Computational Multispectral Endoscopy

Geoffrey Alan Jones

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Centre for Medical Image Computing

Computer Science

University College London

February 19, 2020

I, Geoffrey Alan Jones, confirm that the work presented in this thesis is my own.

Where information has been derived from other sources,

I confirm that this has been indicated in the work.

My thesis' central issue,
is assessing the blood in tissue.
If you have the right light,
using mathematical slight,
you can acquire these data in situ.

Abstract

Minimal access surgery (MAS) is increasingly regarded as the de-facto approach in interventional medicine for conducting many procedures; this is due to the reduced patient trauma and consequently reduced recovery times, complications and costs. However, there are many challenges in MAS that come as a result of viewing the surgical site through an endoscope and interacting with tissue remotely via tools, such as lack of haptic feedback; limited field of view; and variation in imaging hardware. As such, it is important best utilise the imaging data available to provide a clinician with rich data corresponding to the surgical site. Measuring tissue haemoglobin concentrations can give vital information, such as perfusion assessment after transplantation; visualisation of the health of blood supply to organ; and to detect ischaemia. In the area of transplant and bypass procedures measurements of the tissue perfusion/total haemoglobin (THb) and oxygen saturation (SO_2) are used as indicators of organ viability, these measurements are often acquired at multiple discrete points across the tissue using with a specialist probe. To acquire measurements across the whole surface of an organ one can use a specialist camera to perform multispectral imaging (MSI), which optically acquires sequential spectrally band limited images of the same scene. This data can be processed to provide maps of the THb and SO_2 variation across the tissue surface which could be useful for intra operative evaluation. When capturing multispectral (MS) image data, a trade off often has to be made between spectral sensitivity and capture speed.

The work in thesis first explores post processing blurry MS image data from long exposure imaging devices. It is of interest to be able to use these MSI data because the large number of spectral bands that can be captured, the long capture

times, however, limit the potential real time uses for clinicians. Recognising the importance to clinicians of real-time data, the main body of this thesis develops methods around estimating oxygenated haemoglobin (HbO) and deoxygenated haemoglobin (Hb) concentrations in tissue using only monocular and stereo RGB imaging data.

Impact Statement

The impact of the work presented herein is already present in contributions made to scientific 3 conference papers, a journal paper, as well as a patent (refereed in subsection 1.3.1), with the potential for further works as outlined in the future work section in Chapter 7. This work builds on a greater body of work produced within the Surgical Robot Vision group, at University College London, for the purposes of using computational methods in place of hardware solutions to solve vision and imaging tasks in MAS. The motivation for using predominantly computation methods is: the translation to real clinical use is likely to be much more accelerated compared to novel physical devices, which may often be more difficult to rapidly prototype, as well as being subject to greater regulatory control.

The potential output of the work presented here from a clinical mindset is a framework that provides intraoperative methods for evaluating tissue viability, via measuring the concentrations of oxy- and deoxy- haemoglobin. This has potential value to a clinician; as during a procedure they may need to make a decision based on how well perfused tissue is or its degree oxygen saturation. Existing methods that can provide such information intraoperatively have various limitations such as requiring contact with tissue; only measuring at individual spots; or requiring use of a dedicated port during laparoscopic procedures in place of other devices. These limitations prevent the wider exploration and study of how this data could be beneficial to patient safety and outcomes.

The work developed thought this thesis delivers a flexible framework that can, with minimal set-up, be translated to clinical study for making preliminary evaluation of scenarios where having this information may be of use. The greater ease of

translating software based analysis tools to clinical practice also means that there are greater opportunities for rapid commercial exploitation of this work as the capital expenditure for development is likely to be lower than that of physical devices.

Acknowledgements

Over the years which have brought me to this point I have been very fortunate to have had the help and support of many wonderful people. The person who has probably been most exposed to my decision to embark on this project is my wife Anna. Anna has been supportive and encouraging through all the deadlines and hard times in my research, as well as putting up with me pulling all nighters and out of social engagements to work on the PhD. She is a constant inspiration to try harder and has shown great patience and has celebrated milestones with me along the way, it would have been very different attempting this alone. My supervisor Dan Stoyanov who floated the idea originally has consistently been an excellent motivator, helping me freely explore ideas while maintaining focus. Simon Arridge, my second supervisor, has always helped me think more analytically about problems and I have always left our discussions with a much better understanding of the fundamentals. In the early part of my study Dan Elson and Neil Clancy frequently made me welcome in their labs at Imperial College London and have been fantastic collaborators and mentors, it has also been great to get a multiple institution view on academia. Neil has continued to be fantastic touch stone for ideas throughout the rest of my studies, help me see the wood for the trees on may occasion. My research group at UCL has been a wonderful collection of people and has grown while I've been here from just Dan, Max, Xiaofei, Martin and I to its current form with multiple post-doc's PhD and masters students. It has constantly been a open and constructive forum for developing ideas, asking questions, and sometimes getting answers! Finally I thank my parents, they are an incredible duo and have been behind me from the start and throughout have given their unwavering support and confidence.

Contents

Figures

Tables

Abbreviations and Notation

Abbreviations used in this document:

MAS Minimal access surgery

CAI Computer aided interventions

MIC Medical image computing

MS Multispectral

VIS Visible spectrum

MSI Multispectral imaging

RGB Red green blue - often refering to three chanel filtered imaging centered on red green and blue wavelengths in the visible spectrum

SPD Spectral power distribution

CT Computed tomography

MRI Magnetic resonance imaging

PET Positron emision tomography

fMRI Functional magnetic resonance imaging

NIR Near infrared

US Ultra sound

LCTF Liquid crystal tunable filter

AOTF Acousto-optical tunable filter

PhA Photo acoustic

PPG Photoplethysmogram

iPPG Photoplethysmogram imaging

MC Monte Carlo

MMC Mesh-based monte carlo

MCML Monte carlo for multi-layered media

HbO Oxygenated haemoglobin

Hb Deoxygenated haemoglobin

THb Perfusion/total haemoglobin - this would be the sum of both oxygenated and deoxygenated haemoglobin

SO_2	Oxygen saturation - this is the ratio of oxygenated haemoglobin to total haemoglobin	**CoD**	Coefficient of determination
PSF	Point spread function	**MAP**	*MAXIMUM A POSTERIORI*
SVPSF	Spatially varying point spread function	**LuT**	Look up table
		MSE	Mean squared error
		SNR	Signal to noise ratio

Mathematical notation and shorthand use in this document:

- I The image, for a multi-channel image a subscript is used to denote the channels I_{RGB} corresponds to an image with three channels named R G and B.
- $\boldsymbol{x}$ A point in image space (x,y)
- k The point spread function
- Ω The index space of the image
- μ The mean distance light travels within tissue
- λ Wavelength
- G Illumination loss due to scattering out of angle of the detector
- ξ Absorption due to a specific chromophore
- α Concentration of chromophores, this is also given as $\boldsymbol{hb}$ for concentrations of both the oxy and deoxy haemoglobin, unless other wise stated this is in units of g/litre

Chapter 1

Introduction

1.1 Minimal Access Surgery and Imaging

Minimal access surgery (MAS) has revolutionised the way we perform medical interventions, drastically reducing patient trauma in many procedures. By making minimal incisions in the patient to insert tools and cameras it has become possible to work directly on the pertinent tissue with limited damage to the surrounding matter. The benefit to the patient has been seen in long term improvement in physical function and bodily pain [11]; reduced hospital stay and blood loss during surgery [12] as well as fewer respiratory complications [13]; and lower postoperative infection rate without compromising on safety margins for resected tissue [14]. Early

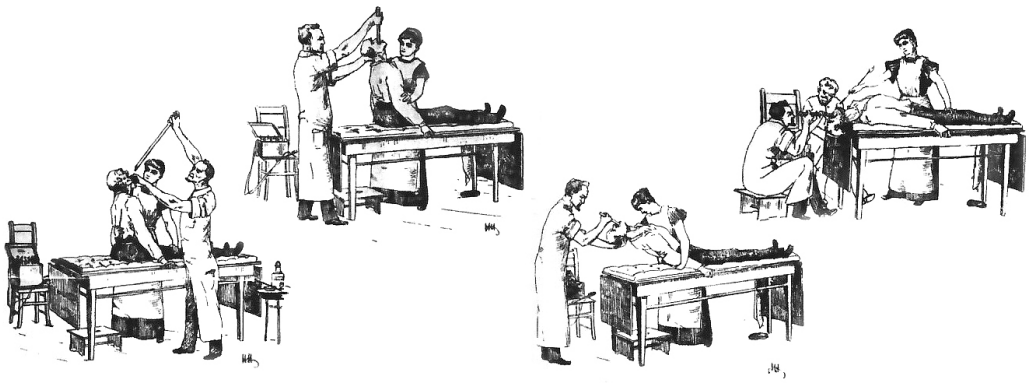

Figure 1.1: Reproduced from [1]: Hacker describing the methods performing esophagoscopy pioneered by Kussmaul who perform the first human esophagoscopy in 1868 on a sword swallower using a rigid endoscope, the illumination is provided by a gas lamp positioned externally to the patient.

endoscopes, first of which developed for gastroscopy by Aldof Kussmaul [15, 16] Figure 1.1, were often simple rigid tubes in design with a single hollow channel. These devices opened up a new potential field in surgical medicine allowing inspection and intervention of parts of the human body without invasive and traumatic procedures.

Modern MAS imaging systems have eliminated many of the restrictions of these early devices integrating sophisticated optics [17] and with multiple working channels for illumination imaging and passage for tools. It is possible to mount a variety imaging hardware at the proximal end, although recently there has been a growth in exploring uses for distally mounted imaging, in particular for flexible or very fine endoscopic devices. While there are manifold benefits to using MAS over open surgery new issues have arisen surrounding surgeon stress and fatigue [18] relating to remote operation, reduced sensory feedback, and restricted manoeuvrability. The maturing of this field has introduced robotic systems (see Figure 1.2) to hold tools, providing a more ergonomic operating posture as well as allowing multiple surgeons to easily switch between operating and assistant roles, giving opportunity to rest during long procedures. Robotic systems are also capable of scaling large hand motions down, giving fine scale accuracy when performing micro surgical tasks [19] within the restricted manoeuvrability environments of MAS. However, despite the many advances in MAS making procedures less invasive and reducing patient trauma there remains a lack of haptic feedback with most frequently used systems. This tactile sensory input can be of great use to clinicians for the purposes of localising and identifying particular tissue structures [20]. The areas of computer aided interventions (CAI) and intra-operative medical image computing (MIC) hope to assist surgical teams in overcoming some of these hurdles by providing real time information and feedback to clinicians. A large body of this work for intraoperative use has focused on problems such as localisation, for instance overlaying pre-operative imaging data with the live surgical view.

Intra-operative imaging using digital cameras and visible spectrum (VIS) wavelengths the primary sensing modality used to guide surgical actions during MAS.

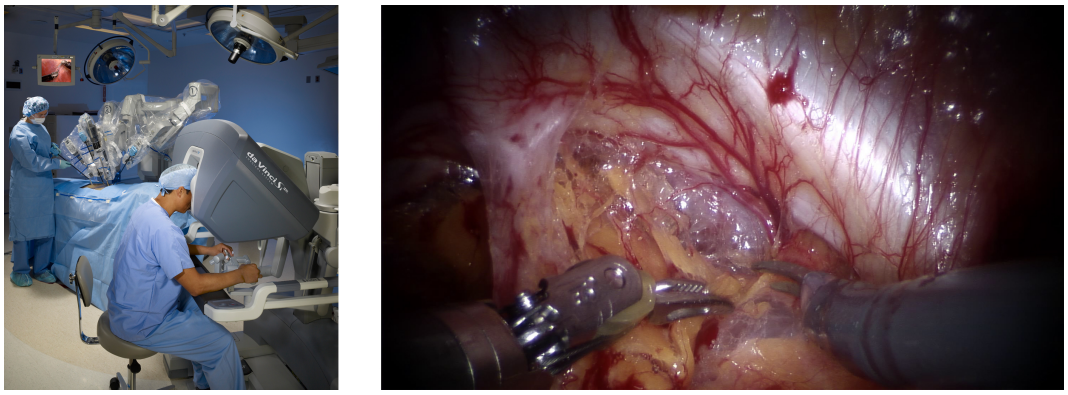

Figure 1.2: Left the da Vinci surgical robotic system for remote operation of endoscope and tools via a console ©2017 Intuitive Surgical, Inc. Right a typical endoscopic view from one of the stereo cameras as seen by a surgeon with such a device, when viewing through the console the surgeon would see the stereo pair to aid with depth perception.

In fact, MAS is largely feasibly today due to the evolution of digital imaging technologies that make it possible to deliver and harvest light from the surgical site through an endoscopic device [21]. Currently, surgeons are primarily reliant on optical images acquired in the VIS wavelength range using three channel red green blue (RGB) sensors [22] either mounted at the tip of the endoscope or attached to the distal end and linked via rod lenses. The endoscopic system then provides real-time images of the tissue surface that are used to guide diagnostic and therapeutic procedures. However access to the surgical site is restricted by the size of entry ports [23], and as a result it is challenging to develop different MAS imaging modalities for interrogating the structure, shape and function of tissues during surgery. This makes it difficult for surgeons to assess tissue viability, for example to evaluate post-transplantation revascularisation, or to characterise tissue types for diagnostic decisions. While there have been developments towards providing haptic feedback to surgeons via the control consoles, instrumenting tools for forces applied is non-trivial and further inflates the cost of devices, due to the increased complexity of devices.

To better assist with surgical guidance, a large body of work has grown around how to augment intraoperative imaging with pre-operative diagnostic imaging [24], such as computed tomography (CT), magnetic resonance imaging (MRI) and ultra sound (US). These modalities are capable of localising subsurface structures

or lesions. However for use in endoscopy they often require sophisticated 3D reconstruction and tracking of tissue to align with endoscopic VIS imaging [25]. In order to do so this requires the introduction of extra hardware to the operating theatre, increasing clutter, as well as necessitating calibration protocols that need to be completed before surgery or between channing endoscopes. When performing laparoscopic procedures there is a premium on the number of ports available for tools and during robotic procedures the space above the paitient is often restricted by space needed for the robotic arms, see Figure 1.2. While methods for localising subsurface features of interest continue to become more robust, the issue of registering diagnostic imaging data to the endoscopic view remains [26]. This is especially acute for abdominal procedures where the majority of tissue is highly deformable, and can be integral to determining patient outcomes for the case of interventional treatments of epilepsy, where it is highly desirable to avoid eloquent brain regions in spite of brain shift. As such fusing, multiple intraoperative imaging modalities to provide rich real-time patient data can be beneficial. Subsurface structural information is readily acquired intraoperatively via US probe; to overcome the added cognitive load in mentally fusing the US data and endoscopic view, optical tracking [27] from the same endoscopic view can be used to integrate to two modalities.

Purely optical imaging based visualisations can be achieved using fluorescence agents with near infrared (NIR) imaging to either highlight submerged structures (up to 10mm) or indicate tissue function such as perfusion mapping. NIR imaging of markers such as indocyanine green can be used to characterise renal cysts [28]; due to this fluorophore being confined to the vascular system it is possible to infer perfusion and degree of vascularisation beneath the tissue surface. Optical approaches such as those fluorescence imaging used intraoperatively have the advantage that the fluorescence images are implicitly registered to the endoscopic view, provided the camera sensor is capable of switching between that and normal white light imaging. Limitations for the use of fluorescence include the need for exogenous dyes, their transient use as the dyes are removed by normal bodily functions, as well as not being sensitive to oxygenation. An optical technique that is sensi-

tive to haemoglobin concentrations in tissue is narrow band imaging that utilises a small discreet set of narrow (20-30nm) band-limited measurements [29] to image tissue chosen in such a way that vascular structure is visible with a high contrast from surrounding tissue. Developing from these ideas multispectral imaging (MSI) that captures a spanning set of narrow-band measurements over the VIS wavelength range is beginning to be used in surgical applications, from tissue classification to measurement of oxygen saturation.

1.2 Multispectral Imaging

Multispectral imaging captures reflectance image data over many discrete wavelength bands such that the spectral response can be measured for each pixel of the image. Knowing the tissue reflectance spectra is useful due to the way in which they are effected by the concentrations of chromophores within tissue; these are molecules present in a volume that will absorb light of different wavelengths in characteristic ways, for instance in human tissue haemoglobin is responsible for significant absorption of light with in the VIS spectrum. MSI data if often described as cubes (see Figure 1.3) comprising the two spatial dimension corresponding to the extents of the field of view of the camera and a third dimension corresponding to spectral power distribution measured over the VIS wavelength range. This multispectral (MS) data can be reduced to a synthetic RGB view of the image by taking the inner product of the spectral response of each pixel with a known camera response curve spanning the same wavelength range to generate the red green and blue pixel data.By modelling the interaction of light and the tissue the spectral features of interest can be observed and used to infer information about the sample. This can be used to detect chromophores such as haemoglobin [30, 31], melanin and water [32] in order to perform tissue characterisation and functional interrogation of the surgical site. Examples include the diagnosis of mesenteric ischaemia, assessment of bowel anastomosis to identify the risk of anastomotic dehiscence, bowel ischaemia surgery and transplanted organ viability or visualisation of the bile duct [17]. These different method can be used to evaluate tissue viability; for na-

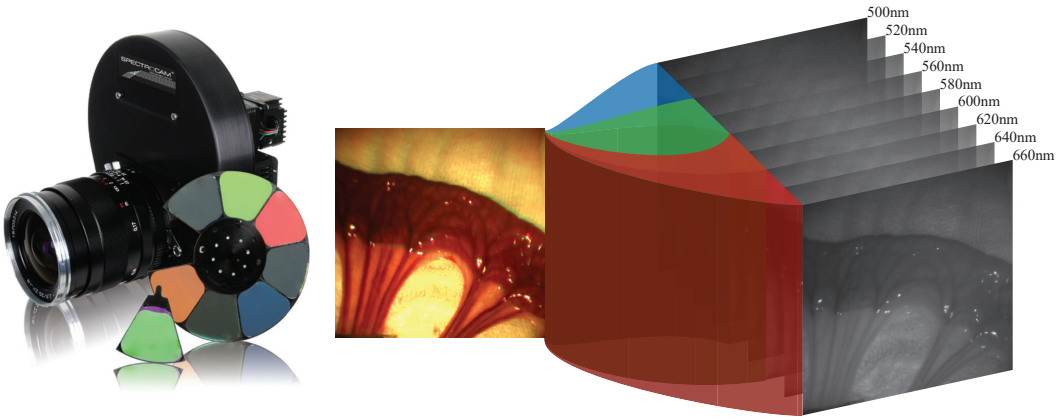

Figure 1.3: a) Commercially available SpectroCam™MS camera exhibiting a filter wheel design. b) Multispectral information (left) comprises multiple spectrally distinct measurements each imaging tissue response in a narrow wavelength range often of only few nanometres. Shown here a MS cube formed from 20nm wide bands centred at the wavelength indicated. A corresponding RGB image (right) is also shown, since the three channels of the RGB are mixtures of the underlying spectral response from tissue the appearance of the RGB image will vary depending on the camera response curve while the intrinsic MS data remains constant.

vigation and guidance; see the review by [33] for an overview of general medical applications.

However, the down side of performing MSI is that most devices have to make the trade off between spectral and temporal fidelity. The design of devices can be grouped into the categories of; snapshot: capturing a small number of discrete band-filtered images simultaneously; spectral scan: sequentially capturing band limited images over a spectral range; and spatial scan: when a spectral response is measured a point or line which gradually scans over region of interest.Devices that are able to capture multiple band limited images simultaneously are able achieve high frame rates, but the lower spectral discrimination limits their flexibility for use in varied sensing scenarios. Snapshot capture devices tend to have a high monetary cost due to needing non-standard sensor fabrication, similarly spatial scanning devises can also be prohibitively expensive for general use. Spectral scanning also introduces issues surrounding registering the separate band limited images together even when there is moderate normal tissue motion [9] this can cause image displacements of over 35% of the diagonal length image plane. Both types of MS imaging configu-

rations can be subject to blurring for very narrow band measurement as the radiant flux at the sensor can get very low, requiring longer exposure times. To perform MSI during MAS these trade-offs need to be balanced depending on the application, but still in most cases require the use of a distinct imaging system to those of the dominant modality: white light RGB imaging. Since the dominant imaging modality in MAS is endoscopic RGB[23], imaging capture techniques tend to be grouped into integrating a separate MSI channel into the endoscope; substituting the RGB camera for a hybrid imaging device capable of imaging in RGB and MS, often achieved through additional capture hardware or modified lighting environment; or in computational techniques to perform spectral unmixing of the RGB. In Chapter 2 a deeper inspection of state of the art as well as notable outstanding research questions is made covering devices and computational methods in this field.

1.3 Thesis Overview

The work presented in this thesis is motivated by the use of endoscopic MSI for the end purpose of estimating haemoglobin concentrations across the surface of imaged tissue. Initially I investigate the blurring problem arising from using narrow spectral bands in MSI, and discuss how the use of hybrid endoscopic imaging systems comprising white light RGB and MS could be of use to deconvolve MS data in a reasoned way. In Chapter 3 I look at how to use optical flow information gleaned from the white light imaging channels to estimate the spatially varying blur kernels to deconvolve MS view. However, if the desired end-point is to estimate the haemoglobin concentrations then it might be possible to do so from a lower spectral fidelity measurement then that of MSI. As such, the following chapters address leveraging white light RGB imaging for the same end, to produce a technique to rapidly and accurately estimate haemoglobin concentrations from the same endoscopic view that the surgeon is using. Throughout these chapters I investigate methods of estimating MS image data from white light RGB image data, which utilises knowledge of the the camera or imaging system's spectral response, as well as priors on the expected chromophore composition in tissue. Lastly, I look at ways

of composing multiple haemoglobin concentration estimation techniques in such a way that this estimation can be performed at high frame rates with minimal loss of accuracy.

As the anticipated clinical use for this would be for MAS, the guiding principles has been to achieve this without introducing new hardware to an already congested surgical environment. As already highlighted, working in an MAS environment brings many constraints with respect to what is practical to integrate into the surgical work flow. By letting this principle steer thinking, the result is more translatable and has the potential to subsequently be used more widely as the addition cost of use would be minimal when compared to adding additional hardware based solutions.

1.3.1 Contributions

The following contributions are made within this thesis:

- Chapter 3: a method is developed for utilising higher speed imaging, captured in parallel via another optical channel in the same laparoscope, to estimate blur kernels and demonstrates how these can be incorporated with an existing deblurring scheme.

- Chapter 4: a method for performing oxygenated haemoglobin (HbO) and deoxygenated haemoglobin (Hb) estimation from RGB VIS imaging is put forward, it is shown that smoothness constrained estimation of the latent MS data is reasonable to assume and sufficient for perfusion/total haemoglobin (THb) estimation.

- Chapter 5: a Bayesian expression for better precision in estimating MSI concentrations is developed. This has improved accuracy of HbO and Hb estimation notably for the computation of oxygen saturation (SO_2).

- Chapter 6: Haar decomposition of the image data is shown to enable a fusion of the two previous methods, such that computational performance and estimation accuracy are not degraded.

The work presented in this thesis has also contributed to the following publications:

1. Geoffrey Jones, Neil T. Clancy, Simon Arridge, Dan S. Elson, and Danail Stoyanov. Deblurring multispectral laparoscopic images. In *Information Processing in Computer-Assisted Interventions*, pages 216–225. Springer International Publishing, 2014

2. Geoffrey Jones, Neil T. Clancy, Simon Arridge, Daniel S. Elson, and Danail Stoyanov. Inference of tissue haemoglobin concentration from stereo RGB. In *Medical Imaging and Augmented Reality*, pages 50–58. Springer International Publishing, 2016

3. Geoffrey Jones, Neil T. Clancy, Yusuf Helo, Simon Arridge, Daniel S. Elson, and Danail Stoyanov. Bayesian estimation of intrinsic tissue oxygenation and perfusion from RGB images. *IEEE Transactions on Medical Imaging*, pages 1–1, 2017

4. Geoffrey Jones, Neil T. Clancy, Xiaofei Du, Maria Robu, Simon Arridge, Daniel S. Elson, and Danail Stoyanov. Fast estimation of haemoglobin concentration in tissue via wavelet decomposition. In *Lecture Notes in Computer Science*, pages 100–108. Springer International Publishing, 2017

Further the contributions of this thesis comprise substance of the novel material component for the 2017 United States of America patent No. 16/310,286; G. Jones, D. Stoyanov, D. Elson and N. Clancy *Method and Apparatus For Estimating The Value Of A Physical Parameter In a Biological Tissue*.

Chapter 2

Multispectral Imaging in Surgical Applications

This review focuses on three area of MSI specifically targeted for the context of MAS. Topics discussed include hardware MSI devices; computational techniques for estimation and processing of multispectral data; and methods of evaluation for these*ex vivo*, and *in vivo*. Papers for this review were drawn from the PubMed Database and Google Scholar. There is a comprehensive general review of of medical MAS found in [33]. The search criteria used to automatically filter papers are those listed in the key words section variants of these key words used were:

- *Minimal[ly] Invasive Surgery / minimal access surgery / key[-,]hole surgery*
- *Multispectral [imaging] / hyperspectral [imaging]*

This chapter is divided in two core sections; the first section looks at developments in MSI hardware, which is divided by approaches that filter light arriving at the camera and approaches that filter light leaving the light source(s); the second section looks at computational techniques of estimating and processing MS data and within the later explores the contrasting requirements for different domains of task such as computer vision, and biophotonic tissue analysis. Finally this review concludes with an exposition on the key hurdles for clinical translation in MAS of MSI and data processes.

2.1 Multispectral Imaging Hardware

Several physiological variables such as HbO and Hb concentration and the diameter of cell nuclei, affect how light of different wavelengths is transmitted through tissue. These can be measured by detecting their characteristic attenuation of light in the visible wavelength range [37, 38]. This can be achieved using multispectral imaging where multiple narrow band images are acquired , resulting in a three-dimensional dataset, as shown in Figure 1.3. The contribution of individual absorbing and scattering elements in the tissue can be extracted from the reflectance spectrum recorded at each pixel. This information can then be further used to generate concentration maps that show spatial variation in the tissue. For example, visualisation of tissue oxygenation has been used in various procedures throughout the body to provide information on the oxygenation of tissue in and around the surgical site [39, 17, 31].

MSI hardware used to record this spatial and spectral information must store this three-dimensional data on two-dimensional sensors. The methods used to do this fall into three broad categories: spectral scanning, spatial scanning and snapshot acquisition [33].

Spectral scanning devices record a sequence of images, each under illumination from a narrow range of wavelengths. The sample is illuminated with broadband light and a series of optical bandpass filters are placed in front of the camera. This may be done using a filter wheel [40, 41]or tuneable filter [42, 43].The technique may also be implemented by scanning the illumination source via dispersive optics [44], filter wheels [45]or multi-wavelength LEDs [46]Spatial scanning systems work by acquiring a single line of spatial information in one dimension and dispersing the spectral information in the other. The line is then scanned across the field of view to obtain the second spatial dimension [47, 48, 49].Snapshot systems aim to collect both spatial and spectral information on a single sensor. This may be done using prisms and dichroic filters to generate multiple, spatially-distinct and wavelength-resolved images on the sensor, or slicing the image into spectrally-dispersed strips [50].More recently sensors with integrated filters attached directly onto the chip have provided a low-footprint option.[51, 52].

