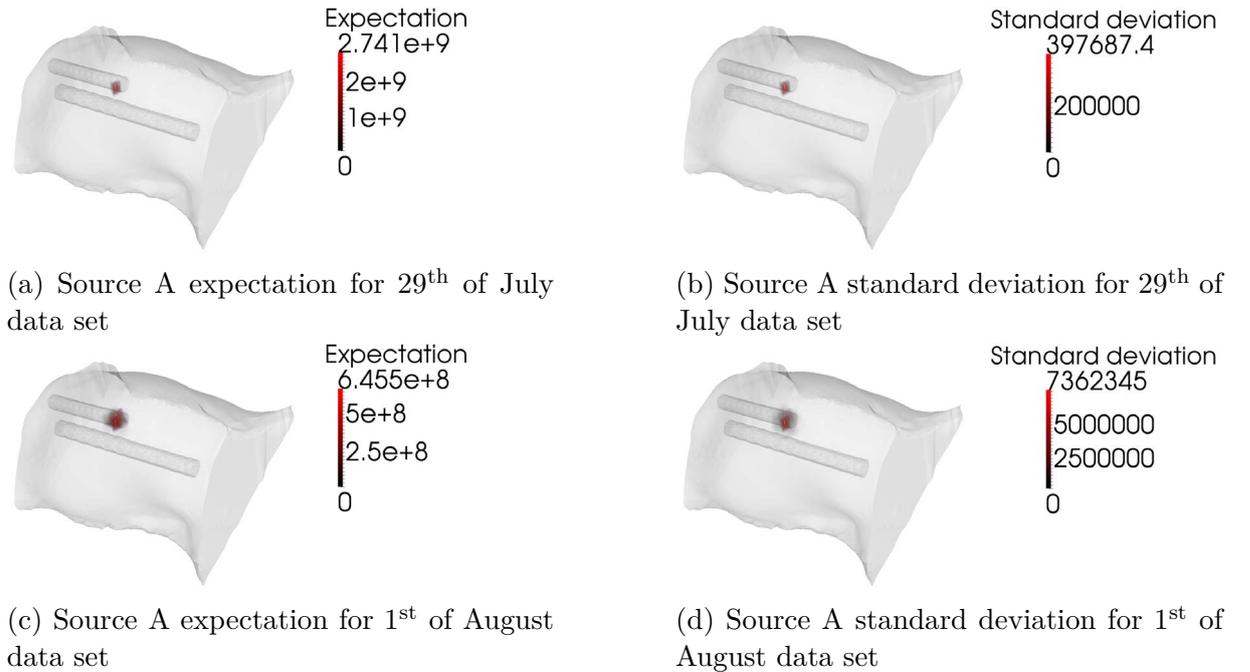


Figure 5.1: Expectations and standard deviations estimated from posterior samples for source A. The results are given in units of photon/s. Figures 5.1a and 5.1b correspond to the data set taken on the 29th of July, and figs. 5.1c and 5.1d correspond to the data set taken on the 1st of August.

Figure 5.2 shows posterior distributions and standard deviations for source B calculated for two data sets. The characteristics of this data set are similar to those for source A. The source expectation lies near to the tunnel exit but slightly above, and the source

is smaller or of a similar size to that of the tunnel exit itself. The expectation for the 1st of August data set is again wider than that of the 29th of July data set, and the standard deviation is larger. The standard deviations are again on the order of 0.01% and 1.0% of the expectation value for the 29th of July and 1st of August data sets respectively. The similarities in the results for both sources may indicate a systematic error in some part of the imaging process, as the data for sources A and B were acquired separately on each day.

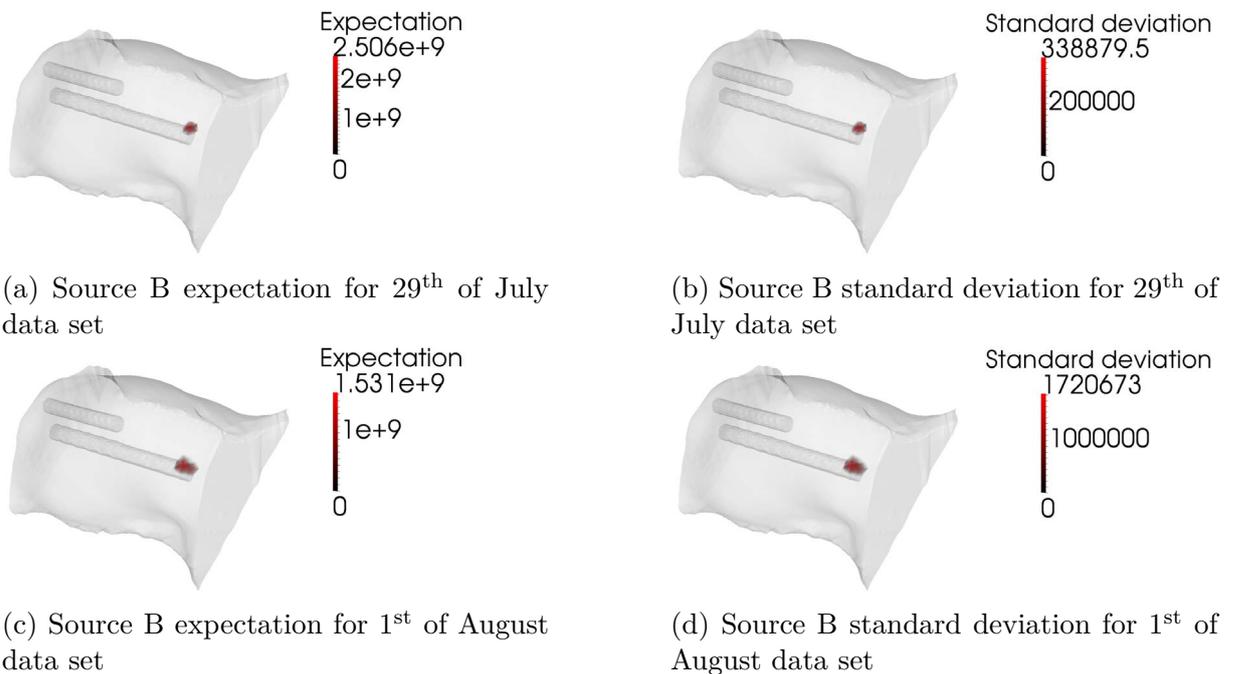


Figure 5.2: Expectations and standard deviations estimated from posterior samples for source B. The results are given in units of photon/s. Figures 5.2a and 5.2b correspond to the data set taken on the 29th of July, and figs. 5.2c and 5.2d correspond to the data set taken on the 1st of August.

Figure 5.3 shows posterior distributions and standard deviations for source C, calculated for two data sets. The expectation and standard deviation for the 29th of July data set do not show the presence of two sources. Instead, the source visible in fig. 5.3a lies on the bottom surface of the XPM-2 phantom and is significantly more intense to compensate for the increased depth. This result is likely a consequence of the MCMC algorithm failing to traverse the entire distribution during the sampling period. Large differences in probability can trap the algorithm within a region of the probability distribution unless it

is executed for a sufficiently long time. The use of multiple Markov Chains with random initialisation locations can mitigate this to an extent but the use of 4 Markov Chains was insufficient in this example. The results for the 1st of August data set however show a successful reconstruction. The two sources were reconstructed in appropriate locations with similar intensities. The standard deviation of the poor reconstruction for the 29th of July data set is an order of magnitude larger than that for the 1st of August data set, but crucially does not indicate uncertainty in the region of source A as a result of the samples being confined to the region around the expectation estimate.

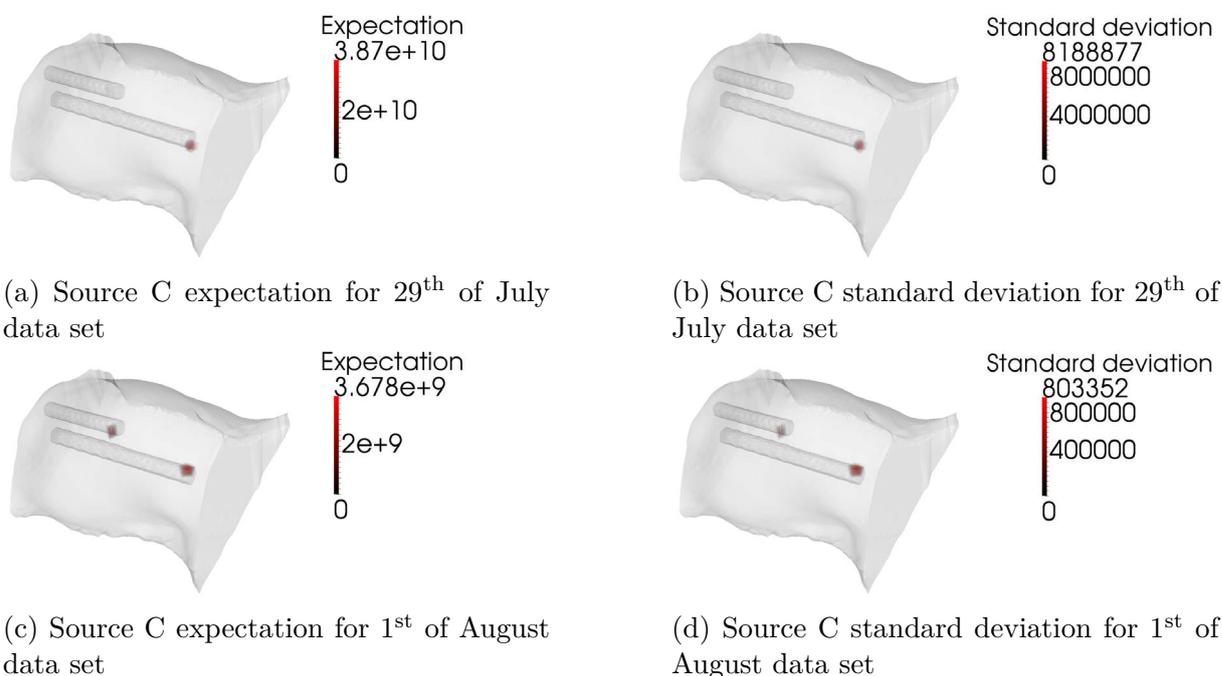


Figure 5.3: Expectations and standard deviations estimated from posterior samples for source C. The results are given in units of photon/s. Figures 5.3a and 5.3b correspond to the data set taken on the 29th of July, and figs. 5.3c and 5.3d correspond to the data set taken on the 1st of August.

Table 5.1 shows the total reconstructed source intensities for both data sets. The reconstructions are within two orders of magnitude for all sources. The true intensities of sources A and B are similar, as are the reconstructions. The reduction in reconstructed intensity compared to the true values may result from systematic errors in instrument calibration or phantom optical property measurement.

	Source A	Source B	Source C
29 th of July, 2013	2.7×10^9	2.5×10^9	4.1×10^{10}
1 st of August, 2013	2.26×10^9	3.05×10^9	6.44×10^9

Table 5.1: Total reconstructed intensity of XPM-2 sources A, B, and A and B together in units of photon/s. The ideal values are 8.09×10^{10} photon/s, 8.4×10^{10} photon/s, and 1.649×10^{11} photon/s.

5.5 Conclusion

The presented work shows that it is possible to reconstruct the locations of the luminescent sources within the XPM-2 phantom using a Bayesian framework and MCMC sampling. The reconstructed locations and intensities are comparable to reconstructions of the data using the CSCG algorithm with measurements compressed via random projection and using column normalisation (see section 4.4) but with superior reconstruction of total intensity and without the large surface artefacts associated with reconstruction using column normalisation.

The reconstruction on the data acquired on the 29th of July of source C is an exception to this, and demonstrates that this algorithm can still reconstruct unsuccessfully without an obvious indication that the reconstruction process has failed. It was hoped that in this case the standard deviation of the reconstruction would reflect great uncertainty but this has not occurred in practice. Both issues can be explained by the large range of probabilities due to the sheer number of measurements involved in the likelihood calculation, and the resulting steepness of the basins of attraction around the local maxima of the posterior distribution. The Markov Chain getting trapped (for finite time) in a single basin of attraction is a known concern with MCMC algorithms. The incorporation of the prior distribution into the proposal distribution allowed the algorithm to make large jumps but in this case that proved to be insufficient. An increase in the number of samples or in the number of independent Markov Chains may mitigate this in future. Alternatively, intelligent methods of initiating the Markov Chain may address this issue (Freiberger et al. 2011). However, the applicability or implementation of the MCMC algorithm is brought

into question in this case because the samples produced by the algorithm should be independent and one of the purposes of the burn-in period is to remove the dependence on the initialisation position.

Another possible factor in the failure of the 29th of July reconstruction of source C lies in the parametrisation of the bioluminescence distribution in terms of one or two sources. The parametrisation was justifiable in this case because a lot of information is known about the sources present in the XPM-2 phantom, but this parametrisation restricts the space of possible bioluminescence distributions greatly and so restricts the directions of movement of the MCMC algorithm. The shape of bioluminescence distributions is not known in general in biological subjects and so future work should involve the design of a more general prior distribution.

The MCMC algorithm required approximately 11 hours to acquire samples from the four Markov Chains using a four core desktop computer containing 8GB of RAM (see Appendix A.6). In comparison, the CSCG algorithm required less than 1 hour to perform a reconstruction of the same data when executed on the same computer. However, each Markov Chain of the MCMC algorithm can be executed in parallel on separate computers, reducing the total time to under 3 hours at the expense of additional computer hardware. Each Markov Chain can also be parallelised to a small extent as multiple potential steps can be evaluated simultaneously, although these represent multiple possible execution paths of which only one is ultimately chosen, and the dependence of the state of the chain at time n on the state of the chain at time $n - 1$ limits the potential for parallelisation.

An alternative method of reducing computational time would be to reduce the number of samples per chain and increase the number of parallel chains. However, this would not reduce the time required for the burn-in period which would still need to be executed independently for each chain.

The MCMC algorithm uses matrix-vector multiplications and vector inner product calculations similarly to CSCG. These operations can be parallelised and MATLAB does so automatically when suitable hardware is available.

The Jacobian matrix is typically the largest object that must be stored in memory. The samples produced by the Markov Chains must also be stored in memory, but the Gaussian parameterisation greatly reduces the storage cost to 48B per sample for a bioluminescence distribution consisting of one Gaussian, and 96B per sample for a bioluminescence distribution consisting of two Gaussians, using a standard double precision floating-point data format.

Reconstruction using sampling in a Bayesian framework allows estimation of properties of the posterior distribution beyond the most probable bioluminescence distribution or the expectation of the distribution. It is hoped that estimation of the standard deviation of the posterior distribution will provide information on the reconstruction reliability that is not accessible when using standard reconstruction techniques, and that this information will improve interpretation of images of bioluminescence distributions. There are situations in which the standard deviation can be non-trivial to interpret, such as in the case where the posterior distribution is strongly multi-modal, and future work would examine potential alternative metrics of reconstruction reliability.

CHAPTER 6

OPTIMISATION OF BIOLUMINESCENCE TOMOGRAPHY MEASUREMENT WAVELENGTHS USING MUTUAL INFORMATION AS A METRIC OF INFORMATION CONTENT

6.1 Introduction

A Bayesian representation of the relationship between bioluminescence distributions and instrument measurements allows reconstruction of a most probable bioluminescence distribution for a particular set of measurements and calculation of measures of reconstruction confidence. However the representation as a probability distribution also allows the measurement process in general to be explored by treating it as a communication channel.

In this perspective the bioluminescence distribution \mathbf{x} is a random variable which is distributed as the prior distribution. The bioluminescence distribution associated with a particular imaging subject is a sample from that distribution, and is the message that must be transmitted. The measurement process is the communication channel and transforms the message deterministically via the free space transfer matrix and Jacobian matrix and stochastically via processes including the sources of noise. The result of the measurement process is the observed signal \mathbf{c} which is itself a random variable. The goal is then to reconstruct the original message \mathbf{x} from the observed signal \mathbf{c} . If the observed signal \mathbf{c}

can be used to select \mathbf{x} then the two random variables must be dependent. Information theory provides a means of measuring the degree of this dependence through a quantity called mutual information (MI).

MI measures the mutual dependence of \mathbf{x} and \mathbf{c} rather than the dependence of \mathbf{c} on a particular instance of \mathbf{x} , and so provides a metric of expected imaging performance in general. This metric is independent of the choice of reconstruction algorithm subsequently used to calculate \mathbf{x} from \mathbf{c} , and so is a more objective measure of measurement process quality than reconstructed image quality-based metrics. However, it should be noted that MI measures the dependence of \mathbf{c} on \mathbf{x} and vice versa but that does not necessarily mean that a chosen reconstruction algorithm is capable of making use of that information. The measurement and reconstruction processes cannot be separated if a pre-selected algorithm is to be used for image reconstruction and in such a case it may be more appropriate to use a reconstructed image quality-based metric.

MI $I(X; C)$ can be calculated from the joint probability distribution $P(\mathbf{x}, \mathbf{c})$, and the probability distributions $P(\mathbf{x})$ and $P(\mathbf{c})$, where $P(\mathbf{x})$ is commonly called the prior distribution:

$$I(X; C) = \int_{\mathbf{x} \in X} \int_{\mathbf{c} \in C} P(\mathbf{x}, \mathbf{c}) \log \left(\frac{P(\mathbf{x}, \mathbf{c})}{P(\mathbf{x})P(\mathbf{c})} \right) d\mathbf{x}d\mathbf{c} \quad (6.1)$$

In this case the form of the joint probability distribution is not known directly, but the conditional probability $P(\mathbf{c}|\mathbf{x})$ is known. The conditional probability distribution $P(\mathbf{c}|\mathbf{x})$, also known as the likelihood, can be used to calculate the joint probability distribution in conjunction with the prior distribution:

$$P(\mathbf{x}, \mathbf{c}) = P(\mathbf{c}|\mathbf{x})P(\mathbf{x})$$

The distribution $P(\mathbf{c})$ can also be calculated from $P(\mathbf{c}|\mathbf{x})$ and $P(\mathbf{x})$ (see eq. (5.3)). Con-

sequently the MI can be calculated via the equation:

$$I(X; C) = \int_{\mathbf{x} \in X} \int_{\mathbf{c} \in C} P(\mathbf{c}|\mathbf{x})P(\mathbf{x}) \log \left(\frac{P(\mathbf{c}|\mathbf{x})}{\int_{\hat{\mathbf{x}} \in X} P(\mathbf{c}|\hat{\mathbf{x}})P(\hat{\mathbf{x}})d\hat{\mathbf{x}}} \right) d\mathbf{x}d\mathbf{c} \quad (6.2)$$

The distribution $P(\mathbf{c}|\mathbf{x})$ contains the model of the measurement process, and the distribution $P(\mathbf{x})$ contains any prior knowledge of the distribution of bioluminescence distributions. The MI can be calculated if both distributions have been specified.

