

University of Southampton Research Repository
ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON

**Imaging Through A Scanner, Darkly:
Spectral Imaging for Sentinel Lymph
Node Biopsies**

by

Jack D. O'Sullivan

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the
Faculty of Physical and Applied Sciences
Optoelectronics Research Centre

April 2012

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
OPTOELECTRONICS RESEARCH CENTRE

Doctor of Philosophy

**Imaging Through A Scanner, Darkly: Spectral Imaging for Sentinel Lymph
Node Biopsies**

by Jack D. O'Sullivan

Breast cancer is the single most prevalent form of cancer in the UK today, accounting for around 16% of all diagnoses, and around 31% of diagnoses in women. The survival rates are good, however the prognosis is heavily dependent on the stage to which the cancer has progressed at diagnosis. In order to help accurately determine this stage, the sentinel lymph node of patients undergoing tumour resection surgery is removed and examined cytologically for the presence of cancerous cells. This examination of the lymph node is currently the rate-limiting step in the operation as a whole. There is evidence in the literature to suggest that cancerous tissue has a measurably different infrared spectrum from healthy tissue owing to chemical and morphological differences in the tissue. There is further evidence to suggest that in the visible and near infrared region, the spectra of healthy lymph node tissue is different from that of cancerous tissue. This thesis details a project, performed in collaboration with a surgical team at St Mary's Hospital, Newport, Isle of Wight, to analyse spectral images taken in the visible and near infrared, of biopsied lymph node tissue. In the course of the project, an unsupervised classification technique, based on an extension to the well established 'spectral angle', was developed to analyse the spectral images.

Psoriasis affects 2-3% of the UK population causing itchy and/or painful plaques on the skin. One of the main treatments for psoriasis is UV phototherapy, exposure to which is a risk factor for burning and the development of cancers. This thesis details an investigation into the possibility of developing a targeted UV phototherapy system based on spectral imaging to delineate the plaques and a proposed new UV laser for treatment.

Contents

Acknowledgements	xviii
Abbreviations and selected terms	xxi
1 Introduction	1
1.1 Brief Project Overview & Summary of Prior Work	1
1.2 Thesis Structure	3
1.3 Photon Interactions with Biological Tissue	3
1.3.1 Absorption	3
1.3.2 Scattering	5
1.3.3 Key Chromophores	8
1.4 Spectral Imaging	11
1.4.1 What Is Spectral Imaging?	11
1.4.2 Technologies and Applications	13
1.4.3 The Curse of Dimensionality	17
2 Skin Damage Imaging and Treatment	21
2.1 Background	21
2.1.1 A Brief Introduction to Psoriasis	22
2.1.2 Medical Review	23
2.1.3 UV Laser Generation	25
2.2 System Requirements	30
2.2.1 Imaging System	30
2.2.2 Laser System	33
2.3 Discussion	35
3 Pixel Classification	37
3.1 Image Processing Concepts	37
3.1.1 Difference Mapping	37
3.1.2 Ratio Mapping	38
3.2 Measuring Spectral Similarity	39
3.2.1 Distance Measures	40
3.2.2 Comparing Spectral Similarity Measures	43
3.3 Extending the Spectral Angle	49
3.3.1 Limitation of the Spectral Angle	50
3.3.2 Numerical Degeneracy - The Other Curse of Dimensionality	50
3.3.3 Azimuthal Angles in Three Dimensions	51
3.3.4 The Spectral Position Map in Many Dimensions	52