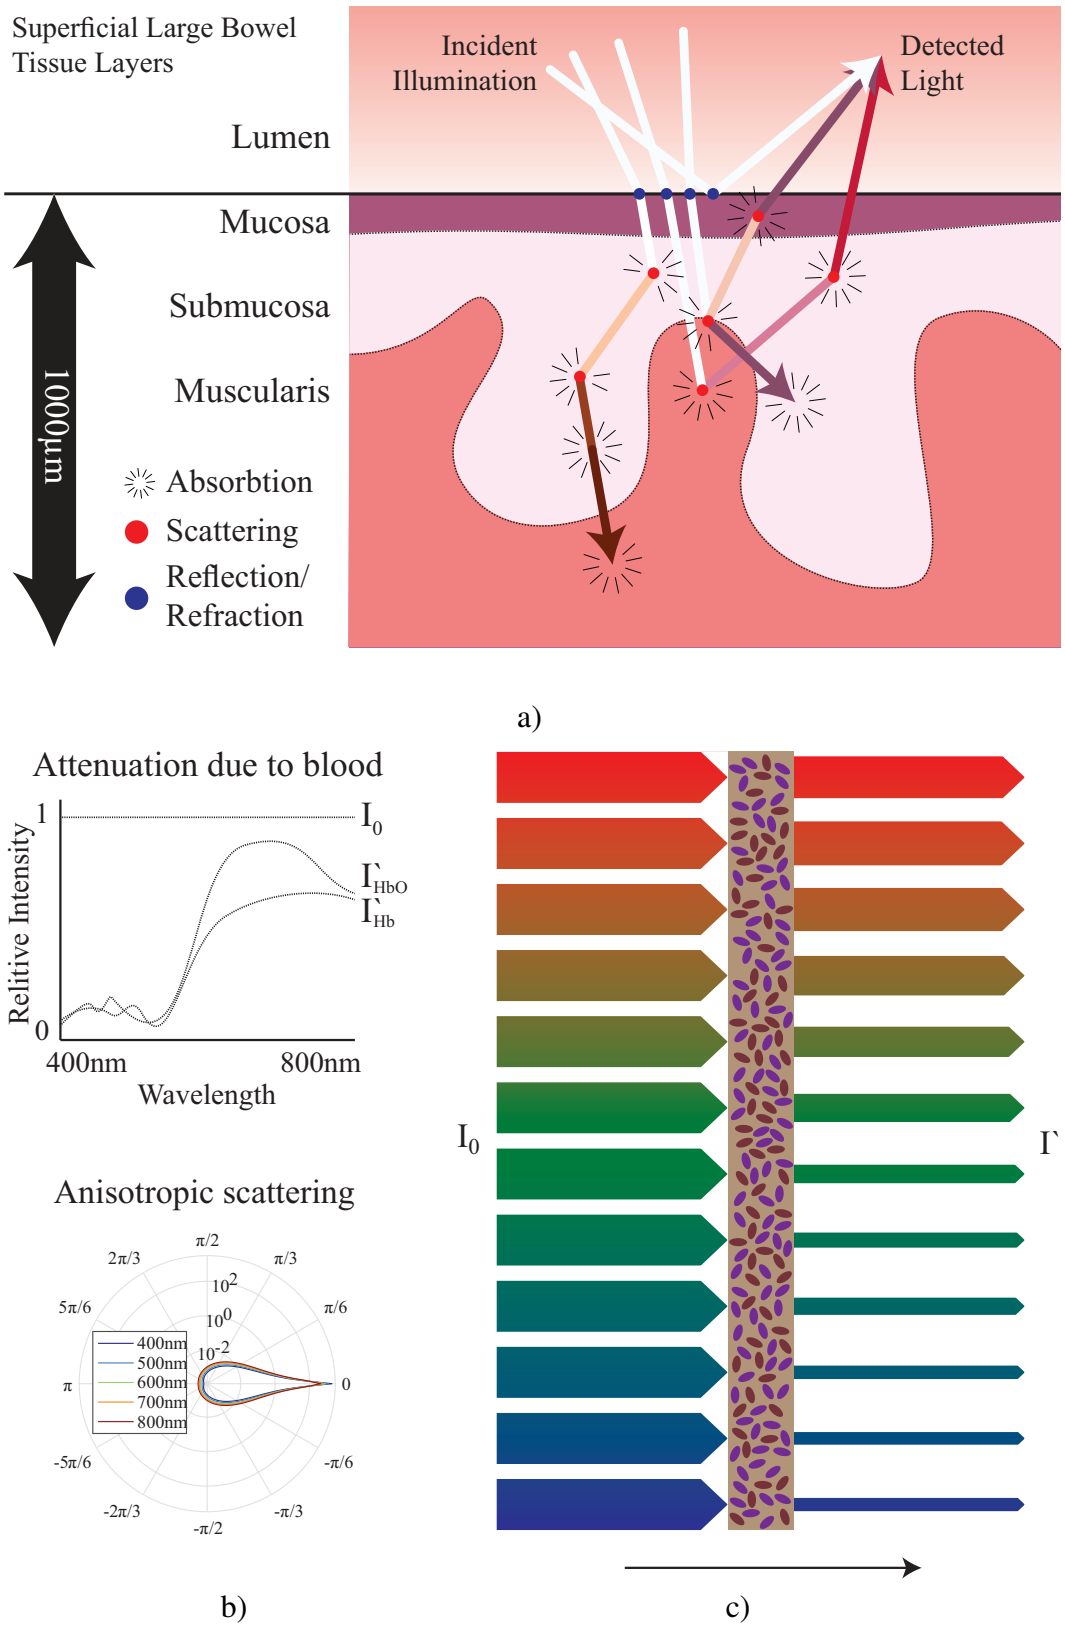

Figure 2.1: a) The interaction of light in the visible wavelength range iterating with the superficial tissue layer. b) top: the transmission of light through HbO or Hb, bottom: the anisotropic scattering of light due to HbO over wavelength. c) light passing through tissue is more greatly attenuated at the blue end of the visible spectrum.

Spatial scanners can achieve high spectral resolution (5 nm) at up to hundreds of wavelengths but require a secondary scanning mechanism. Spectral scanning mechanisms can also use high quality filters to achieve high spectral resolution, but at fewer wavelengths. The sequential nature of the acquisition means that both methods are vulnerable to motion-induced errors. An additional drawback of filtering methods is that they significantly reduce the amount of light reaching the camera sensor necessitating either specialised high sensitivity cameras or long exposure times. The main drawback of modified illumination methods, is that they can be visually disruptive to an operator or interfere with other concurrent imaging modalities such as RGB imaging. Snapshot systems avoid this problem but this comes at the expense of spatial resolution and the number of available wavelengths as everything must be captured on a single sensor. These trade-offs between spatial, temporal and spectral resolution must be considered when choosing the type of mechanism used.

A way to avoid the slow imaging of a high spectral resolution MS data set is to only image the spectral bands that are most informative to the clinical task at hand: as demonstrated by [53] for skin imaging; [54] for fluorescence imaging; and [55] for MAS based haemoglobin concentration estimation. These techniques, while sufficiently rapid to provide real time clinical information, all suffer from the same draw back: they are specific to a single imaging task, and it is desirable in MAS that these functions could be performed from a single imaging device. An ultra-high frame rate (2800fps) MSI device proposed by [56] uses passive optical elements including a periscope array within a spectral shaper, splitting the image into multiple bands and offsetting each to a unique part of the sensor. This methods imaging rate is only limited by the read-out speed of the sensor and as such is able to achieve very high frame rates with a sufficiently fast sensor. Further, due to the optical configuration and all band limited images being projected onto the same sensor, this is able to image without distortion. The rapid imaging comes at the sacrifice of resolution as the image as the separated band-limited images are all projected onto the same sensor, however it may be possible to mitigate by using a

higher resolution sensor.

Multispectral imaging has a number of attractive qualities compared to other modalities. For example, it has been used to make *in situ* assessment of ischaemia in heart tissue [57] without ionising radiation, compared to imaging via CT or positron emision tomography (PET). In other use-cases, MSI can be used to minimise photo toxicity [58] by only imaging at salient wavelengths. In MAS, imaging the oxygenation of tissue has also been used to identify malignant tissue [59] due to the increased vascularisation of malignant tissue resulting in a local increase in the THb within tissue. The use of MSI to assess the success of a uterine transplant procedure has been shown by [38], making inference of the SO_2 and THb in the transplanted organ. In these works the multispectral imaging is performed using a liquid crystal tunable filter which separately captures images across multiple narrow bands. While this method of capture is very versatile, its temporal resolution is inversely proportional to the spectral resolution, and the sequential capture protocol can result in displacement between images taken at different wavelengths. To compensate for displacements, [38] includes a registration step to correct for tissue motion relative to the camera in between captures for each spectral band.

In neuro-surgical applications, the disturbance from camera and tissue motion can be minimised; this allows for the use of MSI[60] to image the change in oxygenation in the exposed upper cortex.The key observation is that after controlling for tissue and camera motion, temporal variation in the scene will be due to changes in concentration of constituent chromophores. As a result the temporal oxygenation signal can be used to make inference about the physical response in the brain to various stimuli, without requiring injection of any staining agents into cortex itself. This work make use of variable illumination strobed light emitting diodes tuned to particular wavelengths, to allow a higher speed capture of the dynamic tissue response; however in do so places a potentially onerous demand on the surgical environment by requiring such a unique lighting configuration.

While differing in practical configuration, the methods of [60] and [9, 55], share a similar theoretical basis. Measurements of the attenuation of illumination

across multiple non-overlapping spectral bands, are made and utilising the Beer-Lambert equation estimation of chromophore concentration (described in Section 2.2). In the case where there is no fluorescence, these techniques are analogous as they make spectrally distinct measurements of the scene. The chromophore concentrations can then be used to estimate the intrinsic SO_2 and THb in Figure 2.2 2.3 and 2.4 the use of various SO_2 mapping applications are illustrated.

Capturing a full multispectral image of a surgical site at sufficient spectra resolution to perform inference on can be slow, especially if using a filter based approach as a significant amount of the illuminating light is filtered out necessitating long exposure times. Rapid capture through illumination [60] is achieved by high speed strobe illumination: typically utilising a carefully chosen set of LEDs that illuminate the surgical site with different narrow band filtered light sources. Rapid filter based approaches such as that of [55] use a reduced number but maximally discriminative filter set, allowing for rapid capture of the full image set required to make an estimation of the haemoglobin distribution and oxygenation.

The method is very successful, however it is less flexible due to the reliance on learning the characteristic reflectance spectral for a specific tissue type. In [45] a combination of multiple spectrally distinct bandpass filters with an RGB sensor are used to rapidly sample the power distribution of light in the visible wavelengths returning from tissue. This is achieved by using triple band pass filters mounted in a filter wheel to make simultaneous and spectrally distinct measurements, recoding a set of 18 band filtered measurements in six captures. This framework accounts for spatially varying geometry and achieves a frame rate of 15fps by using hardware accelerated (GPU) linear matrix inversion to estimate SO_2 and THb distributions.

An alternative approach for delivering filtered MS illumination is to use digital mirror device (DMD)[61] to select spectral bands from standard wideband illumination source such as a typical surgical xenon arc lamp (covered in greater depth in [33]). While originally conceived for open procedures, due to the complexity of delivering illumination in this way it is possible to adapt it for use in MAS [44]. Here it is again used to characterise the tissue SO_2 levels during, in this case, partial

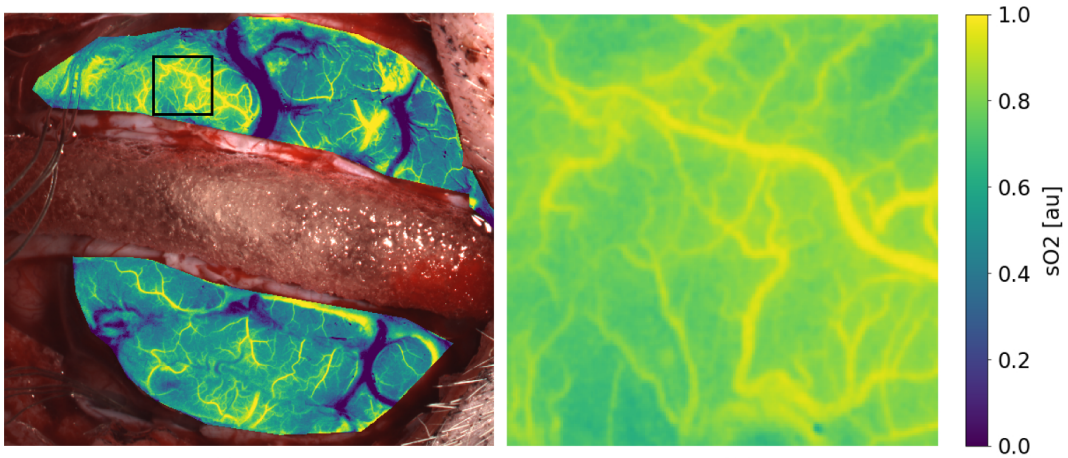

Figure 2.2: Oxygenation map of a porcine brain generated by spectral sweep 8-band MSI device. a The full brain. b Cropping of the region marked in a. SO_2 generated using invertible neural networks as part of work to evaluate MS imaging hardware selection for estimation of HbO and Hb. Reproduced with permission from [2]

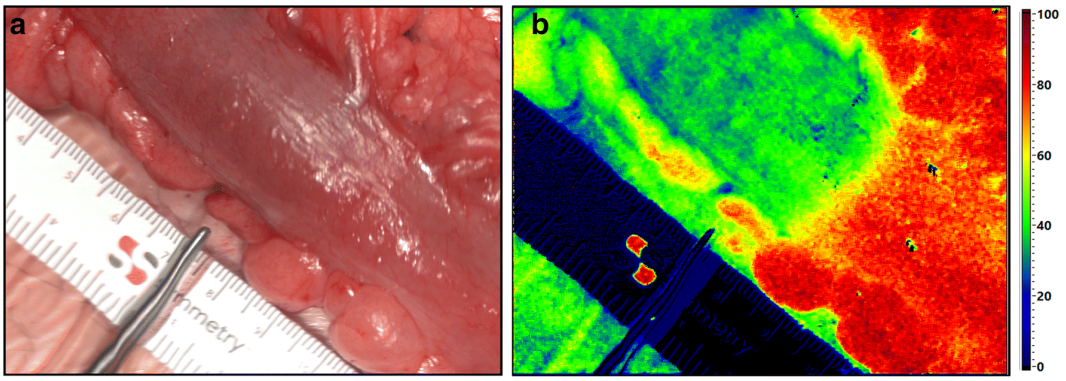

Figure 2.3: Snapshot imaging of the bowel visualising SO_2 from MSI, visible is the boundary between sections of a devascularised ischaemic bowel (left) and healthy bowel (right). Reproduce with permission from [3].

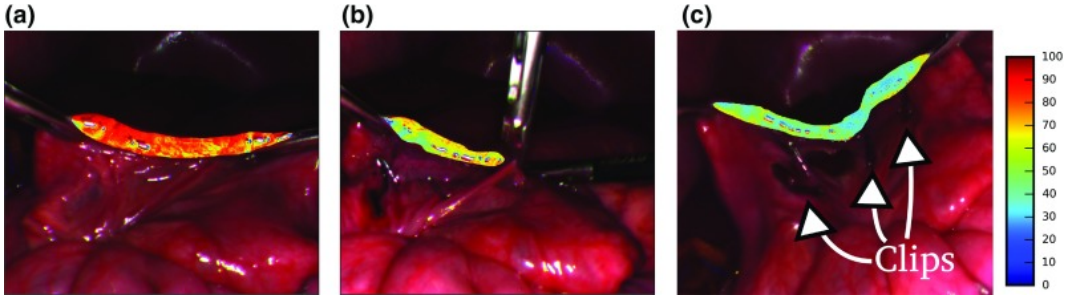

Figure 2.4: Visualization of small bowel SO_2 estimation segmented and shown as percentage oxygen saturation. Data captured using an 8-band spectral sweep MSI device. a) Before clipping. b) Setting of the fourth clip. c) Nine minutes after applying first clip. Reproduce from [4] creative commons cc-by.

nephrectomy procedures.

A MAS compatible structure light MS illumination probe has been developed using the output of super-continuum laser[62] dispersed via a prism onto an incoherent fibre bundle creating a randomised spectrally encoded pattern. They demonstrate that this pattern can be used for surface reconstruction, observed a measurable attenuation in the reflectance of spot data for different types of tissue potentially giving a dispersed MS reflectance imaging signal across the surface. The main draw-back of this system is the sensitivity in detecting the spot, pattern especially in parts of the spectrum where there is large overlap in the RGB sensor response curve. Combining this approach with a 8 channel MS filter wheel based camera in replacement of the RGB sensor it is possible to more accurately detect the spectrally encoded pattern[63].

2.2 Computational Multispectral Techniques in MAS

When full MSI information is available, often the straight forward linear solution is sought, to the Beer-Lamber law for many independent attenuation data for different wavelengths. The apparent absorption at different wavelength is then compared to known extinction coefficients thought to be present in the observed tissue. This linear fitting deduces relative concentrations of these chromophores by making use of the Beer-Lambert law; this states that for a given wavelength λ the intensity of light leaving tissue $I(\lambda)$ can be expressed in terms of the incident illumination intensity $I_0(\lambda)$

$$I(\lambda) = I_0(\lambda) e^{-\mu \Sigma^n \cdot \alpha_n \cdot \xi_n(\lambda)} \tag{2.1}$$

where $\xi_n(\lambda)$ is the n^{th} attenuation coefficient at concentration α_n and μ is the width of the cuvette holding the sample imaged, as illustrated in Figure 2.5. This is originally posed as a transmission context with the sample is placed in between the light source and the detector. For intra operative use case, however, the light returning

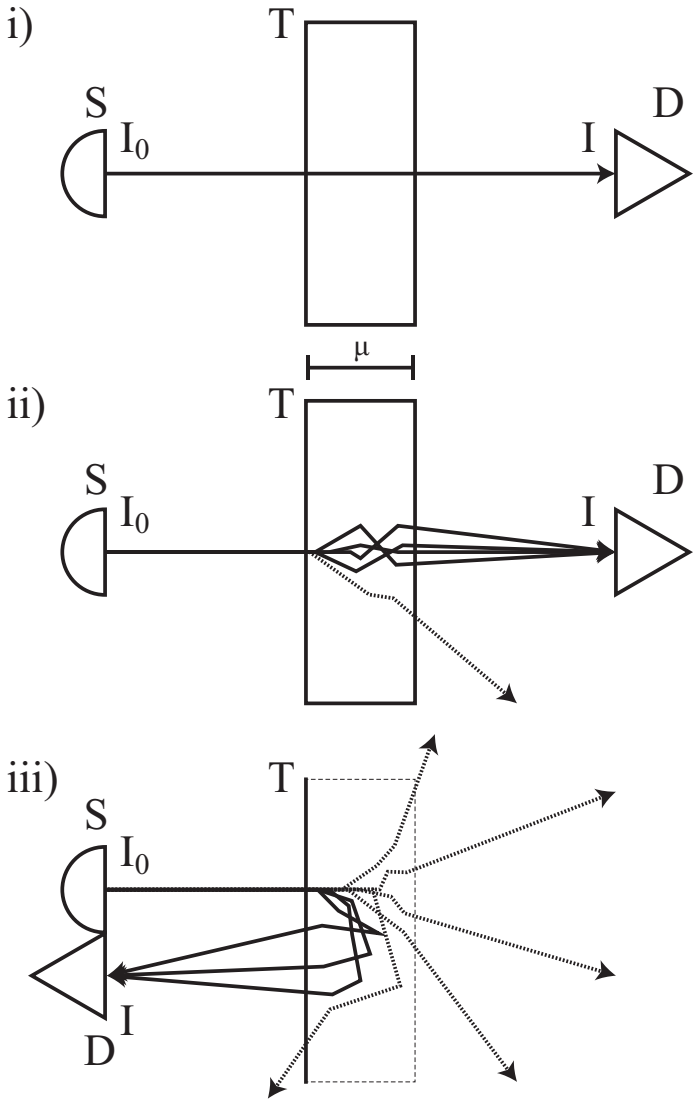

Figure 2.5: The variants on attenuation Beer-Lambert law as developed for scattering media with light emitted from source S passing through tissue T and arriving at detector D. i) Classical Beer-Lambert attenuation where initial illumination with spectra I_0 passes through a thin (μ wide) non-scattering medium with the resultant detected spectra I; ii) the modified Beer-Lambert equation introducing a DPF to account for scattering induced longer path length within tissue as well as geometric factor to account for light lost due to being scattered away from the detector; iii) Backscatter imaging, most of the incident illumination that enters tissue is lost due to absorption before the compound scattering deflect the light back out of tissue towards the detector

Table 2.1: Comparison of recent MSI and processing methods applicable to MAS

Method	MS method[1]	Hardware[2]	MS resolution (nm)	frame rate	MAS contexts[3]
[64]	Df	c	50		G
[62]	Dp	f	5		F,H,G
[9]	Df	c	10		F,H,G
[44]	Df	h	4	3fps	F,H,G
[65]	R	a	-	15fps	H,G
[66]	Df	e	13	.5fps	F,H,G
[55]	Dp	d	20-25	400ms	H,G
[67]	Dp(NrwBnd)	a	25/60	>30fps	H,G
[56]	Dp	[4]	25-50	2800fps	H,G
[45]	Df	a+d	15	15pfs	F,H,G
[63]	Dp	f,d	-		F,H,G
[68]†	I	a	-	30fps [12s delay]	H
[69]	Dp	g	150		G
[70]	Dp	b	12-20	2fps	G
[71]	Dp	d	20-25	400ms	G
[4]	Dp	d	20-25	400ms(+200ms)	H
[8]	R	a	-	1fps [>30fps: LUT]	H,G
[72]†	R	a	10nm	<1fps	G

[1] MS imaging method: [D]irect [f]ull or [p]artial MS imaging; [E]xplicit MS reconstruction; [I]mplicit reconstruction [2] Method's core imaging hardware: a) RGB b) acousto-optical tunable filter (AOTF) c) LCTF d) Filter Wheel e) Scanning/Multiplex Illumination f) Structured Illumination g) Photoacoustic h) DMD [3] Viable for [F]lourescence imaging, [H]aemodynamic imaging, [G]eneral imaging and guidance [4] Uses a periscope array in a spectral shaper to simultaneously separate images of multiple bands † not yet validated on endoscopic data

to the camera is prominently surface reflections and backscattered light. A plot of light scattering in HbO is shown in Fig. 2.1 which illustrates the highly anisotropic forward scattering profile, also notable is variation between light of shorter wavelengths and that of longer red wavelengths. When imaging from the same side of the tissue as the illumination, the effect of scattering has a significant impact on the resultant spectra of returning light. This is further compounded by the fact that the amount of absorption scales exponentially with distance travelled through the tissue, and the wavelength varying scattering profiles add another degree of bias. As such, for the application of imaging highly a scattering medium like human tissue, this algorithm was augmented by [73] as:

$$I(\lambda) = I_0(\lambda) e^{G - \mu \sum^n \mathrm{DPF}_n(\lambda) \cdot \alpha_n \cdot \xi_n(\lambda)} \tag{2.2}$$

Here, still imaging in a transmission context, a constant geometric factor G represents the losses due to scattering the path of a photon out of the detection region of the detector; further the differential path length factor (DFP) represents the increased path length within tissue, that scattered photons will have travelled (see Figure 2.5). Studies such as [7] have collated data from many sources to accurately estimate the attenuation and scattering coefficients of whole blood. More generally with respect to haemoglobin levels found in humans: [10] note that the average THb concentration in blood is variable from person to person and across gender as well as age, they utilise an approximation of $150 g/L$ as general value for THb in whole blood. This approximate is in line with values stated by [74] (Table 18-2) suggesting $140 g/L$ for women and $155 g/L$ for men. However these are estimates for the total haemoglobin as measured in whole blood in where the haematocrit may be between 40-45%, in other tissue this may reasonably be expected to not be the same. Indeed when considering THb concentrations in different tissue types [10] note that there can be significant variance as noted in table 3 of their paper and summarised in Table 2.2.

When switching to an imaging context where both light source and detector are on the same side of the tissue, it is necessary to rigorously incorporate the effects

Table 2.2: Varying expected concentrations of haemoglobin in tissue of the human body summarising [10] table 3

Region of the body	g/L
Brest	100 - 219
Abdomen	81
Skin	7.5 - 52.5
Forearm	754.5
Calf	541.5
Head	502.5
Prostate	1,386

of the scattering on attenuation as the main way of light returning from subsurface is via a series of scattering events. While this could be accurately modelled as a diffusion process [75], it results in the need to use non-linear solvers which would become impractical for high-definition video data. As such it can be desirable to preserve the linear solutions as given by Equation 4.1 but modifying the attenuation coefficients appropriately to accommodate for the departure from a transmission imaging context. It is possible to derive an effective attenuation coefficient μ_{eff} from diffusion theory [76] as follows:

$$\mu_{\text{eff}} = \sqrt{3\mu_a(\mu_a + \mu_s')} \tag{2.3}$$

where μ_a is the attenuation and the transport scattering coefficient $\mu_s' = \mu_s(1-g)$ comprises the scattering coefficient μ_s and anisotropy g. This can be used to compute the fluence rate at depth within tissue however this does rely on a homogeneous tissue composition and a semi-infinite tissue sample, as such numerical methods can be used instead to generate single attenuation parameter $\tilde{\xi}$ for each chromophore. These new parameters can be determined via Monte Carlo simulation [77] using a digital tissue phantom and accumulating the resulting backscattered fraction of simulated photons entering the tissues. The model used by [77] comprises 9 layers of homogeneous tissue parametrised by concentration of blood; oxygen saturation;

water content; melanin fraction as well as a few others shown in table 1 of their paper. Illustrated in their work (Figure 3) are the effects of absorption due to other chromophores such as water and melanin, clearly shown is how over VIS the different forms of melanin are most dominant in terms of absorption however in the absence of melanin haemoglobin is the dominant absorber over water.

For imaging where the detector (camera) is on the same side of the tissue as the light source, the effect of scattering is even more crucial: if the tissue is homogeneous, then the only way that light will return to exit the tissue back in the direction of the light source will be due to compound scattering.

$$I(\lambda) = I_0(\lambda) e^{G - \Sigma^n \alpha_n \cdot \tilde{\xi}_n(\lambda)} \tag{2.4}$$

The compute attenuation coefficient $\tilde{\xi}$ combines the travel distance factors μ and DPF, which are dependant on scattering, with the attenuation. In the work presented in this paper simulation input variables such as the absorption, scattering and scattering anisotropy were taken from [7] tabulated at the end of the paper. Differential path length factor were computed by mean photon distance travelled through simulated tissue to give a net attenuation for light of a given wavelength passing through tissue of that type.

All of these methods result in a system of linear equations fitting concentrations α_n to the negative log ratio of incident illumination to detected light across all wavelengths:

$$-\log\left(\frac{I}{I_0}\right) = A\boldsymbol{\alpha} \;\; (-G) \tag{2.5}$$

Where A is the matrix of total combined attenuation, this is possible due to the high absorption in human tissue of light in the VIS wavelength range meaning light leaving tissue will have most likely only probed a very small volume.

Within the field of pulse oximetry an alternative approach has been to take multiple measurements over time, making the assumption that while structure in the tissue may be unknown, it is constant, and so any variation in measurement will be due to changes in blood volume and oxygenation. The typical way in which this

is performed is to take measurements at two wavelengths and then analysing the normalised ratio of the two measurements as described by [78].

$$S(t) = \frac{AC(t)_{\lambda_1}/DC_{\lambda_1}}{AC(t)_{\lambda_2}/DC_{\lambda_2}} \tag{2.6}$$

Where $AC(t)_\lambda$ is the instantaneous measurement and DC_λ is the rolling average over some time window corresponding to the baseline total blood volume within tissue. This provides a normalised estimation $S(t)$ for the SO_2 as a time series. The key requirement for this is that the tissue being sampled does not move or change due to other factors. A further detailed introduction to current material on light tissue interaction and modelling can be found in [79, 80]

2.2.1 Narrow Band Measurement

The evaluation of MS data for tissue parameter estimation (Hb, HbO etc.) by the Beer-Lambert law has limitations while being fast (outlined at the beginning of section 2.2). While extra terms can be added to account for factors such as scattering, in order to maintain the nice linear formulation, approximations are often substituted such as determining via Monte Carlo (MC) simulation a single value for the DPF irrespective of the tissue type. As such, fully model-based approaches invert the MS data in a non-linear fashion allowing for variable tissue physical properties as well as chromophore concentrations. While accurate, these techniques do not scale computationally to the task of real-time processing of over mega-pixel sized MS image data. The method of [4] proposes an innovative middle way, obtaining the high accuracy as well as rapid computation even for mega-pixel image data. This is achieved via training a random forest regressor on MC simulated MS data using a porcine bowel tissue model. In this manner they were able to learn a very high degree of possible tissue model parameter variation in the training data. Whilst multi-day training is required initially, prediction is naturally much faster with the imaging and regression speeds at 400 ms and 200 ms respectively. The core limitation is that the random forest regressor is limited to the specific tissue type on which it was trained, however due to the success of tissue classification form MS

data it is reasonable that this could be mitigated by having many trained random forest regressors and picking the particular one for the classified tissue type.

[66] uses a variable optical filter, which transmits a narrow band of wavelengths the centre of this band varies along the length of the filter. By translating the filter in front of an aperture the desired waveband can be selected from the illumination source, self calibrating by adjusting exposure and iris aperture from white reflectance standard. The hardware implemented calibration for ambient illumination has the benefit over software calibrated methods which typically resort to multiplying wave bands at the edges of their sensitivity by large scalers amplifying noise in the measurements.

Another recent method for MSI by varying illumination is demonstrated by [81, 70] where they use an AOTF to create band pass filters for the illumination source. They show this to have long term robustness with minimal drift in the calibration over half a year of clinical use over all 6 bands. This is a noteworthy point because it potentially shows how such a device could be routinely used in clinic without the need for an expert to re-calibrate the imaging setup for every use. The imaging device can image MS data at $>$2fps and operates either in MS imaging mode or RGB white light imaging. They show that the MS data is sufficient to perform basic *in-vivo* detection of stomach carcinoma.