It is possible to examine the effect of changes to certain aspects of the measurement process via the metric of MI. These aspects may include:

- the central wavelengths of the imaging filters.
- the bandwidth of the imaging filters.
- the number of imaging filters.
- the position of imaging mirrors.
- the location of the camera.
- the integration time.

The optimum central wavelengths of the imaging filters are examined in this work.

Further, it is also possible to examine the efficacy of the measurement process in different imaging studies via the modification of $P(\mathbf{x})$. For example, the bioluminescence distributions associated with a biomedical study may be predominantly shallow as a result of initial bioluminescent cells being injected subcutaneously. In another biomedical study initial bioluminescent cells may be injected into the bloodstream and so the bioluminescence distributions may be distributed and deep in tissue. The measurement process design can be tailored to each study to provide the maximum amount of information taking into account all prior knowledge.

The dependence of MI on the measurement process was explored in this work in the context of the number and choice of measurement wavelengths. An *in silico* two dimensional imaging subject and measurement process was created and the MI was estimated for a sets of measurement wavelengths consisting of a single wavelength, two wavelengths, and three wavelengths.

6.2 Survey of measurement optimisation methods in optical tomographic imaging

Optimal measurement strategies for BLT have not been an area of recent research, perhaps because of the limitations of using luminophores with narrow and fixed luminescence spectra, the inability to control excitation, and the predominance of camera-based imaging. Measurement strategy optimisation has however been explored for the closely related imaging modality of FMT where it is possible to control excitation, and optical fibre-based light collection is common. Space and cost constraints can limit the number of sources and detectors that can be incorporated into an optical fibre-based system to a greater degree than if using projection-based excitation and camera-based measurement, and so source and detector placement becomes more important. As a consequence a number of studies have examined optimal placement of sources and detectors using a number of metrics. This work examines optimal wavelengths for BLT rather than detector placement but the metrics used for FMT source and detector placement can also be used to optimise other aspects of the imaging process such as imaging wavelengths. Consequently it is worth examining some representative examples.

Graves et al. (2004) presented an application of singular value (SV) analysis to FMT using *in silico* and *in simulacra* examples. SVDs were calculated for Jacobians for different source and detector fields of view. The number of singular values above an empirically determined cut-off of 1×10^{-4} of the value of the largest singular value was used as a metric of measurement quality. The authors found that there were optimal fields of

view which were dependent on the number and arrangement of sources and detectors. Reconstructions of *in silico* and *in simulacra* data were consistent with the trends in the SV analysis.

These results are promising and can be applied to examine the effect of measurement wavelength, but have limitations in that it is not clear how to apply prior knowledge of the bioluminescence distributions such as the expected size and concentration, and knowledge of system noise characteristics. This information has the potential to affect optimal measurement configurations, particularly in the case where camera-based measurement is used and the exposure times of all measurements at a given wavelength are then linked. Use of the system noise characteristics may allow a more principled choice of SV cut-off, or the replacement of the cut off with a superior criteria. Lasser and Ntziachristos (2007); Zhang and Badea (2009) also explored FMT measurement configurations using SV analysis.

SV analysis has also been used for DOT to perform wavelength optimisation (Boas et al. 2004; Brendel and Nielsen 2009; Corlu et al. 2003, 2005; Eames et al. 2008; Mazhar et al. 2010) through the use of the condition number of the Jacobian, which is the ratio of the largest and smallest singular values of the Jacobian. A second metric called the residual norm (Corlu et al. 2005) was also used, but this is related to chromophore separability and is dependent on the spectral extinction properties of the chromophores. Consequently it is not applicable to BLT. BLT and FMT Jacobians are constant under the assumption that tissue optical properties are known, whereas tissue optical properties are the unknown of interest in DOT and so the Jacobian for non-linear DOT reconstruction is not known and changes during the optimisation process. The effectiveness of Jacobian analysis-based measurement configuration metrics may not transfer well between DOT, and BLT and FMT.

Freiberger and Scharfetter (2009) presented a method of reducing the number of sources and detectors in FMT while maximising measurement set information content. Rows of the Jacobian correspond to measurements, which are associated with a partic-

ular source and detector. The cosine of the angles between two rows of the Jacobian is a measure of the similarity between the two measurements associated with those rows. It is possible to calculate measures of the similarity between two sources or between two detectors by calculating averages of the cosines of measurements associated with the two sources or detectors. It is also possible to calculate measures of the similarity between a single source or detector and all other sources and detectors in the same manner. This measure was used to iteratively reduce the number of sources and detectors by removing the sources and detectors with greatest similarity to the remaining sources and detectors. This process minimised the loss of information at each iteration.

The authors demonstrated the effect of this method on an *in silico* cylindrical problem with rings of alternating sources and detectors around the curved surface of the cylinder. The Jacobian matrix was biased towards certain regions through column weighting and the optimised source and detector locations were spread out but placed near the region of interest. Detectors were spread out further than sources. The resulting Jacobians were examined using SV analysis, and the condition number of the Jacobian was used as a measure of the measurement redundancy. Large condition numbers indicate that some measurements are largely redundant. Low condition numbers indicate good conditioning, which can lead to good performance of conventional reconstruction algorithms. The condition numbers of the optimised Jacobians were significantly smaller than that of the original Jacobian.

This method produced results that were consistent with those that would be intuitively expected, and the SV analysis supported the authors' conclusions. The redundancy-based optimisation method was also computationally inexpensive in comparison to the SV analysis on the test problem.

There are several limitations of this approach. Firstly, the final number of sources and detectors must be chosen before the optimisation process. This could be remedied by setting a stopping criteria based on a measure of the information content of the Jacobian as the authors suggested. Secondly, the method of biasing the Jacobian towards certain

regions is crude and it is not clear how to translate more nuanced prior knowledge of source distributions into an appropriate bias on the Jacobian. Thirdly, this method assumes that each of the detectors has the same level of noise which is not true in general, especially for camera-based detectors where every detector pixel must be exposed for the same amount of time regardless of signal strength.

Freiberger et al. (2010) presented an improvement on the technique described in Freiberger and Scharfetter (2009). The information theoretic entropy for a set of measurements and MI between sets of measurements were derived for an FMT instrument, under the assumption that bioluminescence values at each node are independent and normally distributed. Evaluation of the derived expression requires calculation of inverse matrices and the authors stated that this method is computationally infeasible as a result.

The redundancy-based method in Freiberger and Scharfetter (2009) was described and the use of a geometric average rather than an arithmetic average was suggested. A redundancy measure was derived using the geometric average and the authors suggested that the measure using the geometric average is more similar to the MI than the feature using the arithmetic average. The iterative reduction process in Freiberger and Scharfetter (2009) was applied to the measure incorporating the geometric average, and the same region biasing was applied.

The optimisation methods incorporating arithmetic and geometric averaging were applied to the same cylindrical problem as in Freiberger and Scharfetter (2009). Identical sources were chosen, and similar detector positions were chosen. Jacobian condition numbers were also very similar. Example bioluminescence sources were reconstructed and reconstruction error was compared for the two redundancy-based optimisation procedures and an intuitively designed configuration with evenly spaced sources and detectors. The reconstructions were compared against a reconstruction using data from all sources and detectors. The method using geometric averaging produced superior results.

This work showed largely the same performance and benefits as Freiberger and Scharfetter (2009) and could potentially also be applied to optimisation of other imaging factors

such as wavelength. This work also shares the same set of limitations as Freiburger and Scharfetter (2009), in terms of the pre-selection of an final number of sources and detectors, the method of biasing the Jacobian towards certain regions, and the assumptions with regard to measurement noise and the probability distribution for the bioluminescence distribution. The prior probability distribution for the bioluminescence distribution could be used to create a more principled method of biasing the Jacobian, but it is unclear how the derivation of the MI could be extended to incorporate more complex prior distributions for the bioluminescence, and if that were possible it is also unclear whether the MI would retain the same similarity to the redundancy measure using geometric averaging.

Leblond et al. (2010) presented an analysis of SV-based methods of FMT measurement optimisation in the context of measurement noise. A varying number of regularly spaced source and detector placements were examined, unlike Freiburger et al. (2010); Freiburger and Scharfetter (2009). An *in silico* test problem created from CT data of a mouse head was used. Measurements were generated in the presence and absence of 10% Gaussian noise, for an optically homogeneous mesh and for an optically heterogeneous mesh to simulate uncertainty in the physical model. SV analysis was used to predict optimal measurement types. SV analysis was not found to predict correctly the order of the different source and detector configurations, particularly in the presence of measurement noise and model error. This indicates the importance of incorporating measurement noise, and model error if possible, into metrics used to compare different measurement configurations.

6.3 Methods

The Nirfast software package was used to create a two dimensional circular mesh and perform the modelling of light propagation through tissue. Two dimensions were used instead of three because of the computational demands of the MI calculation. Tissue labels associated with the Digimouse mouse atlas (Dogdas et al. 2007; Stout et al. 2002)

were used to calculate optical properties for a two dimensional slice through the phantom and these properties were interpolated onto the two dimensional mesh.

MI was investigated for pairs of wavelengths. Prior and likelihood distributions were designed to emulate compact radially symmetric sources imaged by an ideal (shot noise dominated) contact imaging system in order to allow MI estimation.

6.3.1 Design of the prior distribution

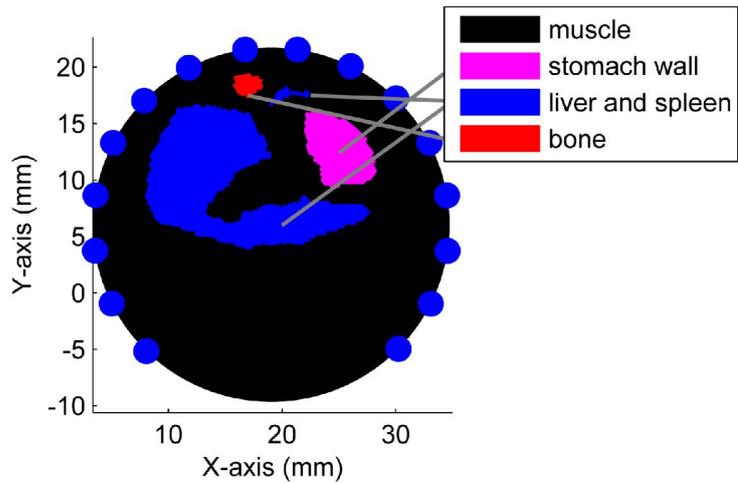
The prior distribution was designed to restrict the space of possible bioluminescence distributions to those following a mixture of Gaussian distributions. These Gaussians were also radially symmetric around their centre, which implies that they possess identical variance in each dimension and zero covariance. A radially symmetric Gaussian source can be defined in terms of a centre location \mathbf{p}_0 , a variance σ^2 , and an intensity I :

$$g(\mathbf{p}) = I \exp\left(-\frac{\|\mathbf{p} - \mathbf{p}_0\|_2^2}{2\sigma^2}\right) \quad (6.3)$$

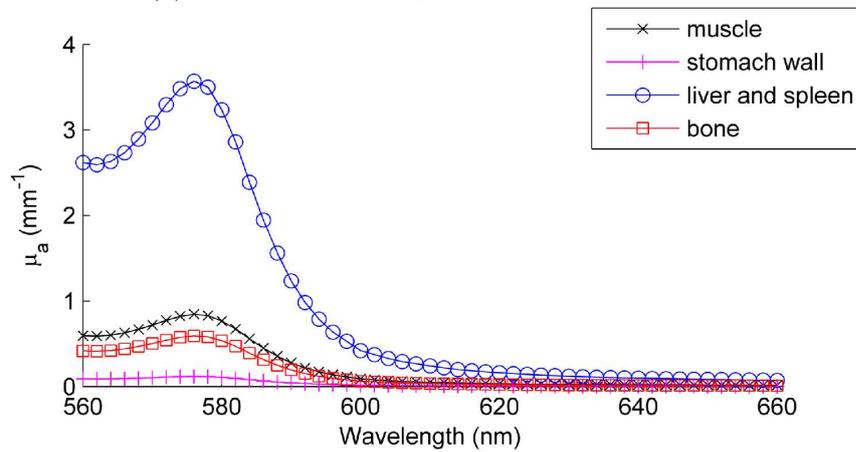
where \mathbf{p} is a point in two dimensions. The bioluminescence distributions that were examined consisted of one, two, or three of these Gaussian sources.

The bioluminescence distributions were parametrised in terms of \mathbf{p} , σ , and I . These parameters were assumed to be independent and the prior distributions for each parameter were designed separately:

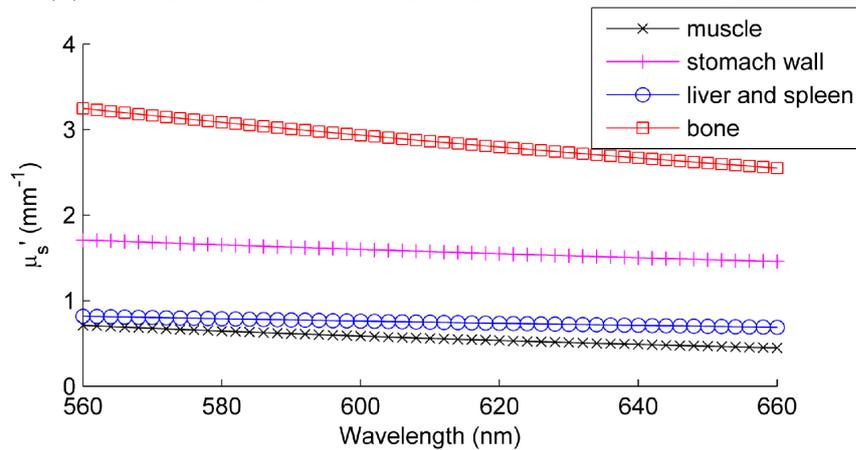
- The prior distribution for \mathbf{p} was chosen to be completely uninformative, and so uniform. This implies that the source centre locations were equally likely to be found anywhere within the imaging region.
- The prior distribution for σ was chosen to be a folded normal distribution, with a mean of 3mm and a standard deviation of 0.5mm.
- The prior distribution for I was chosen to be a folded normal distribution, with a mean of 10 photon/s and a standard deviation of 1.67 photon/s.



(a) Anatomical configuration of test problem



(b) Absorption spectra of organs present within test problem



(c) Scattering spectra of organs present within test problem

Figure 6.1: Test problem used for mutual information estimation. Tissue types were adapted from Dogdas et al. (2007) and tissue optical properties were adapted from Alexandrakis et al. (2005). Images from Basevi et al. (2013a).

For the formula for a folded normal distribution, see eq. (5.13).

The distributions for \mathbf{p} , σ , and I were chosen to be independent, and so the probability distribution for a bioluminescence distribution consisting of N Gaussian bioluminescence sources takes the following form:

$$P(\mathbf{p}_1, \dots, \mathbf{p}_N, \sigma_1, \dots, \sigma_N, I_1, \dots, I_N) = \prod_{i=1}^N P(\mathbf{p}_i)P(\sigma_i)P(I_i) \quad (6.4)$$

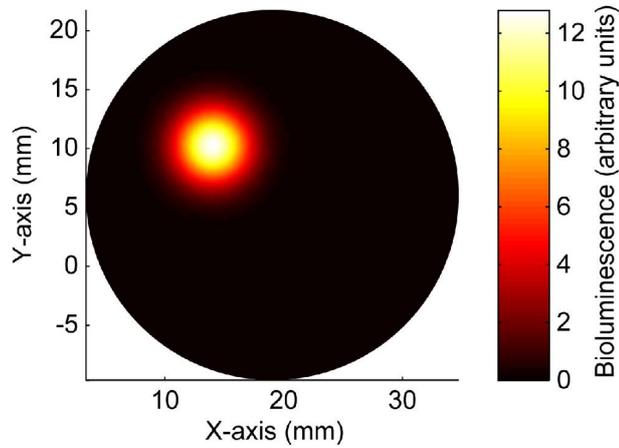


Figure 6.2: A sample from the prior distribution. Image from Basevi et al. (2013a).

6.3.2 Design of the likelihood distribution

The likelihood distribution was designed to represent a system consisting of optical fibres in perfect contact with tissue coupled to perfect 16-bit detectors. Measurement noise was assumed to consist of shot noise and a normally distributed background noise of the form $\mathcal{N}(0, 6)$, and so shot noise dominates. The background noise represents the noise present in a scientific instrument such as that described in Appendix A in ideal conditions. All detectors were exposed for the same time at each wavelength. No constraints were placed on maximum exposure time and this was chosen at each wavelength so that the largest measurement at each wavelength took the value of 90% of the capacity of the detector in the absence of noise. It was also assumed that transmission spectra of the filters used at each wavelength took the form of delta functions.

In order to simplify the mathematics and improve the computational tractability of the problem, the discrete Poisson distribution of shot noise was approximated by the continuous distribution of Gaussian noise.