3.4	Established Alternatives - K-Means Classifiers and Principal Components Analysis	54
3.4.1	The K-Means Classifier	54
3.4.2	Principal Components Analysis	55
3.5	Conclusions	56
4	Testing and Validating the Spectral Position Map	59
4.1	A Demonstration Using Specially Generated Images	59
4.2	A Demonstration in Three Dimensions - RGB Images	68
4.3	A Demonstration with Spectral Images	75
4.4	A Comparison to Established Alternatives	83
4.5	Conclusions	85
5	Towards an Optical Sentinel Lymph Node Biopsy	89
5.1	Background	90
5.1.1	Breast Cancer and the Sentinel Lymph Node Biopsy	90
5.1.2	Medical Review	94
5.1.2.1	The Efficacy of the SLNB	94
5.1.2.2	Spectroscopy and Cancer	96
5.1.2.3	Spectroscopy and the Lymph Node	102
5.2	System Development	105
5.2.1	Image Capture Systems	106
5.2.1.1	Multispectral Imager	107
5.2.1.2	Hyperspectral Imager	109
5.2.1.3	System Set-Up	111
5.2.2	Data Processing Suite	117
5.2.2.1	Pre-Processing Routines	117
5.2.2.2	Real Processing Routines	120
5.3	Results	125
5.3.1	Test Hyperspectral Results	125
5.3.2	Real Hyperspectral Results	134
5.3.3	Multispectral Results	147
5.4	Conclusions	153
6	Discussion	163
6.1	A Review of Procedure	164
6.1.1	Illumination Issues	164
6.1.1.1	Use of an AC Lamp	164
6.1.1.2	Fluorescent Lighting Artefacts	166
6.1.2	Image Calibration	167
6.1.3	Appraisal of the System with the Analysis Algorithm	168
6.2	The Extended Spectral Angle Metric and Classification System	169
6.2.1	Future Directions	170
6.3	The Optical Sentinel Lymph Node Biopsy	172
6.3.1	Future Directions	173
A	Publications	177
A.1	CLEO 2011 - Baltimore, Maryland, May 2011	179

A.2 OSA/SPIE European Conferences on Biomedical Optics - Munich, Germany, May 2011	182
A.2.1 Conference Proceedings Paper for ECBO, Munich, May 2011	182
A.2.2 Poster Presented at ECBO, Munich, May 2011	190
B Further Hyperspectral Images	191
C High Dimensional Image Formats	215
C.1 HDI	215
C.1.1 File Header	215
C.1.2 Table of Contents	216
C.1.3 Image Data Block	217
C.1.4 Metric Data Block	217
C.1.5 Class Data Block	217
C.1.6 Mask Data Block	218
C.1.7 Channel Data Block	218
C.2 HDIx	218
C.2.1 File Header	218
C.2.2 Table of Contents	219
C.2.3 Image Data Block	220
C.2.4 Metric Data Block	220
C.2.5 Class Data Block	220
C.2.6 Mask Data Block	221
C.2.7 Channel Data Block	221
Bibliography	223
CD containing source code for all software	Enclosed

List of Figures

1.1	The spectral responsivity curve for a Sony Ex-View CCD	2
1.2	Relative intensities of scattered light in various directions for Rayleigh scattering, Mie scattering and Mie scattering from large particles	5
1.3	Possible photon paths through a scattering and absorbing material	7
1.4	Energy level diagram of scattering events	7
1.5	Absorption spectrum for haemoglobin	8
1.6	Absorption spectrum for water	9
1.7	Absorption spectrum for lipids	10
1.8	Absorption spectrum for melanin	10
1.9	Bayer pattern for CCDs	12
1.10	An illustrative example of the difference between multi- and hyper- spectral imaging	14
1.11	Focal Plane Scanning schematic	15
1.12	Line Scanning schematic	16
1.13	The ratio of the volume of a hypersphere to a hypercube in n dimensions	19
2.1	A <i>psoriasis vulgaris</i> plaque	22
2.2	Human erythema action curve, 280-400nm	25
2.3	Process flow diagram for treatment	32
2.4	Potential layouts of imaging system	33
3.1	Geometry of an arc and the Euclidean distance between unit vectors . . .	43
3.2	Graphic representation of the vectors from table 3.1	45
3.3	Graphic representation of the vectors from table 3.2	46
3.4	A graph showing the change in spectral angle, spectral information divergence and Bhattacharyya co-efficient with systematically deformed test vectors	47
3.5	A graph showing the derivatives of the plots from figure 3.4	48
3.6	A vector \mathbf{R} in three dimensions with a cone of vectors \mathbf{T} subtending an angle θ around it	52
3.7	The scene from figure 3.6 after the co-ordinate space has been rotated .	53
4.1	An image showing how the points in co-ordinate space correspond to the spatial pixels in the 3D test image, oriented such that the reference vector is orthogonal to the page	60
4.2	An RGB representation of a computer generated three plane multispectral test image	61