Illumination artefacts can cause numerous issues across the study of imaging and reconstruction and models will often resort to treating these as outliers to the acquisition process, leaving holes in the data. The methods of [67] and [82] make use of cross polarisation to remove the surface reflections and short wavelength depolarisation metric (DoLP - Degree of linear polarisation) corresponding to regions of high tissue profusion. Methods like these that are able to eliminate the superficial reflections from the surface may be of particular interest in the exploration of imaging techniques that try to detect or image subsurface features such as haemoglobin concentration or detecting fluorescence markers.

Given an MS data set [71] demonstrate high quality tissue classification using a combination of intra-waveband multi resolution rotationally invariant local binary

pattern (LBP) descriptors. [71] observe that most current MS technique acquire wavebands sequentially necessitating a registration step to combine individual band limited images into an MS cube, and as such the reliability of using the spectral data at a location is limited by the accuracy of the registration. To integrate the MS data they compute an inter-waveband average spectrum in a local patch sacrificing some resolution with the gain of robustness to noise and small mis-alignment of the registration. The LBP for every band and the patch based average spectrum are concatenated as a single descriptor vector.

With a two band MS light source [69] demonstrate the feasibility of a device to image MS photo acoustic (PhA) images using a hollow needle to inject the light stimulus beneath the tissue surface. The use of a MS light source gives greater tissue specificity. Further due to the light source being injected beneath the surface, the device is capable of capturing MS data at much greater depths than surface illuminated MS PhA devices or optical MS imaging, because the attenuation of ultrasound being much less than VIS light in human tissue. In this device's current form, however, there are concerns with the light fluence at the needle tip which is currently above accepted safety margins for human tissue.

Recent rapid capture techniques have focused on the use of RGB cameras to capture data at near real-time speeds, by capturing parallel measurements of the surgical site in each channel. The work of [65] reconstructs intrinsic SO_2 and THb from RGB images through a two step approach: firstly Weiner filtering based reconstruction of the expected spectrum; and secondly Beer-Lambert based of estimation haemoglobin characteristics.

2.2.2 Spectral Estimation from RGB

Methods that make direct use of the RGB camera feed from endoscopic cameras already in place for VIS spectrum human compatible imaging are of particular value as they often require little or no modification to existing hardware. The lack of modification to hardware and lighting environments for imaging means that these techniques are often able to be performed simultaneously with existing VIS imaging, potentially enabling the fusion of multiple vision techniques; and also there is

a significantly lower barrier for clinical translation. These techniques often rely on an intermediate step where a spectrally distinct estimation of the MS data is made from the spectrally correlated RGB data. The use of RGB data to evaluate concentrations of constituent chromophores is explored by [77], validating the correspondence between simulated CIE L*a*b* [83] and measured MS data for a variety of body parts. While their presented method does not elaborate to inverting RGB data converted to CIE L*a*b* it does illustrate the perceivable difference in RGB images of tissue simulated to have a range of chromophore concentrations, potentially enabling the creation of an atlas of RGB values and corresponding concentrations for a given illuminant and camera tri-stimulus.

A Wiener filtering technique is made use of by [65] to estimate the MS data of skin measurements. The MS data is then processed as above using the Beer-Lambert relation with MC determined backscattered attenuation coefficients for HbO, Hb and melanin. The measurement process is posed as a linear iteration $\boldsymbol{I}_c = \boldsymbol{r}_c \boldsymbol{I}$ where $\boldsymbol{I}_c$ corresponds to sensor measurement in channel c; $\boldsymbol{r}_c$ is the combined illumination spectra and camera spectral sensitivity in each colour channel; and $\boldsymbol{I}$ is the reflectance spectra corresponding to the location of a pixel measurement in the image. As is usual for Wiener filtering based reconstruction method seeks to find an estimation matrix $\tilde{\boldsymbol{I}} = \boldsymbol{W}\boldsymbol{I}_c$ that minimises the squared error between the estimated ($\tilde{\boldsymbol{I}}$) and real MS data:

$$e = \left\langle \left(\boldsymbol{I} - \tilde{\boldsymbol{I}}\right)^t \left(\boldsymbol{I} - \tilde{\boldsymbol{I}}\right) \right\rangle \tag{2.7}$$

The derivation of the Wiener estimation matrix results in the formulation:

$$\boldsymbol{W} = \left\langle \boldsymbol{I}\boldsymbol{I}^t \right\rangle \boldsymbol{r}^t \left(\boldsymbol{r} \left\langle \boldsymbol{I}\boldsymbol{I}^t \right\rangle \boldsymbol{r}^t\right)^{-1} \tag{2.8}$$

Where the $\langle \cdot \rangle$ denotes an ensemble averaging operator. To determine the factor $\langle \boldsymbol{I}\boldsymbol{I}^t \rangle$ [65] collate many observation of skin $\boldsymbol{I}$ using a spectrometer and compute the auto correlation of the all of these measurements. This results in very accurate estimation of the MS data and in turn chromosphere concentrations, provided there is low sensor measurement noise and tissue is at the correct imaging distance. This

method is a local pixel-wise method and as such does not include and neighbour information also due to the lack of consideration for sensor measurement noise this method may be less well applied to the lighting environment of endoscopic imaging. Further more, the need to learn the spectral reflectance auto correlation matrix $\langle \boldsymbol{I}\boldsymbol{I}^t \rangle$ for each tissue type would necessitate a tissue segmentation and classification pre-processing step to select the correct Wiener filter to use for each image pixel. This prerequisite scene processing may not be a significant draw-back as classification and segmentation can be performed rapidly via various machine learning approaches.

The work of [72] also utilises known response curves from multiple cameras to reconstruct the latent MS reflectance signal for the scene, as well as the illumination spectra. This is achieved through an alternating least squares optimisation solving the bilinear system comprising the surface reflectance and illumination models. Integrated into this approach are constraints for positivity and smoothness of the estimated spectra at every iteration. This method assumes the multi-view registration problem is solved prior to performing the MS estimation. While in their work this is less of an issue, they are able to image planer targets thus having a simple homography mapping; for adaptation to surgery, rapid dense stereo matching would be required. Further, while impressive in the accuracy of reconstruction the method has a high computational cost, producing MS acquisition times equivalent to that achieved with liquid crystal tunable filter (LCTF) based MS set-ups. In their work they propose an optimisation by making the assumption of uniform illumination to only compute the illumination at small number of locations (30) and then keep this fixed for the reflectance estimation.

Techniques that attempt to measure a photoplethysmogram (PPG) from non-contact RGB measurements also have the potential to be applied to endoscopic assessment of tissue. These typically make measurements over a given time period extracting periodic pathological signals; as such, motion artefacts can have detrimental effects on the analysis. The presence of significant tissue and camera motion in endoscopic procedures would preclude the use of PPG techniques that do

not adequately compensate for motion as discussed by [84]. The remote photoplethysmogram imaging (iPPG) method of [68] monitors RGB video data imaging tissue extracting periodic signals to effectively track relative and absolute oxygen saturation changes within tissue. This method combines dynamic region of interest selection and temporal averaging to increase robustness to motion. To perform iPPG the original ratio-of-ratios approach of [78] is used, however [68] re-derive the formulation generalising for use with wideband RGB sensor measurements.

$$\begin{aligned} S(t) &= \frac{\hat{R}_1}{\hat{R}_2} \\ &= \frac{\sum_{\lambda_1} \hat{I}_0(\lambda) r_c(\lambda) m(\lambda, \overrightarrow{x}) b_{AC}(\lambda, t)}{\sum_{\lambda_1} \hat{I}_0(\lambda) r_c(\lambda) m(\lambda, \overrightarrow{x}) b_{DC}(\lambda, t)} \cdot \\ &\qquad \frac{\sum_{\lambda_2} \hat{I}_0(\lambda) r_c(\lambda) m(\lambda, \overrightarrow{x}) b_{DC}(\lambda, t)}{\sum_{\lambda_2} \hat{I}_0(\lambda) r_c(\lambda) m(\lambda, \overrightarrow{x}) b_{AC}(\lambda, t)} \\ &= \frac{AC_1/DC_1}{AC_2/DC_2} \end{aligned} \tag{2.9}$$

Where R_n is the RGB camera response in the n^{th} channel

The core contribution of this work is the many checks a controls that [68] introduce to ensure that a valid measurement of the AC is being made discarding estimates where there is too greater subject motion or specular reflection or low signal to noise ratio (SNR) of patient heart rate. While this dynamically re-evaluates the many regions of interest throughout the process, in the presence of significant motion it requires them to be re-initialised necessitating the use of a non-rigid tissue tracking technique such as that of [85]. The other consideration would be to pre-process with a suitable rapid reconstruction step in order to accurately normalise the DC estimate by distance from the camera and surface normal direction.

2.3 Discussion

General MSI devices currently available have the limitations of, complexity of design; requiring expert set up; high unit cost; a trade off between acquisition time and number of spectral bands imaged. There are, however, tangible uses cases for being able to perform MSI during surgical procedures including such as navigation,

tissue classification and analysis as well as for better fluorescence imaging.

Minimal waveband MSI devices have shown to be a viable trade off of spectral discrimination for elevated capture speed. Using these devices, it is possible to gather sufficient information to perform a specific computer vision task; specialised tissue analysis; or targeted fluorescence imaging. However, this reduces the versatility of the system, often limiting such imaging devices to a single imaging function as the set of wavebands used is often tune for one specific task. One interesting method that makes inroads to multifunction MSI is demonstrated by [63], combining surface reconstruction via structured light with a minimal 8 filter MS camera. While combining both MS illumination of tissue with MSI camera could enable concurrent MS imaging/analyses via a single device, it makes such a system arbitrarily more complex in construction and calibration, which would likely negatively impact the translation of such a system for clinical use.

MS reconstruction methods that try to estimate the full MS cube from one or more RGB cameras have the potential to provide a general and cheeper method for acquiring MS data. These techniques are still at a fairly early stage in development in the domain of MAS, and as such there still is a lack of high accuracy methods that are also high speed in computation. The computation speed barrier may not be as significant factor as it initially seems since the direction of compute hardware is currently making huge improvements in concurrent code execution, which plays well to the parallel nature of MS reconstruction algorithms. One of the potential pitfalls of MS estimation from sparse or low fidelity measurements is often the requirement for using priors on tissue spectral reflectance. The risk here is that a general reconstruction technique would only be able to be used in a instances where the patient pathology is congruent with the priors or learning data set, as such the careful selection of these is important.

From currently available techniques there appear to be two promising areas of development: either augmenting and compensating deficiencies in existing MSI techniques for more generally and better performance in surgery; or developing robust MS estimation methods that utilise existing non-specialised imaging hardware

commonly used in MAS. In Chapter 3, the first of these two approaches is investigated. Here the issues of large scale non-rigid motion blur in MSI of dynamic tissue is tackled, using tissue tracking to estimate scene blur. However, despite developing methods that would be capable of deconvolving individual channels of the MS data the practical applications remain limited by the capture speed of the MSI device, necessitating a minimum delay of several seconds before sufficient band limited data has been acquired to perform inference on. As such, the real-time applications of this approach are not viable unless the underlying MSI protocol was able to capture at a faster rate, although in this case the presence of motion blur in the scene would be greatly reduced. Thus for Chapter 4 through Chapter 6 the second approach is explored by investigating robust methods for estimation of MS from endoscopic RGB video data. The methods developed in these chapters make use of priors on the spectral reflectance of tissue to build a methods capable of estimating MSI data robustly enough to calculate tissue oxygenation and perfusion, at speeds fast enough to enable real time applications.

Chapter 3

Deblurring for Multispectral Endoscopic Images

When adapting imaging techniques for use in an MAS context there are a host translational issues that arise which are either negligible or easily mitigated in other use cases. One such issue that is a significant factor across many modalities of medical imaging is the highly dynamic component of imaging human tissue. Motion can originate from macroscopic patient movement; bodily functions such as breathing, peristalsis, and heartbeat; and in some cases from the clinician performing the imaging i.e. not holding the imaging device still. In addition the operational distance for MAS is often very small, causing even small motions to become large displacements in the camera's field of view, having the effect of amplifying the impact of these various sources of motion.

As discussed, MSI systems come in various forms, many of which are based on a single image channel laparoscope [86] and either rely on switching of the illumination source, or use an additional scope for visible white light visualisation. Devices that are LCTF based are of interest as they are capable of acquiring full frame image data with spatially resolved spectral response data of tissue and the subject matter. In human tissue, light in the VIS wavelength range is strongly absorbed, however, it is also this range that is of interest for measuring haemoglobin concentrations. Further the transmission properties of LCTF devices are rather poor: less tan 50% of the light is transmitted though even in optimal configurations. Depending on the

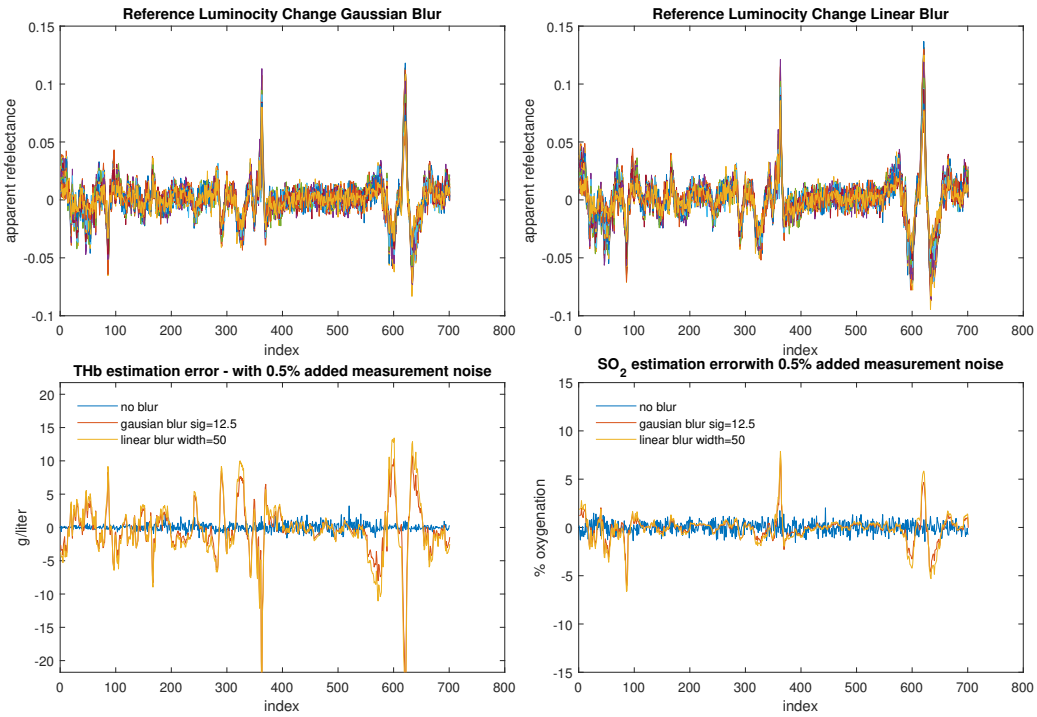

Figure 3.1: An illustration of the effects of estimating haemoglobin concentrations from blurred data. Using synthetic blurred data with either linear or gaussian blurring, it can be seen that the estimation THb from blurry data under estimates in areas of higher luminosity and vice versa, where as the saturation estimation shows the opposite correlation with less sensitivity.

working distance and the concentration of absorbers, acquisition of a single band image within an MS stack of images may take over hundred milliseconds, making its recording highly susceptible to motion blurring.

Blurring due to motion or other factors can have significant impact on the results of subsequent processing of acquired MS data. One such issues is large scene displacement between bands that can arises as a product of using a sequential band sweeping imaging protocol, as used by LCTF based devices. Registration is required to align the separate bands [39, 9, 85], and this is made more difficult in the presence of significant blurring as it obfuscates tissue surface detail used for the alignment. Another issue, one that this chapter aims to address, is that blurring in an image causes the apparent brightness of pixel data, corresponding to a specific location in an image, to be shifted toward the mean brightness of the clear image as a whole. As a result, one can no longer treat the MS data as a localised spectral

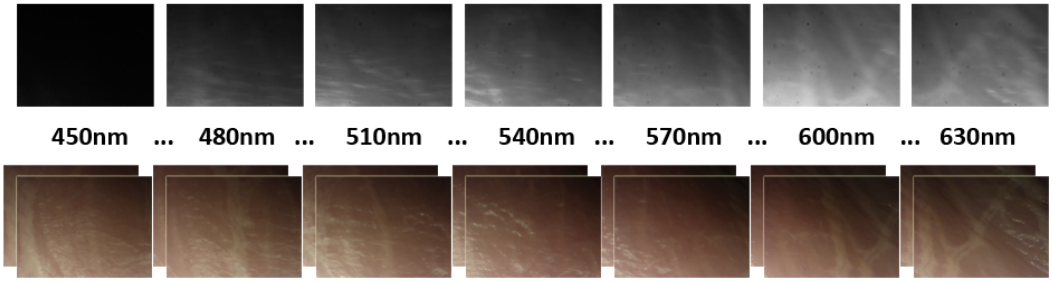

Figure 3.2: Example data for a simultaneous capture ofspectral sweep multispectral (top) stack, with corresponding stereo VIS RGB (bottom). The MSI data is visibly degraded by the tissue motion, an artifact barely perceptible in the RGB data.

response measurement as it is now an average over the point spread function (PSF) of the blur. This makes it more difficult to evaluate spatial variation in derived metrics such as estimating HbO and Hb from blurry MS. To demonstrate this one can take a scan-line of a multispectral data set and inspect the luminosity change at each point of the scan line after applying a linear blur in the scan line direction, Figure 3.1. If one were to then compare the estimation error between computing HbO and Hb from the clear and blurred MS data the results is naturally similarly effected. Interestingly as THb is computed as the sum of HbO and Hb it's error is negatively correlated with the luminosity change. The estimation error for SO_2 however, would be positively correlated with the luminosity change, and this is to be expected as SO_2 is computed as the fraction of HbO to THb. Illustrating as such, it is apparent that scene motion that can cause linear blur in long exposure imaging can have a pronounced effect on the accuracy of chromophore concentration estimates.

3.1 Problem formulation

The field of image deconvolution is very mature research area and typically divides protocols into blind, semi-blind and non-blind, corresponding to how much is known about the shape of the convolution to which image data were subjected from the perspective of the camera. The Richardson-Lucy (RL) algorithm [87, 88] is one of the most established approaches requiring a known PSF. Semi-blind methods incorporate priors [89] however often this makes optimisation non-convex. Blind

deblurring using natural image statistics is also possible [90], [91] but a major challenge for *in vivo* multispectral images is the lack of constraints to anatomical structural correspondence. Further reason as to why it may be inadvisable to apply some purely blind methods is that the statistics of images, such as total variation, acquired during MAS differs from natural images these techniques were designed for such as photography or astronomy. Stochastic deconvolution [92] allows for better correspondence utilising a known spatially varying point spread function (SVPSF), as well as being able to include arbitrary regularisation without the need for complex optimisation, at the cost of a prohibitively slower execution time.

Multi-camera methods are able recover the information from one or more cameras and use it to formulate the SVPSF in another [93]. This can be achieved using or optical flow to estimate a discrete SVPSF to use as a seed for iterative optimisation [5]. Existing algorithms have difficulty with scenes that are undergoing non-rigid deformations as in MAS and which exhibit complex reflectance functions. However, there are methods that have been explicitly designed for tracking in MAS[85] which can be adapted for this use case, and this method can generalise to allow for dense, non-rigid scene deformation. Within this context, this chapter will explore semi-blind deconvolution, using synthetic data, as well as, data acquired using a rigid endoscopic system that comprises two VIS RGB channels and a third LCTF driven MS channel. This integrated device's advantage for clinical use is tht it can deliver a variety of optical imaging through a single trocar. Additionally the LCTF is capable of achieving a high spectral resolution ($<$10nm) and can be electronically-controlled, allowing on-demand access to wavelengths of interest. The use of this device enables the estimation of a dense SVPSF for the scene which can be then used in the deconvolution of the MSI data. By combining both feature and intensity stereo tracking the method of [85] propagates a mesh over tissue and is robust to many issues facing tissue tracking *in-vivo* such as specularities and sparse feature identification. The resultant mesh deformations can be then used to approximate scene motion and ultimately estimate a SVPSF in another camera view.

3.2 Extending Richardson-Lucy Deconvolution

In the following section 3.2.1 the procedure for modifying Richardson-Lucy deconvolution for use with SVPSF. The choice of a non-frequency domain method to develop from is due to the difficulties presented in parametrising a SVPSF for use with a frequency domain method, in particular when the SVPSF arises in the most part from motion blurring. In section 3.2.2 the methods used to validate the proposed technique are outlined comprising both synthetic and *ex-vivo* experimental configurations.

3.2.1 Deconvolution Methods

Semi-blind deconvolution is performed by using the stereo RGB images to estimate a deformation field in the image plane of the multispectral camera. The deconvolution method is an extension of the RL algorithm which can be expressed using the convolution operator [94, 95] as follows:

$$I^{(n+1)} = \left\{ \left\{ \frac{B}{I^{(n)} \otimes k} \right\} \otimes \hat{k} \right\} I^{(n)} \tag{3.1}$$

where n is the number of the iteration and B the blurred image and such that for a 2D PSF k with an index space $\Omega \subseteq \mathbb{R}^{2+}$, the inverse point spread function $\hat{k}$ is

$$\hat{k}(\mathbf{x}) = k(\max(\Omega) - \mathbf{x}), \mathbf{x} \in \Omega.$$

$\mathbf{x}$ being the index into the PSF, so if k has dimensions 2x2 then then $\hat{k}$ is just k but indexed as $k(x_0 - 1, x_1 - 1)$.

The formulation (3.1) can be extended to incorporate an SVPSF by creating an individual k_i for every pixel in the image. However, it has been observed that while it is possible to generate per-pixel kernels from deformation fields, standard RL deconvolution will become unstable for SVPSFs as seen in Fig. 3.5 [6][96]. This stems from the non symmetrical nature of the *forward* blurring kernel compared to the *inverse* error collection kernel. To rectify this, a generalisation is made of the RL algorithm, and propose a new non-symmetric blur method that directly utilises

the source deformation field, via successive re-sampling.

The generalised Richardson-Lucy algorithm can be expressed as the error metric and the iteration update, with the blur model independent of the algorithm. This is observed from (3.1) which articulates the blur model as the convolution operator, equivalent to original summation based expression of the RL algorithm.

```
def blurImage(image, inverse=False):
    # the optical flow defines the spatially varying displacement
    # to invert the blur just negate the optical flow vector field
    if inverse:
        deformation = optical_flow * -1
    else:
        deformation = optical_flow

    # the output blurred image is accumulated by integrating w.r.t
    time
    #  t0  = no displacement the initial positsion
    #  t1  = the fully displaced position
    image_blurred = Zeros()
    for time in range(0, 1, delta_t):
        displacment    = deformation.at(time)
        image_blurred += sample(image, offsets=displacment)

    return image_blurred

with observed_blurred_image as B:
  I_est = B
  for i in range(0, i_max):
    # calculate an error image from the current estimate
    B_est = blurImage(I_est)

    # element wise division
    err   = B / B_est

    # accumulate the error to contributor locations
    err_b = blurImage(err, inverse=True)

    # element wise multiplication
    I_est = err_b * I_est
```

Listing 3.1: Generalised Richardson Lucy Algorithm: psudo code.

This generalisation is also similar to to way in which [96] incorporate their accumulative blur model with the RL framework. However the crucial distinction is that [96] make the assumption that the scene is under piecewise projective deformation and the method presented here does not, allowing a continuous per-pixel treatment of the deformation instead of dividing the scene into regions of discrete deformations.

Re-sampling blur evolves a sampling grid by repeated re-sampling the grid

locations deformed by a spatially varying deformation field. The sampling grid is then used to accumulate samples from a source image to form the blurred image. The intuition of this method follows from the PSF generation method of [6], where blur kernels are rendered by drawing line segments of a motion path weighted by the relative time duration the segment corresponds to. Using the same approach, a sampling point s_n is propagated through a deformation field D, and instead of drawing the trace of this sampling point into a kernel it is used it to accumulate successive samples from the source image. So to create the blurred image B from input image I, the value at $b_i \in B$ is calculate as:

$$b_i = \Delta \sum_{n=0}^{N} I(s_n) \tag{3.2}$$

where $\Delta = \frac{1}{N}$, and s_n is the sampling position calculated recursively as

$$s_{n+1} = s_n + \Delta t D(s_n) \tag{3.3}$$

with s_0 initialised as the coordinate location of b_i, and t is the length of exposure. B-Spline sampling is used to sample the deformation $D(s_n)$ and input $I(s_n)$ for each location s_n. In order to integrate this blur model with the generalised RL algorithm an inverse blurring function can elegantly be expressed by negating the deformation fields. To avoid zigzag drift, as shown in Fig. 3.5 as an artefact of rotational blur, the update for s_n is modified to use a weighted filter over a neighbourhood centred at s_n. So for filter f and neighbourhood offsets d, indexed over Ω, the update equation for s_n becomes

$$s_{n+1} = s_n + \sum_{x}^{\Omega} f(x) \Delta t D(s_n + d(x)) \tag{3.4}$$

The method described here is illustrated in Listing 3.1 as pseudo code.

3.2.2 Validation methods

Synthetic data are generated from blurring ground truth images for known deformation fields so as to be able to assess the effectiveness of this method against standard

RL under perfect conditions. Various deformation fields could then be applied to the images to explore the stability of the algorithm under different simulated conditions. Figure 3.3 illustrates an example of a respiratory like blur to create a synthetic blurred image.

The experimental set-up used for *ex-vivo* validation is shown in Figure 3.4. The three-channel (trinocular) rigid endoscope delivers colour stereo images at a resolution of 1024×768 pixels using two IDS Imaging, uEye 2230-C cameras. The central channel of the scope is a LCTF (Varispec, CRI, Inc) attached to a monochrome camera (Thorlabs DCU 223M). The LCTF has a spectral range running from 400-720 nm with a resolution of 10 nm. The multispectral camera was synchronised with the LCTF so that a given wavelength range, decomposed into contiguous non-overlapping 10 nm bands, is captured such that each image corresponds to a single band. A 25.4 mm diameter, 75 mm focal length was used to provide magnification in the central channel in order observe the same site as the white light channels. For the *ex-vivo* experiment, the integration time and gain were set to 1000ms and 25 respectively for the multispectral camera, and 90 ms and 20 for the stereo cameras[1]. *Ex vivo* experiments were carried out using a translation stage, to simulate spatially varying motion. By positioning the plane of translation non-parallel to the imaging plane of the camera a large spatially varying projective motion could be generated. Example data from this type of capture as show in Figure 3.2 illustrates how the multispectral capture exhibits significant motion blur when compared to the image data from the VIS RGB channels. The multispectral stacks were registered and points of interest identified, before surrounding patches then being extracted. The multispectral stacks for each patch were processed, extracting the reflectance spectra of the tissue of interest, in order to calculate haemoglobin concentration [9].

3.3 Results

For synthetic blur using scene displacement typical of respiratory motion, as shown in Figure 3.3, it is plainly visible that this method contains fewer artefacts arising

[1]Further details of scope design and configuration are found in [9].

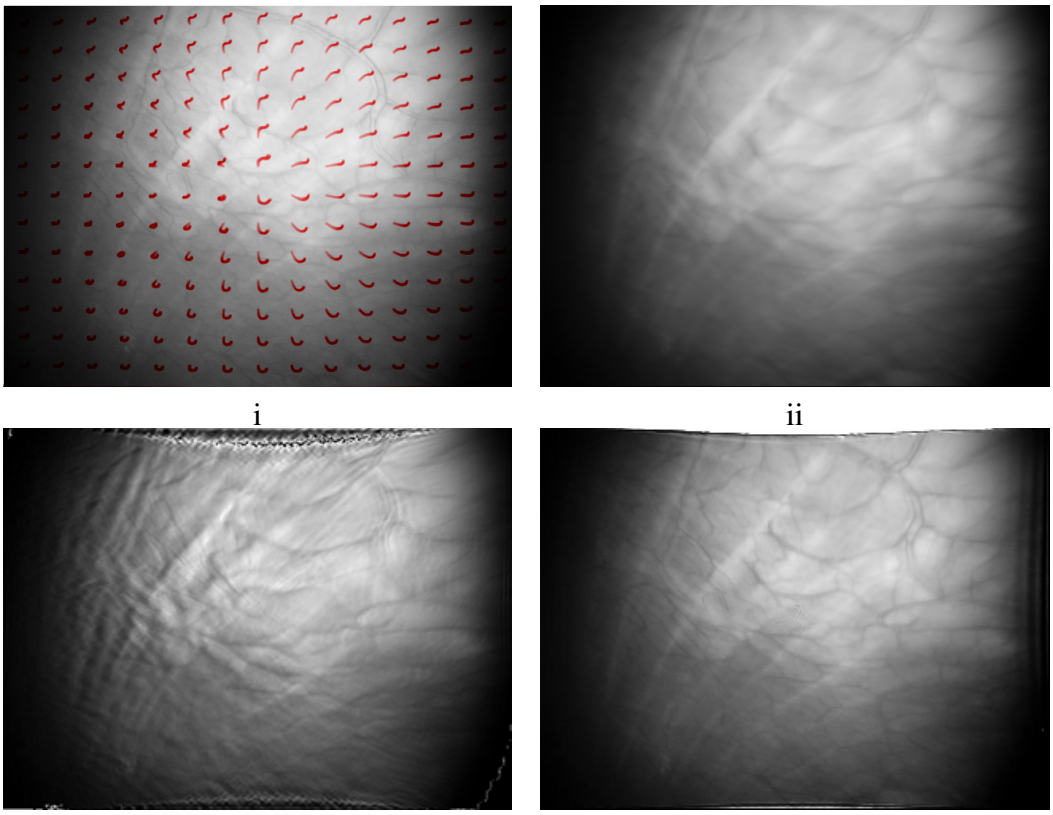

Figure 3.3: i) Kernels sampled form an SVPSF overlaid on a clear image. ii) The clear image blurred by the SVPSF iii) Deblurring using typical [5] spatially varying extension of the Richardson Lucy algorithm. iv) Deblurring using our proposed adapted method.