The noise associated with each measurement was chosen to be independent. The likelihood $P(\mathbf{c}|\mathbf{x})$ is dependent on \mathbf{x} and so \mathbf{x} was used to determine the measurement standard deviations that were applied to the likelihood calculation. The standard deviations, σ_i , of measurement likelihoods were determined as follows:

$$\sigma_i = \sqrt{\mathbf{J}_i \mathbf{x}} \quad (6.5)$$

where \mathbf{J}_i is the Jacobian row corresponding to measurement i . The likelihood takes the following form:

$$\begin{aligned} P(\mathbf{c}|\mathbf{x}) &= \prod_{i=1}^N P(c_i|\mathbf{x}) \\ &= \prod_{i=1}^N \frac{1}{\sigma_i \sqrt{2\pi}} \exp\left(-\frac{(\mathbf{J}_i \mathbf{x})^\top \mathbf{J}_i \mathbf{x}}{2\sigma_i^2}\right) \end{aligned} \quad (6.6)$$

6.3.3 Mutual information calculation

The MI can be calculated from eq. (6.2), and involves the solution of three nested integrals. This calculation is not tractable in general due to the complexity of possible priors and likelihoods. These integrals however can be estimated using MC sampling. Two MC sampling steps are necessary. The value of $P(\mathbf{c})$ can be estimated by applying MC sampling to eq. (5.3), which is found in the denominator of the logarithm in eq. (6.2). The result approximates $P(\mathbf{c})$:

$$P(\mathbf{c}) \approx \frac{1}{N} \sum_{i=1}^M P(\mathbf{c}|\mathbf{x}_i) \quad (6.7)$$

This can then be substituted into eq. (6.2) which enables the MI to be estimated as follows:

$$I(X; C) \approx \frac{1}{M} \sum_{i=1}^M \log \left(\frac{P(\mathbf{c}_i | \mathbf{x}_i)}{\frac{1}{N} \sum_{j=1}^N P(\mathbf{c}_i | \mathbf{x}_j)} \right) \quad (6.8)$$

where \mathbf{c}_i and \mathbf{x}_i are joint samples from $P(\mathbf{c}, \mathbf{x})$, \mathbf{x}_j is a sample from $P(\mathbf{x})$, and M and N are the size of sample sets. This is still computationally expensive but is more tractable. MI was estimated using this method for all the results presented here.

6.3.4 Experiments

MI was estimated for imaging of the specified subject using a single wavelength, two wavelengths, and three wavelengths. Wavelengths ranging from 560nm to 660nm in 10nm steps were examined.

MI was estimated using eq. (6.8). A total of 2.5×10^7 samples were acquired for wavelength sets consisting of a single wavelength, where $M = 5000$ and $N = 5000$ samples were acquired for the outer and inner Monte Carlo sampling operations respectively. A total of 7.5×10^7 samples were acquired for wavelength sets consisting of two wavelengths, where $M = 15000$ and $N = 5000$ samples were acquired for the outer and inner Monte Carlo sampling operations respectively.

The condition numbers of the Jacobian matrices were also calculated in order to compare the optimal wavelength sets using the MI metric and using the condition number metric.

6.3.5 Validation

Validation of these results was performed by extracting samples from the posterior distribution $P(\mathbf{x} | \mathbf{c})$ for a number of measurement sets including a single wavelength and measurement sets including two wavelengths using MCMC sampling. 500 test bioluminescence distributions were sampled from $P(\mathbf{x})$ for each set of measurement wavelengths, and noisy measurement sets were sampled from $P(\mathbf{c} | \mathbf{x})$ for these bioluminescence distribu-

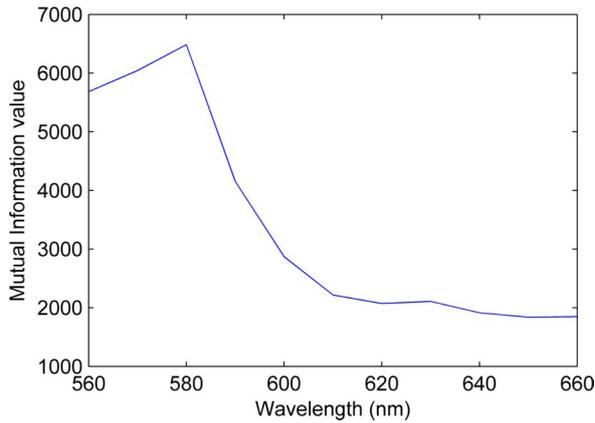
tions. MCMC sampling was then used to generate 2000 bioluminescence samples from the posterior distribution for each of these bioluminescence distribution-noisy measurement set pairs. The differences between the locations, source widths, and source intensities was calculated for each test bioluminescence distribution and associated set of bioluminescence samples. The standard deviation of these differences was also calculated and used as a measure of the width of the posterior distribution for the test bioluminescence distribution.

6.4 Results and discussion

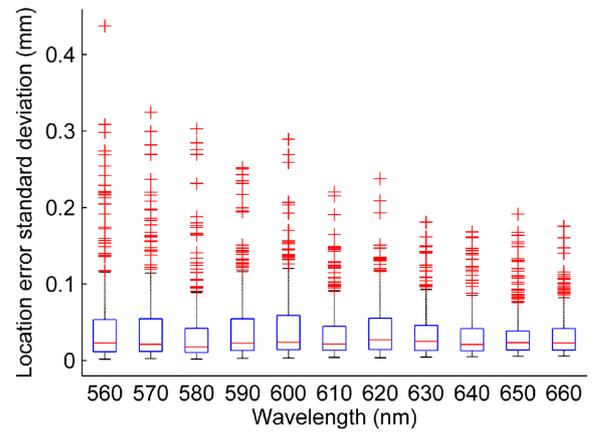
Figure 6.3 shows the MI and validation for the test problem using a single wavelength. There is a peak in the MI at 580nm and a rapid drop-off in the MI between 580nm and 610nm at which point a further increase in the measurement wavelength does not affect MI greatly.

This distribution mirrors that of the absorption coefficient and implies that information content may be dependent on the absorption, or the ratio of absorption to scattering. The distribution of estimated standard deviations for the validation examples indicates that there is no correlation between the MI and the spread of the posterior distribution with respect to location (see fig. 6.3b). Correlations are however visible in the spread of the posterior distribution with respect to intensity (see fig. 6.3d) and in particular to width (see fig. 6.3c). There is a clear reduction in the spread around 580nm for width and a much smaller reduction for intensity. This implies that the additional information present at 580nm relative to a wavelength above 610nm relates more to the bioluminescence distribution width than to the intensity, and more to the intensity than to the location.

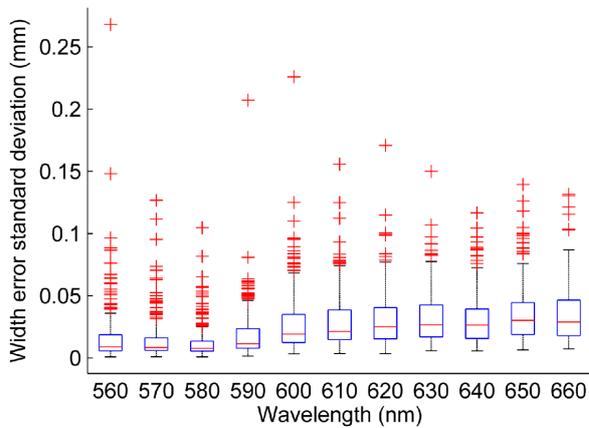
This information is useful, as knowledge of the type of benefit provided by optimising the measurement wavelength can allow a user to make an informed decision about whether the benefit provided by using an optimal wavelength outweighs the cost. Here the optimal



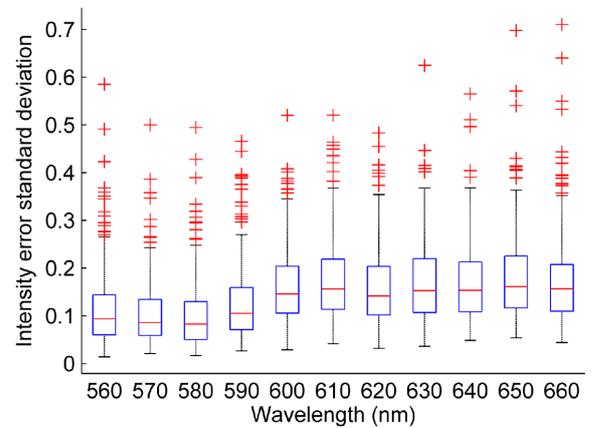
(a) Mutual information for measurements acquired at a single wavelength



(b) Sample location difference standard deviation



(c) Sample width difference standard deviation



(d) Sample intensity difference standard deviation

Figure 6.3: Mutual information and validation for single wavelengths. Images from Basevi et al. (2013a).

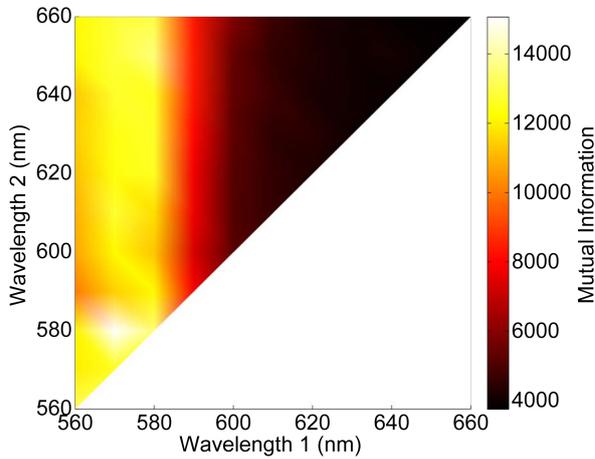
wavelength is 580nm. Potential costs involved in its use include the expense of purchasing a suitable filter and in particular the increase in imaging time necessary to acquire a satisfactory SNR due to the maximum of tissue absorption at 580nm (see fig. 6.1b). It may be best to forgo the increase in width information if source width is unimportant relative to source location and intensity, and minimising imaging time is important due to a limit on sedation duration or metabolism of the bioluminescence substrate.

Figure 6.4 shows the MI for measurement sets consisting of two wavelengths. The optimal measurement set consists of 570nm and 580nm using the criteria of MI. However, the validation data is not consistent with this unlike in the single wavelength case.

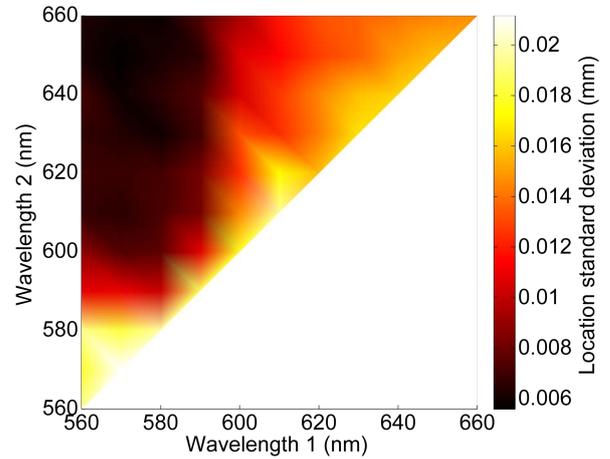
The spread of the posterior distribution is minimised with respect to location, width, and intensity for wavelength sets consisting of one wavelength below 590nm and one wavelength above 600nm. This implies acquiring measurements at a wavelength with relatively high tissue absorption and also acquiring measurement at a wavelength with relatively low tissue absorption (see fig. 6.1b). The MI is in agreement with the validation data that measurement sets consisting of one wavelength below 590nm and one wavelength above 600nm provide more information than wavelength sets with two wavelengths above 600nm, but is in disagreement in the region where both wavelengths are below 600nm.

Figure 6.5 shows a summary of the top twenty wavelength sets consisting of three wavelengths in the context of MI. The MI for these data sets is graphed in fig. 6.5a and the frequency of individual wavelengths is plotted in fig. 6.5b. Eighteen of these wavelength sets contain two wavelengths between 560nm and 580nm, which is consistent with the trend exhibited for dual wavelength sets. The third wavelength is roughly uniformly distributed over all of the higher wavelengths.

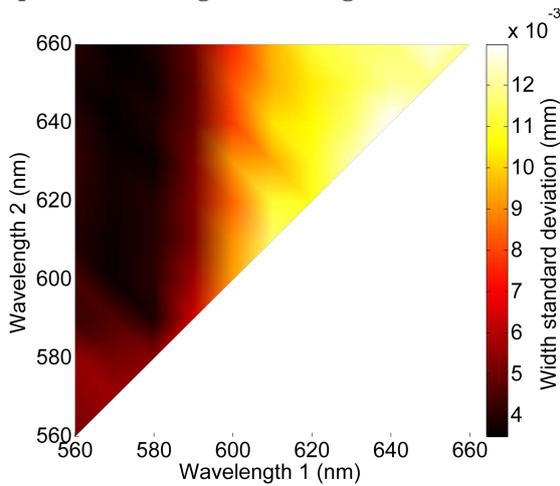
Figure 6.6 shows a condition number analysis of the Jacobians used in the MI calculation. Low condition numbers are optimal and the best single wavelengths according to condition number are the same as those according to MI. Similar trends are observed for measurement sets consisting of two wavelengths in that sets consisting of two high wavelengths perform poorly and good wavelength sets include at least one low wavelength.



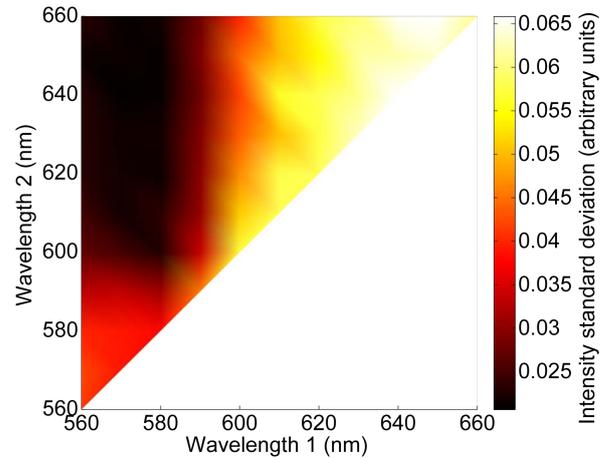
(a) Mutual information for measurements acquired at a single wavelength



(b) Sample location difference standard deviation



(c) Sample width difference standard deviation



(d) Sample intensity difference standard deviation

Figure 6.4: Mutual information and posterior distribution standard deviation estimates for wavelength sets consisting of two wavelengths. The optimal wavelength set consists of 570nm and 580nm using the mutual information metric. Images adapted from Basevi et al. (2013a).

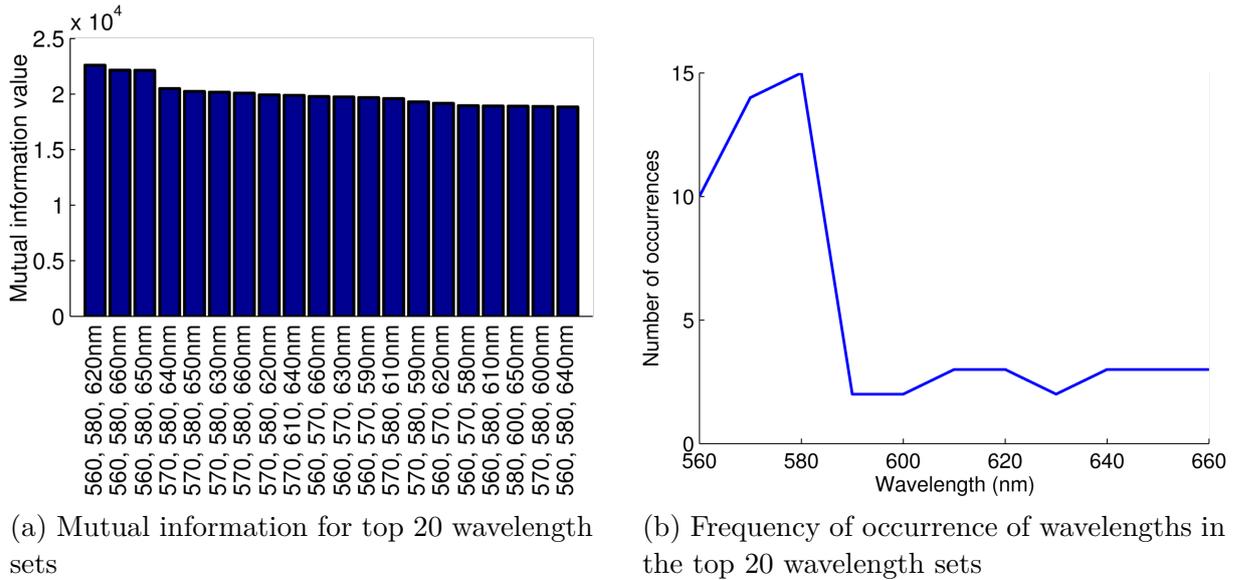


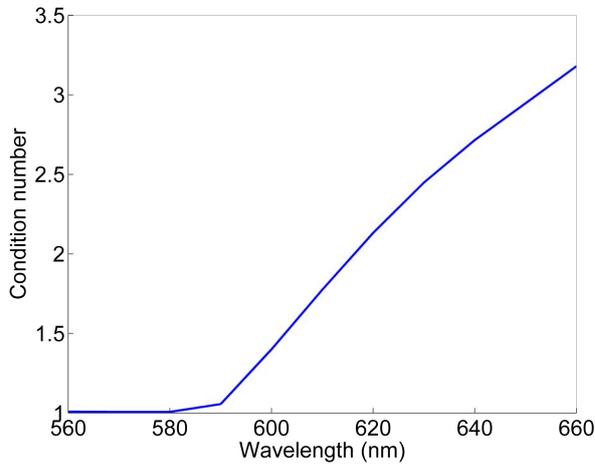
Figure 6.5: Mutual information and wavelength frequencies for the top 20 wavelength sets consisting of three wavelengths.

However, a larger range of low wavelengths are optimal than for the MI metric. Similar trends are again observed in the case of wavelength sets consisting of three wavelengths, where wavelengths below 600nm occur more frequently than wavelengths above 600nm. However, the MI metric prioritises 570nm and 580nm, while the condition number metric prioritises 590nm.