4.3	An image showing the spatial pixels in the test image as points in a 3D space with axes rotated to demonstrate the shape of the distribution of the points	62
4.4	An image showing the spatial pixels in the test image as points in a 3D space with axes rotated such that the reference vector is orthogonal to the plane of the page	63
4.5	The SAM for the 3D test image	64
4.6	The SPM for the 3D test image	64
4.7	Scale bar for SPMs	65
4.8	The SPM for the 5D test image	66
4.9	The SPM for the 6D test image	67
4.10	Intensity profiles of the planes of the SPM of a 5D test image	67
4.11	Intensity profiles of the planes of the SPM of a 6D test image	68
4.12	Lena - A staple image processing target	69
4.13	The SAM for the Lena image	69
4.14	Histogram of spectral angles from the Lena image	70
4.15	Gaussian peaks automatically fitted to the histogram of spectral angles from the Lena image	71
4.16	Spectral angle classified version of the Lena image	72
4.17	Scale bar showing class numbers and the colours assigned to them	72
4.18	The planes of the SPM of the Lena image	73
4.19	RGB image of some snooker balls	73
4.20	Spectral angle classified version of the snooker balls image	74
4.21	The planes of the SPM of the snooker balls image	74
4.22	Pseudo-colour image of a countryside landscape spectral image	76
4.23	Spectral angle classified version of the landscape image	77
4.24	2D histograms of spectral angle against each plane of the SPM for the landscape image	78
4.25	2D histogram of spectral angle against the first plane of the SPM for the landscape image	79
4.26	Partially classified version of the landscape image, classified using the spectral angle and SPM	80
4.27	Pseudo-colour image of a colour wheel target spectral image	81
4.28	Spectral angle classified version of the colour wheel target image	81
4.29	2D histogram of spectral angle against the first plane of the SPM for the colour wheel image	82
4.30	Partially classified version of the colour wheel image, classified using the spectral angle and SPM	83
4.31	ENVI K-Means classified version of the landscape image	84
4.32	The first principal component of the landscape image	85
4.33	The second, third and fourth principal component of the landscape image	86
4.34	ENVI K-Means classification of the principal components of the landscape image	87
5.1	The cut surface of a butterflied lymph node	93
5.2	the cut surface of a butterflied lymph node	93
5.3	Photograph of the mark I spectral imaging system	112
5.4	Schematic diagram of the mark I spectral imaging system	113