Figure 3.4: Left: The front of the laparoscope and the three camera multiplex; Right the experimental configuration: the scope is positioned perpendicular to the tissue surface as mounted on a translation stage. The axis of motion of the translation stage is normal to the view direction.

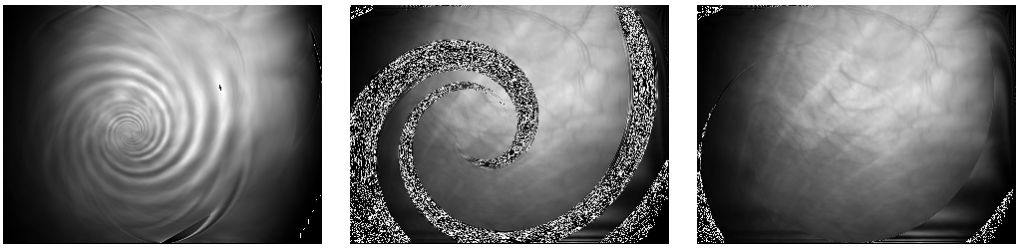

Figure 3.5: Comparative stability of different deblurring algorithms on a particularly problematic synthetic spatially varying blur of large rotation. From right to left; deblurring using standard RL deconvolution with pixel wise kernels[6] after only 10 iterations; deconvolution using RL and re-sampled accumulation at 500 iterations; deconvolution using RL and 3x3 filtered re-sampled accumulation at 500 iterations.

from incorrect error accumulation. This method was also stress tested on particularly taxing blur such as large rotations for which the performance of standard RL is very unstable. Notably this method can be seen to be robust even for many iterations Figure 3.5. The rotational experiment also clearly demonstrated how improving the tracking of the sampling grid translates directly to an improvement in the stability of this algorithm.

Making comparisons in the frequency domain (Figure 3.6) it is seen that these deblurred results have a similar frequency profile to that of the ground truth, where as the blurred data has significantly more weighting to the lower frequencies as expected. Furthermore, by looking at how the SNR for the frequency domain varies, over scale space, it can be seen how this method improves on the standard RL, and blurred results, by recovering more fine image structure at the trade off of less large scale accuracy. Frequency domain comparisons were chosen because they would be more sensitive to the errors often created by RL deconvolution such as ringing artefacts introduced during deblurring.

Figure 3.7 shows two selected feature points observed in the data and compares the results of deblurring and deblurring with denoising (pre-filtering bilaterally). There are visible improvements to localisation of structure in the total haemoglobin analysis, for the deblurred and in particular where the data is also initially de-noised.

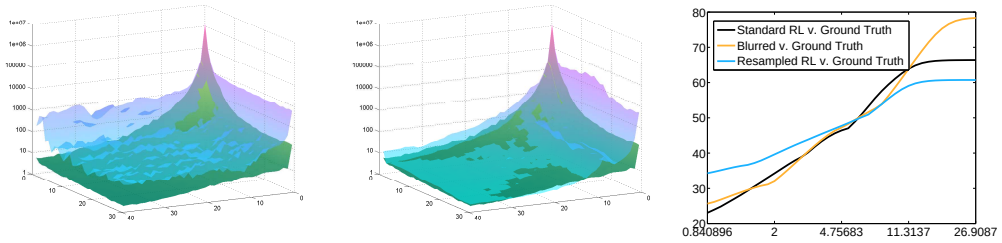

a) b) c)

Figure 3.6: A frequency domain comparison of, a) the blurred image and b) deblurred (in blue) against the ground truth (in green). c) Comparison of SNR of in the frequency domain over scale space.

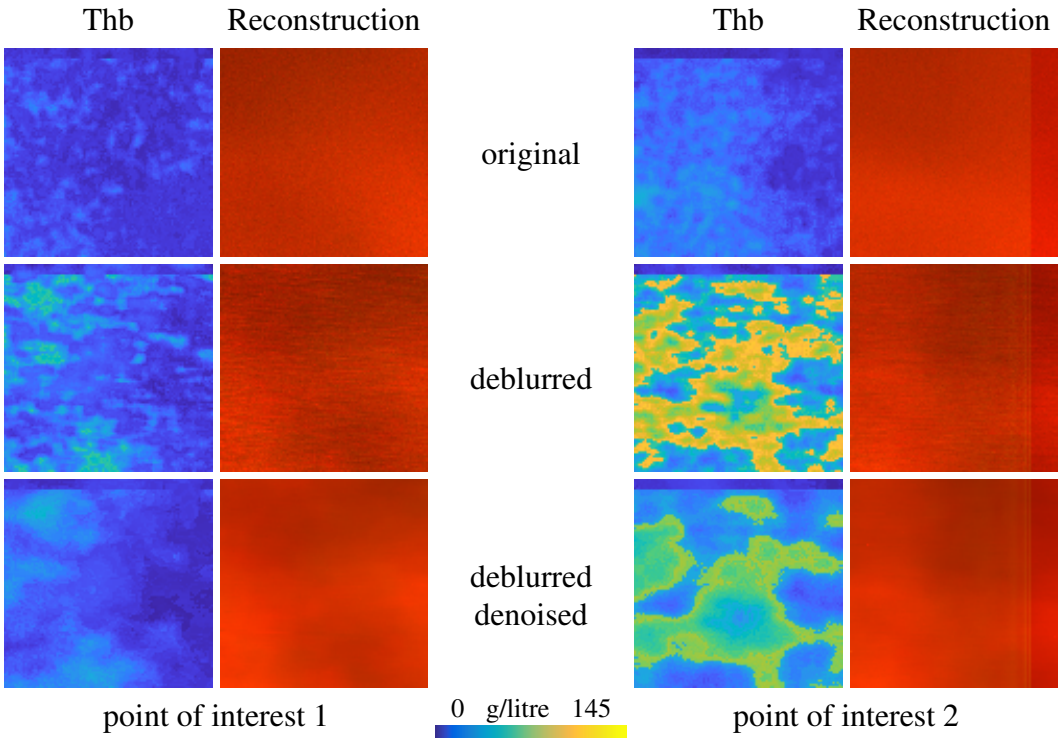

Figure 3.7: Multispectral regression results for two points of interest, in the *ex-vivo* data set (upper of Figure 3.2), showing total haemoglobin (Thb) and colour reconstruction. The reconstruction are predominantly red because most of the shorter wavelength bands are not available.

3.4 Discussion

Incorporating this spatially varying blur model with the generalised RL algorithm provides an accurate and computationally efficient deblurring of multispectral images even with non-uniform blur. Bilateral filtering prior to deconvolution has a positive effect down-stream on the results of structural cohesion visible in total haemoglobin analysis. The marked improvement from denoising the data prior to deblurring is likely related to the low transmission of the LCTF at shorter VIS wavelengths causing a lower SNR in this range; these wavelengths are more useful at discriminating HbO and Hb. Characteristic artefacts of the RL approach appear when very hard edges are visible in the image, such as near surface vessels, but it may possible to dampen the expression of these artefacts [97]. The atomic design of this algorithm allows for parallelisation and, while for our *ex vivo* data with 100 pixels PSFs deblurring times of around five minutes, with a well optimised implementation near real-time performance is possible. Memory usage is efficient since deformation is stored as a field instead of per-pixel kernels. This is important because uncompressed per-pixel floating point 40×40 kernels would require over 6GB for all kernels for a one mega-pixel image.

While this approach makes headway towards improving the results of processing the MSI data via deblurring, the net result is still some way from being able to provide real-time imaging of haemoglobin concentrations to a clinician. Even with further optimisation to the deblurring algorithm and implementation, the limiting factor of the MS camera's acquisition rate would remain. As such, to provide potential real-time imaging of haemodynamic concentrations within tissue for instance to track ischaemia in organs during a MAS procedure, alternative acquisition methods may be more appropriate. To this end, the next chapter will look at ways of inferring the MS data from wide band RGB measurements in the VIS spectrum to achieve potential real-time *in-vivo* imaging of haemoglobin concentrations.

Chapter 4

Smoothness Constrained Estimation of Multispectral Data

As discussed in Chapter 2 there are many different imaging modalities available to clinicians for the purposes of guidance navigation and assessment, MSI having the potential to deliver a wide range of beneficial insights. In Chapter 3 methods were proposed to reduce spatially varying motion blur induced artefacts as seen in some sequences of *in-vivo* MSI data. A related challenge is that many existing techniques fall short of being able to deliver rich real-time spectral imaging, limiting their ability to be deployed clinically where it may be desirable to perform rapid assessment of tissue; or track dynamic changes in haemoglobin concentrations [60]. A cross-section of potential rapid MSI solutions are inspected in Chapter 2 however many of these demandthe use of non-standard imaging sensors, scanning mechanisms or specialist, tunable illumination sources. Many of these systems are not concurrent compatible with standard RGB imaging of the tissue and therefore require and additional, dedicated entry port into the abdomen (for MAS). Since the primary form of optical imaging used during MAS procedures comprises VIS spectrum RGB imaging, it makes this modality a natural avenue to explore to investigate how it can be adapted for use in estimating haemoglobin concentrations.

Notable advantages of co-opting the RGB data are: first it removes the acquisition bottleneck, potentially enabling video rate imaging; and second it would remove then need to modify a current surgical imaging setup. However, the chal-

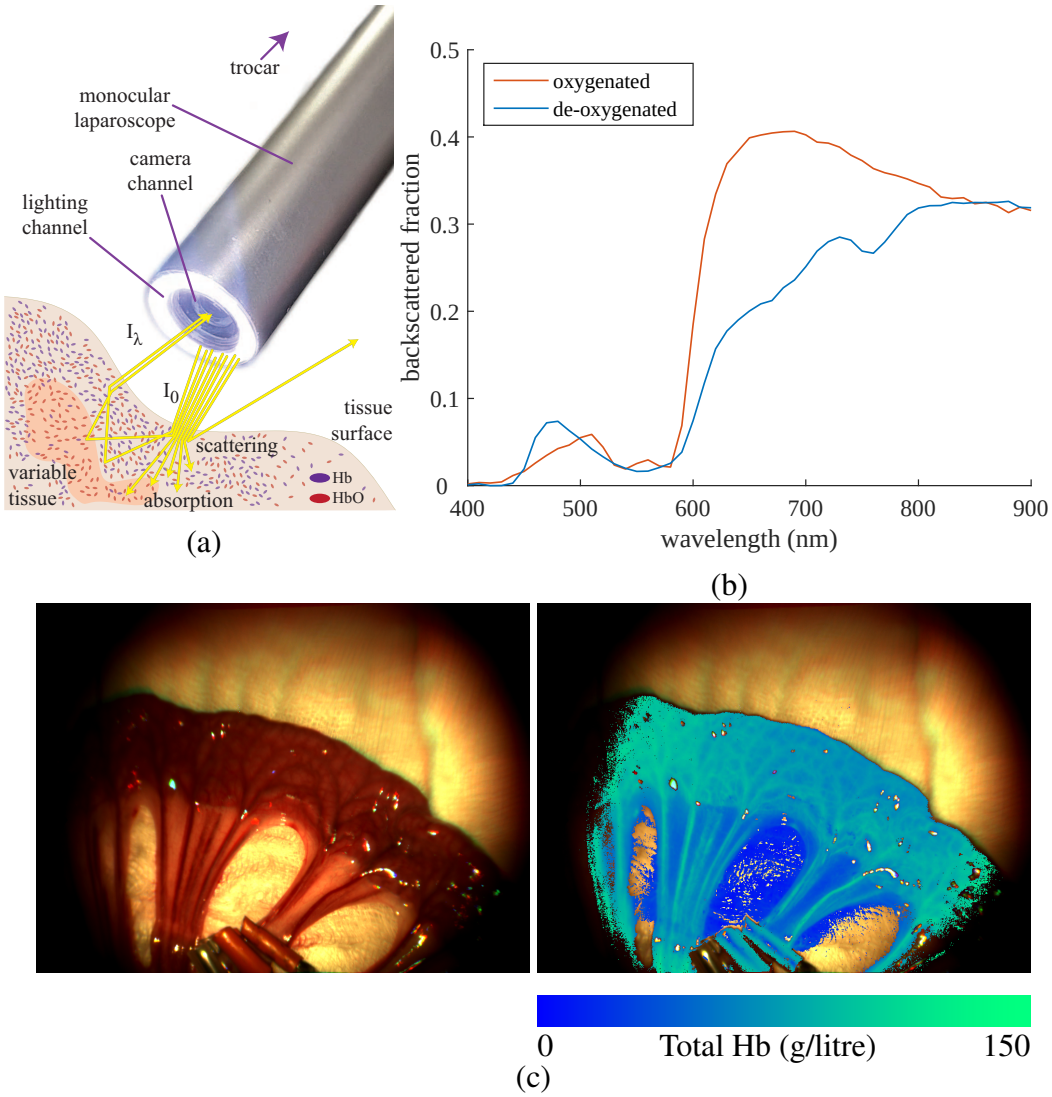

Figure 4.1: (a) The interaction between light and tissue comprises both scattering and absorption and the inhomogeneity in the concentration of chromophores results in a spatially varying distribution of detected light. The effective attenuation of the measured light in an endoscopic image is a combination of these effects. (b) The backscattered fraction of light for oxygenated and de-oxygenated blood under uniform illumination generated by MC simulation [7]. This provides an aggregate attenuation coefficient ($-log(\cdot)$) that combines the attenuation due to absorption and scattering as well as the mean travel distance within tissue. (c) The total haemoglobin (THb) is a measure formed by the sum of concentrations of these two chromophores and saturation is the ratio of oxygenated blood to total haemoglobin. These can be superimposed on endoscopic images to provide the surgeon with functional information during an intervention.

lenges with RGB imaging are that the bands for each channel are wide (approximately 100nm full-width at half maximum) and have a high degree of overlap, as well as inconsistencies between devices in terms of each channels specific spectral response. The low spectral resolution of each channel, as is shown later, causes uncertainty firstly in estimating the underling spectral reflectance of the imaged tissue, and in turn deriving quantities such as THb and SO_2. In this chapter a method for estimating THb and SO_2 is developed that uses the RGB sensors from a stereo laparoscope. The calibrated sensors' response curves can be considered as defining the mapping of the latent MS data onto RGB space. This mapping can be inverted using a Tikhonov regularisation scheme to preserve smoothness, and in this chapter it is illustrated how this is a reasonable constraint to use. Performing such an inversion reduces the haemoglobin estimation problem into a two step process: the first step having a closed form solution, enabling rapid processing of full frame stereo data; and the second step being the linear fitting of haemoglobin absorption characteristics to the MS data estimated in the first step.

4.1 Estimation of Latent Multispectral Data

This section describes the theory on the interaction of visible light with tissue and how RGB image data can be used to invert the model to estimate the latent MSI data used to derive underlying tissue properties (Section 4.1.1). A computational experiment, using a numerical MSI simulation, to generate synthetic test data is then described (Section 4.1.3). Finally, a porcine surgical experiment is outlined, which was used to obtain *in-vivo* data (Section 4.1.4).

Light arriving at the endoscopic camera (Figure 4.1) is assumed to have been back scattered and attenuated by passing through tissue containing spatially varying concentrations of chromophores. By inspecting the light-tissue attenuation model it is possible to make smoothness generalisations on the spectrum shape of MS data imaged leaving tissue. This generalisation enables the constrained inversion of RGB imaging data to estimate the MSI data.

4.1.1 Light-Tissue Attenuation Model

The dominant chromophores in tissue are assumed to be the two oxygenation states that haemoglobin can take, HbO and Hb ($\boldsymbol{hb}$), corresponding to oxygenated and de-oxygenated respectively. At a given x, y image location ($\boldsymbol{x}$) the concentration of chromophores $\boldsymbol{hb}(\boldsymbol{x})$, within the corresponding tissue location, can be estimated using the modified Beer-Lambert relationship [98]. This defines the model for how the concentration of all chromophores present in a sample (the tissue corresponding to image position $\boldsymbol{x}$) attenuate the incident light $\boldsymbol{I}_{0,\boldsymbol{\lambda}}(\boldsymbol{x})$ for a set of illumination wavelengths $\boldsymbol{\lambda}$. The resulting intensity of light leaving the scene at that position, $\boldsymbol{I}_{\boldsymbol{\lambda}}(\boldsymbol{x})$ can then be calculated as:

$$\boldsymbol{I}_{\boldsymbol{\lambda}}(\boldsymbol{x}) = \boldsymbol{I}_{0,\boldsymbol{\lambda}}(\boldsymbol{x}) e^{\left(-\boldsymbol{G}(\boldsymbol{x}) - \boldsymbol{hb}(\boldsymbol{x}) \boldsymbol{\mu}_{\boldsymbol{hb},\lambda} \boldsymbol{\xi}_{\boldsymbol{hb}}\right)} \tag{4.1}$$

From the original Beer-Lambert equation, the attenuation can be expressed as due to absorption $\boldsymbol{\xi}_{\boldsymbol{hb},\boldsymbol{\lambda}}$ over the wavelength range $\boldsymbol{\lambda}$ for each chromophore. The modified form of the Beer-Lambert relationship introduces two additions. The first new variable is the mean path length of light travelling through the tissue $\boldsymbol{\mu}_{\boldsymbol{hb},\lambda}$. This is assumed to be longer than the shortest distance between the source and detector because each photon will have been scattered along the way. The assumption is appropriate for endoscopic systems where the illumination and the camera are modelled at the tip of the scope and also since each photon will take a slightly different path the value of $\boldsymbol{\mu}$ represents the mean path length. The second extension is to include a geometric factor $\boldsymbol{G}(\boldsymbol{x})$ to describe the effect of loss due to scattering such that the resultant photon path never intersects the detector. These are reduced to a single term representation, but this geometry factor can also be expanded as a constant and linear term with the latter being dependant on wavelength [99]. Rearranging the terms of (4.1) it is possible to formulate a linear relationship:

$$\mathrm{A} = -\log\left(\frac{\boldsymbol{I}_{\lambda}(\boldsymbol{x})}{\boldsymbol{I}_{0,\lambda}(\boldsymbol{x})}\right) = \boldsymbol{G}(\boldsymbol{x}) + \boldsymbol{hb}(\boldsymbol{x}) \boldsymbol{\mu}_{\boldsymbol{hb},\lambda} \boldsymbol{\xi}_{\boldsymbol{hb}}. \tag{4.2}$$

The negative log ratio of the incident to exit intensity is the total attenuation,

A. When the wavelength range is narrower, for example 500nm to 600nm, the constant offset plus linear scattering term can be used for $\boldsymbol{G}(\boldsymbol{x})$ [99]. However, when using an RGB sensor, over the wide spectral range of each measurement, the scattering parameters for whole blood containing HbO and Hb are not completely linear. This observation is illustrated in the scattering and attenuation coefficients data as compiled by [7], especially at shorter wavelengths, and varies slightly between each chromophore. For this formulation, instead of analytically modelling the wavelength dependent absorption and scattering attenuation for each chromophore, an aggregate attenuation factor $\hat{\hat{\boldsymbol{\xi}}}_{\boldsymbol{hb}}$ is empirically derived by using MC simulation [60], [100]. Given the incident illumination is known, the chromophore concentration parameters $\boldsymbol{hb}(\boldsymbol{x})$ on the right can be found by making many spectrally distinct measurements of the tissue and solving the linear equation:

$$\mathrm{A} = \hat{\hat{\boldsymbol{\xi}}}_{\boldsymbol{hb}} \boldsymbol{hb}(\boldsymbol{x}) \tag{4.3}$$

Approaches where many MS narrow bands can be obtained will typically use a form of (4.2) or (4.3) either with or without the scattering component, however, for the estimation method developed here, the above light-tissue interaction relationship is used. One can solve Equation 4.3 using a non-negative least squares solver: the fast non-negative least squares method of [101] is utilised here to constrain the estimation from MS data. The need to constrain the estimation to be non-negative is to restrict the fitting to physically plausible values as it is possible otherwise to produce negative values for the concentration when solving Equation 4.3 via an unconstrained fitting. Ultimately, however, this type of chromophore concentration estimation is only applicable when the conditions of the Beer-Lambert equation are satisfied; however, the wide band color channel measurement from an RGB camera would violate one of these conditions.

An indirect approach is needed when the measurement is actually from RGB cameras comprising channels that each having significant spectral overlap and spanning a wide spectrum. In order to preserve the conditions of the Beer-Lambert equation, to solve this directly from stereo RGB measurements it would be necessary to

pose the solution as the minimisation of 6 variably weighted sums of exponential terms. For example an individual channel measurement would be:

$$I_N = C_{N,\lambda_1} I_{0,\lambda_1} \exp^{\xi_{\lambda_1} \boldsymbol{hb}} + C_{N,\lambda_2} I_{0,\lambda_2} \exp^{\xi_{\lambda_2} \boldsymbol{hb}} + \ldots C_{N,\lambda_n} I_{0,\lambda_n} \exp^{\xi_{\lambda_n} \boldsymbol{hb}} \tag{4.4}$$

Here I_N is the measured light intensity at the N^{th} channel and C_{N,λ_n} is the weighting corresponding to wavelength λ_n for that channel i.e. a discretisation of that channels response curve. Solving Equation 4.4 for the concentration $\boldsymbol{hb}$ directly would not have a concise closed form linear solution, so instead the estimation of $\boldsymbol{hb}$ is divided into two steps. In the first step the MS intensity I_{λ_n} is estimated for each wavelength λ by solving across all N channels

$$\begin{pmatrix} I_1 \\ \vdots \\ I_N \end{pmatrix} = \begin{pmatrix} C_{1,\lambda_1}\tilde{I}_{\lambda_1} + C_{1,\lambda_2}\tilde{I}_{\lambda_2} + \ldots C_{1,\lambda_n}\tilde{I}_{\lambda_n} \\ \vdots \\ C_{N,\lambda_1}\tilde{I}_{\lambda_1} + C_{N,\lambda_2}\tilde{I}_{\lambda_2} + \ldots C_{N,\lambda_n}\tilde{I}_{\lambda_n} \end{pmatrix} \tag{4.5}$$

from this estimate I_λ, the concentrations can be estimated as normally would be done using Beer-Lambert methods.

The naïve approach of estimating the MS data from the measured RGB image data I_{RGBs} is to invert the spectral response curve of the camera C solving the linear system though a least squares minimization such as:

$$\hat{I}_\lambda = \arg\min_{I_\lambda} \|CI_\lambda - I_{RGBs}\|^2 \tag{4.6}$$

where $||\cdot||$ is the L^2 norm. This yields a poor estimation of the true MS data due to the problem being vastly under constrained. I_λ can have an order of magnitude more entries then I_{RGBs}, thus the naïve solution, while in sense correct, is often a metamerism of the true MS data for the given C, as shown in Figure 4.2. To constrain the estimation a prior (Γ) on I_λ is imposed: this is done using a Tikhonov regularisation,

$$\hat{I}_\lambda = \arg\min_{I_\lambda} \|CI_\lambda - I_{RGBs}\|^2 + \|\gamma\Gamma I_\lambda\|^2 \tag{4.7}$$

To prevent the prior dominating the term $\gamma = 0.01$ is used, this value of γ was determined as effective experimentally via linear search on synthetic data sets. As γ affects the expression of the prior it is important that its value is not too small otherwise when solving Equation 5.9 the left hand side will be ill conditioned. Similarly, too great a value will cause over fitting to the prior estimation ignoring the RGB measurements. Equation 4.7 is typically solved implicitly as:

$$I_\lambda = (C^T C + \Gamma^T \Gamma)^{-1} C^T I_{RGBs} \tag{4.8}$$

where Γ is can be constructed as the identity matrix to minimise the overall magnitude of I_λ values. However, for this use case, a Laplacian matrix is used for Γ to penalise non-smooth I_λ, the Laplacian matrix is formed of ones on the leading diagonal and negative half on the first super and sub diagonals. This choice of Γ is made because the MS data is expected to be similar to that predicted by the Beer-Lambert relationship, which is mostly smooth across the visible wavelength range for $\boldsymbol{\xi}$ comprising attenuation due to oxy and de-oxy haemoglobin.

This method can also be applied to monocular imaging by reducing the number of columns in C and the length of I_{RGBs}. Doing so allows concentration estimation to either happen jointly, utilising data from 2 or more cameras, or independently for each camera.

4.1.2 RGB Camera Spectral Calibration

This method is dependant on having access to an accurate spectral response calibration for the camera sensor(s); protocols do exist for capturing very high quality calibration with negligible error using a monochromatic light source [102]. For this work estimation of the spectral response curve calibration measurement was performed using the method described in [9]. To measure the spectral sensitivity of the spectral response in each channel an LCTF was positioned between the camera and a uniform white reflectance standard. The LCTF was then used to sweep across the VIS spectrum in 10nm tuning intervals accumulating a measure of the spectral response, intensity recorded by the sensor, for each channel at each given band centre.

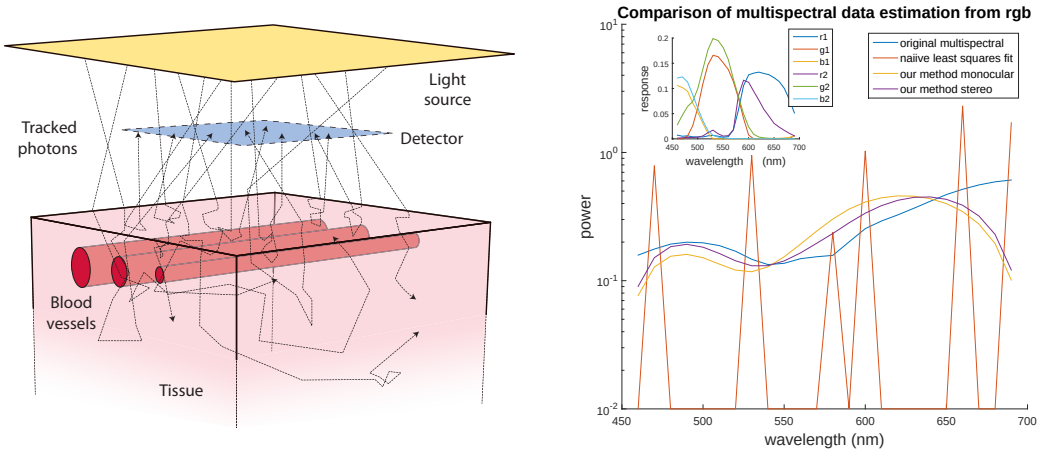

Figure 4.2: L) The synthetic model used for MC simulation of test data comprising variable diameter blood vessels embedded within soft tissue. R) Three estimates of the MS data from RGB. Stereo sensor response inset.

This was used to capture an ideal scenario calibration of the RGB camera response, however, such lengthy and refined calibration may not always be available. An alternative calibration method is performed via imaging coloured patches of known reflectance [103]; although this type of calibration has the limitation of being less accurate and at high sensor-noise levels. The result of multiple patch measurements can then be combined via linear fitting to enabled accurate estimation of the RGB camera sensor response. In the next chapter the implications of the types of issues arising from mis-calibrated RGB sensor response are explored further (section 5.2.2).

4.1.3 Synthetic Evaluation

One of the difficulties in validating technique that estimate full frame HbO and Hb concentrations is that it can be difficult to source high quality ground truth data to test against.