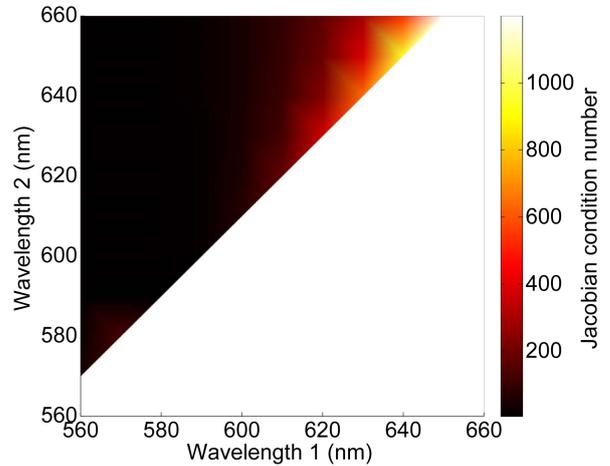
6.5 Conclusions

The use of MI for measurement protocol optimisation shows promise. This type of analysis is advantageous over conventional methods in that it allows many varied types of information to be integrated into a single quantitative measure of measurement performance. This measure is independent of any choice of reconstruction algorithm and so is well suited to imaging systems for which a reconstruction algorithm has not yet been chosen or for which a number of reconstruction algorithms are compatible.

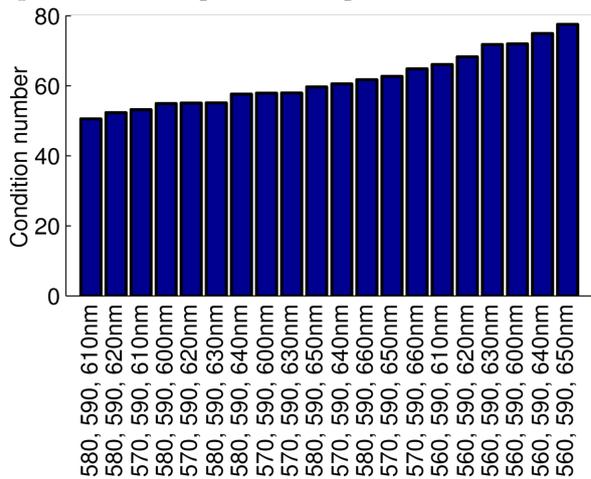
The MI for a single wavelength indicates that 580nm is the optimum wavelength of those examined, which is initially counter-intuitive as 580nm is also the wavelength with the greatest absorption. The distribution of MI for a single wavelength follows in general



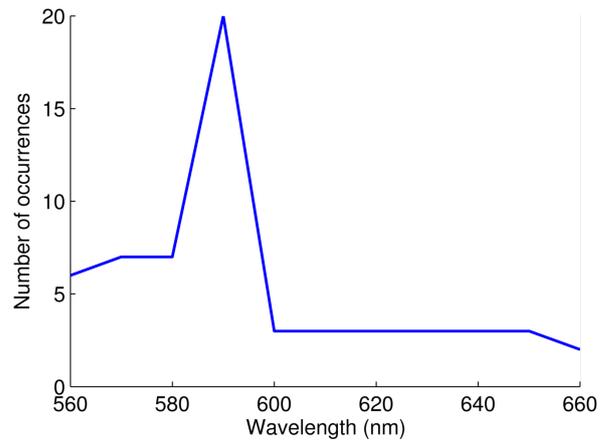
(a) Condition numbers for measurements acquired at a single wavelength



(b) Condition numbers for measurements acquired at two wavelengths



(c) Condition numbers for top 20 wavelength sets consisting of three wavelengths



(d) Frequency of occurrence of wavelengths in the top 20 wavelength sets consisting of three wavelengths

Figure 6.6: Condition numbers of Jacobians for measurement sets consisting of a single wavelength (fig. 6.6a), two wavelengths (fig. 6.6b), and three wavelengths (figs. 6.6c and 6.6d). Low condition numbers are ideal. The optimal single wavelength is 570nm. The optimal set of two wavelengths consists of 580nm and 600nm. The optimal set of three wavelengths consists of 580nm, 590nm and 610nm.

the spectral distribution of tissue absorption. However, it was assumed for this test case that it was always possible to integrate until one measurement reached 90% of the detector capacity. In this case the maximum magnitude and SNR is fixed and the large absorption results in a greater range between the highest and lowest measurements. The attenuation due to absorption is exponential and so the large absorption favours more direct paths between source and detector.

The validation method also suggests that the additional information at 580nm with respect to the higher wavelengths is associated with source width over intensity, and intensity over location. This both supports the trends in MI and provides another source of potentially useful information. Knowledge that the benefit of using a wavelength of 580nm is predominantly in terms of information with regard to source width over intensity and location would allow a practitioner to decide whether these advantages are worth the increase in imaging times necessitated by the large tissue absorption.

The MI for sets of two wavelengths is less supported by the validation data. The MI suggests that measurement sets consisting of one wavelength below 590nm and one wavelength above 600nm provide more information than wavelength sets with two wavelengths above 600nm which is supported by the validation data. However it also suggests that two measurements below 590nm provide further information in contrast to the validation data. This may result from an inadequate number of MC samples used to estimate the MI integrals, but may also result from a characteristic of the posterior distribution that is not measured by the validation data.

Calculation of MI as performed here via nested MC estimation is computationally expensive. Estimating the MI for a single wavelength required approximately 1 hour and estimating the MI for a pair of wavelengths required approximately 6 hours using a four core desktop computer containing 8GB of RAM (see Appendix A.6). It should be noted that the MI estimation for a pair of wavelengths made use of three times the number of samples (7.5×10^7) that were used to perform MI estimation for a single wavelength (2.5×10^7). This examination was conducted on a two dimensional situation

and computational time would increase significantly if the technique were to be applied to a three dimensional situation. The Jacobian matrix is typically the largest object that must be stored in memory and it is not necessary to store all of the samples in memory simultaneously.

There is potential to reduce the computational time of the algorithm through the use of parallelisation. The MI for each wavelength or pair of wavelengths can be calculated independently. Further, within the MI estimation calculation for a wavelength or pair of wavelengths, the MI for each pair of measurements, c_i , and bioluminescence, x_i , can be calculated independently. Consequently the MI estimation process would scale very well to the number of available computer processors and is suitable for execution on a distributed grid of computers. Memory requirements at each node would also be reduced by distributing the computation across a grid of computers.

A major benefit of the use of MI is the principled integration of a number of different types of information. This allows for great flexibility in analysis. For example, it would be possible to integrate a limitation on maximum imaging time if enforced by a particular practical constraint. This may change the situation to the extent where the maximum information is not found at a wavelength of 580nm. It may be that the region of interest changes to a different part of the animal, and MI would allow the integration of the new distribution of internal organs and optical properties. It may be that the bioluminescent cell line is changed, and as a result has different growth properties. MI would allow the integration of this change into the prior distribution without affecting the implementation of application of the technique as a whole. This technique can potentially aid imaging in a number of situations.

CHAPTER 7

CONCLUSIONS

The first aim of this work was to enable BLT image reconstruction via the design and implementation of a surface metrology system to measure the surface geometry of imaging subjects. Surface geometry is necessary to create a physical model of light propagation within the imaging subject. Accurate surface geometry is essential because differences in surface shape of a few millimetres can greatly change recorded light intensity due to the attenuation of tissue.

The work presented in chapter 2 successfully provided this information. The surface metrology subsystem required the addition of two digital projectors, and did not interfere with the configuration of existing optical components. A general optical model was used which allowed the projectors to be placed in any position, further simplifying the process of hardware integration. This general optical model also allowed the two existing 45° mirrors to be integrated into the surface metrology imaging process, which benefited the imaging process in two ways. Firstly, the mirrors extended the field of view of the system, and allowed the surface geometry to be measured in locations visible only through the mirrors. This allowed the creation of a more accurate physical model. Secondly, the surface metrology system also enabled the positions of the mirrors to be calibrated using surface metrology data and this allowed bioluminescence measurements acquired through the mirrors to be mapped onto the physical model. This information also aided the modelling of free space light propagation using the model described in Guggenheim et al.

(2013b).

The system measured surface locations to an accuracy of less than 0.2mm in 82% of cases, when comparing the measured surface of an XPM-2 phantom to manufacturer-provided surface points. A small number of points were associated with much larger measurement inaccuracies of up to 1.2mm. This is likely to be a result of mesh creation inaccuracies due to incomplete surface measurement coverage. The geometry around the phantom ears is complex which prevents complete measurement coverage in those regions. In addition the positions of the optical components and size of the mirrors prevents measurement of the entire length of the phantom and prevents measurement of the underside of the phantom including the lower jaw. Measurements could be extended to the entire length of the phantom by using longer mirrors, but it would not be possible to measure the underside of the phantom without significant changes to the imaging configuration such as by using a glass stage and placing a mirror underneath the phantom as in Chaudhari et al. (2005). This would however necessitate the enhancement of the surface metrology system to allow indirect projection through the mirrors to the underside of the subject as the integration of an additional projector would be difficult.

The completed surface metrology system was used to acquire surface geometry for the phantom which was used in chapters 4 and 5. The qualified success of these reconstructions indicates that this system meets the necessary imaging requirements. The most significant current system limitation lies in its treatment of mirrors, as manual segmentation is necessary to distinguish pixels that correspond to direct imaging and pixels that correspond to imaging through a mirror. This process could be automated via the use of an object to block light from the object reaching the mirrors, but this would require two separate sets of images where the light block is present and absent. A more powerful solution would involve visible markers on the edges of the mirrors which would allow automatic segmentation of the four corners and which would be sufficient to roughly calculate the mirror position and orientation. The projection could then be masked to prevent projection through the mirrors and the mirror calibration could then be refined through

the surface metrology process. This would also allow projection through the mirrors to illuminate more of the imaging subject surface.

The system is also currently limited by the focal depth of the camera and projectors. Gaussian blurring of a sinusoidal pattern does not alter the frequency component but reduces the signal strength, so can be tolerated to an extent. However, blurring can limit the maximum sinusoid frequency that can be used in practice. This problem can be partially addressed through the use of a laser-based projection method which scans a laser beam using mirrors. These projectors do not use lenses and so do not go out of focus, although at large distances the scan lines can become separated or overlapping depending on the dispersion of the laser and the angular resolution of the scanning mirrors. Projections on surfaces at high angles to the camera or projector are subject to a similar challenge in that aliasing of the projected signal occurs. A potential solution to both of these problems would be to use an adaptive scheme where the pattern frequency is spatially varying. This would allow high resolution measurement in regions of the surface that are in focus and at a shallow angle to the camera and projector, and lower resolution but robust measurement in regions for which this is not the case.

The surface metrology method and mirror design in Lanman et al. (2009) may provide better surface coverage and would simplify the treatment of mirrors, but would necessitate a significant redesign of the imaging system to accommodate mirrors and Fresnel lenses for the projectors. This would be worth pursuing if greater surface coverage was necessary but the system performs well for the current application.

The second aim of the project was to enable effective bioluminescence reconstruction in a principled manner through the use of prior information including the surface geometry. The CSCG algorithm was developed to accomplish this aim.

The CSCG algorithm produced promising results, with certain caveats. The purpose of the algorithm was to allow reconstruction of underdetermined BLT problems through the use of prior knowledge about the spatial distribution of the bioluminescence, without requiring user interaction to determine algorithm parameters. The CSCG algorithm

uses techniques from CS to apply prior knowledge about distribution sparsity and uses a measurement noise-based stopping criterion to determine final algorithmic parameters. Reconstruction is also made more efficient by initialising the algorithm for each set of parameters at a position which is a local optimum for the algorithm using a different but similar set of parameters. This potentially initialises the algorithm in the neighbourhood of the new optimum. Conceptually the process of finding solutions for multiple regularisation weights is similar to the L-curve method, but uses a system-noise constraint rather than using the relationship between the measurement discrepancy term and the regularisation term. It is also not necessary for the CSCG algorithm to perform reconstruction using sparsity weightings that are higher and sparsity weightings that are lower than the eventually chosen weighting.

The CSCG algorithm behaviour was consistent with that of a CS algorithm when tested using a typical *in silico* CS reconstruction problem, validating the algorithm design and implementation. The algorithm behaved differently when tested using an *in silico* two dimensional BLT reconstruction problem, but this was expected given the manner in which the properties of the BLT measurement process differ from those of CS measurement processes. Reconstruction performance was limited, with less obvious trends and no sharp algorithm performance transition. The use of Jacobian column normalisation also had a great influence on CSCG performance.

The measurement noise stopping criterion did not influence the algorithm performance as expected when applied to *in silico* CS measurements. The effect on reconstruction quality was small but negative when using a stopping criterion such that the algorithm terminated when measurement discrepancies were an average of one standard deviation away from the measured value. The effect on reconstruction quality was large and negative when this criterion was set to two standard deviations. The metrics most influenced were pixel error, sensitivity and specificity. These effects were not observed for *in silico* BLT measurements with a stopping criterion of one standard deviation, but changes were observed for a stopping criterion of two standard deviations, particularly for nodal er-

ror, standard deviation, sensitivity, and total bioluminescence in the case where column normalisation was employed. The cause is not clear but indicates that column normalisation should be used with care. It also suggests that the reduction in measurement overfitting due to the stopping criteria may not have the desired effect on reconstruction performance.

Further testing of the performance of CSCG on an *in silico* three dimensional BLT problem produced positive results. The *in silico* problem was chosen to replicate imaging of the XPM-2 phantom using a Dartmouth College instrument. As the noise characteristics of the instrument were not known, the noise characteristics of simulations including noise were not provided to the CSCG algorithm and a measurement discrepancy-based stopping criterion was not used. CSCG was compared against two benchmark reconstruction algorithms: Tikhonov regularised GN and Tikhonov regularised NNLS algorithms. A single regularisation weight for the benchmark reconstruction problems was chosen empirically from simulation results and then fixed to replicate the lack of prior knowledge of these regularisation weights in real situations. The results indicate that CSCG provides good reconstruction quality for measurements that contain no noise, 1% noise, and 5% noise, without requiring the choice of a regularisation weight. It should be noted that no measurement noise-based stopping criteria was used due to the absence of knowledge of instrument noise statistics, and so the CSCG solution overfit for noise to no obvious detrimental effect. NNLS displayed similar localisation ability and lack of artefacts to CSCG, but with reduced compactness. The localisation ability can be explained by the fact that the conventional NNLS algorithm can behave like a CS algorithm (Foucart and Koslicki 2014), and the reduction in compactness is a result of the regularisation used.

CSCG reconstruction of XPM-2 phantom data acquired using a Dartmouth College instrument was performed in the absence of detailed knowledge about the noise characteristics of the measurement instrument. Consequently the reconstruction was performed without the use of a covariance matrix and a measurement noise-based stopping criterion was not applied. The CSCG algorithm outperformed the other two algorithms in

reconstructing and resolving the two luminescent sources, which demonstrates that the algorithm is capable of reconstructing bioluminescence from real measurements. The reconstruction from NNLS was less compact than that of CSCG and the reconstruction of the lower source was inferior. GN did not successfully reconstruct either of the two sources. It should be noted that NNLS was also capable of resolving both sources and the reconstructions from CSCG and NNLS both contained artefacts.

The behaviour of CSCG in the *in silico* experiments was interesting. In particular, the behaviour of the algorithm differed from that expected of a CS algorithm in the case where measurement noise was present and a measurement discrepancy-based stopping criteria was used. CS theory states that there is a value of the regularisation weight such that the measurement discrepancies in the final reconstruction are accounted for by the statistics of the measurement noise model, and this is a solution which minimises the contribution of measurement error to the reconstruction result. However, the CSCG algorithm showed decreased performance in terms of pixel error, sensitivity, and specificity in both the *in silico* CS experiments and in the *in silico* BLT experiments. If this was the case only in the BLT experiments then the trend could potentially be explained by the difference in the properties of the BLT measurement process compared to classical CS problems. However, the observation of this effect in the CS experiments suggests that this may be a property of the algorithm design and further investigation may be illuminating.

The case study in chapter 4 demonstrates that it is possible to reconstruct the same two sources using surface geometry measurements from the system detailed in chapter 2 and the CSCG reconstruction algorithm. This imaging system uses a CCD camera to acquire measurements and interference filters to provide spectrally resolved information resulting in a large number of measurements. Consequently, it was also necessary to integrate measurement compression techniques into this work in order to allow reconstruction using a computer with limited memory capacity. Reconstruction was successful for source A, source B, and source C (sources A and B switched on simultaneously), but reconstruction performance was heavily dependent on the choice of measurement compression technique

and whether or not column normalisation was employed.

Reconstructions in the case where a measurement summation compression technique was used were uniformly inferior to those where a random projection compression technique was used. Reconstructions without column normalisation produced very diffuse sources, whereas reconstructions using column normalisation produced more compact sources. However, this was also associated to a large increase in the intensity of surface artefacts which typically were more intense than the true source. It was necessary to use random projection compression, column normalisation, and discard the surface nodes of the reconstruction in order to produce results of similar compactness to the reconstructions in chapter 3. This is despite the fact that neither measurement compression nor column normalisation were used in chapter 3.

This can be explained by the difference in measurement coverage in the two imaging systems. The Dartmouth College system has superior measurement surface coverage and individual measurement efficiency but inferior measurement density. The benefit of the extra surface coverage, specifically encompassing the underside of the phantom, may account for the difference in performance. However, a benefit of using the Birmingham system is that the characteristics of the measurement components is known and so it was possible to calibrate the physical model and provide quantitative reconstructions. The sizes of the sources are unknown but the intensities of the sources are known, and the reconstructions were consistently within two orders of magnitude of the true values. This may seem a large discrepancy, but is understandable in the context of the influence of geometry errors on surface intensities, and the use of ideal manufacturer-provided imaging component calibration data.

The surface artefacts were located far away from measurement locations, and so signal from these artefacts would appear as nearly uniform at the detector locations due to scattering. It is also possible that the measurement calibration process did not account fully for an unknown background signal or systematic measurement error, which may affect the measurements in a similar manner to the surface artefacts.

The attenuation coefficient of the phantom is on the order of 1mm^{-1} , and so small changes in the measured surface can lead to large changes in the modelled surface intensity. The sum-millimetre accuracy of the surface system (see fig. 2.9c) should minimise this discrepancy. The manufacturer-provided calibration data is also likely to deviate from the true values because the components are now several years old and do degrade over time. The free space propagation model also does not take into account the angular dependence of the interference filters, which includes a shift in central wavelength and efficiency and is not trivial to model. A natural step in improving the reconstruction quality would be to perform a calibration of the optical components, measure the angular dependence of the interference filters, and integrate all of this into the free space model.

The CSCG algorithm proved capable of producing good results and resolving two separate internal bioluminescent sources. However, several parameters with no physical justification had a strong effect on the performance of the algorithm, such as the use of column normalisation. The second aim of the project was not fully fulfilled as a result. It was desirable to design a reconstruction method that either did not include such parameters or was insensitive to the choice of these parameters. It was also desirable to produce a metric of reconstruction uncertainty which would ideally allow principled segmentation of the true bioluminescent sources from features like surface artefacts.

Chapter 5 addressed the second and third aims of the project. The third aim was to provide measures of the reconstruction reliability to aid interpretation of results. This chapter detailed an effort to avoid the use of algorithm parameters with no natural physical meaning by integrating any available prior information on the problem into a universal Bayesian representation. This universal representation means that problem analysis and reconstruction can be performed in the same manner regardless of the degrees and types of prior information available. The influence of arbitrary non-physical factors is then minimised in the reconstruction and analysis processes. An additional benefit of the Bayesian representation is that it contains additional information to the MAP solution which is accessible and this can aid in the interpretation of the MAP solution or distribution ex-

pectation. It should be noted that one potentially important source of uncertainty was not integrated into the Bayesian representation: the uncertainty present in the physical models of light propagation. This is a limitation of the presented work.

A MCMC technique was applied to obtain the bioluminescence distribution expectation and standard deviation. The standard deviation of the distribution is a measure of the uncertainty in the parameter estimation, and addresses the third aim of the project. This technique operates on the solution representation and the posterior distribution directly, and so minimises the influence of non-physical parameters. The solution representation was that of a truncated Gaussian distribution in this case, which is justifiable because of the known construction of the sources in the XPM-2 phantom. However, this representation does not generalise well to arbitrary bioluminescence distributions and so one area of potential future work would be to formulate a general bioluminescence distribution which incorporates biological and chemical information to elegantly constrain the solution space.

The reconstructions produced distributions with good localisation and intensity, except for the reconstruction of source C for the 29th of July data set. The reconstruction of source C for the 1st of August data set successfully reconstructed both sources with similar intensities. These results were superior to the CSCG reconstructions in the previous chapters (chapters 3 and 4) with the previously mentioned exception and show that it is possible to obtain high quality reconstructions with the involvement of a minimum of unphysical parameters. Interestingly, the reconstructed source intensities were similar to those in chapter 4 and underestimated the source intensities by more than an order of magnitude, suggesting that the underestimation in chapter 4 may have been a result of the physical model and not the CSCG algorithm.

The exception in this chapter was the source C expectation for the 29th of July data set. This was a consequence of the number of measurements and resulting large gradients in the posterior distribution. The MCMC algorithm is by necessity a local technique in the absence of a good approximation of the posterior distribution that can be used as a proposal distribution. The MCMC algorithm can then require a large number of iterations

to leave a local optimum as a consequence of the steepness of the well of attraction. This is a known issue with MCMC algorithms and certain types of distributions. A non-local component was added to the proposal distribution and four independent Markov chains were used in order to compensate for this challenge, but were insufficient in this case. Simple solutions to this problem would be to use a larger number of independent Markov chains, which would each in effect sample a different region of the posterior distribution, or to run the existing Markov chains for a much longer period of time and discard more samples to reduce correlations. However, doing so would require an unknown increase in computational power or execution time and this may not be practical.

The standard deviations of the posterior distributions provided by the algorithm gave estimates of reconstruction reliability. These were observed to be related to the success of the reconstruction, with the blurrier reconstructions of sources A and B using the 1st of August data sets producing larger standard deviations than the reconstructions of the 29th of July data sets, and the failed reconstruction for source C producing standard deviations ten times the size of that of the successful reconstruction. However, these standard deviations are unrealistically small with respect to the expectations. This may result from the parametrisation of the bioluminescence distribution acting as an excessively strong prior, but may also result from the lack of consideration of the uncertainty in the physical models themselves. Incorporating physical model uncertainty into the Bayesian representation would form an interesting and potentially important piece of future work.

The final piece of work presented here addressed the fourth aim of the project, which was to use prior knowledge to improve the measurement process itself, in addition to improving the way that measurements are used. It explored the potential of the Bayesian representation to go beyond image reconstruction and aid in instrument design, or potentially in optimisation of imaging protocols. This was performed by examining and quantifying the dependence between the space of bioluminescence distributions and the space of bioluminescence measurements in the form of mutual information. This provides an objective metric of comparison between two instrument designs that incorporates mea-

surement noise characteristics and information on the types of bioluminescence distribution that will be imaged, in addition to the measurement sensitivities. Crucially, this metric is also independent of the use of any reconstruction algorithm. It gives a measure of the amount of available information. This is an upper bound in a sense, and how much of that information is accessible to the reconstruction process then depends on the reconstruction algorithm chosen to perform the eventual reconstruction.

The work presented in chapter 6 examined the effect of measurement wavelength on mutual information. The width of the posterior distributions was also estimated using MCMC techniques from chapter 5 and used as a validation mechanism. If the posterior is assumed to be a unimodal distribution then the mutual information should be negatively correlated to the width of the distribution. However, this assumption breaks down when the distribution consists of two narrow optima that are greatly separated as the mutual information here would be high but the distribution width would also be high.

In the case of a single wavelength, the results showed a strong preference for the wavelength of 580nm, which was reflected in the standard deviations of bioluminescence intensity and width of the posterior distribution. Interestingly the correlation with source location was weaker, suggesting that the choice of wavelength is less important if measuring source locations is most important. The condition number metric produced similar results, and favoured the wavelength of 590nm.

In the case of two wavelengths there was broad agreement between the mutual information metric and posterior standard deviations on good pairs of wavelengths, but disagreement in the region where both wavelengths were below 600nm. The cause is unknown, and an exploration of the reasons behind this would be an important piece of future work. It is possible that there are non-linear correlations between the bioluminescence location, intensity, and width that the standard deviation cannot capture. The mutual information metric favoured the wavelength set of 580nm and 590nm, which is surprising because intuitively one would expect two very different wavelengths to provide more information than two very similar wavelengths. Here again the condition number

was in broad agreement, and favoured the wavelength set consisting of 580nm and 600nm.

An interesting pattern emerged in the case of three wavelengths for the best sets according to the mutual information metric, where two wavelengths were almost always chosen from the region between 560nm and 580nm, and the third wavelength was chosen almost uniformly from the remaining higher wavelengths. This suggests that two low measurement wavelengths are sufficient to preserve most of the information obtainable from that region, which then makes measurement from the high wavelength region attractive. A similar trend was also observed for the condition number metric in terms of a low wavelength region with two measurement wavelengths and a high wavelength region with one wavelength, although 590nm was preferred above the other wavelengths, while the mutual information metric favoured 580nm and did not favour 590nm.

The similarity between the mutual information metric and the established metric of the condition number suggests that the mutual information metric is working as hoped and picking out good wavelengths. However, this method incorporates a range of prior information on the measurement instrument and the bioluminescence distribution and so can potentially be used to probe a variety of questions without any changes to the underlying method, whereas the condition number metric is far less flexible.

However, the mutual information estimation process was computationally expensive and this limited the practical application of this implementation of the technique to a two dimensional problem. This limitation could be remedied through an increase in computational power or through improvements/optimisations to the estimation method. Doing so would form an important strand of future work as it could be crucial to expanding the application and impact of this technique.

The work in this chapter has only scratched the surface of the information available through analysis of the Bayesian representation of the problem. The effect of measurement locations, maximum exposure time limitations, tissue properties, bioluminescence emission spectra, internal optical property distributions, bioluminescence injection sites and other variables can be investigated using this technique. The current major limitation

is the computational expense of mutual information estimation using the method in its present form. Optimisation of the method and integration of analytical techniques where possible are immediate directions of future work.

In summary, the work presented in this thesis has taken steps towards the three aims of this project. It has enabled surface geometry measurement for an imaging system which allows the use of physical models of light propagation in tissue, and these physical models have been integrated with prior knowledge about the imaging system and about the space of bioluminescence distributions through reconstruction and analysis algorithms to enable reconstruction of bioluminescence distributions and estimation of other properties of interest. This is an advance in the state of the art in the field of BLT. The use of a Bayesian framework enabled principled regularisation which minimised the use of unphysical algorithmic parameters, and allowed the quantification of the information content in bioluminescence imaging instruments and protocols. It has the potential to change the way that BLT imaging and studies are conducted if issues of computation cost can be resolved. It is hoped that the adoption of a Bayesian framework and description of the biological and physical imaging problem may eventually enable efficiencies for all aspects of the process of a BLT study, from the biological experiment design to the imaging process, and to the way in which results are analysed and conclusions drawn. These efficiencies will allow investigators to learn more with fewer imaging subjects.

Appendices

APPENDIX A

THE MULTI-MODALITY IMAGING SYSTEM USED IN THIS WORK

A.1 Contributions to this work

The multi-modality DOT, BLT, and surface metrology instrument detailed here was primarily designed, constructed, and programmed by James Guggenheim (Guggenheim 2014). The surface metrology subsystem was designed, integrated and implemented by the author jointly with James Guggenheim. For more details see chapter 2. It should be noted that *in simulacra* imaging was also performed using the system detailed in Leblond et al. (2011) in chapter 3. The measurement and calibration processes was not performed by the author.

A.2 Introduction

A variety of imaging instruments exist that are capable of performing surface metrology measurements, and BLT measurements. The primary instrument (Guggenheim et al. 2013a) used in this work is a multi-modality imaging system capable of spectral continuous wave DOT, spectral BLT, and non-contact structured light surface metrology. This system was designed to collect sufficient information to reconstruct bioluminescence distributions corresponding to small animal subjects.

Surface metrology measurements are necessary to define the surface geometry of the subject, which itself is used in physical modelling of light propagation and to calculate the surface light distribution from camera measurements. Calculation of the surface light distribution requires knowledge of the optical properties, locations, and orientations of the optical components within the instrument and a physical model of light propagation from the surface of the imaging subject to the camera sensor.

DOT measurements are necessary to perform physical modelling of the propagation of light through the tissue of the imaging subject. These measurements enable reconstruction of the tissue optical properties of the imaging subject in terms of absorption and scattering. These optical properties may be further divided into chromophore concentrations and scatterer properties using an approximation of Mie theory (Dehghani et al. 2008; Mourant et al. 1997) if the DOT measurements are spectrally resolved.

The primary imaging instrument used in this work is described here. The system hardware is detailed and the optical properties of the components are shown. The software used to control the hardware and perform imaging is described, and the hardware used to reconstruct bioluminescence images is detailed. The optical and luminescence properties of the optical phantom used to perform *in simulacra* imaging are also shown.

A.3 Instrument hardware

A labelled image of the system can be seen in fig. A.1. The system contains a number of components that are modality-specific, but some components are used for all modalities. The most significant of these is the C9100-14 ImageEM Electron Multiplying Charge Coupled Device (EM-CCD) camera (Hamamatsu Photonics K.K., Hamamatsu City, Japan). The 1 megapixel CCD is capable of a number of imaging and binning modes and these are chosen to best enable each application. Binning modes are employed to reduce integration time for BLT and DOT due to the low signal. The Electron Multiplying (EM) mode of the camera was not utilised during this work as the signal to noise ratio for the EM mode is

inferior to that of the normal mode above a certain signal level, and the use of a phantom allowed the integration times to be increased to compensate for the low signal of optical tomographic methods. This is also the case for the intended final imaging subjects of the system as they will be sedated and hence still.

A sample stage was mounted to a L490MZ Motorised Lab Jack (Thorlabs, Ely, United Kingdom) which was fixed to the bottom of the system. This enabled the subject to be raised and lowered so as to place the camera's focal plane in an optimal location.

A FW102c Motorised Filter Wheel (Thorlabs, Ely, United Kingdom) was also utilised for all modalities. This contained up to 5 changeable Thorlabs 10nm full-width-half-maximum (FWHM) interference filters in the region between 500nm and 850nm. The filters were primarily used in the optical tomography modalities, but the 580nm filter was also used for the surface metrology subsystem as an attenuation mechanism due to the strength of the associated projector light sources.

A NT59-871 25mm Compact Fixed Focal Length visible-near infra-red (VIS-NIR) Lens (Edmund Optics, York, United Kingdom) was connected to the filter wheel and is used to form the image on the camera CCD.

The final optical components common to all of the imaging modalities are the N-BK7 75mm Enhanced Aluminum Coated Right Angle Mirrors (Edmund Optics, York, United Kingdom), which have greater than 95% reflectance between 500nm and 700nm and are used to increase both the coverage and the signal-to-noise ratio of measurements. The use of mirrors placed a number of restrictions on the surface metrology system and precipitated the development of a robust and effective method of measuring surface geometry both directly, and indirectly through the mirrors.

Both the surface metrology and DOT systems use specialised components in the form of projectors. In the case of the surface metrology system two Pocket Projector MPRO120 (3M United Kingdom, Berkshire, United Kingdom) liquid crystal-based projectors were used to generate structured light patterns. These were fixed in place above and on either side of the imaging subject and were positioned so as to provide maximum coverage of

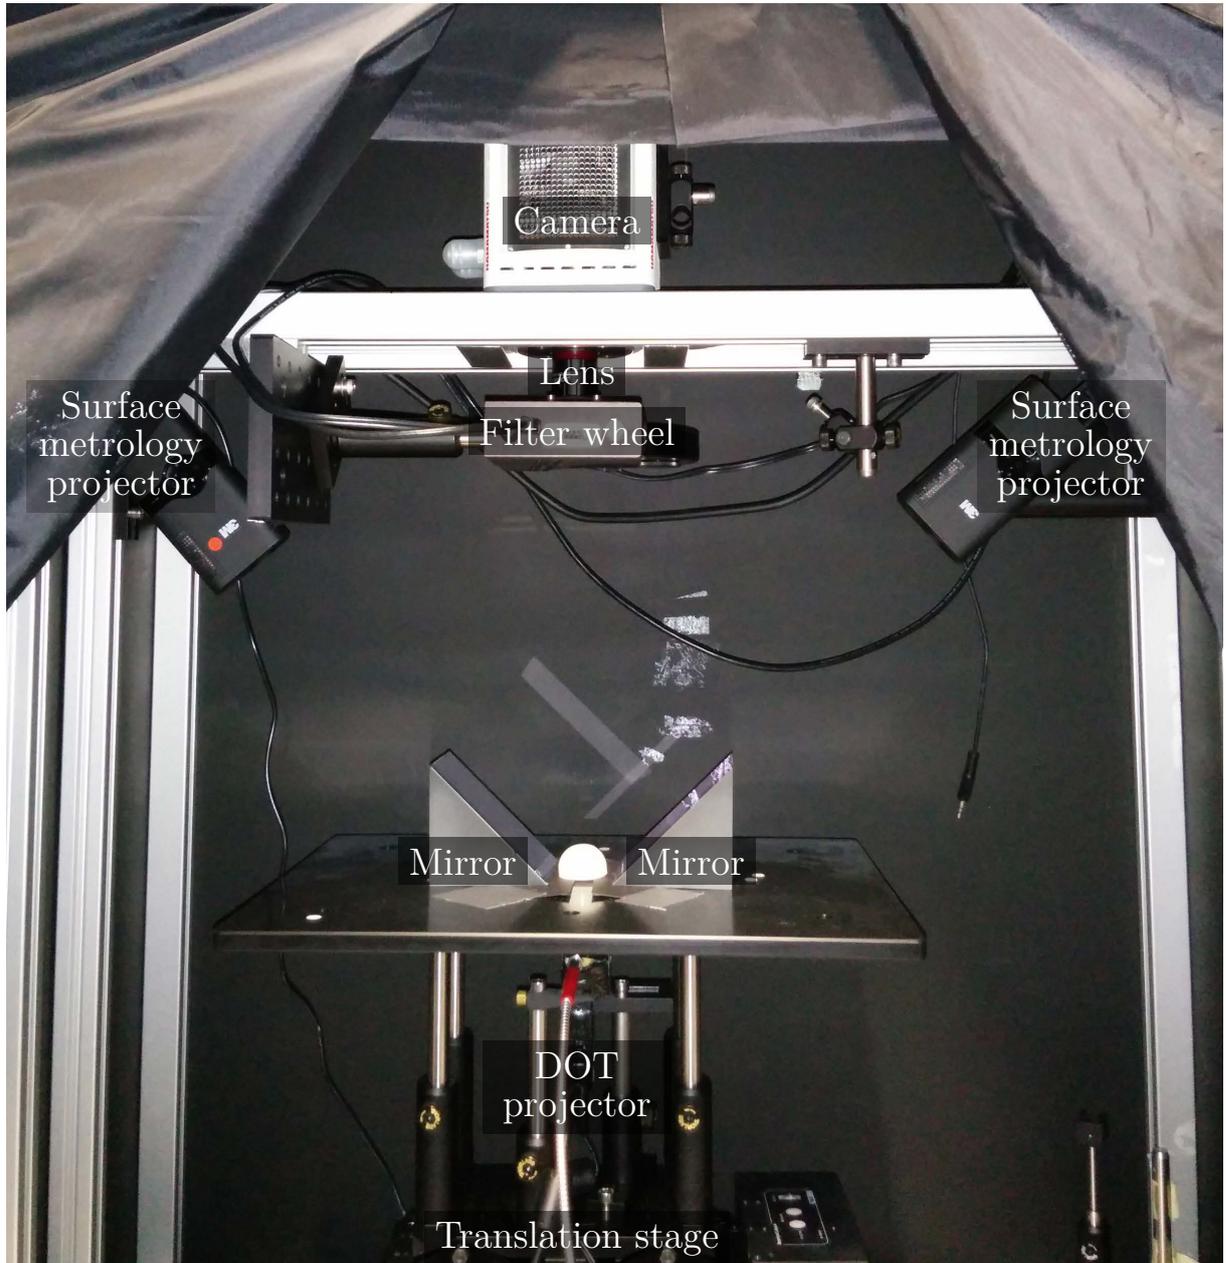


Figure A.1: The multimodal imaging instrument developed by James Guggenheim. It consists of a EM-CCD camera, lens, filter wheel, 10nm bandwidth filters, two 45° aluminium-coated mirrors, a translation stage, two digital projectors for surface metrology, and a digital projector coupled to a halogen light source for DOT. The back and side walls of the system are painted black, and the front of the system is covered by a black curtain. The components visible here are connected to a desktop computer (not visible) running Microsoft Windows, and the system is controlled by software written in LabVIEW and MATLAB.