For simulating the interaction of light and tissue, monte carlo for multi-layered media (MCML) simulation [104] has been an established method for simulating light transport through layered turbid media. The mesh-based monte carlo (MMC) simulation framework of [105] is used here as it allow as convenient means of simulating 3D structures, such as blood vessels, as well as refraction and reflection across

boundaries of different refractive indices. MSI data are simulated using MMC and wavelength-resolved physiologically realistic, values of tissue optical absorption and scattering. These data are subsequently filtered by applying known RGB camera spectral responses to create RGB images. The digital phantom, shown in Figure 4.2, comprised a homogeneous block of soft tissue, bowel submucosa, with three superficial vessels containing either oxygenated or de-oxygenated blood. The three blood vessels had different uniform diameters of 2 mm, 1 mm and 0.5 mm and the top edges of each vessel were at 0.1 mm below the surface of the tissue. This depth was chosen to balance the obfuscating effect of the surrounding soft-tissue volume while maintaining reasonable detectability of the vessels. While this may seem shallow, the high degree of scattering makes it increasingly difficult to detect vessels embedded more than 2 mm beneath the surface. However, this is typical for MSI which is known to be only capable of measuring a superficial layer of tissue. The optical characteristics of blood and colonic sub-mucosa were compiled from [7] and [106] respectively. The mesh was made denser near the blood vessel boundaries for smoother approximation of the curved vessel walls.

During MMC simulation, photons were generated at intervals of 10nm across the visible range between 400 nm to 900 nm. To detect the backscattered light, the final direction of travel for each photon was recorded leaving the bounds of the meshed region. Photons that did not exit through the side of the mesh that was illuminated were discarded as were photons leaving at angles to the surface too oblique to be detected by a detector placed at 10 mm above the illuminated surface. The simulated detector was placed close to the tissue surface because the light exiting the tissue is highly diffuse and placing the detector further from the tissue would have an exponential penalty in simulation time.

Image data for each wavelength band were then generated at the detector location by photo counting. A total of 2×10^9 photons were generated uniformly distributed across the illumination surface with random initialisations. The number of photons was chosen by following MMC examples and documentation, as well as from the practical constraint of the number of photons that could be simulated in

5 weeks of simulation time on available hardware. To simulate MSI data, photons were first filtered into spectral bands and then used the incident angle at the detector to further filter photons according to the camera model. For computational reasons, the detector in the simulation environment had to be close to the tissue surface, so an orthographic camera model was used. The orthographic camera model also avoids angular bias found in projective cameras, where light at the edge of the image exits the tissue at a much shallower angle than that at the centre. simulated RGB images were generated by filtering the MS data with the response curves corresponding to RGB cameras, a stereo response curve is inset in Figure 4.2. Noise was added to the MS data by adding zero-mean normally-distributed vales to each channel, however for the RGB data, noise was generated correlated based on the response curve of the camera.

4.1.4 In-vivo Validation

Multispectral data from a porcine study [43][1] were used, in which MS imaging was performed on tissue undergoing periods of artificially induced restricted blood flow [43]. The MS imaging hardware is of the previously described LCTF MS imaging design. It comprised a 30° (Karl STORZ); a LCTF (Varispec; CRi); 50mm focal length achromatic imaging lens (Thorlabs) and monochrome camera (DCU223M; Thorlabs Ltd., UK). Illumination was from a Xenon surgical light source (Xenon 300; Karl STORZ). The camera and laparoscope were held by adjustable arms whose pose could be locked in place to avoid camera shake. MS data M were captured serially in non-overlapping bands 10nm wide from 500nm to 620nm, the integration time at the sensor was 200ms per band limited capture, requiring in a total of approximately 7 seconds for a full capture, after including filter switching and data readout times (specific and further details may be found in paper related to the original acquisition [43]). To compensate for the tissue motion during from the long capture times for each complete MS capture, all band limited image were registered after acquisition [39].

[1] The original procedure was undertaken with the assistance of Northwick Park Institute for Medical Research. The surgical procedure was conducted under UK Home Office personal animal licence (PIL) No. 70/24843 and project licence (PPL) No. 8012639.

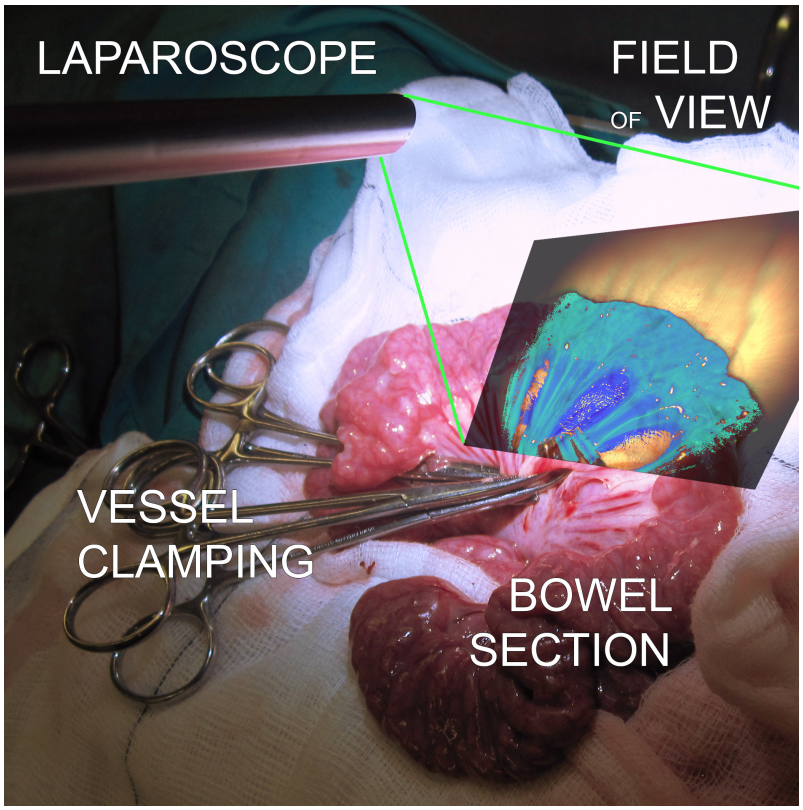

a)

b)

Figure 4.3: a) *In vivo* experimental configuration for imaging a section of porcine bowel, overlaid is a reconstructed synthetic RGB image (from captured MSI stack and RGB response) with illustrative heat map of THb concentrations. b) A sample MS stack captured by the device imaging at 490nm to 710nm at 20nm increments, from left to right top to bottom. Full capture stack includes 10nm increments and ranges form 450nm to 720nm.

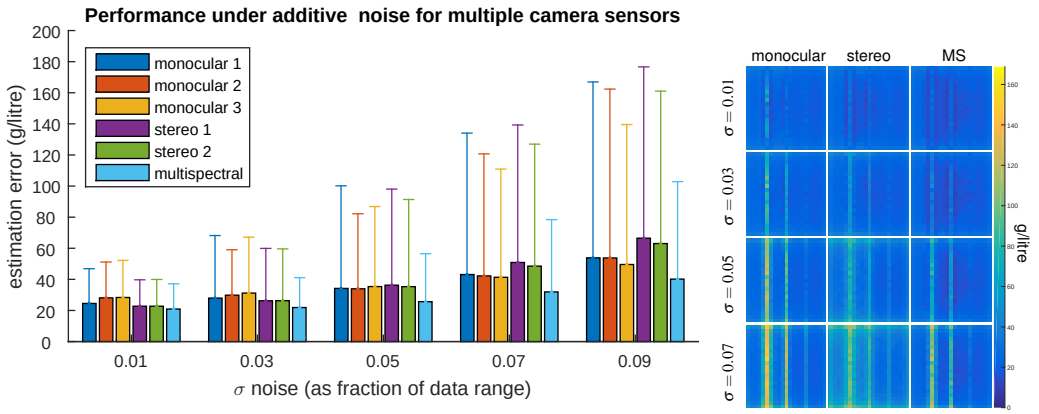

Figure 4.4: L) The absolute concentration estimation error for HbO and Hb combined across both test cases, showing mean absolute error and one standard deviation. For reference typical total haemoglobin concentration for whole blood in an adult male is approx. 145g/litre. Noise σ is shown generally as for floating point image data, in a typical 8 bit sensor $\sigma = 0.09$ corresponds to $\sigma_{8\text{bit}} = 23.04$. R) The mean concentration estimation error for monocular, stereo and MS at four noise levels.

The experiment imaged a section of porcine small bowel (45kg domestic white pig) under 4 different restricted blood flow configurations induced by occluding vessels for up to 15 minutes with rubber shod clamps. Blood flow was potentially restricted at two sites and 58 MS captures were made with the blood flow either fully restricted (both sites clamped) partially restricted (one site clamped) or unimpeded. Using the method of [43] a ground truth total haemoglobin estimation was created for each MS data set M. A subset $\hat{M}$ of the original MS data was masked where the coefficient of determination (CoD) was over 0.5, creating an MS data set where there is high confidence in the ground truth concentration. To create the input RGB data the MS data were filtered into a three channel image using a known RGB camera response curve. The imaging configuration can be seen in Figure 4.3, a section of bowel is exposed and some vessels are clamped to restrict blood flow.

4.2 Results

4.2.1 Robustness to Noise Synthetic Data

The performance of this method as seen in Figure 4.4 is close to the estimation from full MS data as the noise level increases. Further, at low noise levels the stereo (6 channel) version of this method outperforms the monocular (3 channel) version.

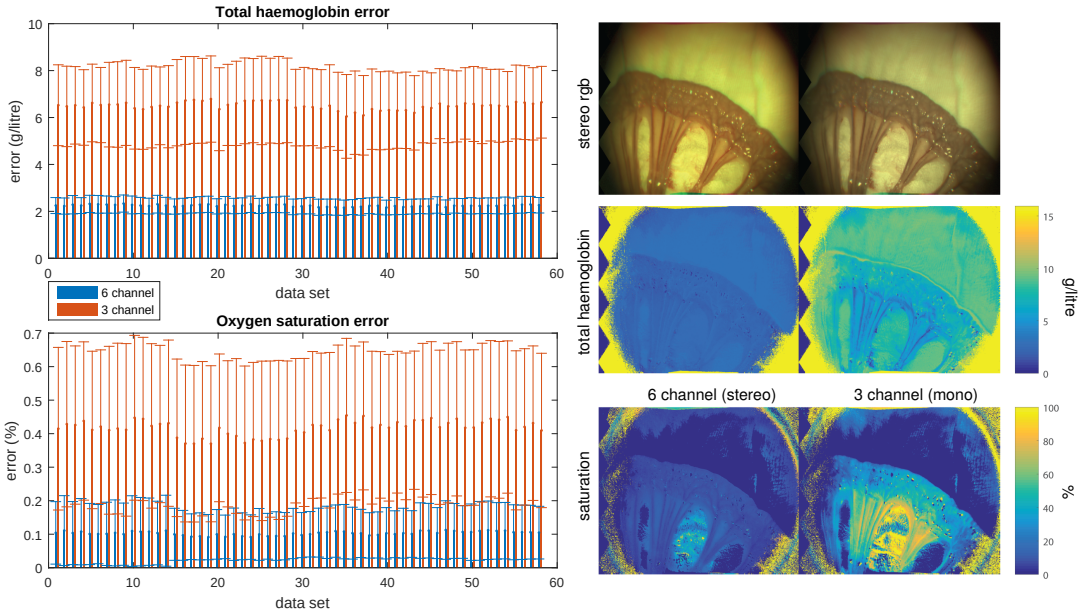

Figure 4.5: L) Estimation error for each of the 58 data sets comparing stereo (blue) and monocular (red) against MS derived ground truth. Clearly visible is the improved precision of the stereo method especially for the correct evaluation of the oxygen saturation. R) A registered stereo view of a surgical site with. Maps of the estimation error for saturation and total haemoglobin computed from stereo and monocular data.

However, at high levels of noise, the stereo version under-performs due to the increased likelihood of over or under saturated pixel data in the six channels compared to the three of the monocular version. The impact of over or under saturated measurement data is more significant in this method compared to an MS approach because each channel in this method corresponds to wide wavelength range, which in the presence of the smoothness prior on the I_λ causes large global under or over chromophore concentration estimation. The presence of a few saturated outliers has less impact on the MS method as it is directly fitting against the MS data, and the effect of a saturated outlier is localised to an individual wavelength band.

4.2.2 Estimation of Tissue Parameter In-vivo

Both the monocular and stereo versions of this method were run on $\hat{M}$ and comparison was then made against the MS derived ground truth. Concentrations of Hb and HbO were then converted into the total haemoglobin and oxygen saturation measures since these are the markers that would then be used to clinical evaluate a surgical site. As shown in Figure 4.5, the use of stereo significantly improves on the estima-

tion of total haemoglobin with an overall mean absolute error of less than 3g/litre when using 6 channels from two cameras compared to over 6g/litre for a 3 channel monocular approach. The standard deviation of the error for each method remains close, yet is slightly lower for the 6 channel variant. Given that THb concentration of whole blood is normally in the region of 145g/litre[7], the relative small magnitude of the THb estimation error indicates a high degree of accuracy when imaging for the purposes of perfusion mapping. SO_2 estimation shows the most marked improvement when using stereo over monocular with an overall saturation estimation error of 10.27% down from 41.71% for the stereo and mono variants respectively. In evaluating the *in vivo* performance, results that corresponded to a THb concentration greater than 200g/litre were considered outliers this enumerated as less than 0.1% of the results being rejected as outliers.

In both cases this method produces highly similar results to those from MS inference, and the error is typically located in areas not corresponding to vasculature. This is illustrated in Figure 4.5 where a section of bowel is exposed on a gauze background. For the stereo case the error in THb estimation is very low across the view however the SO_2 estimation performs less well in areas of low THb. This is to be expected as the oxygen saturation is a ratio of Hb and HbO, and when both are at low concentration small errors in estimation of either become amplified in the aggregate saturation measure.

4.3 Discussion

The novelty of this described approach comes from its ability to estimate concentrations of haemoglobin directly from existing laparoscopic RGB imaging devices that may currently be already in use in theatre. The method provides greater accuracy when applied to stereoscopic data as typically found in robotic assisted MAS. This is to be expected as the combined six, RGB channel, sensor responses are likely to encode for more spectral variation then three when using monocular data, however it is also reasonable that such an improvement would be unlikely if both cameras were identical or having very similar per channel sensor responses. It is shown that

the method performs well on synthetic data and is comparable to the result from raw MSI data acquired using modified imaging hardware such as an LCTF camera. This method's only requirement is to have a calibration of the laparoscopic sensors and light source to capture the response curve of each channel. This makes the technique very applicable to a wide range of MAS procedures and easily to integrate in the operating theatre. The success of this stereo method is going to be strongly linked to the quality of the registration of the two camera views, while this remains an open problem effective techniques specifically targeted at registration for MS inference have been developed [85] [9]. This method also requires imaging to be at a constant distance from the tissue surface, however, it is possible that integrating the stereo acquired depth information may provide was to normalise for these global changes in irradiance.

The key limitation of this method is the accuracy of the estimation result. While including more RGB measurements from cameras with distinct channel spectral responses does improve the estimation of THb to a degree, the estimated SO_2 remains in the realms of the unusable. It is likely that this limitation arises from over smoothing the latent MS data estimate which would obfuscate differences between attenuation spectra corresponding to either HbO or Hb. As SO_2 is a ratio of these values it would be more sensitive to poor estimation of their concentrations. In contrast THb being a sum of the respective chromophores would scale proportionally with the magnitude of the sensor measurement which is likely the reason for good THb estimation yet poor SO_2. In the next chapter, a more robust method for jointly estimating the haemoglobin concentrations with the latent MS reference data is developed to improve upon the final estimation accuracy.

Chapter 5

Robust Estimation of Haemoglobin Concentration via a Bayesian Prior on Latent Multispectral Data

This chapter focuses on refining the HbO and Hb estimation accuracy from RGB data by developing a more rigorous constraint of the distribution of latent spectral data. The technique presented in the previous chapter illustrated how haemoglobin concentrations can be estimated from RGB image data by making an intermediate estimate of the corresponding MS reflectance of tissue. However while the smoothness constraint allows for rapid computation, the Tikhonov method is limited in its accuracy, especially in the estimation of SO_2. This is because there are subtle differences between the attenuation spectra of HbO and Hb and these are lost when a strong smoothness prior is applied to the MS data. The assumption put forward that justifies the use of a smoothness prior on the latent MS reflectance is that the dominant chromophores in the image are haemoglobin [107], however it would be possible to provide a more rigorous and general prior that simply stated the the distribution of MS data must follow that which is likely to be produced by a combination of known chromophores in the field of view. The strength of this more constrained prior is that it preserves potential variation in the MS data estimate allowing for better separation of chromophores. The improved estimation accuracy that this brings enables the use of monocular RGB imaging to achieve better results

then multichannel estimation via Tikhonov regularisation.

An iterative Bayesian approach is developed from modeling the sensor's pixel measurements as weighted sums over a mixture of Poisson distributions and optimising the variables SO_2 and THb to maximise the likelihood of the observations. To validate the technique, synthetic images are generated from MC physics simulation of light transport through soft tissue containing sub-surface blood vessels, as done in the previous chapter. The synthetic data sets are also used to evaluate the robustness of the Bayesian method to a poor initial calibration of the RGB sensor. *In vivo* validation is also performed by comparing the method to an MSI dataset acquired with a LCTF based imaging device.

5.1 An Improved Prior on Latent Spectral Data

In the previous chapter it was shown that a simplistic modeling of the latent MS data corresponding to an RGB measurement could be sufficient to estimate total haemoglobin concentrations. This arose from the observation that since the attenuation spectrum of HbO and Hb are approximately smooth in VIS wavelength range the corresponding transmission spectra as predicted by the Beer-Lambert relation would also be smooth. The use of a smoothness constraint was insufficient to model the MS data and resultantly was unable to account for variation between chromophores. To better estimate the individual concentrations of HbO and Hb this section will assemble a more descriptive model for the latent MS data. This is done by taking a deeper consideration of how light interacts with the scene and camera (Section 5.1.1); then showing how the MS data and RGB data and hemoglobin concentrations can be be composed to give a Bayesian *maximum a posteriori* (MAP) estimate on the MS data (Section 5.1.2). Discussion is also presented on the optimisation of such a formulation (Section 5.1.3) alternating between estimating either the concentrations of haemoglobin or the spectral data while keeping the other fixed. The resultant method is validated on synthetic data to assess performance under measurement noise as well as for mis-calibration of the sensor, also *in-vivo* validation is performed as outlined in Section 5.1.4.

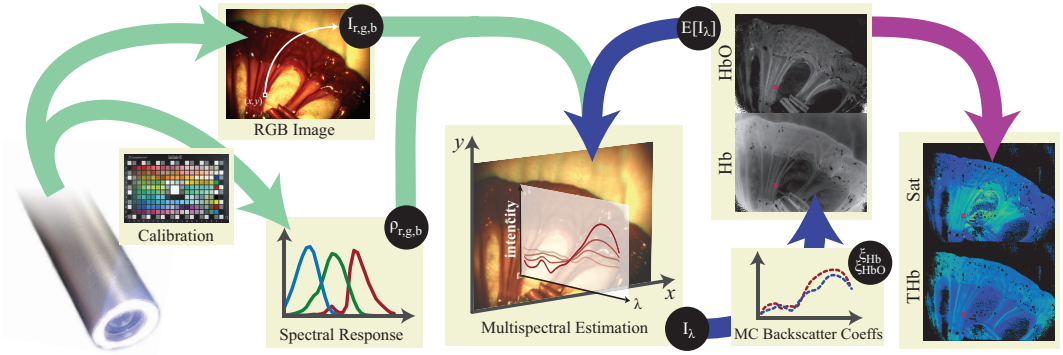

Figure 5.1: A spectral calibration was captured for the endoscopic sensor by using a colour chart illuminated with standard lighting. When an RGB image from the scope is captured, at each pixel alternately estimates are made of the MS signal from the RGB camera calibration and the current concentrations estimate (detail in Figure 5.2). Then the concentrations estimate is updated from the current MS estimate and the MC computed backscatter data.

5.1.1 Light-Tissue Reflectance and Sensor Models

The first component of the MAP estimate is the prior on the MS data. The prior is derived using the Beer-Lambert relation as a forward model to generate expected tissue transmission spectra from the current chromophore concentrations estimate. This light-tissue model are as discussed in Section 4.1.1.In this chapter the predicted reflectance spectra of light within the scene is used to yield a more specific shape prior on the MS data when estimating from RGB data. Again, the Beer-Lambert equation is used to fit and estimate for hemoglobin concentrations.

The second component of the MAP estimate is to model the interaction of light with the camera sensor. The probability for a sensor to detect photons is assumed dependent on two random variables. The first is the probability that a photon left the scene and arrived at the sensor $\boldsymbol{I}(\boldsymbol{x}, \boldsymbol{hb}(\boldsymbol{x}))$ and the second describes whether the sensor will detect it ($\boldsymbol{\rho}$) based on its spectral sensitivity as shown in see Figure 5.1. It is also assumed that $\boldsymbol{\rho}$ is constant throughout the imaging process but may vary dependent on wavelength. The resulting formulation becomes a counting problem and the probability distribution of the sensor becomes $\text{Poisson}(\rho_\lambda \cdot \text{E}[I_\lambda(\boldsymbol{x}, \boldsymbol{hb}(\boldsymbol{x}))])$ for any given wavelength, where $\text{E}[\cdot]$ represents the expectation of the random variable.

For a wide band detector, such as an entire channel (c) of an RGB sensor, the weighted sum over the sensitive range of wavelengths has weights corresponding to the sensitivity $\rho_{c,\lambda}$ at that particular wavelength. The random variable of the channel measurement $\boldsymbol{I}_c(\boldsymbol{x}, \boldsymbol{hb}(\boldsymbol{x}))$ can be written as the inner product over all wavelengths:

$$\boldsymbol{I}_c(\boldsymbol{x}, \boldsymbol{\xi}(\boldsymbol{x})) = \langle \boldsymbol{\rho}_c, \boldsymbol{I}(\boldsymbol{x}, \boldsymbol{hb}(\boldsymbol{x})) \rangle_\lambda \tag{5.1}$$

This is a straightforward sum because the contributions from different wavelengths are independent when conditioned on the chromophore concentration $\boldsymbol{hb}(\boldsymbol{x})$. The distribution of this random variable is also Poisson distributed with a mean and variance:

$$\mu_c = {\sigma_c}^2 = \sum^{\lambda} \rho_{c,\lambda} \mathrm{E}[I_\lambda(\boldsymbol{x})] \tag{5.2}$$

The random variables for sensor measurements are also conditional on the unknown concentration values $\boldsymbol{hb}(\boldsymbol{x})$ at a given position.

5.1.2 Reflectance Spectra Estimation

Since Equation 5.1 is a linear model, the natural way to recover $\boldsymbol{I}_\lambda(\boldsymbol{x}, \boldsymbol{hb}(\boldsymbol{x}))$ is to use regularised linear inversion. Regularisation in this context is needed to account for Equation 5.1 being ill conditioned. An attractive method to do so is Tikhonov regularisation, as it provides a closed form regularised estimation of the MS reflectance [8, 108]. Solving at point $\boldsymbol{x}$ (omitted for brevity[1]) as a linear system with a regularisation parameter γ:

$$\boldsymbol{I_\lambda} = \left(\boldsymbol{\rho}^{\mathrm{T}}\boldsymbol{\rho} + \gamma * Id\right)^{-1} \boldsymbol{\rho}^{\mathrm{T}} \cdot \boldsymbol{I_c} \tag{5.3}$$

Adopting normal notation Id is the identity matrix. While using Tikhonov regularisation is computationally efficient, it often trades stability of the inversion at the cost of over smoothing the result, which introduces a bias in the computed haemoglobin concentrations. While it may be possible to determine a heuristic to

[1] From here on, spatial variable ($\boldsymbol{x}$) that describes a position in the image is omitted because all calculations are made local to a point without neighbourhood information.

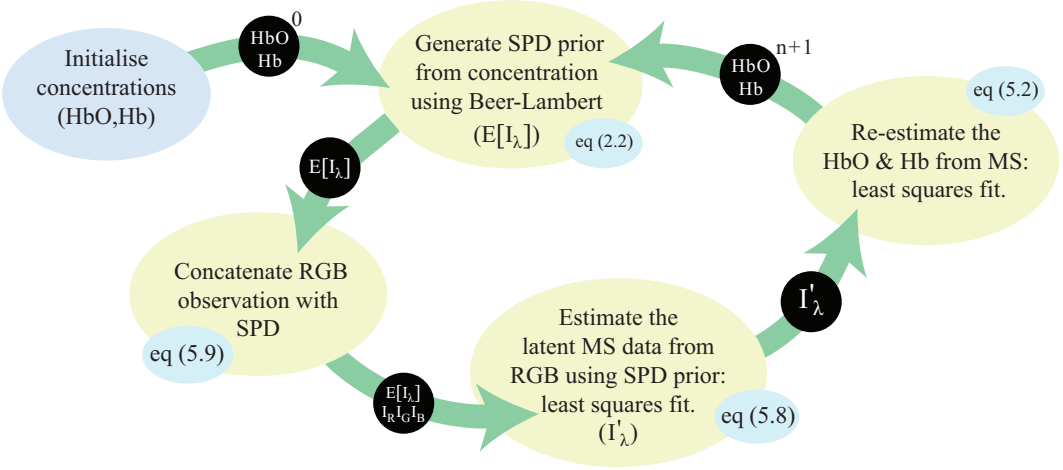

Figure 5.2: Iterative Bayesian estimation of the concentrations alternately estimates the latent MS information $\boldsymbol{I'_\lambda}$ regularised by the current concentration estimate, and subsequently updates the concentration estimate by fitting against $\boldsymbol{I'_\lambda}$.

correct for the resultant estimation bias, by fitting computed to known concentration, this would limit the flexibility and robustness of the method and would be unique for every sensor. To avoid having to calculate sensor specific correction factors, as well as to have greater resilience to sensor noise, an alternative approach is to perform estimation using Bayesian inference.

For an RGB sensor, Bayes' rule gives the relationship:

$$P(\boldsymbol{hb}|I_r,I_g,I_b) = \frac{P(I_r,I_g,I_b|\boldsymbol{hb})P(\boldsymbol{hb})}{P(I_r,I_g,I_b)} \tag{5.4}$$

where $P(\cdot)$ is a probability. Due to the channel response being independent and conditioned on the chromophore concentrations,the right hand side of Equation 5.4 can be formulated as:

$$\frac{P(I_r|\boldsymbol{hb})P(I_g|\boldsymbol{hb})P(I_b|\boldsymbol{hb})P(\boldsymbol{hb})}{P(I_r,I_g,I_b)} \tag{5.5}$$

This gives the product of measurement probabilities for individual channels and a prior on the concentrations $\boldsymbol{hb}$.

Since each channel accumulates over a wide wavelength range, the mean counts will be large and hence the Poisson distribution of count values in each channel can be approximated as normally distributed $\mathcal{N}(\mu,\sigma)$. Because the deno-

minator of Equation 5.5 does not depend on $\boldsymbol{hb}$, it can be ignored for the purposes of maximisation, and that Equation 5.4 will be maximised for:

$$\arg\max_{\xi} \left\{ P(\boldsymbol{hb}) \prod^{c} \mathcal{N}(\boldsymbol{\mu}_{c,\boldsymbol{hb}}, \boldsymbol{\sigma}_{c,\boldsymbol{hb}}) \right\} \tag{5.6}$$

The only prior on $\boldsymbol{hb}$ are that all elements are non-negative because it is not possible to have negative concentrations. While this may seem too relaxed a constraint allowing concentrations to be potentially infinite, the cases when this might occur are limited to degenerate situations where the count for each channel of the sensor is zero or very close to zero. These degenerate cases are an inevitability of the discretisation of the sensor measurements.

5.1.3 Parameter Optimisation

Since the modified Beer-Lambert model is only suitable for measurements that are non-overlapping and narrow banded, first the spectrum of incidental light $\boldsymbol{I}_\lambda{}'$ arriving at the RGB sensor is estimated. Secondly the concentrations of HbO and Hb from $\boldsymbol{I}_\lambda{}'$ solving (4.2) is estimated by minimising the sum of square difference (SSD). This is performed iteratively alternating between estimating the MS data $\boldsymbol{I}_\lambda{}'$ and then concentrations $\boldsymbol{hb}$.