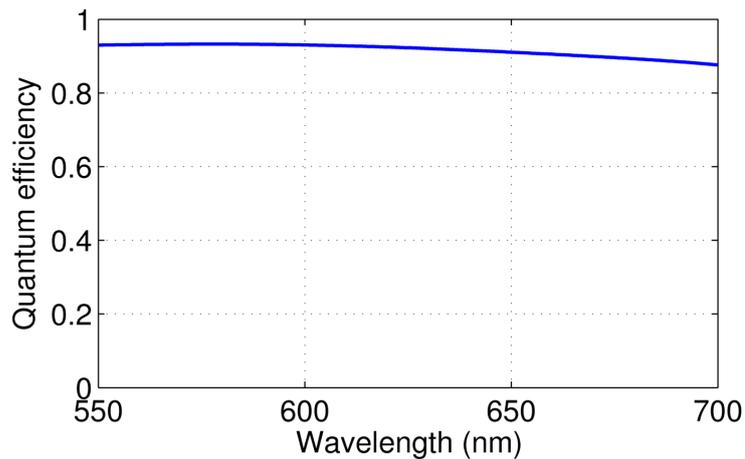


Figure A.2: C9100-14 ImageEM-1L EM-CCD camera quantum efficiency

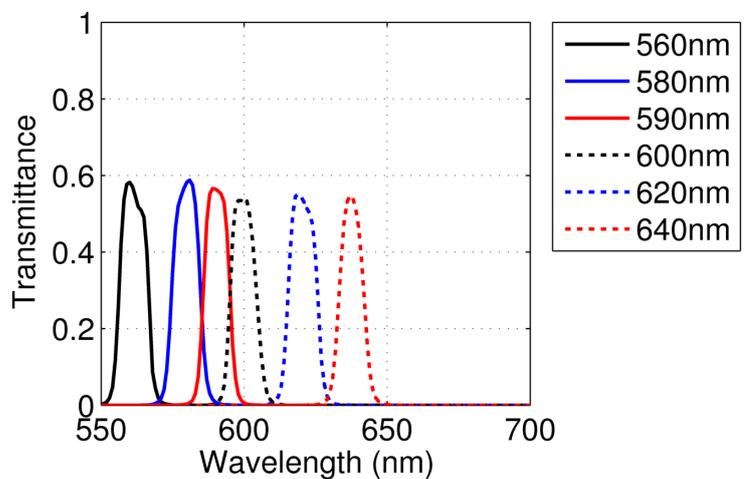


Figure A.3: 10nm bandwidth interference filter transmittances

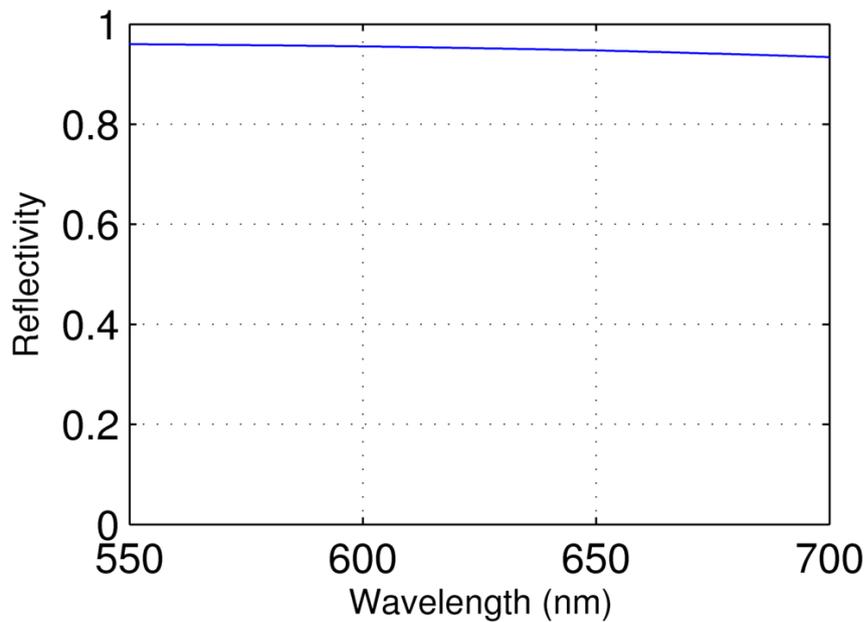


Figure A.4: Mirror reflectivity

the subject. White light patterns were used, and projection intensities were calibrated to produce a linear response to image intensity.

The diffuse optical tomography subsystem used a modified (Chen et al. 2010; Venugopal et al. 2010) Pico PK-102 Pocket Projector (Optoma, London, United Kingdom) mounted underneath the platform. This was coupled to a HL-2000-FHSA Tungsten Halogen Light Source (Ocean Optics, Oxford, United Kingdom). No filter was placed between the light source and the projector, or between the projector and the imaging subject.

A desktop computer was used to operate the various hardware components and perform processing. Four video outputs were required to operate the computer and provide signal to the three projectors present in the system, and multiple video cards were used to provide this capability. The computer hardware included an Intel i7 processor and 16GB of RAM.

A.4 Instrument software

The Microsoft Windows 7 operating system was installed in the desktop computer, and the instrument hardware was operated using the LabVIEW (National Instruments, Texas, United States of America) and MATLAB (MathWorks, Massachusetts, United States of America) software packages, in conjunction with drivers and software bundled with the individual hardware components. LabVIEW provided the ability to manipulate hardware to individually or sequentially acquire data in surface metrology, DOT, and BLT imaging modes, whilst MATLAB was used to perform the calculations necessary to manipulate projected illumination and to save data.

The Nirfast software package (Dehghani et al. 2008) was used to model the propagation of light through tissue. The implementation of the forward model within Nirfast was modified using the reciprocity principle (Dehghani et al. 2008) to allow efficient calculation of measurement sensitivity matrices for BLT. The reciprocity principle typically reduces sensitivity matrix calculation computational cost in the case of BLT because the number

of unknowns that parameterise the bioluminescence distribution is normally much larger than the number of measurements.

A.5 The free space light propagation model

The process of mapping camera measurements onto the imaging subject surface geometry is performed using the surface metrology data and a free space model of light propagation (Guggenheim et al. 2013b). This model allows the integration of measurements from multiple views into a single reconstruction and enables quantitative reconstruction.

Surface metrology data provides the surface geometry of the imaging subject in terms of the three dimensional coordinates of a set of surface points which are visible to the camera directly or through one of the two system mirrors. It should be noted that surface points may be visible both directly and through one of the mirrors, effectively providing two measurements for the same location if the differences in imaging sensitivity due to differing imaging distances and angles can be accounted for. These surface points are provided in the coordinate system of the camera regardless of whether the surface points are imaged directly or through one of the mirrors because the position of the mirrors is calibrated during the surface metrology imaging process.

The measured set of surface points can be converted into a surface or volume mesh using existing algorithms and software (Cignoni 2014; Kazhdan et al. 2006; Schoberl 2014). The creation of a surface or volume mesh allows estimation of the surface normal at the points of measurement which in turn allows correction of the camera measurements to account for the angle of the surface with respect to the camera.

Other aspects of the system affect the intensity of camera measurements. The instrument lens collects rays of light originating from the subject surface and focuses them onto the camera CCD. The lens has a fixed location and focal length and so light from surface points in the focal plane of the system may map to single CCD pixels but light from surface points outside of the focal plane may map to a number of pixels. The lens is also

of finite radius, and so can collect only a small portion of the surface light. The optical filters and mirrors are also of finite size and so also reduce the portion of the surface light collected. This effect is not uniform across the CCD and manifests as vignetting in the simplest case.

The free space light propagation model described in Guggenheim et al. (2013b) was used to calibrate CCD measurements for these effects. This model produces a free space transfer matrix which maps surface light onto CCD measurements, taking into account the sizes, positions, and orientations of the system optical components as well as the shape, position, and orientation of the imaging subject. This matrix does not include the spectral dependence of the system, which can be accounted for using a multiplicative factor.

The spectral dependence of the system is actually a function of the angle of the incident rays upon the optical filters. This angle affects both the transmission efficiency and central wavelength of the filter. This is a consequence of the use of interference filters (Lissberger and Wilcock 1959; Pollack 1966). However, this has not been accounted for within the free space model due to the complexity of modelling this dependence.

The free space transfer matrix \mathbf{T} provides a mapping between surface light \mathbf{y} and camera measurements \mathbf{c} of the form $\mathbf{c} = \mathbf{T}\mathbf{y}$, ignoring other contributions to system sensitivity. The measurements \mathbf{c} are acquired by the system, and \mathbf{y} is required to reconstruct bioluminescence. It is necessary to reconstruct \mathbf{y} from \mathbf{c} but this task may be complicated by the properties of \mathbf{T} , which may be underdetermined. Regularisation may be used to reconstruct \mathbf{c} to an unknown degree of accuracy, or \mathbf{T} can be integrated into the description of the bioluminescence reconstruction problem directly.

A.6 Image reconstruction hardware

Data acquired by the imaging instrument were processed primarily using a Microsoft Windows-based desktop computer containing an Intel i7-860 processor, 8GB of RAM,

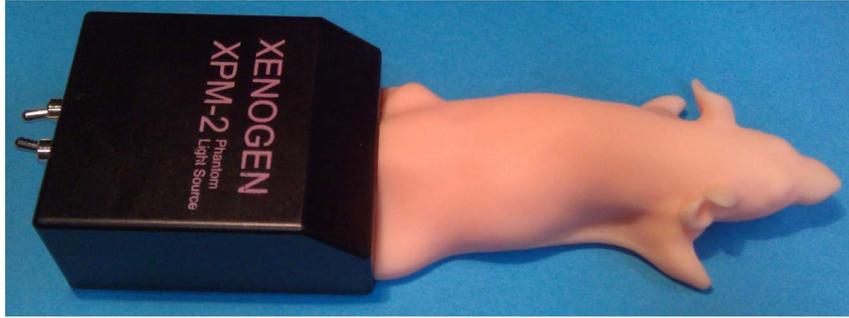
and an nVidia GeForce 570 Graphics Processing Unit (GPU). Reconstructions for chapters 3 and 4, and MCMC calculations for chapter 5 were also executed on this desktop computer. Validation calculations for chapter 3 and MI estimation for chapter 6 were executed on a Linux-based cluster within the School of Computer Science of the University of Birmingham.

A.7 Image reconstruction software

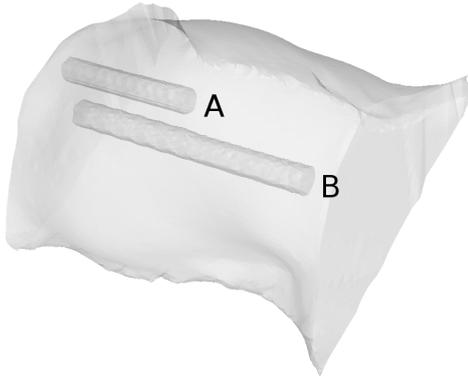
Image reconstruction software was primarily written in MATLAB. MATLAB was used because of the reduced developer time investment compared to languages like C which execute rapidly but require many more lines of code than MATLAB to execute the same task. Certain MC sampling operations were rewritten in CUDA and executed on a GPU to enhance performance.

A.8 XPM-2 Phantom Mouse

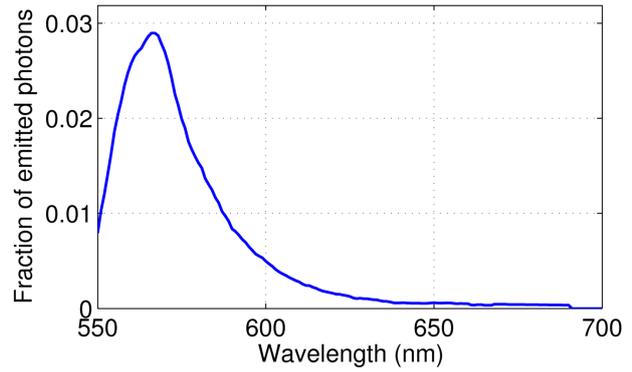
In simulacra experiments for the majority of this work were conducted using the XPM-2 Phantom Mouse (PerkinElmer, Massachusetts, United States of America). The XPM-2 is an optical phantom shaped like a mouse that is sold for calibration of the IVIS preclinical imaging system (PerkinElmer, Massachusetts, United States of America). It is made of a polymer resin that contains pigments that replicate the absorption and scattering of tissue to a limited extent. It also contains two optical fibres that are fed by LEDs, which approximate two bioluminescent point sources within tissue. The locations and strength of the light sources, and the shape of the surface, have been provided by the manufacturer. The phantom is 83mm long, and 30mm wide and 23mm high at the chest.



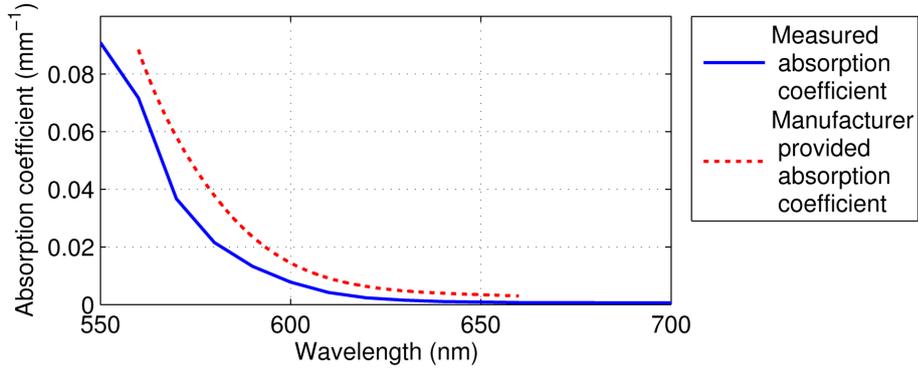
(a) XPM-2 Phantom Mouse



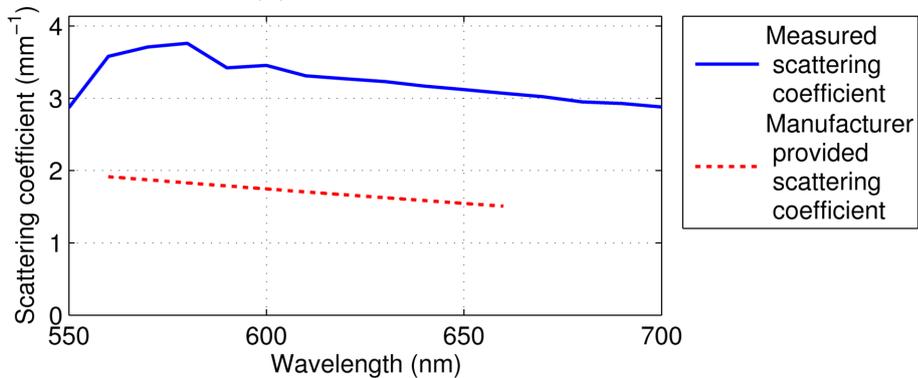
(b) XPM-2 source optical fibre tunnels



(c) XPM-2 source emission



(d) XPM-2 absorption coefficient



(e) XPM-2 scattering coefficient

Figure A.5: XPM-2 Phantom Mouse photograph and optical properties. The optical properties were measured by James Guggenheim (Guggenheim 2014). Image from Basevi et al. (2013b).

Source	Intensity (photon/s)
A	8.09×10^{10}
B	8.4×10^{10}
C (A and B simultaneously)	1.649×10^{11}

Table A.1: Intensity of XPM-2 phantom sources, as provided by the manufacturer. Source A is the higher source.

A.9 Conclusions

The system detailed in this chapter is capable of surface metrology, DOT, and BLT imaging. It reuses components for multiple modalities where possible, which reduces the occupied space, complexity, and cost of the system. The use of the system camera for surface metrology, DOT, and BLT provides an additional advantage as the camera calibration and surface geometry measurements in the coordinate system of the camera can be directly applied to DOT and BLT imaging.

LIST OF REFERENCES

- Ahn, S., A. J. Chaudhari, F. Darvas, C. A. Bouman, and R. M. Leahy (2008). “Fast iterative image reconstruction methods for fully 3D multispectral bioluminescence tomography”. In: *Physics in Medicine & Biology* 53.14, pp. 3921–3942.
- Alexandrakis, G., F. R. Rannou, and A. F. Chatziioannou (2005). “Tomographic bioluminescence imaging by use of a combined optical-PET (OPET) system: a computer simulation feasibility study”. In: *Physics in Medicine & Biology* 50.17, p. 4225.
- Arridge, S. R., M Schweiger, M Hiraoka, and D. T. Delpy (1993). “A finite element approach for modeling photon transport in tissue”. In: *Medical Physics* 20, pp. 299–309.
- Arridge, S. R. (1999). “Optical tomography in medical imaging”. In: *Inverse Problems* 15.2, R41.
- Arridge, S. R. and J. C. Hebden (1997). “Optical imaging in medicine: II. Modelling and reconstruction”. In: *Physics in Medicine & Biology* 42, pp. 841–853.
- Arridge, S. R. and W. R. B. Lionheart (1998). “Nonuniqueness in diffusion-based optical tomography”. In: *Optics Letters* 23.11, pp. 882–884.
- Arridge, S. R. and M. Schweiger (1997). “Image reconstruction in optical tomography.” In: *Philosophical Transactions of the Royal Society B* 352.1354, p. 717.
- Baraniuk, R. G. (2007). “Compressive sensing”. In: *IEEE Signal Processing Magazine* 24.4, pp. 118–120,124.
- (2008). “Single-pixel imaging via compressive sampling”. In: *IEEE Signal Processing Magazine* 25.2, pp. 83–91.

- Basevi, H. R. A., J. A. Guggenheim, H. Dehghani, and I. B. Styles (2012a). “Application of Compressive Sensing to Bioluminescence Tomography”. In: *Biomedical Optics and 3-D Imaging*. Optical Society of America, BSu3A.64.
- Basevi, H. R. A., K. M. Tichauer, F. Leblond, H. Dehghani, J. A. Guggenheim, R. W. Holt, and I. B. Styles (2012b). “Compressive sensing based reconstruction in bioluminescence tomography improves image resolution and robustness to noise”. In: *Biomedical Optics Express* 3.9, pp. 2131–2141.
- Basevi, H. R. A., J. A. Guggenheim, H. Dehghani, and I. B. Styles (2013a). “Information-theoretic method for wavelength selection in bioluminescence tomography”. In: *European Conference on Biomedical Optics*. Optical Society of America, p. 879909.
- Basevi, H. R. A., J. A. Guggenheim, H. Dehghani, and I. B. Styles (2013b). “Simultaneous multiple view high resolution surface geometry acquisition using structured light and mirrors”. In: *Optics Express* 21.6, pp. 7222–7239.
- Basevi, H. R. A., J. A. Guggenheim, S. Taylor, H. Dehghani, and I. B. Styles (2014). “Towards Bayesian Reconstruction and Analysis in Bioluminescence Tomography via Markov Chain Monte Carlo Techniques”. In: *Biomedical Optics*. Optical Society of America, BM3A–45.
- Baum, K., R. Hartmann, T. Bischoff, J. O. Oelerich, S. Finkensieper, and J. T. Heverhagen (2012). “Three-dimensional surface reconstruction within noncontact diffuse optical tomography using structured light”. In: *Journal of Biomedical Optics* 17.12, pp. 126009–126009.
- Berryman, F., P. Pynsent, and J. Cubillo (2003). “A theoretical comparison of three fringe analysis methods for determining the three-dimensional shape of an object in the presence of noise”. In: *Optics and Lasers in Engineering* 39.1, pp. 35–50.
- Bingham, E. and H. Mannila (2001). “Random projection in dimensionality reduction: applications to image and text data”. In: *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, pp. 245–250.

- Biospace Lab (2014a). *3D Module*. <http://www.biospacelab.com/m-90-3d-module.html>.
Last visited on 21/09/2014.
- (2014b). *4-View Module*. <http://www.biospacelab.com/m-89-4-view-module.html>.
Last visited on 21/09/2014.
- (2014c). *PhotonIMAGERTM*. <http://www.biospacelab.com/m-31-optical-imaging.html>. Last visited on 21/09/2014.
- Bioucas-Dias, J. M. and G. Valadão (2007). “Phase unwrapping via graph cuts”. In: *IEEE Transactions on Image Processing* 16.3, pp. 698–709.
- Boas, D., J Culver, J Stott, and A Dunn (2002). “Three dimensional Monte Carlo code for photon migration through complex heterogeneous media including the adult human head”. In: *Optics Express* 10.3, pp. 159–170.
- Boas, D. A., A. M. Dale, and M. A. Franceschini (2004). “Diffuse optical imaging of brain activation: approaches to optimizing image sensitivity, resolution, and accuracy”. In: *Neuroimage* 23, S275–S288.
- Brendel, B. and T. Nielsen (2009). “Selection of optimal wavelengths for spectral reconstruction in diffuse optical tomography”. In: *Journal of Biomedical Optics* 14.3, pp. 034041–034041.
- Bruning, J. H., D. R. Herriott, J. E. Gallagher, D. P. Rosenfeld, A. D. White, and D. J. Brangaccio (1974). “Digital wavefront measuring interferometer for testing optical surfaces and lenses”. In: *Applied optics* 13.11, pp. 2693–2703.
- Burgos, J. S., M Rosol, R. A. Moats, V Khankaldyyan, D. B. Kohn, M. D. Nelson Jr, and W. E. Laug (2003). “Time course of bioluminescent signal in orthotopic and heterotopic brain tumors in nude mice”. In: *Biotechniques* 34.6, pp. 1184–1188.
- Candes, E. and J. Romberg (2007). “Sparsity and incoherence in compressive sampling”. In: *Inverse Problems* 23.3, pp. 969–985.
- Candes, E. J. (2008). “The restricted isometry property and its implications for compressed sensing”. In: *Comptes Rendus Mathematique* 346.9, pp. 589–592.

- Cao, L., M. Breithaupt, and J. Peter (2010). “Geometrical co-calibration of a tomographic optical system with CT for intrinsically co-registered imaging”. In: *Physics in Medicine & Biology* 55.6, p. 1591.
- Chamorro-Servent, J., J. Aguirre, J. Ripoll, J. J. Vaquero, and M. Desco (2011). “Feasibility of U-curve method to select the regularization parameter for fluorescence diffuse optical tomography in phantom and small animal studies”. In: *Optics Express* 19.12, pp. 11490–11506.
- Chaudhari, A. J., F. Darvas, J. R. Bading, R. A. Moats, P. S. Conti, D. J. Smith, S. R. Cherry, and R. M. Leahy (2005). “Hyperspectral and multispectral bioluminescence optical tomography for small animal imaging”. In: *Physics in Medicine & Biology* 50, p. 5421.
- Chen, J., V. Venugopal, F. Lesage, and X. Intes (2010). “Time-resolved diffuse optical tomography with patterned-light illumination and detection”. In: *Optics Letters* 35.13, pp. 2121–2123.
- Chen, X., D. Yang, Q. Zhang, and J. Liang (2014). “L1/2 regularization based numerical method for effective reconstruction of bioluminescence tomography”. In: *Journal of Applied Physics* 115.18, p. 184702.
- Cheong, W.-F., S. A. Prahl, and A. J. Welch (1990). “A review of the optical properties of biological tissues”. In: *IEEE Journal of Quantum Electronics* 26.12, pp. 2166–2185.
- Cignoni, P. (2014). *MeshLab Home Page*. <http://meshlab.sourceforge.net/>. Last visited on 21/09/2014.
- Cong, A., W. Cong, Y. Lu, P. Santago, A. Chatziioannou, and G. Wang (2010). “Differential evolution approach for regularized bioluminescence tomography”. In: *IEEE Transactions on Biomedical Engineering* 57.9, pp. 2229–2238.
- Contag, C. H. and M. H. Bachmann (2002). “Advances in in vivo bioluminescence imaging of gene expression”. In: *Annual Review of Biomedical Engineering* 4.1, pp. 235–260.

- Corlu, A., T. Durduran, R. Choe, M. Schweiger, E. Hillman, S. R. Arridge, and A. G. Yodh (2003). “Uniqueness and wavelength optimization in continuous-wave multispectral diffuse optical tomography”. In: *Optics Letters* 28.23, pp. 2339–2341.
- Corlu, A., R. Choe, T. Durduran, K. Lee, M. Schweiger, S. R. Arridge, E. Hillman, and A. G. Yodh (2005). “Diffuse optical tomography with spectral constraints and wavelength optimization”. In: *Applied Optics* 44.11, pp. 2082–2093.
- Cox, B. T., J. G. Laufer, and P. C. Beard (2009). “The challenges for quantitative photoacoustic imaging”. In: *SPIE BiOS: Biomedical Optics*. International Society for Optics and Photonics, pp. 717713–717713.
- Culver, J. P., R. Choe, M. J. Holboke, L. Zubkov, T. Durduran, A. Slemple, V. Ntziachristos, B. Chance, and A. G. Yodh (2003). “Three-dimensional diffuse optical tomography in the parallel plane transmission geometry: evaluation of a hybrid frequency domain/continuous wave clinical system for breast imaging”. In: *Medical Physics* 30.2, pp. 235–247.
- Custo, A., W. M. Wells III, A. H. Barnett, E. Hillman, and D. A. Boas (2006). “Effective scattering coefficient of the cerebral spinal fluid in adult head models for diffuse optical imaging”. In: *Applied Optics* 45.19, pp. 4747–4755.
- Dehghani, H., S. C. Davis, and B. W. Pogue (2008). “Spectrally resolved bioluminescence tomography using the reciprocity approach”. In: *Medical Physics* 35.11, pp. 4863–4871.
- Dehghani, H., M. E. Eames, P. K. Yalavarthy, S. C. Davis, S. Srinivasan, C. M. Carpenter, B. W. Pogue, and K. D. Paulsen (2008). “Near infrared optical tomography using NIRFAST: Algorithm for numerical model and image reconstruction”. In: *Communications in Numerical Methods in Engineering* 25.6, pp. 711–732.
- Deliolanis, N., T. Lasser, D. Hyde, A. Soubret, J. Ripoll, and V. Ntziachristos (2007). “Free-space fluorescence molecular tomography utilizing 360° geometry projections”. In: *Optics Letters* 32.4, p. 382.

- Dogdas, B., D. Stout, A. F. Chatziioannou, and R. M. Leahy (2007). “Digimouse: a 3D whole body mouse atlas from CT and cryosection data”. In: *Physics in Medicine & Biology* 52.3, p. 577.
- Donoho, D. and J. Tanner (2009). “Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367.1906, pp. 4273–4293.
- Dorrington, A. A., M. J. Cree, A. D. Payne, R. M. Conroy, and D. A. Carnegie (2007). “Achieving sub-millimetre precision with a solid-state full-field heterodyning range imaging camera”. In: *Measurement Science and Technology* 18.9, p. 2809.
- Dothager, R. S., K. Flentje, B. Moss, M.-H. Pan, A. Kesarwala, and D. Piwnica-Worms (2009). “Advances in bioluminescence imaging of live animal models”. In: *Current Opinion in Biotechnology* 20.1, pp. 45–53.
- Eames, M. E., J. Wang, H. Dehghani, and B. W. Pogue (2008). “Wavelength band optimization in spectral near-infrared optical tomography improves accuracy while reducing data acquisition and computational burden”. In: *Journal of Biomedical Optics* 13.5, pp. 054037–054037.
- Edinger, M., Y. A. Cao, M. R. Verneris, M. H. Bachmann, C. H. Contag, and R. S. Negrin (2003). “Revealing lymphoma growth and the efficacy of immune cell therapies using in vivo bioluminescence imaging”. In: *Blood* 101.2, pp. 640–648.
- Epstein, E., M. Granger-Piché, and P. Potilin (2004). “Exploiting Mirrors in Interactive Reconstruction with Structured Light”. In: *Vision, Modeling, and Visualization 2004*. IOS Press, p. 125.
- Feng, J., K. Jia, C. Qin, G. Yan, S. Zhu, X. Zhang, J. Liu, and J. Tian (2009). “Three-dimensional Bioluminescence Tomography based on Bayesian approach”. In: *Optics Express* 17.19, pp. 16834–16848.

- Feng, J., C. Qin, K. Jia, S. Zhu, K. Liu, D. Han, X. Yang, Q. Gao, and J. Tian (2012). “Total variation regularization for bioluminescence tomography with the split Bregman method”. In: *Applied Optics* 51.19, pp. 4501–4512.
- Foucart, S. and D. Koslicki (2014). “Sparse Recovery by Means of Nonnegative Least Squares.” In: *IEEE Signal Processing Letters* 21.4, pp. 498–502.
- Freiberger, M., C. Clason, and H. Scharfetter (2010). “Adaptation and focusing of optode configurations for fluorescence optical tomography by experimental design methods”. In: *Journal of biomedical optics* 15.1, pp. 016024–016024.
- Freiberger, M. and H. Scharfetter (2009). “Sensor optimization for fluorescence optical tomography by experimental design methods”. In: *SPIE Proceedings*. Vol. 7174, p. 717419.
- Freiberger, M., A. Laurain, M. Hintermüller, A. Köstinger, and H. Scharfetter (2011). “Using the topological derivative for initializing a Markov-chain Monte Carlo reconstruction in fluorescence tomography”. In: *European Conferences on Biomedical Optics*. International Society for Optics and Photonics, 80881Q–80881Q.
- Gao, H. and H. Zhao (2010). “Multilevel bioluminescence tomography based on radiative transfer equation Part 2: total variation and l1 data fidelity”. In: *Optics Express* 18.3, pp. 2894–2912.
- Geng, J. (2011). “Structured-light 3D surface imaging: a tutorial”. In: *Advances in Optics and Photonics* 3.2, pp. 128–160.
- Geng, Z. J. (2007). “Diffuse optical tomography system and method of use”. Pat. 7242997.
- Gorthi, S. S. and P. Rastogi (2010). “Fringe projection techniques: whither we are?” In: *Optics and Lasers in Engineering* 48.2, pp. 133–140.
- Graves, E. E., J. P. Culver, J. Ripoll, R. Weissleder, and V. Ntziachristos (2004). “Singular-value analysis and optimization of experimental parameters in fluorescence molecular tomography”. In: *Journal of the Optical Society of America A: Optics, Image Science, and Vision* 21.2, pp. 231–241.

- Guggenheim, J. A. (2014). “Multi-modal diffuse optical tomography and bioluminescence tomography system for preclinical imaging”. PhD thesis. University of Birmingham.
- Guggenheim, J. A., H. R. A. Basevi, J. Frampton, I. B. Styles, and H. Dehghani (2013a). “Multi-modal molecular diffuse optical tomography system for small animal imaging”. In: *Measurement Science and Technology* 24.10, p. 105405.
- Guggenheim, J. A., H. R. A. Basevi, I. B. Styles, J. Frampton, and H. Dehghani (2013b). “Quantitative surface radiance mapping using multiview images of light-emitting turbid media”. In: *Journal of the Optical Society of America A: Optics, Image Science, and Vision* 30.12, pp. 2572–2584.
- Guo, X., X. Liu, A. Wang, F. Tian, F. Liu, B. Zhang, G. Hu, and J. Bai (2010). “A combined fluorescence and microcomputed tomography system for small animal imaging”. In: *IEEE Transactions on Biomedical Engineering* 57.12, pp. 2876–2883.
- Halko, N., P.-G. Martinsson, and J. A. Tropp (2011). “Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions”. In: *SIAM Review* 53.2, pp. 217–288.
- Han, W. and G. Wang (2007). “Theoretical and numerical analysis on multispectral bioluminescence tomography”. In: *IMA Journal of Applied Mathematics* 72.1, pp. 67–85.
- Hansen, P. C. and D. P. O’Leary (1993). “The use of the L-curve in the regularization of discrete ill-posed problems”. In: *SIAM Journal on Scientific Computing* 14.6, pp. 1487–1503.
- Hastings, W. K. (1970). “Monte Carlo sampling methods using Markov chains and their applications”. In: *Biometrika* 57.1, pp. 97–109.
- Häusler, G. and W. Heckel (1988). “Light sectioning with large depth and high resolution”. In: *Applied Optics* 27.24, pp. 5165–5169.
- He, X., J. Liang, X. Qu, H. Huang, Y. Hou, and J. Tian (2010). “Truncated total least squares method with a practical truncation parameter choice scheme for bioluminescence tomography inverse problem”. In: *Journal of Biomedical Imaging* 2010, p. 12.

- He, X., Y. Hou, D. Chen, Y. Jiang, M. Shen, J. Liu, Q. Zhang, and J. Tian (2011). “Sparse regularization-based reconstruction for bioluminescence tomography using a multilevel adaptive finite element method.” In: *International Journal of Biomedical Imaging* 2011, p. 203537.
- Hebden, J. C., S. R. Arridge, and D. T. Delpy (1997). “Optical imaging in medicine: I. Experimental techniques”. In: *Physics in Medicine & Biology* 42.5, p. 825.
- Jiang, X., L. Cao, and J. Peter (2014). “Surface reconstruction from multiview projection data employing a microlens array-based optical detector: a simulation study”. In: *Optical Engineering* 53.2, pp. 023104–023104.
- Jiang, X., L. Cao, W. Semmler, and J. Peter (2012). “A surface recognition approach for in vivo optical imaging applications using a micro-lens-array light detector”. In: *Biomedical Optics*. Optical Society of America.
- Judge, T. R. and P. J. Bryanston-Cross (1994). “A review of phase unwrapping techniques in fringe analysis”. In: *Optics and Lasers in Engineering* 21.4, pp. 199–239.
- Kazhdan, M., M. Bolitho, and H. Hoppe (2006). “Poisson surface reconstruction”. In: *Proceedings of the Fourth Eurographics Symposium on Geometry Processing*.
- Kilmer, M. E., P. C. Hansen, and M. I. Espanol (2007). “A projection-based approach to general-form Tikhonov regularization”. In: *SIAM Journal on Scientific Computing* 29.1, pp. 315–330.
- Kinell, L. (2004). “Spatiotemporal approach for real-time absolute shape measurements by use of projected fringes”. In: *Applied Optics* 43.15, pp. 3018–3027.
- Klose, A. D. and E. W. Larsen (2006). “Light transport in biological tissue based on the simplified spherical harmonics equations”. In: *Journal of Computational Physics* 220.1, pp. 441–470.
- Koenig, A., L. Hervé, A. Da Silva, J. Dinten, J. Boutet, M. Berger, I. Texier, P. Peltié, P. Rizo, V. Jossierand, and J. Coll (2007). “Whole body small animal examination with a diffuse optical tomography instrument”. In: *Nuclear Instruments and Methods*

- in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 571.1, pp. 56–59.
- Kraus, K. (2007). *Photogrammetry: Geometry from images and laser scans*. de Gruyter.
- Kumar, A. T. N., S. B. Raymond, A. K. Dunn, B. J. Bacskai, and D. A. Boas (2008). “A time domain fluorescence tomography system for small animal imaging”. In: *IEEE Transactions on Medical Imaging* 27.8, pp. 1152–1163.
- Kuo, C., O. Coquoz, T. L. Troy, H. Xu, and B. W. Rice (2007). “Three-dimensional reconstruction of in vivo bioluminescent sources based on multispectral imaging”. In: *Journal of Biomedical Optics* 12.2, p. 024007.
- Kuthirummal, S. and S. K. Nayar (2006). “Multiview radial catadioptric imaging for scene capture”. In: *ACM Transactions on Graphics*. Vol. 25. 3. ACM, pp. 916–923.
- Lanman, D., D. Crispell, and G. Taubin (2009). “Surround structured lighting: 3-D scanning with orthographic illumination”. In: *Computer Vision and Image Understanding* 113.11, pp. 1107–1117.
- Lasser, T. and V. Ntziachristos (2007). “Optimization of 360° projection fluorescence molecular tomography”. In: *Medical Image Analysis* 11.4, pp. 389–399.
- Lasser, T., A. Soubret, J. Ripoll, and V. Ntziachristos (2008). “Surface Reconstruction for Free-Space 360° Fluorescence Molecular Tomography and the Effects of Animal Motion”. In: *IEEE Transactions on Medical Imaging* 27.2, pp. 188–194.
- Laufer, J., E. Zhang, G. Raivich, and P. Beard (2009). “Three-dimensional noninvasive imaging of the vasculature in the mouse brain using a high resolution photoacoustic scanner”. In: *Applied Optics* 48.10, pp. D299–D306.
- Lawson, C. L. and R. J. Hanson (1995). *Solving least squares problems*. Vol. 15. Society for Industrial Mathematics.
- Leblond, F., K. M. Tichauer, and B. W. Pogue (2010). “Singular value decomposition metrics show limitations of detector design in diffuse fluorescence tomography”. In: *Biomedical Optics Express* 1.5, pp. 1514–1531.