Estimating the MS data from RGB measurements by linearly solving using the camera response matrix results in a very poor approximation even with the positivity constraint, as shown in Chapter 4. Typically this is because the formulation results in a sparse estimate with high values in a small number of wavelengths due to being highly under constrained. A Bayesian MAP estimation derived from (5.6) can be utilised to pose this step in a constrained way as:

$$\boldsymbol{I}'_\lambda = \arg\max_{\boldsymbol{I}_\lambda} \left\{ \left(\prod^{c} \mathcal{N}(\boldsymbol{\mu}_{c,\boldsymbol{I}_\lambda}, \boldsymbol{\sigma}_{c,\boldsymbol{I}_\lambda}) \right) P(\boldsymbol{I}_{\boldsymbol{\lambda}}) \right\} \tag{5.7}$$

The prior is generated from the previous estimation of $\boldsymbol{\xi}$ using (4.1) to estimate the absorption. Each per channel variable is maximised when the difference is minimised between $\langle \boldsymbol{\rho}_c, \boldsymbol{I}_\lambda \rangle$ and the measurements made in each channel. The effect of the prior is to penalise the per wavelength deviation of $\boldsymbol{I}_\lambda$ from the the current

estimate hence Equation 5.8 can be posed as the minimisation:

$$\boldsymbol{I}'_{\lambda} = \underset{\boldsymbol{I}_{\lambda}}{\arg\min} \left\{ \sum^{c} (\langle \boldsymbol{\rho}_{c}, \boldsymbol{I}_{\lambda} \rangle - I_{c})^{2} + \sum^{\lambda} (\boldsymbol{I}_{\lambda} - \mathrm{E}[I_{\lambda}])^{2} \right\} \tag{5.8}$$

The expected value of the spectral power distribution (SPD) of $\boldsymbol{I}_{\lambda}$, written $\mathrm{E}[\boldsymbol{I}_{\lambda}]$, is derived generatively from Equation 4.1 using the previous iteration's concentration estimates. It is important to note that the approximation of the Poisson distribution by the Normal distribution for large counts could potentially allow negative count values for each wavelength. To avoid this, a non-negativity constraint is imposed on $\boldsymbol{I}'_{\lambda}$ to prevent physically impossible values. The full estimation and optimisation process can be summarised as shown in Figure 5.1.

Notes on Implementation

While Poisson distribution is utilised in the construction of the method shown above this is for correctness and in implementation the approximation to the normal is used. As mentioned this approximation is valid when counts are large as would be expected in an optical imaging context. The optimisation of Equation 5.8 is possible via a single linear expression solving for $\boldsymbol{I}'_{\lambda}$ on the left hand side below:

$$\begin{bmatrix} \gamma & & 0 \\ & \ddots & \\ 0 & & \gamma \end{bmatrix} \boldsymbol{\rho}_{r,g,b} \, \boldsymbol{I}'_{\lambda} = \begin{bmatrix} \gamma \mathrm{E}[\boldsymbol{I}_{\lambda}] \\ I_r \\ I_g \\ I_b \end{bmatrix} \tag{5.9}$$

where I_r, I_g and I_b are the sensor measurements in each channel and γ is a regularisation constant. The choice of the scalar $\gamma = 0.01$ is as in section 4.1.1, and similarly reasoned to make solving Equation 5.9 not over fit or be ill conditioned. Since the left hand side of Equation 5.9 is constant over all iterations, a computationally efficient approach is to reuse decompositions of the left hand side matrix at each iteration. The two step estimation process of alternatively estimating the latent MS data and then the chromophore concentrations is illustrated in Figure 5.2.

A fast non-negative least squares method is used to solve Equation 5.9 and

also for the subsequent parameter estimation from the estimated emission MS data [101]. Direct computation using this method with on-the-fly optimisation yields frame rates of 2.5 frames per second (720p), using a Surface Pro 3 2.3GHz 8GB RAM. However, the optimisation is deterministic for a given illumination spectrum and camera sensor response and real-time performance can easily be achieved by using a precomputed look up table (LuT) of all possible RGB values. Using a LuT consisting of all possible 24bit RGB colours (approx. 16 million) it is possible to process endoscopic images at video rate (over 30fps) using a serial implementation requiring 270MB of RAM.

5.1.4 Validation methods

Synthetic and *in-vivo* MS data were used to validate the proposed method using the same MC and porcine experiments, respectively, as described in Chapter 4. Further to this the effect of poor quality or mis-calibrated sensor response is explored using synthetic data

In lieu of high quality *in-vivo* ground truth data, the Bayesian technique is first validated using synthetic data generated from measured optical characteristics of blood [7] and sub mucosa [106]. This validation is performed in the same way as in section 4.1.3 to enable cross comparison of the two methods and the process of data generation is covered in greater detail there. The same digital phantom comprising three sub surface vessels was used to compare performance of additive sensor noise. Zero mean normally distributed noise was added to the input RGB and MS data at various levels to observe the performance deterioration in the presence of measurement or sensor noise. For the synthetic RGB data the variance of the noise in each channel was correlated to the spectra response data being used for each. The results of Bayesian estimation as described in this chapter are compared to a Tikhonov regularised estimation from Chapter 4, ([8]) as well as a full MS analysis [9]. In the top of Figure 5.3 there are shown example synthetic RGB channel responses corresponding to the increasing sensor noise. The estimation error for each method is also shown in Figure 5.3 as well as tabulated numerically in Figure 5.4.

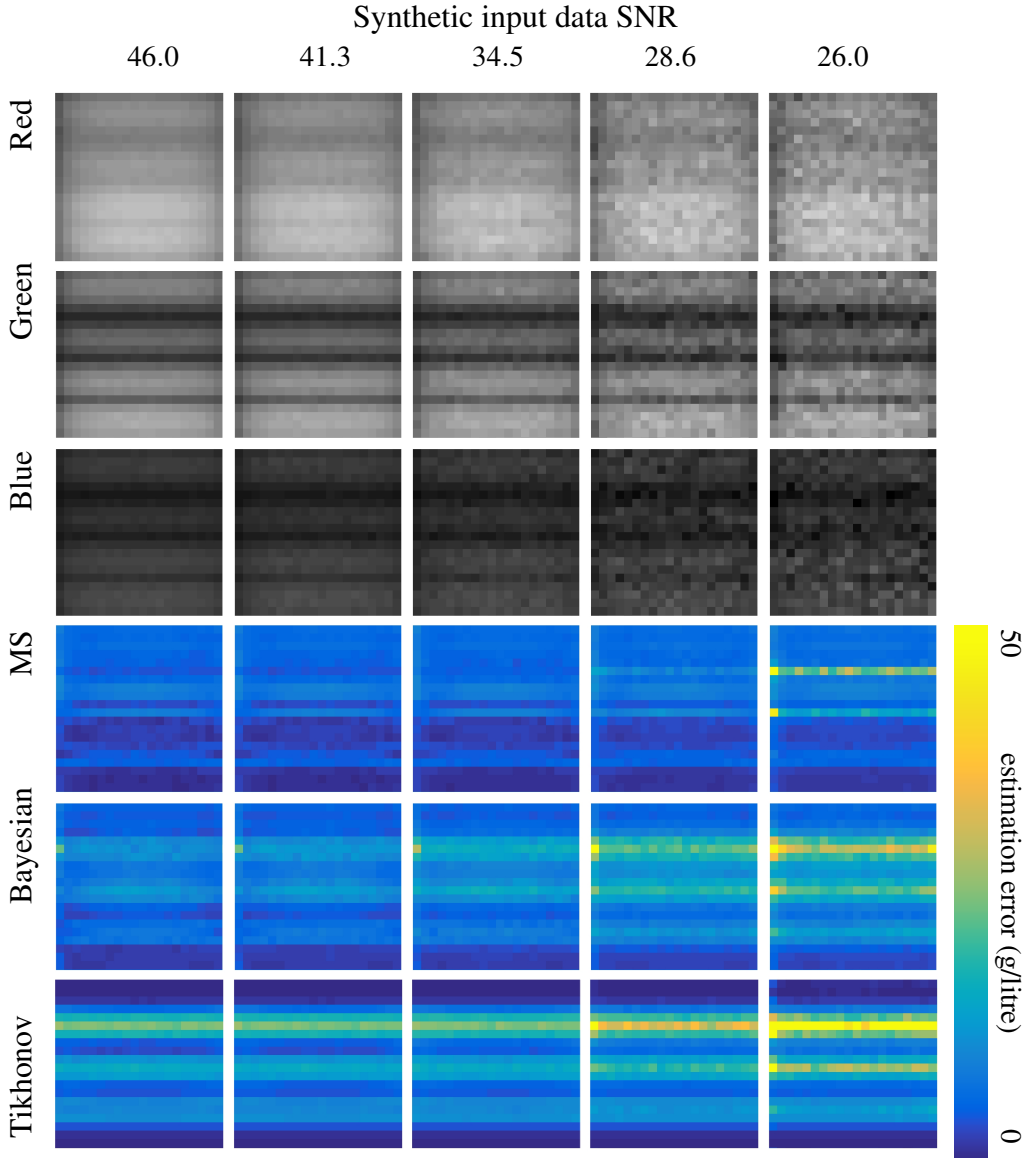

Figure 5.3: Top three rows: example synthetic RGB responses for increasing levels of added noise, SNR calculated with respect to zero mean added noise data. Lower three rows: visualisation of the mean absolute error in chromophore concentration for the Bayesian method, Tikhonov regularised estimation [8] and a full MS approach [9].

5.2 Results

5.2.1 Synthetic Results

For the estimation of individual concentration values for HbO and Hb, the Bayesian method performs close to the hardware MSI method [9], and better than the Tikhonov method presented in the previous chapter. This constitutes a high degree of equivalent accuracy in estimating THb. The SO_2 estimation is particularly sensitive for Tikhonov estimation and the Bayesian method offers a notable improvement.

	Source	HbO Estimation Error (g/litre)			Hb Estimation Error (g/litre)		
σ	PSNR	[9]	Ours	[8]	[9]	Ours	[8]
0	-	0.23	0.44	0.82	0.94	1.23	3.71
0.005	46.0	0.23 ± 0.2	0.51 ± 0.3	0.82 ± 0.6	0.94 ± 0.6	1.36 ± 1.3	3.79 ± 3.1
0.006	44.2	0.24 ± 0.2	0.55 ± 0.3	0.83 ± 0.6	0.95 ± 0.6	1.42 ± 1.4	3.82 ± 3.2
0.009	41.3	0.24 ± 0.2	0.64 ± 0.4	0.84 ± 0.6	0.95 ± 0.6	1.57 ± 1.5	3.83 ± 3.2
0.013	37.9	0.25 ± 0.2	0.85 ± 0.6	0.90 ± 0.7	0.97 ± 0.6	1.95 ± 1.9	3.86 ± 3.2
0.019	34.6	0.27 ± 0.2	1.20 ± 1.0	1.06 ± 0.8	1.04 ± 0.7	2.40 ± 2.3	3.88 ± 3.3
0.027	31.4	0.31 ± 0.3	1.77 ± 1.5	1.72 ± 0.5	1.18 ± 0.8	3.04 ± 3.0	3.92 ± 3.3
0.037	28.6	0.48 ± 0.4	2.71 ± 2.4	3.92 ± 3.8	2.16 ± 1.7	3.84 ± 3.8	4.61 ± 4.1
0.050	26.0	1.83 ± 1.7	4.66 ± 4.3	11.3 ± 11	7.76 ± 7.1	6.66 ± 6.6	9.17 ± 8.8

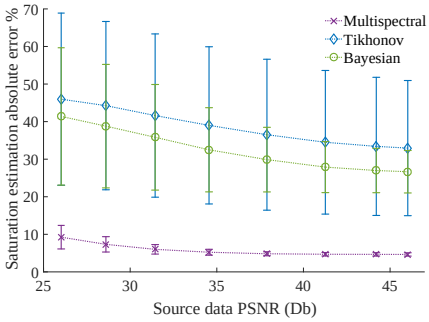

Figure 5.4: Left: MSE in estimating HbO and Hb concentrations (g/litre) under additive noise on synthetic data. NB numbers have been truncated for display but the variance is less than the mean error. Right: mean standard deviation in absolute error against the synthetic ground truth for increasing noise levels. Comparing this method (Bayesian) with [8] (Tikhonov) and [9] (MS).

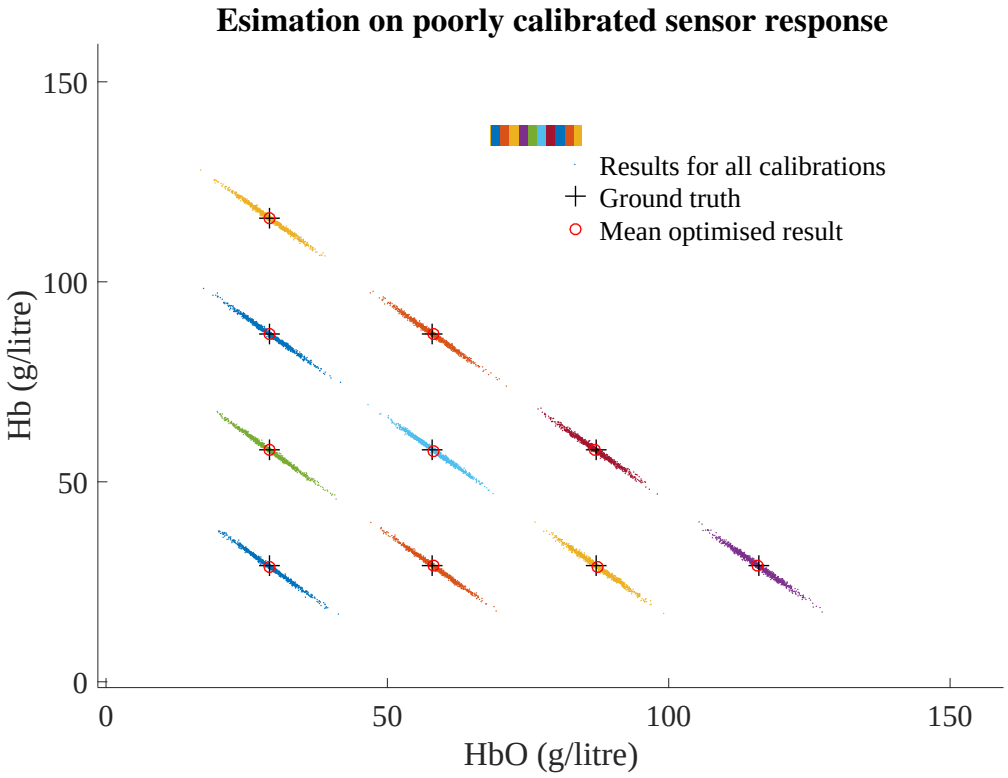

Figure 5.5: The impact of noisy sensor calibration on Bayesian estimation is shown in this figure. The optimisation result using 1000 randomly sampled noisy sensor calibrations for 10 concentration pairings, here the noise level is 39.9 dB compared to ground truth.

While the mean improvement in SO_2 estimation error is close to 10%, at low noise levels there is a large reduction in the variance of this error. There is still a sizeable short fall in the accuracy compared to the hardware MSI. While not as capable as a multi-band approach the Bayesian method at least seems to be generally indicative of tissue saturation and such relevant information may still be pertinent clinically.

The simulated imaging process to generate MS data using MC sampling is inherently noisy as it is the product of the passage of a finite number of photons, although the MC framework of [105] will eventually converge to the noise free result. Further as mentioned in [106], it is difficult to eliminate the contribution of haemoglobin entirely when measuring the optical properties of tissue such as those for the sub-mucosa used in the synthetic model. As such it is possible that the physical parameter absorption and scattering data for submucosa used is likely actually a combination of the true values for submucosa as well as small fraction of those corresponding to haemoglobin. This can result in overly idealised ground truth haemoglobin distribution map, for which the submucosa was assumed to be entirely vacant of HbO and Hb. These are likely the dominant factors for why the reference method [43] does not achieve perfect results on the synthetic data.

5.2.2 Results in the Presence of Calibration Noise

The proposed framework is fundamentally dependent on the accuracy of the endoscopic sensor's spectral sensitivity calibration. In practice, a high quality calibration can be performed using a monochromator and this would need to be repeated infrequently to maintain an accurate calibration, dependent on sensor characteristics and degradation over time [102]. Such a calibration can even be performed during manufacturing and remain accurate even for devices that are in regular clinical use [70]. The method of performing a high quality calibration via the means of an LCTF is discussed in Section 4.1.2. However, in the case where such a technique is unavailable, it is worth investigating the effect of a poor calibration for a more complete numerical understanding of our method's characteristics.

For the purposes of numerical investigation the performance under poor initial calibration was observed via simulated data. Practically these could arise if

calibration is performed by using low-cost solutions such as colour matching on Macbeth colour charts with known reflectance. Images of such objects can be used to estimate the endoscopic sensor's response similarly to how geometric calibration is performed [103]. To assess the impact of poor calibration, reflectance spectra for 24 colour swatch chart [109] were used to simulate an RGB camera response under additive zero mean Gaussian noise. For each of the 24 swatches a 256x256 patch measurement was generated and then the calibration run on the noisy data. These synthetic calibrations were performed 10e4 times to numerically determine the distribution of calibration errors over wavelength, using this distribution realistic mis-calibrated sensor response curves were generated. MS data for 10 concentration pairings were generated upon which the Bayesian method was run using 10e3 randomly sampled mis-calibrated sensor response data.

The estimation for the 10 ground truth concentration pairings is shown in Figure 5.5 for the calibration noise level of 39.9dB. This was chosen as it was the point at which errors noise began to be observable in the case of measurement noise. Visible in the results is the distribution of the error such that for any given ground truth value the errors are negatively correlated when comparing HbO and HbO. This means that an incorrect calibration would produce results that are still fairly reliable for measuring the THb, as this measure is based on the sum of the two values. However, the SO_2 estimation is going to be very sensitive to poor calibration data as the ratio of HbO to HbO varies considerably.

5.2.3 In-vivo Results

The Bayesian method described in this chapter is compared with the Tikhonov regularised estimation (from Chapter 4) using the MS results as ground truth, and is illustrated in Figure 5.6 and 5.7. A strong correspondence is observed between the MS ground truth and both the Bayesian and Tikhonov regularised estimation, however, the Bayesian method is more accurate in correctly estimating the ratio between the concentration of HbO to Hb as shown in Figure 5.6. The bias of the Tikhonov estimation is shown in the corresponding error histogram of Figure 5.6 as twin peaks either side of zero. There are similar, much smaller peaks visible in the error

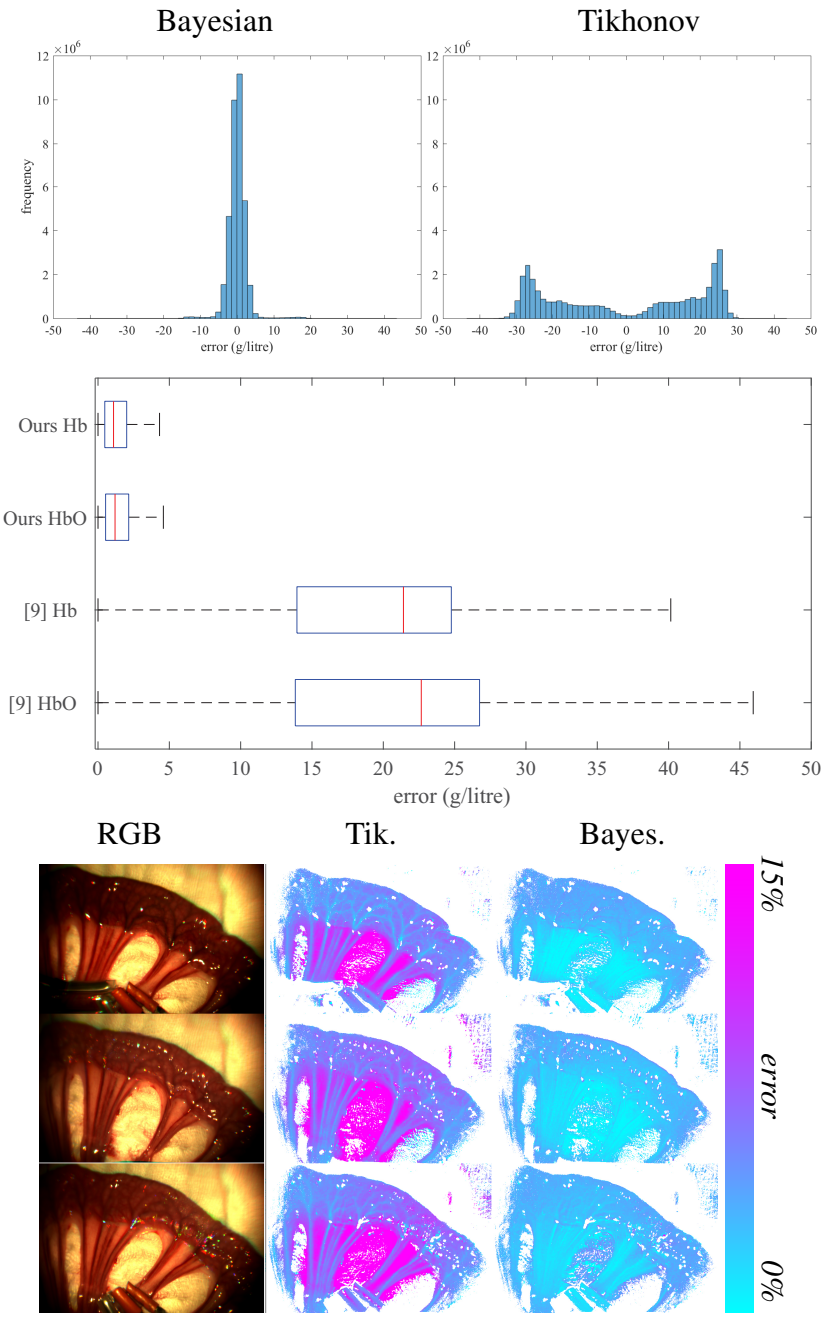

Figure 5.6: Top: Error histograms for the Bayesian and Tikhonov methods against [9] (used as ground truth) combined for HbO and Hb. The Bayesian method's error is more closely grouped around zero. Centre: Comparison of the absolute error for HbO to HbO concentration estimation, on *in-vivo* data for the Bayesian and Tikhonov methods, against the MS ground truth. Bottom: example absolute error maps for the saturation estimation with RGB view of the surgical site for reference, visible is the lower error in areas corresponding to vasculature (higher perfusion).

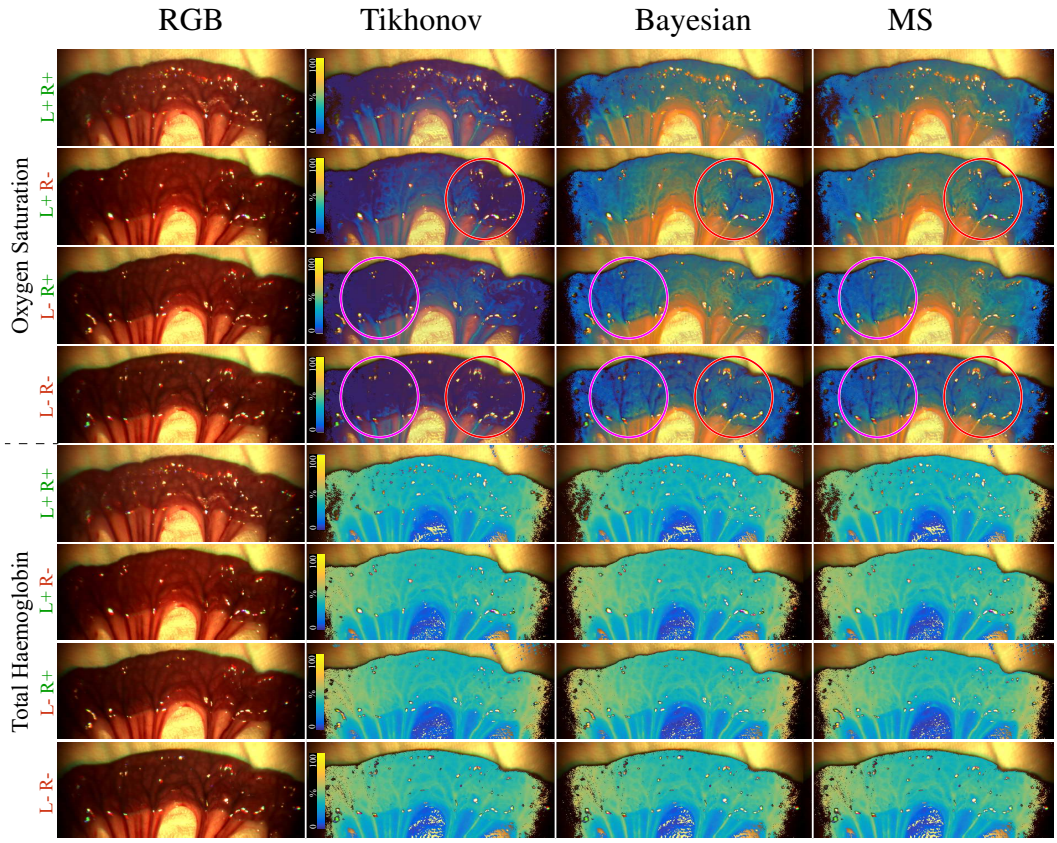

Figure 5.7: SO_2 and THb maps corresponding to the four configurations of the *in vivo* experiments, clamping vessels leading to either the left or the right of the bowel as viewed, green (+) corresponds to clamp open and red (−) corresponds to clamp closed. The Bayesian method is shown with the MS method of [9] and Tikhonov regularised estimation from Chapter 4. Maps of SO_2 distribution are displayed at the top and THb maps are displayed to the bottom, both overlaid on a reference RGB view as visible in the left most column. For the SO_2 maps they have been overlaid by alpha blending using the THb as an alpha value, regions are also highlighted corresponding to the part of the bowel with restricted circulation.

histogram for the Bayesian method, however most of that method's errors are concentrated around zero. Numerically, the Bayesian method Tikhonov methods are evaluated against the result from [43] as the motivation is to closely approximate the result from direct MS imaging. The SNR is measured for the Bayesian method as 45.9±8dB for THb estimation, 28.6±3dB for SO_2 estimation and 35.8±4dB for the overall raw Hb and HbO concentration values. In contrast, the Tikhonov regularised estimation yielded a SNR of 38.2±2dB for THb estimation, 8.2±1dB for SO_2 estimation and 16.8±1dB for the overall estimation. As observed in the synthetic experiments, a marked improvement in the accuracy of saturation estimation can be

seen for the Bayesian method over the Tikhonov regularised approach. The improvement in the THb estimation is less significant when comparing against the MS result as ground truth. This contrast in performance is shown in Figure 5.7. Pearson correlation coefficients[2] were calculated for the haemoglobin concentrations estimated from the Bayesian and Tikhonov approaches against the result from direct measurement of MSI data[9]. For the Bayesian method a correlation coefficient of 0.996 was calculated and for the Tikhonov a correlation coefficient of 0.766. For both, the confidence level on the correlation was above 99.9% and sample sizes were greater than 3×10^7 data points. The stronger correlation for the Bayesian method explains why a smaller error in saturation estimation is observed for that method in comparison to the Tikhonov regularised estimation.

A comparison of the estimation errors for THb and SO_2 is shown in Figure 5.6; illustrated here is how the SO_2 estimation error seems to be lower in areas corresponding to vasculature, and worse in regions where there is expected to be a lower THb value. This is to be expected because when there are low concentrations of both chromophores; slight errors in estimation of either can lead to a significant change in the ratio of one to the other. The THb estimation for the Bayesian method showed a small improvement compared to that of the Tikhonov regularised estimation. For both the raw concentration values (HbO and Hb) and the composite measures (THb and SO_2), the Bayesian method produced a lower mean absolute error and the estimation error for our method also had a lower variance, when comparing against the Tikhonov approach.

It is possible to observe similar localisation patterns in the saturation of blood within tissue in our results as to that of hardware MSI analysis. Highlighted in Figure 5.7 the parts of the small bowel that correspond to the tops of blood vessels that were either restricted or unimpeded. In the cases where blood flow is restricted a noticeable drop in the oxygenation of these regions can be seen. While it has been shown that the result of the Bayesian method is numerically similar to the result from MSI, in Figure 5.7 it can be seen how the visualisation of the result from our

[2] calculated using MATLAB *corrcoef* function

method is also visually similar to the visualisation of the MSI result.

5.3 Discussion

In this chapter, a framework has been developed and evaluated for estimating haemoglobin concentrations from RGB image data that bypasses the need for direct MS imaging yet which produces comparable results. The results shown are commensurate with those obtained through hardware MSI and while the Bayesian technique does not achieve the same accuracy on simulated data, its potential for providing additional information within existing endoscopic imaging systems could be clinically valuable. Additionally, while MSI using higher numbers of wavelength limited acquisitions achieves better spectral discrimination, it suffers from misalignment and blurring in the data cube when imaging dynamic tissue that the Bayesian and Tikhonov methods are much less susceptible to.

A limitation of the Bayesian method presented here is that it is not a real-time algorithm, insofar that it is not capable of processing images at a typical video rate of 24-30fps. To achieve real-time use of estimation afforded by this method one could use a precomputed LuT. The downsides of using a LuT is that the computational requirements scale exponentially with the number of channels or cubically with the sensor colour depth (for three channels). Another limitation of the Bayesian method is that it requires the camera and illumination source to be at a fixed normal to the tissue surface because this assumption is used in order to simplify the estimation term. The method is, however, robust to changes in imaging distance between the camera and the tissue surface up to a point. This is achieved by adding a constant term to the estimation of chromophores, assuming that the loss due to greater distance from the tissue to the camera would be constant across wavelengths. For normal to surface imaging it is assumed that the variance across wavelengths for dispersion over distance would be less noticeable compared to acute imaging angles where the variable refraction may start to play a more significant role. If surface reconstruction could be performed then it may be possible to modify the constant term in such a way to compensate for the potential attenuation due to off

normal imaging angle. An important note about the Bayesian framework is that it can be extended to incorporate additional constraints, although dynamic adaptation to things such as surface to imaging axis angle would prohibit the use of a precomputed LuT. The Bayesian method developed in this chapter operates on a per-pixel basis but it is apparent, in endoscopic images, neighbouring points on the tissue surface are related and represent physiological structures like blood vessels or tissues from different organs. Combining information, for example from image segmentation, within the framework could allow better regularisation and estimation with multiple pixel measurements rather than the current single measurement approach.