- Leblond, F., S. C. Davis, P. A. Valdés, and B. W. Pogue (2010). “Pre-clinical whole-body fluorescence imaging: Review of instruments, methods and applications”. In: *Journal of Photochemistry and Photobiology B: Biology* 98.1, pp. 77–94.
- Leblond, F., K. M. Tichauer, R. W. Holt, F. El-Ghoussein, and B. W. Pogue (2011). “Toward whole-body optical imaging of rats using single-photon counting fluorescence tomography”. In: *Optics Letters* 36.19, pp. 3723–3725.
- Lei, S. and S. Zhang (2009). “Flexible 3-D shape measurement using projector defocusing”. In: *Optics Letters* 34.20, pp. 3080–3082.
- Li, C., G. S. Mitchell, J. Dutta, S. Ahn, R. M. Leahy, and S. R. Cherry (2009). “A three-dimensional multispectral fluorescence optical tomography imaging system for small animals based on a conical mirror design”. In: *Optics Express* 17.9, p. 7571.
- Li, E., X Peng, J. Xi, J Chicharo, J. Yao, and D Zhang (2005). “Multi-frequency and multiple phase-shift sinusoidal fringe projection for 3D profilometry”. In: *Optics Express* 13.5, pp. 1561–1569.
- Lissberger, P. H. and W. L. Wilcock (1959). “Properties of all-dielectric interference filters. II. Filters in parallel beams of light incident obliquely and in convergent beams”. In: *Journal of the Optical Society of America A: Optics, Image Science, and Vision* 49.2, pp. 126–128.
- Liu, J., D. Chen, X. Li, X. Ma, H. Chen, W. Fan, F. Wang, X. Qu, J. Liang, F. Cao, and J. Tian (2010). “In vivo quantitative reconstruction studies of bioluminescence tomography: effects of peak-wavelength shift and model deviation”. In: *IEEE Transactions on Biomedical Engineering* 57.10, pp. 2579–2582.
- Lu, Y., X. Zhang, A. Douraghy, D. Stout, J. Tian, T. F. Chan, and A. F. Chatziioannou (2009). “Source Reconstruction for Spectrally-resolved Bioluminescence Tomography with Sparse A priori Information”. In: *Optics Express* 17.10, pp. 8062–8080.
- Lu, Y., H. B. Machado, Q. Bao, D. Stout, H. Herschman, and A. F. Chatziioannou (2011). “In vivo mouse bioluminescence tomography with radionuclide-based imaging validation”. In: *Molecular Imaging and Biology* 13.1, pp. 53–58.

- Luker, K. E. and G. D. Luker (2008). “Applications of bioluminescence imaging to antiviral research and therapy: multiple luciferase enzymes and quantitation”. In: *Antiviral Research* 78.3, pp. 179–187.
- Lustig, M., D. Donoho, and J. M. Pauly (2007). “Sparse MRI: The application of compressed sensing for rapid MR imaging”. In: *Magnetic Resonance in Medicine* 58.6, pp. 1182–1195.
- Lustig, M., D. L. Donoho, J. M. Santos, and J. M. Pauly (2008). “Compressed sensing MRI”. In: *IEEE Signal Processing Magazine* 25.2, pp. 72–82.
- Mao, X., W. Chen, and X. Su (2007). “Improved Fourier-transform profilometry”. In: *Applied Optics* 46.5, pp. 664–668.
- Mazhar, A., S. Dell, D. J. Cuccia, S. Gioux, A. J. Durkin, J. V. Frangioni, and B. J. Tromberg (2010). “Wavelength optimization for rapid chromophore mapping using spatial frequency domain imaging”. In: *Journal of Biomedical Optics* 15.6, pp. 061716–061716.
- Meyer, H., A. Garofalakis, G. Zacharakis, S. Psycharakis, C. Mamalaki, D. Kioussis, E. N. Economou, V. Ntziachristos, and J. Ripoll (2007). “Noncontact optical imaging in mice with full angular coverage and automatic surface extraction”. In: *Applied Optics* 46.17, pp. 3617–3627.
- Milstein, A. B., J. J. Stott, S. Oh, D. A. Boas, R. P. Millane, C. A. Bouman, and K. J. Webb (2004). “Fluorescence optical diffusion tomography using multiple-frequency data”. In: *Journal of the Optical Society of America A: Optics, Image Science, and Vision* 21.6, pp. 1035–1049.
- Mourant, J. R., T. Fuselier, J. Boyer, T. M. Johnson, and I. J. Bigio (1997). “Predictions and measurements of scattering and absorption over broad wavelength ranges in tissue phantoms”. In: *Applied Optics* 36.4, pp. 949–957.
- Naser, M. A. and M. S. Patterson (2010). “Algorithms for bioluminescence tomography incorporating anatomical information and reconstruction of tissue optical properties”. In: *Biomedical Optics Express* 1.2, pp. 512–526.

- Naser, M. A. and M. S. Patterson (2011a). “Bioluminescence tomography using eigenvectors expansion and iterative solution for the optimized permissible source region”. In: *Biomedical Optics Express* 2.11, pp. 3179–3193.
- (2011b). “Improved bioluminescence and fluorescence reconstruction algorithms using diffuse optical tomography, normalized data, and optimized selection of the permissible source region”. In: *Biomedical Optics Express* 2.1, pp. 169–184.
- Ng, R., M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan (2005). “Light field photography with a hand-held plenoptic camera”. In: *Computer Science Technical Report* 2.11.
- Nilson, D., M. D. Cable, B. W. Rice, and K Kearney (2007). “Structured light imaging apparatus”. Pat. 7298415.
- Nocedal, J and S. J. Wright (1999). *Numerical optimization*. Springer New York.
- Ntziachristos, V. and J. Ripoll (2011). “Method and system for free space optical tomography of diffuse media”. Pat. 7962200.
- Ntziachristos, V. (2006). “Fluorescence molecular imaging”. In: *Annual Review of Biomedical Engineering* 8, pp. 1–33.
- Ntziachristos, V., J. Ripoll, L. V. Wang, and R. Weissleder (2005). “Looking and listening to light: the evolution of whole-body photonic imaging”. In: *Nature Biotechnology* 23.3, pp. 313–320.
- Park, T. and G. Casella (2008). “The bayesian lasso”. In: *Journal of the American Statistical Association* 103.482, pp. 681–686.
- Payne, A. D., A. P. P. Jongenelen, A. A. Dorrington, M. J. Cree, and D. A. Carnegie (2009). “Multiple frequency range imaging to remove measurement ambiguity”. In: *Multiple frequency range imaging to remove measurement ambiguity*, pp. 139–148.
- PerkinElmer (2014). *IVIS 200 Pre-clinical In Vivo Imaging System*. <http://www.perkinelmer.com/Catalog/Product/ID/IVIS200>. Last visited on 21/09/2014.

- Pollack, S. A. (1966). “Angular dependence of transmission characteristics of interference filters and application to a tunable fluorometer”. In: *Applied Optics* 5.11, pp. 1749–1756.
- Prescher, J. A. and C. H. Contag (2010). “Guided by the light: visualizing biomolecular processes in living animals with bioluminescence”. In: *Current Opinion in Chemical Biology* 14.1, pp. 80–89.
- Rauhut, H. (2010). “Compressive sensing and structured random matrices”. In: *Theoretical Foundations and Numerical Methods for Sparse Recovery*. Vol. 9. Radon Series Comp. Appl. Math. deGruyter, pp. 1–92.
- Rehemtulla, A., L. D. Stegman, S. J. Cardozo, S. Gupta, D. E. Hall, C. H. Contag, and B. D. Ross (2000). “Rapid and Quantitative Assessment of Cancer Treatment Response Using In Vivo Bioluminescence Imaging”. In: *Neoplasia* 2.6, pp. 491–495.
- Reinbacher-Köstinger, A., M. Freiberger, and H. Scharfetter (2011). “Uncertainty analysis for fluorescence tomography with Monte Carlo method”. In: *European Conferences on Biomedical Optics*. International Society for Optics and Photonics, 80881P–80881P.
- Rice, B. W., M. D. Cable, and K. Kearney (2010). “3D in-vivo imaging and topography using structured light”. Pat. 7797034.
- Robert, C. P. (2001). *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation Second Edition*. Springer-Verlag New York.
- Roberts, G. O., A. Gelman, and W. R. Gilks (1997). “Weak convergence and optimal scaling of random walk Metropolis algorithms”. In: *The Annals of Applied Probability* 7.1, pp. 110–120.
- Saldner, H. O. and J. M. Huntley (1997). “Temporal phase unwrapping: application to surface profiling of discontinuous objects”. In: *Applied Optics* 36.13, pp. 2770–2775.
- Salvi, J., J. Pages, and J. Batlle (2004). “Pattern codification strategies in structured light systems”. In: *Pattern Recognition* 37.4, pp. 827–849.

- Sansoni, G., M. Carocci, and R. Rodella (1999). “Three-dimensional vision based on a combination of gray-code and phase-shift light projection: analysis and compensation of the systematic errors”. In: *Applied Optics* 38.31, pp. 6565–6573.
- Scharstein, D. and R. Szeliski (2003). “High-accuracy stereo depth maps using structured light”. In: *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.* Vol. 1. IEEE, pp. I–195.
- Schoberl, J. (2014). *Netgen Mesh Generator*. <http://sourceforge.net/projects/netgen-mesher/>. Last visited on 21/09/2014.
- Schweiger, M. and S. Arridge (2014). “The Toast++ software suite for forward and inverse modeling in optical tomography”. In: *Journal of Biomedical Optics* 19.4, pp. 040801–040801.
- Seliger, H. H. and W. D. McElroy (1960). “Spectral emission and quantum yield of firefly bioluminescence”. In: *Archives of Biochemistry and Biophysics* 88.1, pp. 136–141.
- Shewchuk, J. R. (1994). *An introduction to the conjugate gradient method without the agonizing pain*. <http://www.cs.cmu.edu/~jrs/jrspapers.html>.
- Srinivasan, S., B. W. Pogue, S. Jiang, H. Dehghani, and K. D. Paulsen (2003). “Validation of hemoglobin and water molar absorption spectra in near-infrared diffuse optical tomography”. In: *Biomedical Optics 2003*. International Society for Optics and Photonics, pp. 407–415.
- Srinivasan, V, H. C. Liu, and M. Halioua (1984). “Automated phase-measuring profilometry of 3-D diffuse objects”. In: *Applied Optics* 23.18, pp. 3105–3108.
- Stearns, D. G., B. W. Rice, and M. D. Cable (2010). “Method and apparatus for 3-D imaging of internal light sources”. Pat. 7860549.
- Storn, R. and K. Price (1997). “Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces”. In: *Journal of Global Optimization* 11.4, pp. 341–359.

- Stout, D, P Chow, R Silverman, R. Leahy, X Lewis, S Gambhir, and A Chatziioannou (2002). “Creating a whole body digital mouse atlas with PET, CT and cryosection images”. In: *Molecular Imaging and Biology* 4.4, S27.
- Styles, I. B., H. R. A. Basevi, J. A. Guggenheim, and H. Dehghani (2013). “Random matrix-based dimensionality reduction for bioluminescence tomography reconstruction”. In: *European Conference on Biomedical Optics*. Optical Society of America, 87990J.
- Su, X. and W. Chen (2001). “Fourier transform profilometry:: a review”. In: *Optics and Lasers in Engineering* 35.5, pp. 263–284.
- (2004). “Reliability-guided phase unwrapping algorithm: a review”. In: *Optics and Lasers in Engineering* 42.3, pp. 245–261.
- Takeda, M. and K. Mutoh (1983). “Fourier transform profilometry for the automatic measurement of 3-D object shape”. In: *Applied Optics* 22.24, pp. 3977–3982.
- Taylor, C. J. (2012). “Implementing high resolution structured light by exploiting projector blur”. In: *2012 IEEE Workshop on Applications of Computer Vision (WACV)*. IEEE, pp. 9–16.
- Thorne, N., J. Inglese, and D. S. Auld (2010). “Illuminating insights into firefly luciferase and other bioluminescent reporters used in chemical biology”. In: *Chemistry & Biology* 17.6, pp. 646–657.
- Thorne, S. H. and C. H. Contag (2005). “Using in vivo bioluminescence imaging to shed light on cancer biology”. In: *Proceedings of the IEEE* 93.4, pp. 750–762.
- Tibshirani, R. (1996). “Regression shrinkage and selection via the lasso”. In: *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288.
- Tichauer, K. M., R. W. Holt, F. El-Ghoussein, Q. Zhu, H. Dehghani, F. Leblond, and B. W. Pogue (2011). “Imaging workflow and calibration for CT-guided time-domain fluorescence tomography”. In: *Biomedical Optics Express* 2.11, pp. 3021–3036.

- Tremoleda, J. L., M. Khalil, L. L. Gompels, M. Wylezinska-Arridge, T. Vincent, and W. Gsell (2011). “Imaging technologies for preclinical models of bone and joint disorders”. In: *EJNMMI Research* 1.1, pp. 1–14.
- Venugopal, V., J. Chen, and X. Intes (2010). “Development of an optical imaging platform for functional imaging of small animals using wide-field excitation”. In: *Biomedical Optics Express* 1.1, p. 143.
- Wang, G., Y. Li, and M. Jiang (2004). “Uniqueness theorems in bioluminescence tomography”. In: *Medical Physics* 31.8, pp. 2289–2299.
- Wang, G., H. Shen, K. Durairaj, X. Qian, and W. Cong (2006). “The first bioluminescence tomography system for simultaneous acquisition of multiview and multispectral data”. In: *International Journal of Biomedical Imaging* 2006.
- Wang, L., S. L. Jacques, and L. Zheng (1995). “MCML—Monte Carlo modeling of light transport in multi-layered tissues”. In: *Computer Methods and Programs in Biomedicine* 47.2, pp. 131–146.
- Wang, L. V. (2008). “Tutorial on photoacoustic microscopy and computed tomography”. In: *IEEE Journal of Selected Topics in Quantum Electronics* 14.1, pp. 171–179.
- Wang, L. V. and S. Hu (2012). “Photoacoustic tomography: in vivo imaging from organelles to organs”. In: *Science* 335.6075, pp. 1458–1462.
- Wang, Z., H. Du, and H. Bi (2006). “Out-of-plane shape determination in generalized fringe projection profilometry”. In: *Optics Express* 14.25, pp. 12122–12133.
- Wang, Z., H. Du, S. Park, and H. Xie (2009). “Three-dimensional shape measurement with a fast and accurate approach”. In: *Applied Optics* 48.6, pp. 1052–1061.
- Weber, D. A. and M. Ivanovic (1999). “Ultra-high-resolution imaging of small animals: implications for preclinical and research studies”. In: *Journal of Nuclear Cardiology* 6.3, pp. 332–344.
- Wessels, J. T., A. C. Busse, J. Mahrt, C. Dullin, E. Grabbe, and G. A. Mueller (2007). “In vivo imaging in experimental preclinical tumor research—a review”. In: *Cytometry Part A* 71.8, pp. 542–549.

- White, E. H., E. Rapaport, H. H. Seliger, and T. A. Hopkins (1971). “The chemi- and bioluminescence of firefly luciferin: an efficient chemical production of electronically excited states”. In: *Bioorganic Chemistry* 1.1, pp. 92–122.
- Ye, J. C., K. J. Webb, C. A. Bouman, and R. P. Millane (1999). “Optical diffusion tomography by iterative-coordinate-descent optimization in a Bayesian framework”. In: *Journal of the Optical Society of America A: Optics, Image Science, and Vision* 16.10, pp. 2400–2412.
- Yu, J., F. Liu, J. Wu, L. Jiao, and X. He (2010). “Fast Source Reconstruction for Bioluminescence Tomography Based on Sparse Regularization”. In: *IEEE Transactions on Biomedical Engineering* 57.10, pp. 2583–2586.
- Zappa, E. and G Busca (2008). “Comparison of eight unwrapping algorithms applied to Fourier-transform profilometry”. In: *Optics and Lasers in Engineering* 46.2, pp. 106–116.
- Zhang, Q., H. Zhao, D. Chen, X. Qu, X. Chen, X. He, W. Li, Z. Hu, J. Liu, J. Liang, and J. Tian (2011). “Source sparsity based primal-dual interior-point method for three-dimensional bioluminescence tomography”. In: *Optics Communications* 284.24, pp. 5871–5876.
- Zhang, S, Q. X. Chen, J. W. Luo, and J Bai (2013). “A fast surface reconstruction method for fluorescence molecular tomography based on cross-beam edge back projection”. In: *Measurement* 46.4, pp. 1565–1571.
- Zhang, X. and C. Badea (2009). “Effects of sampling strategy on image quality in non-contact panoramic fluorescence diffuse optical tomography for small animal imaging”. In: *Optics Express* 17.7, pp. 5125–5138.
- Zhang, X., Y. Lu, and T. Chan (2012). “A novel sparsity reconstruction method from poisson data for 3D bioluminescence tomography”. In: *Journal of Scientific Computing* 50.3, pp. 519–535.