While the experiments in this chapter have shown that the Bayesian method is promising in terms of accuracy, providing information comparable to hardware MSI, it is important to note that additional experimentation is necessary to understand both the variance of haemoglobin concentration quantities in tissue and potential further clinical applications that may benefit. This relates to one of the general difficulties alluded to earlier in the validations sections; most works are either validated against synthetic data or estimation using MS data from the best available MSI device. Sometimes point measurements are used to validate specific regions in an image, however in order to explore the other ways in which MSI like devices can be validated they need to be more easily deployed in a clinical context to be able to investigate the practicalities of different *in-vivo* validation techniques. In the next chapter the idea of addressing the hurdles for clinical translation of computation MSI are further explored looking at how to create a hybrid method that combines the best features of its constituent parts.

Chapter 6

Detecting Real-time Haemodynamic Variation in Tissue

The methods developed in Chapter 4 and Chapter 5 both necessitate a trade off between accuracy and computational efficiency, albeit each prioritises a different one. Either, for the Bayesian method, it can achieve good accuracy but the computation speed would be far too slow to perform in real-time ruling out integrating into a live view and monitoring subtle temporal variation; or vice versa for the Tikhonov regularised method, where its value is limited due to sub par accuracy. As such, in this chapter a method is developed that combines the positives of both of these previous methods while minimising the downsides, to formulate a method that minimises the technical limitations for clinical deployment.

The success of this method comes from compressing the RGB image data in such a way that the more accurate estimation can be performed on a small fraction of the data, while the rapid estimation can process the majority. Wavelet decompositions have been widely used in image processing to transform data into domain spaces where certain operations can be applied more efficiently or components of the signal with little entropy can be omitted [110]. In this chapter, a wavelet decomposition with the Haar wavelet is adopted due to the fast discrete decomposition and recomposition, as well as the way it separates the edge data to leave a small final residual layer. By such a decomposition the image data is split into two distinct groupings with distinct properties, this enables the use or separate algorithms

to process different coefficients of the wavelet transformed image data. Due to the Haar wavelet producing large numbers of near-zero coefficients in the transformed data set, intensity information is concentrated in relatively few coefficients which can then be processed effectively. The resulting composite algorithm can utilise RGB data to arrive at a surrogate MSI signal through a dual optimisation approach which is demonstrated effective for fairly smooth signals, such as laparoscopic video data.

6.1 Estimation from Haar Wavelet Coefficients

In this section the rationalisation is given for using two different estimation techniques on two subsets of Haar coefficients of an image (Section 6.1.1). The methods of validating such an approach are then discussed (Section 6.1.2) for the synthetic and *in-vivo* data sets used.

6.1.1 Decomposing the Image for Separate Computation

As with the techniques developed in Chapter 4 and Chapter 5, underpinning the technique proposed in this chapter is the estimation of a signal surrogate to MSI but obtained from RGB video frames. Similarly this also relies on knowledge of the tissue's transmitted attenuation characteristics [7]; backscattered attenuation coefficients are derived by MC simulation [105], to combine attenuation due to absorption and scattering (as demonstrated in Chapter 5), since camera and light source are on the same side of the tissue during laparoscopic imaging. The fitting to surrogate MSI data can be performed to satisfy the Beer-Lambert relation.

Once an RGB image is captured, the 2D discrete Haar transform can be used to decompose the image into four components, three directional approximate derivatives and a residual low pass coefficient [111]. This fact an be exploited for efficient computation since it enables only performing the expensive estimation on less of the data ($1/2^{1+n}$ where n is the level of decomposition). The matrix H can be expressed for an image window $I = \{I_{x,y}, I_{x,y+1}, I_{x+1,y}, I_{x+1,y+1}\}$ as:

$$\text{Haar}(I) = I * H \tag{6.1}$$

given that the 2d Haar matrix is formed as:

$$H = 0.5 \cdot \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & -1 & 1 & -1 \end{bmatrix} \tag{6.2}$$

As this is a just a linear transformation, if each pixel data is grouped in windows of four elements it can be included in the construction of RGB data I_{RGB} from the MS signal I_λ:

$$C \cdot I_\lambda \cdot H = I_{RGB} \cdot H \tag{6.3}$$

for known camera spectral sensitivity matrix C, the least squares solution being:

$$I_\lambda = ((C^{\mathrm{T}}C)^{-1}C^{\mathrm{T}} \cdot I_{RGB} \cdot H) \cdot H \tag{6.4}$$

since $H = H^*$. Similarly for the Tikhonov regularise solution as used in Chapter 4, this method can be applied to the Haar transformed data as

$$I_\lambda \cdot H = (C^{\mathrm{T}}C + \Gamma)^{-1}C^{\mathrm{T}} \cdot I_{RGB} \cdot H \tag{6.5}$$

where Γ is the identity matrix multiplied by a small constant γ. Here the same windowed grouping is used:

$$I_{RGB} = \{I_{RGB,x,y}, I_{RGB,x,y+1}, I_{RGB,x+1,y}, I_{RGB,x+1,y+1}, \} \tag{6.6}$$

corresponding to the MS window:

$$I_\lambda = \{I_{\lambda,x,y}, I_{\lambda,x,y+1}, I_{\lambda,x+1,y}, I_{\lambda,x+1,y+1},\} \tag{6.7}$$

While this demonstrates how Tikhonov based estimation can be used on any of the Haar coefficients, to attain higher accuracy it is desirable to only use this on the sparse directional components. On the low-pass coefficients the Bayesian method of Chapter 5 is utilised, which is more accurate than Tikhonov estimation, enabling better estimation of SO_2. It is possible to do so because the low pass coefficients are analogous to an over-exposed image taken at a lower resolution, and can be evaluated by dusting the baseline illumination I_0 when estimating haemoglobin concentrations from the MS data.

Formally, it can be shown that the method holds by substituting in the maximisation step of the Bayesian method with low pass transformed data. Using $H_{LP} = 0.5 * [1,1,1,1]$ to represent the low pass component of the discrete Haar transform, starting from the concentration estimation $\hat{\boldsymbol{hb}}$ from low pass data:

$$\xi \hat{\boldsymbol{hb}} = -log(I_\lambda H_{LP}) \tag{6.8}$$

where ξ are the backscatter attenuation coefficients for oxy and de-oxy haemoglobin (HbO and Hb) as well as a constant term. The constant term is included to account for global illumination changes that occur equally across wavelength, such as those due to distance from the camera. Using least squares fitting to solve for $\boldsymbol{hb}$,

$$\hat{\boldsymbol{hb}} = (-\xi'\xi)^{-1}\xi' \log(I_\lambda H_{LP}) \tag{6.9}$$

These concentration estimates are then used to generate the expected value for the MS data $E[I_\lambda]$ used to regularise then subsequent fitting iteration, this is defined as,

$$E[I_\lambda] = e^{-\xi \boldsymbol{hb}} \tag{6.10}$$

substituting the concentration estimate from (6.8) and cancelling terms the expected value will be the MS data transformed by H_{LP}

$$E[I_\lambda] = I_\lambda H_{LP} \tag{6.11}$$

This can only apply to the low pass coefficients because they are all positive and so all terms remain real throughout.

The expectation step of Bayesian estimation method is linear and so compatible with working on Haar transformed data. One notable change made to the expectation step is to change the prior on the expected spectrum from a value prior to a shape prior. This can be done by computing the second derivative of $E[I_\lambda]$ and using that to regularise the second derivative of the estimated $\hat{I}_\lambda$, this better regularises the estimation.

6.1.2 Validation methods

The algorithm described above was validated via two different experimental approaches. The first uses RGB views generated from *in-vivo* animal experiment MSI data sets, which allow evaluation comparing the method to a hardware implemented gold standard. Secondly the method was used to process video from a patch of tissue at the base of the tongue, where due to this algorithm being capable of higher frame-rates it is able to detect pulse rate, by tracking THb variation in a patch of tissue over time.

Real *in-vivo* MSI data was used to generate synthetic RGB images using known camera spectral sensitivity curves, calibrated from the spectral response of cameras on a Da Vinci surgical robot [112]. The MSI data comes from experiments in which MSI was used to monitor and evaluate organ viability, by measuring tissue perfusion and oxygen saturation, throughout uterine transplant procedures performed on sheep and rabbits [47].

Video was captured of the base of the tongue of an adult male using a Da Vinci surgical robot's stereo laparoscope as shown in Figure 6.2. The laparoscopic camera was calibrated to estimate the channel sensitivity. A patch of tissue was tracked

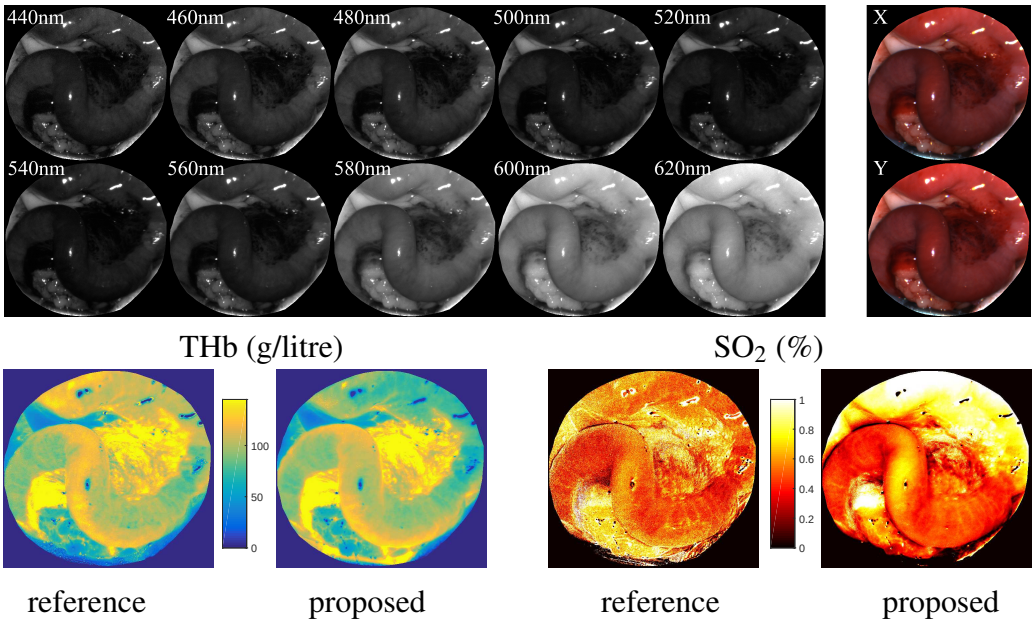

Figure 6.1: Top: Selected bands from a MS datacube and corresponding synthesised RGB views. Bottom left: total haemoglobin. Bottom right: oxygen saturation. Example estimation from this method using RGB data synthesised from MSI and compared with reference estimation directly from the MSI data.

using the method of [85] over time in both left and right cameras, and processed the two views separately. Tracking and image warping was necessary in order to remove residual motion artefacts due to tongue movement allowing visualisation of the spectral and temporal variation of a selected region within a spatio-temporally registered signal.

6.2 Results

Good stability of the method is observed both on the generated RGB data and on the *in-vivo* acquired RGB data. This is conformant to expectations as the method independently combines two methods which either have a closed form solution as used in Chapter 4 or are empirically observed to be stable under noisy conditions as for the Bayesian component derived in Chapter 5.

Table 6.1: HbO and Hb estimation accuracy on synthetic data compared to direct estimation from MS data [9]. All algorithms implemented by myself to ensure consistent baseline for comparison.

Method	Computational architecture	Execution time Hz	mean error g/litre
Tik [8]	CPU	0.3251	60.8951
Tik [8]	GPU	42.9356	60.8952
Bayes [35]	CPU	0.0774	6.6038
Bayes [35]	GPU	3.5626	6.6053
Combined (w1)	CPU	0.1474	6.5367
Combined (w1)	GPU	12.6624	6.5367
Combined (w3)	GPU	14.3833	5.9567

6.2.1 Comparison with Hardware Multispectral Imaging Signals

To evaluate computation speed and accuracy the combined methods is compared against the methods from Chapter 4 and Chapter 5, however, to make fair comparisons of speed on the same architecture both of these methods were also implemented for GPU, using the CUDA programming language. Two variants of the method put forward in this chapter were run, corresponding to single (w1) and multilevel (w3: three levels) Haar decomposition. All methods' accuracy are compared against the results of estimation from directly using the MSI data [43].

As seen in Figure 6.1 the estimation from the composite method bears strong visual similarity to the result of direct estimation from MSI data, with two notable variations. Firstly in the composite method there is much less noise in the estimation, which can be interpreted as a product of the MSI data's band limited images each being moderately noisy. Such noise is inherently smoothed out when synthesising RGB data from the MSI datacube as each RGB band integrates over many MSI bands due to wider sensitivity of the respective colour channels. On the other hand, the RGB technique is less able to robustly estimate haemoglobin concentrations in areas of high luminance; this can be clearly seen in the SO_2 estimation near the top. This is possibly due to the non-linear response of the camera at extreme mea-

surements. In addition, the generally smoother appearance of the combined result means that local variations in the MSI signal are lost due to both the synthesis of the RGB image and subsequently due to the regularisation in the combined estimation approach. However, overall the combined method was closer in accuracy to that of a fully Bayesian estimation than a fully Tikhonov regularised method (shown in Table 6.1), despite Tikhonov estimation being used on the majority of coefficients.

In terms of computational performance, the timed results of haemoglobin estimation were made using mega-pixel sized input images. Timings were calculated including all pre-processing steps, which for the combined method also included the initial Haar decomposition and subsequent recomposition. This illustrates the diminishing returns for the for multi-leveled (w3) variant of the combined method compared to the single decomposition level version (w1), as tabulated in Table 6.1. Interestingly both variations have the same error performance albeit with minor computational variations in terms of speed.

6.2.2 Experiments on In-vivo Data From a Stereo Laparoscope

The comparison of the mean value of total haemoglobin as estimated from each view is shown in Figure 6.2 and a strong similarity was observed in estimation over time. While there appears to be constant offset between the estimate from either view, this could be due to a mis-calibration of either of the cameras' spectral response curve. A large amount of high frequency noise was observed in the time series data, which could be due to specular highlights moving as the tissue surface deformed and moved or to imperfect tracking which is to be expected given the difficult conditions.

Processing the registered video signal after warping using a low pass filter, it was possible to observe a periodic signal in the derivative of the time series which was aligned in both views. This may be representative of physiological signals within the tissue due to cardiovascular activity. By looking at the Fourier transform of the change in THb estimation the power spectrum peaked between 1.2 and 1.3 Hz which corresponds to periodic signal of between 72 and 78 cycles per minute. This corresponds with the heart rate of the subject during the experimental acquisition.

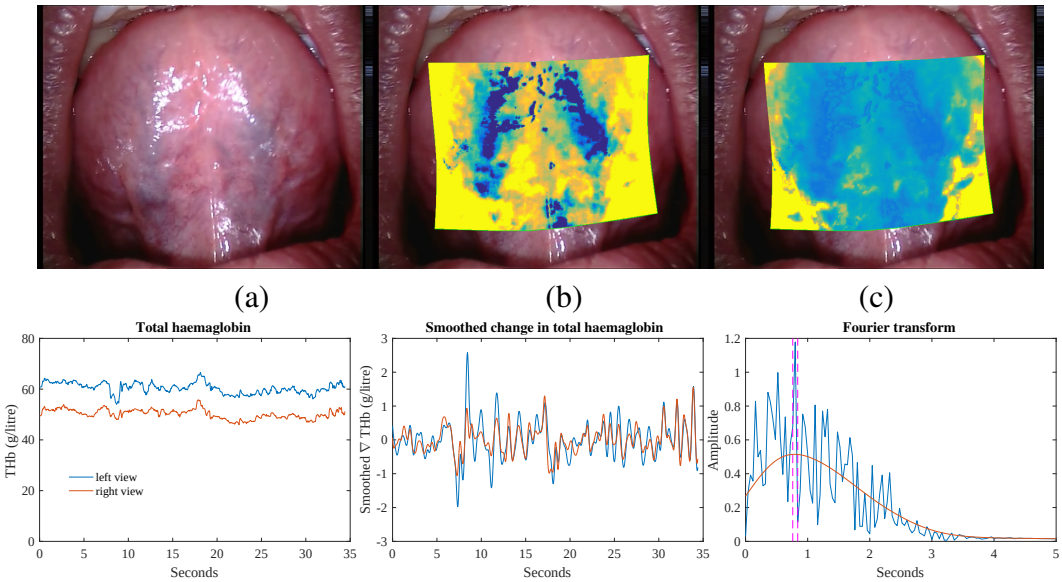

Figure 6.2: Top (a) the original RGB laparoscopic view, (b) SO_2 estimation overlaid on the laparoscopic view, (c) THb estimation overlaid on the laparoscopic view. Bottom left: the tracked THb over time from the left and right camera feeds separately these show a mostly constant THb level throughout, which is reassuring that the the subject is not experiencing an unexplained loss of blood. Bottom centre: the derivative after smoothing the THb trace for each camera view, smoothing is required due to large amounts of high frequency noise in the trace. (roaming high-lights on tissue surface, non-perfect tracking) Bottom right: Frequency analysis of the change in THb over time, fitting a polynomial curve finds the peak between between 1.2 and 1.3 Hz.

6.3 Discussion

The hybrid method presented here, improves on the computational speed with minimal loss in accuracy, with respect to the methods presented in previous chapters. This method's speed improvement, functioning at over 14Hz, is sufficient to enable use in real-time clinical applications, as it is fast enough to resolve variations in the oxygenated and de-oxygenated blood flow within tissue close the exposed surface. During surgical procedure heart rate can become elevated to over 2 Hz [113]; as such, imaging at a rate faster than 4Hz would be required to detect this without aliasing. The precision of the hybrid method is also high enough to evaluate the oxygen saturation within tissue, this is a major improvement on the Tikhonov method presented in Chapter 4. The results on *in-vivo* measurements of haemoglobin concentrations using an existing laparoscopic imaging configuration for RGB imaging are promising, as the hybrid approach presented here actually improves on the estimation accuracy with respect to the purely Bayesian from Chapter 5.

The main limitation of this method is again the accuracy of the sensor response calibration, as such patch based calibration is likely to be insufficient for accurate SO_2 results. A possible solution is to use a monochromator to precisely measure the sensor response or to make use of an LCTF in conjunction with the patch measurements as in [9]. Alternatively it could be a great improvement to further include a joint estimation of the RGB camera's spectral response simultaneously to calibrate the camera response with initial haemoglobin estimation.

Chapter 7

Conclusions

An early motivation was seen to exist in the manifold potential uses for clinical MSI, yet it was also apparent that in so far of it's clinical application and translatability. However, these applications and methods of data capture were relatively new and had many barriers in particular for clinical translation. These hurdles related to early devices having slow capture rates or low spectral specificity; large bulk; limited data processing methods; and sometimes onerous requirements for bespoke illumination or imaging protocols. In spite of these difficulties in performing MSI in an *in-vivo* setting, early clinical investigation had show interest in the capability of measuring perfusion and oxygen saturation across an entire image, rather then limited to a few discrete point measurements. As such the focus of the work contributing to this thesis has centred on: firstly enabling full frame image based measurement of THb and SO_2 using MSI techniques; and secondly doing this in a way that better facilitates the transfer to clinical use.

The opening part of this thesis relates to improving the performance of MSI in experimental conditions where long exposure times are needed and data is corrupted by significant motion blurring (Chapter 3). This work contributes a method of extending existing deblurring approaches, utilising parallel captured higher frame rate imaging to estimate optical flow and blur of the slower frame rate LCTF. Deblurring the MSI sequences enabled restoration of spectral fidelity and recovery of fine spatial details. Apart from MSI, this method could have applications for other optical imaging modalities such as fluorescence imaging where the low signal levels

necessitate long sensor integration times.

Robust and computationally efficient methods for estimating haemoglobin concentration have subsequently presented, making use of estimating latent spectral information from RGB data. The first contribution using RGB sensor data focused on speed of computation (Chapter 4) showing it was possible to rapidly and accurately estimate perfusion in two closed form steps. The first, estimating the MSI data from RGB constrained using a smoothness prior, and then second fitting haemoglobin concentrations to this estimate. This could be computed rapidly and was accurate for estimating the THb but was less precise for the SO_2 estimation. Developing on this design I then present the contribution of a more robust MSI data estimation scheme (Chapter 5) which alternately estimated MSI data and haemoglobin concentrations. This achieved a much higher precision in estimating both the THb and SO_2 due to the use of a Bayesian prior on the distribution of the MSI data, which in turn was derived from the current estimate of haemoglobin concentrations. The key limitation of this method, however, is that it is computationally expensive requiring more time to process the RGB data the to capture it and prohibiting real-time applications. Thus to unify the best aspects of these two approaches I contribute a method (Chapter 6) that elegantly combines both, notably by applying them to separate coefficients of the discrete Haar transformed image. The more accurate Bayesian method operates on the low pass residual coefficients while the faster method could be used on the directional coefficients comprising the majority of the transformed image data. Combining these two methods in such yields the computational speed into the range of real-time application without sacrificing estimation accuracy.

The net result is a framework that can accurately ($<$ 5.96g/litre error) estimate haemoglobin concentrations in real-time ($>$14Hz) and that requires no novel hardware to be installed. The main pre-requisite is a sensor response calibration to be taken of the RGB camera to be used, and a computer equipped with a modest GPU to process the video data. The main limitations of using a framework such as the one developed here are, discussed in the preceding chapters but can be grouped into

two collections. The first group of limitations introduced by this method are that: it requires a fixed illumination environment and assumes that the view direction is approximately normal to tissue; further that the sensor response calibration is accurate; and possibly most importantly it models the tissue reflectance spectra as strongly influenced by haemoglobin concentrations. This group of limitations are then the new ones that are introduced with my framework, distinct from general optical imaging, however in reality these are likely to be of low impact, or potentially solvable in future work. The illumination environment *in-vivo* can be well controlled especially during laparoscopic procedures there are often very few light sources. Where this would fail however, is if the surgeon needed to increase the brightness or other wise change the environment illumination, possibly for some other concurrent optical imaging task. It is possible although likely impractical to keep calibrations for a wide variety of illumination environments that could be switched between where necessary, however a more useful solution would be to develop a dynamic auto-calibration of the sensor and illumination environment using a statistical analysis of the past imaging data. Also acquiring a good sensor response calibration would be a relatively non-arduous limitation as it is can be performed via a variety of methods and is unlikely to change much over time. This could similarly benefit by an auto-calibration extension in the vein of [112] as removing the requirement for a skilled calibration would increase the number of people who could operate such a device. The laparoscopic view direction during surgery is often not normal to tissues especially during navigation phases, however there is already motivation to try to frame regions of interest for ease of operating. Typically it is only at extreme oblique angles that the losses due to internal reflection of the tissue surface would become significant, however further validation would be required to confirm this. Lastly the methods presented are model based methods and as such if there were scenarios that did not fit with that model then estimation would be invalid. While it is reasonable to assume that the dominant chromophores are those relating to haemoglobin, it is also reasonable to not that people who are in need of acute interventional help tend to have something wrong with them that deviates from nor-

mal bodily function. It is concealable that if a patient had a contaminant at high concentrations in their body it is possible it might have a significant effect on the estimation accuracy of this method for imaging those organs. One nice design of the framework presented however is while it only parametrises tissues via haemoglobin it would be straight forward to extend it with another known chromophores. Secondly as an optical method it is affected by the same issues that often trouble other optical imaging methods, such as difficulties with specularities, or sensor over and under exposure related issues. This second group of general limitations of optical imaging modalities, are to be expected with any optical approach and as the methods presented here don't directly mitigate there circumstances they could be expected to fail in the same scenarios.

As such the use a framework such as this would not impose any major requirement on clinicians to modify their work-flow or clinical hardware environment. It is likely to perform within it's limitations under most normal clinical scenarios, with the potential to become extended for greater robustness to changes in the surgical environment and tissue.

7.1 Future Work

Possible ways to expand the available high quality ground-truth data may be from the creation of better digital phantoms to simulate tissue light interaction [105] or from imaging organs *ex-vivo* sustained on artificial perfusion devices[49, 114] capable of controlled adjustment of perfusion and oxygen saturation. However, even with a small data set it could be possible to use transfer learning to train neural-net methods with success, and the work of [115] is an exciting initial step in this direction. For the model based approach, outstanding question is about tackling the calibration of the RGB camera sensor response, which would be a barrier for use outside of a finely controlled clinical setting. If it were possible to bootstrap the calibration of the illumination and camera sensor response from the initial few seconds or minutes of imaging of tissue; then it would open up the uses of ad-hoc imaging devices such as mobile phones, as well as uncontrolled for lighting environments

such as field clinics. This could potentially bring beneficial imaging capabilities to critical locals where standard hardware solutions might not be reachable.

7.2 Closing Remarks

Now more then ever there is the opportunity to measure vast tracts of patient data during procedures. Sources to name but a few such as rich video data streaming from laparoscopic devices; cameras positioned throughout the surgical theatre to track patient and clinician movement alike; chip-on tip instruments unrestricted in range of motion by an absence of optical channels. All of these individually present exciting opportunities to develop techniques to better patient outcomes, and taken together with the rapid development of neural-net based image processing tools could provide real-time meta analysis of a surgical procedure as a whole. At the other end of the spectrum in parts of the world where complex imaging set ups are not available such as in disaster and conflict response situations, there is potentially another rapidly advancing imaging technology that can be found on almost everyone: a smart phone. The improving quality and vast dissemination would make the targeting these types of device to be modified as ad-hoc imaging systems could have a potential significant impact in challenging emergency situations. Thus: with one eye to the highs of integrated imaging in surgical theatres and the other to need and potential of low cost devices in hostile setting, it is an exciting moment to be developing optical imaging solution to clinical questions.

Appendix A

Colophon

This is a description of the tools used to make this thesis. It helps people make future documents.

This document was set in the Times Roman typeface using LaTeX and BibTeX, composed with Visual Studio Code with the various LaTeX related extensions to provide colourisation and code completion. Compilation of the LaTeX sources were managed through the lightweight TeXworks interface supported by a MikTex back-end, later iterations of this document were composed using VScode using the Latex Workshop plugin this provided better multi-file management. The LaTeX style files and skeleton template were created by Ian Kirker[1]. The source text for this document were versioned using the git SCM tool, which provided a mechanism finding and comparing edits made across writing sessions.

Graphs and figures were produced with MATLAB. A particularly useful extension for this was the *tightsubplot.m* MATLAB extension. This extension streamlined the task of laying out plots with less redundant space and greater positional control. More complex figures were composed using a combination of Adobe Photoshop and Adobe Illustrator.

[1] i.kirker+GH@ucl.ac.uk

Bibliography

[1] V. Hacker. *Uber die Technik der Oesophagoskopie*, volume 9, pages 91–94]. Wien, Springer, 1896.

[2] Tim J. Adler, Lynton Ardizzone, Anant Vemuri, Leonardo Ayala, Janek Gröhl, Thomas Kirchner, Sebastian Wirkert, Jakob Kruse, Carsten Rother, Ullrich Köthe, and Lena Maier-Hein. Uncertainty-aware performance assessment of optical imaging modalities with invertible neural networks. *International Journal of Computer Assisted Radiology and Surgery*, March 2019.

[3] Boris Jansen-Winkeln, N. Holfert, H. Köhler, Y. Moulla, J. P. Takoh, S. M. Rabe, M. Mehdorn, M. Barberio, C. Chalopin, T. Neumuth, and I. Gockel. Determination of the transection margin during colorectal resection with hyperspectral imaging (HSI). *International Journal of Colorectal Disease*, 34(4):731–739, February 2019.

[4] Sebastian J. Wirkert, Hannes Kenngott, Benjamin Mayer, Patrick Mietkowski, Martin Wagner, Peter Sauer, Neil T. Clancy, Daniel S. Elson, and Lena Maier-Hein. Robust near real-time estimation of physiological parameters from megapixel multispectral images with inverse monte carlo and random forest regression. *Int J CARS*, 11(6):909–917, may 2016.

[5] Yu-Wing Tai, Hao Du, Michael S Brown, and Stephen Lin. Correction of spatially varying image and video motion blur using a hybrid camera. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(6):1012–1028, 2010.

[6] Moshe Ben-Ezra and Shree K Nayar. Motion deblurring using hybrid imaging. In *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*, volume 1, pages I–657. IEEE, 2003.

[7] Nienke Bosschaart, Gerda J. Edelman, Maurice C. G. Aalders, Ton G. van Leeuwen, and Dirk J. Faber. A literature review and novel theoretical approach on the optical properties of whole blood. *Lasers Med Sci*, 29(2):453–479, Oct 2013.

[8] Geoffrey Jones, Neil T. Clancy, Simon Arridge, Daniel S. Elson, and Danail Stoyanov. Inference of tissue haemoglobin concentration from stereo RGB. In *Medical Imaging and Augmented Reality*, pages 50–58. Springer International Publishing, 2016.

[9] Neil T. Clancy, Danail Stoyanov, David R. C. James, Aimee Di Marco, Vincent Sauvage, James Clark, G. Z. Yang, and Daniel S. Elson. Multispectral image alignment using a three channel endoscope in vivo during minimally invasive surgery. *Biomed. Opt. Express*, 3(10):2567–2578, Oct 2012.

[10] Steven L Jacques. Optical properties of biological tissues: a review. *Physics in Medicine and Biology*, 58(11):R37–R61, May 2013.

[11] V. Velanovich. Laparoscopic vs open surgery. *Surgical Endoscopy*, 14(1):16–21, jan 2000.

[12] Thomas E Ahlering, David Woo, Louis Eichel, David I Lee, Robert Edwards, and Douglas W Skarecky. Robot-assisted versus open radical prostatectomy: a comparison of one surgeon's outcomes. *Urology*, 63(5):819–822, may 2004.

[13] Jim C. Hu. Comparative effectiveness of minimally invasive vs open radical prostatectomy. *JAMA*, 302(14):1557, oct 2009.

[14] Surya SAY Biere, Mark I van Berge Henegouwen, Kirsten W Maas, Luigi Bonavina, Camiel Rosman, Josep Roig Garcia, Suzanne S Gisbertz, Jean HG

Klinkenbijl, Markus W Hollmann, Elly SM de Lange, H Jaap Bonjer, Donald L van der Peet, and Miguel A Cuesta. Minimally invasive versus open oesophagectomy for patients with oesophageal cancer: a multicentre, open-label, randomised controlled trial. *The Lancet*, 379(9829):1887–1892, may 2012.

[15] A. Kussmaul. *Ueber die Behandlung der Magenerweiterung durch eine neue Methode, mittelst der Magenpumpe*. Druck v. Poppen, 1869.

[16] Eric L. Matteson and Friedrich J. Kluge. Think clearly, be sincere, act calmly: Adolf kussmaul (february 22, 1822–may 28, 1902) and his relevance to medicine in the 21st century. *Current Opinion in Rheumatology*, 15(1):29–34, jan 2003.

[17] Neil T. Clancy, Danail Stoyanov, Vincent Sauvage, David James, Guang-Zhong Yang, and Daniel S. Elson. A triple endoscope system for alignment of multispectral images of moving tissue. In *Biomedical Optics and 3-D Imaging*, page BTuD27. Optical Society of America, 2010.

[18] Alfred Cuschieri. Whither minimal access surgery: tribulations and expectations. *The American Journal of Surgery*, 169(1):9–19, jan 1995.

[19] Clément-Claude Abbou, András Hoznek, Laurent Salomon, Leif Eric Olsson, Adrian Lobontiu, Fabien Saint, Antony Cicco, Patrick Antiphon, and Dominique Chopin. Laparoscopic Radical Prostatectomy With a Remote Controlled Robot. *The Journal of Urology*, 165(6):1964–1966, jun 2001.

[20] Sebastian Schostek, Marc O. Schurr, and Gerhard F. Buess. Review on aspects of artificial tactile feedback in laparoscopic surgery. *Medical Engineering & Physics*, 31(8):887–898, oct 2009.

[21] Danail Stoyanov. Surgical vision. *Annals of Biomedical Engineering*, 40(2):332–345, Oct 2011.

[22] Daniel J. Mirota, Masaru Ishii, and Gregory D. Hager. Vision-based navigation in image-guided interventions. *Annual Review of Biomedical Engineering*, 13(1):297–319, Aug 2011.

[23] Ara Darzi and Sean Mackay. Recent advances in minimal access surgery. *British Medical Journal*, 324(7328):31, 2002.

[24] Tobias Sielhorst, Marco Feuerstein, and Nassir Navab. Advanced medical displays: A literature review of augmented reality. *Journal of Display Technology*, 4(4):451–467, dec 2008.

[25] L. Maier-Hein, P. Mountney, A. Bartoli, H. Elhawary, Daniel. S. Elson, A. Groch, A. Kolb, M. Rodrigues, J. Sorger, S. Speidel, and D. Stoyanov. Optical techniques for 3d surface reconstruction in computer-assisted laparoscopic surgery. *Medical Image Analysis*, 17(8):974–996, dec 2013.

[26] Bongjin Koo, Erol zgr, Bertrand Le Roy, Emmanuel Buc, and Adrien Bartoli. Deformable registration of a preoperative 3d liver volume to a laparoscopy image using contour and shading cues. In *Medical Image Computing and Computer Assisted Intervention - MICCAI 2017*, pages 326–334. Springer International Publishing, 2017.

[27] Archie Hughes-Hallett, Philip Pratt, Erik Mayer, Aimee Di Marco, Guang-Zhong Yang, Justin Vale, and Ara Darzi. Intraoperative ultrasound overlay in robot-assisted partial nephrectomy: First clinical experience. *European Urology*, 65(3):671–672, mar 2014.

[28] Scott Tobis, Joy K. Knopf, Christopher Silvers, Edward Messing, Jorge Yao, Hani Rashid, Guan Wu, and Dragan Golijanin. Robot-assisted and laparoscopic partial nephrectomy with near infrared fluorescence imaging. *Journal of Endourology*, 26(7):797–802, jul 2012.

[29] Kazuhiro Gono, Takashi Obi, Masahiro Yamaguchi, Nagaaki Ohyama, Hirohisa Machida, Yasushi Sano, Shigeaki Yoshida, Yasuo Hamamoto, and Takao

Endo. Appearance of enhanced tissue features in narrow-band endoscopic imaging. *Journal of Biomedical Optics*, 9(3):568, 2004.

[30] Michail A. Ilias, Erik Hggblad, Chris Anderson, and E. Gran Salerud. Visible hyperspectral imaging evaluating the cutaneous response to ultraviolet radiation, 2007.

[31] Brian S. Sorg, Owen Donovan, Yiting Cao, Mark W. Dewhirst, and Benjamin J. Moeller. Hyperspectral imaging of hemoglobin saturation in tumor microvasculature and tumor hypoxia development. *Journal of Biomedical Optics*, 10(4):044004–044004–11, 2005.

[32] Michael G. Sowa, Jeri R. Payette, Mark D. Hewko, and Henry H. Mantsch. Visible-near infrared multispectral imaging of the rat dorsal skin flap. *Journal of Biomedical Optics*, 4(4):474–481, 1999.

[33] Guolan Lu and Baowei Fei. Medical hyperspectral imaging: a review. *J. Biomed. Opt*, 19(1):010901, jan 2014.

[34] Geoffrey Jones, Neil T. Clancy, Simon Arridge, Dan S. Elson, and Danail Stoyanov. Deblurring multispectral laparoscopic images. In *Information Processing in Computer-Assisted Interventions*, pages 216–225. Springer International Publishing, 2014.

[35] Geoffrey Jones, Neil T. Clancy, Yusuf Helo, Simon Arridge, Daniel S. Elson, and Danail Stoyanov. Bayesian estimation of intrinsic tissue oxygenation and perfusion from RGB images. *IEEE Transactions on Medical Imaging*, pages 1–1, 2017.

[36] Geoffrey Jones, Neil T. Clancy, Xiaofei Du, Maria Robu, Simon Arridge, Daniel S. Elson, and Danail Stoyanov. Fast estimation of haemoglobin concentration in tissue via wavelet decomposition. In *Lecture Notes in Computer Science*, pages 100–108. Springer International Publishing, 2017.

[37] Brian S Sorg, Benjamin J Moeller, Owen Donovan, Yiting Cao, and Mark W Dewhirst. Hyperspectral imaging of hemoglobin saturation in tumor microvasculature and tumor hypoxia development. *Journal of biomedical optics*, 10(4):044004–044004, 2005.

[38] Neil T. Clancy, Srdjan Saso, Danail, Vincent Sauvage, David J. Corless, Michael Boyd, David E. Noakes, Meen-Yau Thum, Sadaf Ghaem-Maghami, J. R. Smith, and Daniel S. Elson. Multispectral imaging of organ viability during uterine transplantation surgery. In *Advanced Biomedical and Clinical Diagnostic Systems XII*. SPIE-Intl Soc Opt Eng, Feb 2014.

[39] Danail Stoyanov, Aleksandr Rayshubskiy, and Elizabeth Hillman. Robust registration of multispectral images of the cortical surface in neurosurgery. In *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, pages 1643–1646. IEEE, 2012.

[40] F. Wang, A. Behrooz, M. Morris, and A. Adibia. High-contrast subcutaneous vein detection and localization using multispectral imaging. *Journal of Biomedical Optics*, 18(5):050504, 2013.

[41] Neil T. Clancy, K. Gurusamy, G. Jones, B. Davidson, M. J. Clarkson, D. J. Hawkes, and D. Stoyanov. Intraoperative optical characterisation of thermal ablation. In *The Hamlyn Symposium on Medical Robotics*, pages 41–42.

[42] Karel J. Zuzak, Sabira C. Naik, George Alexandrakis, Doyle Hawkins, Khosrow Behbehani, and Edward H. Livingston. Characterization of a near-infrared laparoscopic hyperspectral imaging system for minimally invasive surgery. *Analytical Chemistry*, 79(12):4709–4715, 2007.

[43] Neil T. Clancy, Shobhit Arya, Danail Stoyanov, Mohan Singh, George B. Hanna, and Daniel S. Elson. Intraoperative measurement of bowel oxygen saturation using a multispectral imaging laparoscope. *Biomedical Optics Express*, 6(10):4179, Sep 2015.

[44] Ephrem O. Olweny, Stephen Faddegon, Sara L. Best, Neil Jackson, Eleanor F. Wehner, Yung K. Tan, Karel J. Zuzak, and Jeffrey A. Cadeddu. First place: Renal oxygenation during robot-assisted laparoscopic partial nephrectomy: Characterization using laparoscopic digital light processing hyperspectral imaging. *Journal of Endourology*, 27(3):265–269, mar 2013.

[45] Yasser Fawzy, Stephen Lam, and Haishan Zeng. Rapid multispectral endoscopic imaging system for near real-time mapping of the mucosa blood supply in the lung. *Biomedical Optics Express*, 6(8):2980, 2015.

[46] Y. Ma, M. A. Shaik, S. H. Kim, M. G. Kozberg, D. N. Thibodeaux, H. T. Zhao, H. Yu, and E. M. C. Hillman. Wide-field optical mapping of neural activity and brain haemodynamics: considerations and novel approaches. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 371(1705):20150360, 2016.

[47] Neil T. Clancy, Julian Teare, George B. Hanna, and Daniel S. Elson. Flexible multimode endoscope for tissue reflectance and autofluorescence hyperspectral imaging. In *Biomedical Optics 2016*. The Optical Society, 2016.

[48] R. Kumashiro, K. Konishi, T. Chiba, T. Akahoshi, S. Nakamura, M. Murata, M. Tomikawa, T. Matsumoto, Y. Maehara, and M. Hashizume. Integrated endoscopic system based on optical imaging and hyperspectral data analysis for colorectal cancer detection. *Anticancer Research*, 36(8):3925–3932, 2016.

[49] F. Tetschke, W. Markgraf, M. Gransow, S. Koch, C. Thiele, A. Kulcke, and H. Malberg. Hyperspectral imaging for monitoring oxygen saturation levels during normothermic kidney perfusion. *Journal of Sensors and Sensor Systems*, 5(2):313–318, 2016.

[50] N. Hagen and M. W. Kudenov. Review of snapshot spectral imaging technologies. *Optical Engineering*, 52(9):090901, 2013.

[51] A. S. Luthman, D. J. Waterhouse, L. Ansel-Bollepalli, J. Yoon, G. S. D. Gordon, J. Joseph, M. di Pietro, W. Januszewicz, and S. E. Bohndiek. Bimodal reflectance and fluorescence multispectral endoscopy based on spectrally resolving detector arrays. *Journal of Biomedical Optics*, 24(3):031009, 2018.

[52] E. L. Wisotzky, B. Kossack, F. C. Uecker, S. Arens, P. Dommerich, A. Hilsmann, and P. Eisert. Validation of two techniques for intraoperative hyperspectral human tissue determination. In B. Fei and C. A. Linte, editors, *Proc. of SPIE*, volume 10951 of *Medical Imaging 2019: Image-Guided Procedures, Robotic Interventions, and Modeling*, page 109511Z.

[53] S.J. Preece and E. Claridge. Spectral filter optimization for the recovery of parameters which describe human skin. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(7):913–922, jul 2004.

[54] Tobias C. Wood, Surapa Thiemjarus, Kevin R. Koh, Daniel S. Elson, and Guang-Zhong Yang. *Optimal Feature Selection Applied to Multispectral Fluorescence Imaging*, pages 222–229. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[55] Sebastian J. Wirkert, Neil T. Clancy, Danail Stoyanov, Shobhit Arya, George B. Hanna, Heinz-Peter Schlemmer, Peter Sauer, Daniel S. Elson, and Lena Maier-Hein. *Endoscopic Sheffield Index for Unsupervised In Vivo Spectral Band Selection*, pages 110–120. Springer International Publishing, Cham, 2014.

[56] Kazuki Hashimoto, Hikaru Mizuno, Keiichi Nakagawa, Ryoichi Horisaki, Atsushi Iwasaki, Fumihiko Kannari, Ichiro Sakuma, and Keisuke Goda. High-speed multispectral videography with a periscope array in a spectral shaper. *Optics Letters*, 39(24):6942, dec 2014.

[57] Stephen P. Nighswander-Rempel, R. Anthony Shaw, James R. Mansfield, Mark Hewko, Valery V. Kupriyanov, and Henry H. Mantsch. Regional variations in myocardial tissue oxygenation mapped by near-infrared spectrosco-

pic imaging. *Journal of Molecular and Cellular Cardiology*, 34(9):1195 – 1203, 2002.

[58] Raphael Sznitman, Seth Billings, Diego Rother, Daniel Mirota, Yi Yang, Jim Handa, Peter Gehlbach, Jin U. Kang, Gregory D. Hager, and Russell Taylor. Active multispectral illumination and image fusion for retinal microsurgery. In *Information Processing in Computer-Assisted Interventions*, pages 12–22. Springer Science Business Media, 2010.

[59] Ela Claridge, Džena Hidović-Rowe, Phillipe Taniere, and Tariq Ismail. Quantifying mucosal blood volume fraction from multispectral images of the colon. In *Medical Imaging*, pages 65110C–65110C. International Society for Optics and Photonics, 2007.

[60] Matthew B Bouchard, Brenda R Chen, Sean A Burgess, and Elizabeth Hillman. Ultra-fast multispectral optical imaging of cortical oxygenation, blood flow, and intracellular calcium dynamics. *Optics express*, 17(18):15670–15678, 2009.

[61] Karel J. Zuzak, Robert P. Francis, Eleanor F. Wehner, Maritoni Litorja, Jeffrey A. Cadeddu, and Edward H. Livingston. Active DLP hyperspectral illumination: A noninvasive, in vivo, system characterization visualizing tissue oxygenation at near video rates. *Analytical Chemistry*, 83(19):7424–7430, oct 2011.

[62] Neil T. Clancy, Danail Stoyanov, Lena Maier-Hein, Anja Groch, Guang-Zhong Yang, and Daniel S. Elson. Spectrally encoded fiber-based structured lighting probe for intraoperative 3d imaging. *Biomedical Optics Express*, 2(11):3119, oct 2011.

[63] Jianyu Lin, Neil T. Clancy, and Daniel S. Elson. An endoscopic structured light system using multispectral detection. *Int J CARS*, 10(12):1941–1950, jul 2015.

[64] Irene M. Orfanoudaki, George C. Themelis, Stavros K. Sifakis, Despina H. Fragouli, John G. Panayiotides, Eleftheria M. Vazgiouraki, and Eugenios E. Koumantakis. A clinical study of optical biopsy of the uterine cervix using a multispectral imaging system. *Gynecologic Oncology*, 96(1):119–131, jan 2005.

[65] Izumi Nishidate, Takaaki Maeda, Kyuichi Niizeki, and Yoshihisa Aizu. Estimation of melanin and hemoglobin using spectral reflectance images reconstructed from a digital RGB image by the wiener estimation method. *Sensors*, 13(6):7902–7915, Jun 2013.

[66] Vasileios Kavvadias, George Epitropou, Nikos Georgiou, Fani Grozou, Minas Paschopoulos, and Costas Balas. A novel endoscopic spectral imaging platform integrating k-means clustering for early and non-invasive diagnosis of endometrial pathology. In *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. Institute of Electrical & Electronics Engineers (IEEE), jul 2013.

[67] Neil T. Clancy, Shobhit Arya, Ji Qi, Danail Stoyanov, George B. Hanna, and Daniel S. Elson. Polarised stereo endoscope and narrowband detection for minimal access surgery. *Biomedical Optics Express*, 5(12):4108, nov 2014.

[68] Alessandro R. Guazzi, Mauricio Villarroel, João Jorge, Jonathan Daly, Matthew C. Frise, Peter A. Robbins, and Lionel Tarassenko. Non-contact measurement of oxygen saturation with an RGB camera. *Biomedical Optics Express*, 6(9):3320, aug 2015.

[69] Wenfeng Xia, Daniil I. Nikitichev, Jean Martial Mari, Simeon J. West, Rosalind Pratt, Anna L. David, Sebastien Ourselin, Paul C. Beard, and Adrien E. Desjardins. Performance characteristics of an interventional multispectral photoacoustic imaging system for guiding minimally invasive procedures. *J. Biomed. Opt*, 20(8):086005, aug 2015.

[70] Martin Hohmann, R. Kanawade, F. Klämpfl, A. Douplik, J. Mudter, M. F. Neurath, and H. Albrecht. In-vivomultispectral video endoscopy towardsin-vivohyperspectral video endoscopy. *J. Biophoton*, jul 2016.

[71] Yan Zhang, Sebastian J. Wirkert, Justin Iszatt, Hannes Kenngott, Martin Wagner, Benjamin Mayer, Christian Stock, Neil T. Clancy, Daniel S. Elson, and Lena Maier-Hein. Tissue classification for laparoscopic image understanding based on multispectral texture analysis. In Robert J. Webster and Ziv R. Yaniv, editors, *Medical Imaging 2016: Image-Guided Procedures, Robotic Interventions, and Modeling*. SPIE-Intl Soc Optical Eng, mar 2016.

[72] Seoung Wug Oh, Michael S. Brown, Marc Pollefeys, and Seon Joo Kim. Do it yourself hyperspectral imaging with everyday digital cameras. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

[73] A Villringer. Non-invasive optical spectroscopy and imaging of human brain function. *Trends in Neurosciences*, 20(10):435–442, oct 1997.

[74] Boulpaep EL. *Medical physiology*, chapter 18: Blood, page 448481. Boron F, Boulpaep EL Saunders, Philadelphia, 2 edition, 2009.

[75] Hsing-Wen Wang, Timothy C. Zhu, Mary E. Putt, Michael Solonenko, James Metz, Andreea Dimofte, Jeremy Miles, Douglas L. Fraker, Eli Glatstein, Stephen M. Hahn, and Arjun G. Yodh. Broadband reflectance measurements of light penetration, blood oxygenation, hemoglobin concentration, and drug concentration in human intraperitoneal tissues before and after photodynamic therapy. *Journal of Biomedical Optics*, 10(1):014004, 2005.

[76] B.C. Wilson and S.L. Jacques. Optical reflectance and transmittance of tissues: principles and applications. *IEEE Journal of Quantum Electronics*, 26(12):2186–2199, 1990.

[77] Georgi I. Petrov, Alexander Doronin, Harry T. Whelan, Igor Meglinski, and Vladislav V. Yakovlev. Human tissue color as viewed in high dynamic

range optical spectral transmission measurements. *Biomed. Opt. Express*, 3(9):2154–2161, Sep 2012.

[78] J.G. Webster. *Design of Pulse Oximeters*. Series in Medical Physics and Biomedical Engineering. CRC Press, 1997.

[79] Valery Tuchin. Tissue optics and photonics: Biological tissue structures. *Journal of Biomedical Photonics & Engineering*, 1(1):3–21, 2015.

[80] Valery Tuchin. Tissue optics and photonics: Light-tissue interaction. *Journal of Biomedical Photonics & Engineering*, 1(2):98–134, 2015.

[81] Martin Hohmann, Alexandre Douplik, Janani Varadhachari, Aulia Nasution, Jonas Mudter, Markus Neurath, and Michael Schmidt. Preliminary results for hyperspectral videoendoscopy diagnostics on the phantoms of normal and abnormal tissues: towards gastrointestinal diagnostics. In Nirmala Ramanujam and Jürgen Popp, editors, *Clinical and Biomedical Spectroscopy and Imaging II*. SPIE-Intl Soc Optical Eng, jun 2011.

[82] Anne Pigula, Neil T. Clancy, Shobhit Arya, George B. Hanna, and Daniel S. Elson. Video-rate dual polarization multispectral endoscopic imaging. In Adam Wax and Vadim Backman, editors, *Biomedical Applications of Light Scattering IX*. SPIE-Intl Soc Optical Eng, mar 2015.

[83] I.V. Meglinski and S.J. Matcher. Computer simulation of the skin reflectance spectra. *Computer Methods and Programs in Biomedicine*, 70(2):179–186, feb 2003.

[84] G de Haan and A van Leest. Improved motion robustness of remote-PPG by using the blood volume pulse signature. *Physiol. Meas.*, 35(9):1913–1926, aug 2014.

[85] Xiaofei Du, Neil T. Clancy, Shobhit Arya, George B. Hanna, John Kelly, Daniel S. Elson, and Danail Stoyanov. Robust surface tracking combining

features, intensity and illumination compensation. *International Journal of Computer Assisted Radiology and Surgery*, 10(12):1915–1926, June 2015.

[86] Raimund Leitner, Martin De Biasio, Thomas Arnold, Cuong Viet Dinh, Marco Loog, and Robert P.W. Duin. Multi-spectral video endoscopy system for the detection of cancerous tissue. *Pattern Recognition Letters*, 34(1):85 – 93, 2013.

[87] W. H. Richardson. Bayesian-based iterative method of image restoration. *J. Opt. Soc. Am.*, 62(1):55–59, Jan 1972.

[88] LB Lucy. An iterative technique for the rectification of observed distributions. *The astronomical journal*, 79:745, 1974.

[89] N. Dey, L. Blanc-Feraud, C. Zimmer, Z. Kam, J. C. Olivo-Marin, and J. Zerubia. A deconvolution method for confocal microscopy with total variation regularization. In *Biomedical Imaging: Nano to Macro, 2004. IEEE International Symposium on*, pages 1223–1226 Vol. 2, 2004.

[90] Rob Fergus, Barun Singh, Aaron Hertzmann, Sam T Roweis, and William T Freeman. Removing camera shake from a single photograph. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 787–794. ACM, 2006.

[91] Dilip Krishnan, Terence Tay, and Rob Fergus. Blind deconvolution using a normalized sparsity measure. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 233–240. IEEE, 2011.

[92] James Gregson, Felix Heide, Matthias B. Hullin, Mushfiqur Rouf, and Wolfgang Heidrich. Stochastic Deconvolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, page (to appear), Jun 2013.

[93] Feng Li, Jingyi Yu, and Jinxiang Chai. A hybrid camera for motion deblurring and depth map super-resolution. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.

[94] Timothy J. Holmes. Blind deconvolution of quantum-limited incoherent imagery: maximum-likelihood approach. *J. Opt. Soc. Am. A*, 9(7):1052–1061, Jul 1992.

[95] DA Fish, AM Brinicombe, ER Pike, and JG Walker. Blind deconvolution by means of the richardson–lucy algorithm. *JOSA A*, 12(1):58–65, 1995.

[96] Y.W. Tai, P.Tan, and M.S. Brown. Richardson-lucy deblurring for scenes under a projective motion path. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2009.

[97] Lu Yuan, Jian Sun, Long Quan, and Heung-Yeung Shum. Progressive inter-scale and intra-scale non-blind image deconvolution. *ACM Trans. Graph.*, 27(3):74:1–74:10, Aug 2008.

[98] D T Delpy, M Cope, P van der Zee, S Arridge, S Wray, and J Wyatt. Estimation of optical pathlength through tissue from direct time of flight measurement. *Physics in Medicine and Biology*, 33(12):1433–1442, Dec 1988.

[99] C. E. Thorn, S. J. Matcher, I. V. Meglinski, and A. C. Shore. Is mean blood saturation a useful marker of tissue oxygenation? *AJP: Heart and Circulatory Physiology*, 296(5):H1289–H1295, mar 2009.

[100] Andrew K. Dunn, Anna Devor, Anders M. Dale, and David A. Boas. Spatial extent of oxygen metabolism and hemodynamic changes during functional activation of the rat somatosensory cortex. *NeuroImage*, 27(2):279–290, Aug 2005.

[101] Rasmus Bro and Sijmen De Jong. A fast non-negativity-constrained least squares algorithm. *J. Chemometrics*, 11(5):393–401, Sep 1997.

[102] EMVA standard 1288: *Standard for Characterization of Image Sensors and Cameras*. *European Machine Vision Association*, (3.0), 2010.

[103] Jon Y. Hardeberg, Hans Brettel, and Francis J. M. Schmitt. Spectral characterization of electronic cameras. In Jan Bares, editor, *Electronic Imaging:*

Processing, Printing, and Publishing in Color. SPIE-Intl Soc Optical Eng, Sep 1998.

[104] Lihong Wang, Steven L Jacques, and Liqiong Zheng. Mcmlmonte carlo modeling of light transport in multi-layered tissues. *Computer methods and programs in biomedicine*, 47(2):131–146, 1995.

[105] Qianqian Fang. Mesh-based monte carlo method using fast ray-tracing in plücker coordinates. *Biomed. Opt. Express*, 1(1):165, 2010.

[106] A N Bashkatov, E A Genina, V I Kochubey, V S Rubtsov, E A Kolesnikova, and V V Tuchin. Optical properties of human colon tissues in the 350-2500 nm spectral range. *Quantum Electronics*, 44(8):779–784, Aug 2014.

[107] Jean-Luc Boulnois. Photophysical processes in recent medical laser developments: A review. *Lasers in Medical Science*, 1(1):47–66, 1986.

[108] V.A. Morozov. Regularization of incorrectly posed problems and the choice of regularization parameter. *USSR Computational Mathematics and Mathematical Physics*, 6(1):242–251, Jan 1966.

[109] Danny Pascale. Color checker data., 2012, Accessed: May 2016.

[110] Ingrid Daubechies. *Ten Lectures on Wavelets*. Society for Industrial & Applied Mathematics (SIAM), jan 1992.

[111] Alfred Haar. Zur theorie der orthogonalen funktionensysteme. *Mathematische Annalen*, 69(3):331–371, September 1910.

[112] Maria Ruxandra Robu, Matthieu Leclerc-Chalvet, and Danail Stoyanov. Radiometric calibration for a spectrophotometric analysis pipeline for assessing burns. CRAS, sep 2015.

[113] KASEY P Bensky, Linda Donahue-Spencer, G Erik Hertz, and MI Anderson. The dose-related effects of bolus esmolol on heart rate and blood pressure following laryngoscopy and intubation. *AANA journal*, 68(5):437–443, 2000.

[114] Richard W Laing, Hynek Mergental, Christina Yap, Amanda Kirkham, Manpreet Whilku, Darren Barton, Stuart Curbishley, Yuri L Boteon, Desley A Neil, Stefan G Hübscher, M Thamara P R Perera, Paolo Muiesan, John Isaac, Keith J Roberts, Hentie Cilliers, Simon C Afford, and Darius F Mirza. Viability testing and transplantation of marginal livers (vittal) using normothermic machine perfusion: study protocol for an open-label, non-randomised, prospective, single-arm trial. *BMJ Open*, 7(11), 2017.

[115] Sebastian J. Wirkert, Anant S. Vemuri, Hannes G. Kenngott, Sara Moccia, Michael Götz, Benjamin F. B. Mayer, Klaus H. Maier-Hein, Daniel S. Elson, and Lena Maier-Hein. Physiological parameter estimation from multispectral images unleashed. In Maxime Descoteaux, Lena Maier-Hein, Alfred Franz, Pierre Jannin, D. Louis Collins, and Simon Duchesne, editors, *Medical Image Computing and Computer Assisted Intervention MICCAI 2017*, pages 134–141, Cham, 2017. Springer International Publishing.