5.5	Schematic diagram of the mark II spectral imaging system	114
5.6	Photograph of the mark II spectral imaging system	115
5.7	Flow chart showing the operation of the image acquisition system	116
5.8	Output spectrum from a spectral image of a blank target illuminated by fluorescent strip lights and an argon lamp	119
5.9	SAMs for two dummy spectral images	127
5.10	Histogram of spectral angles for the original dummy image	127
5.11	Histogram of spectral angles for the dummy image with 10% noise level .	128
5.12	Histogram of spectral angles for the dummy image with 20% noise level .	128
5.13	Spectral angle classified version of the dummy images with 5% and 20% noise levels	129
5.14	Graph showing number of misclassified pixels against noise level of the dummy image	129
5.15	Average spectra of normal and metastatic classes of classified dummy images at 5% and 20% noise level	130
5.16	Graph showing the position of the Gaussian peaks fitted to the histogram of spectral angles for each of the dummy images	131
5.17	2D histogram of spectral angle against dominant plane of the SPM of the dummy image with 20% noise level	132
5.18	Classified version of the dummy image with 20% noise level, classified by spectral angle and SPM	132
5.19	Graph showing number of misclassified pixels against noise level of the dummy image (as classified by spectral angle and SPM)	133
5.20	Pseudo-colour image of a hyperspectral image of an SLN	135
5.21	Histogram of spectral angles from the hyperspectral image of an SLN .	136
5.22	Gaussian peaks automatically fitted to the histogram of spectral angles from the hyperspectral image of an SLN	136
5.23	Spectral angle classified version of the hyperspectral image of an SLN .	137
5.24	2D histogram of spectral angle against the dominant plane of the SPM of a hyperspectral image of an SLN	138
5.25	Partially classified image of the hyperspectral image of an SLN, classified using the spectral angle and dominant plane of the SPM	139
5.26	Average spectra of the classes found by classification with spectral angle .	139
5.27	Average spectra of the classes found by partial classification by spectral angle and dominant plane of the SPM	140
5.28	Pseudo-colour image of the hyperspectral image of an SLN with four areas of interest shown	141
5.29	Average spectra of four areas of interest from the hyperspectral image of an SLN	142
5.30	Pseudo-colour image of a hyperspectral image of an SLN	143
5.31	Spectral angle classified version of the hyperspectral image of an SLN .	143
5.32	Average spectra of the classes found by classification with spectral angle .	144
5.33	2D histogram of spectral angle against the dominant plane of the SPM of a hyperspectral image of an SLN	145
5.34	Partially classified image of the hyperspectral image of an SLN, classified using the spectral angle and dominant plane of the SPM	145
5.35	Average spectra of the classes found by partial classification by spectral angle and dominant plane of the SPM	146

5.36	Monochrome image of a multispectral image of an SLN	147
5.37	Histogram of spectral angles from the multispectral image of an SLN . . .	148
5.38	Gaussian peaks automatically fitted to the histogram of spectral angles from the multispectral image of an SLN	149
5.39	Spectral angle classified version of the multispectral image of an SLN . . .	150
5.40	2D histograms for the spectral angle against each plane of SPM for the multispectral image of an SLN	151
5.41	2D histogram showing spectral angle against the first plane of the SPM from the multispectral image of an SLN	152
5.42	Partially classified image of the multispectral image of an SLN, classified using the spectral angle and first plane of the SPM	153
5.43	Monochrome image of a multispectral image of an SLN	154
5.44	Histogram of spectral angles from the multispectral image of an SLN . . .	155
5.45	Spectral angle classified version of the multispectral image of an SLN . . .	156
5.46	2D histograms for the spectral angle against each plane of SPM for the multispectral image of an SLN	157
5.47	2D histogram showing spectral angle against the first plane of the SPM from the multispectral image of an SLN	158
5.48	Partially classified image of the multispectral image of an SLN, classified using the spectral angle and first plane of the SPM	159
5.49	2D histogram showing spectral angle against the first plane of the SPM from the multispectral image of an SLN	160
5.50	Partially classified image of the multispectral image of an SLN, classified using the spectral angle and first plane of the SPM	161
6.2	An image of a car windscreen with and without a polarising filter	167
B.1	Pseudo colour image of a hyperspectral image of an SLN	192
B.2	Histogram of spectral angles from the hyperspectral image of an SLN . .	193
B.3	Gaussian peaks automatically fitted to the histogram of spectral angles from the hyperspectral image of an SLN	193
B.4	Spectral angle classified version of the hyperspectral image of an SLN . .	194
B.5	Average spectra of the classes found by classification with spectral angle .	195
B.6	2D histogram of spectral angle against the dominant plane of the SPM of a hyperspectral image of an SLN	195
B.7	Partially classified image of the hyperspectral image of an SLN, classified using the spectral angle and dominant plane of the SPM	196
B.8	Average spectra of the classes found by partial classification by spectral angle and dominant plane of the SPM	197
B.9	Pseudo colour image of a hyperspectral image of an SLN	198
B.10	Histogram of spectral angles from the hyperspectral image of an SLN . .	199
B.11	Gaussian peaks automatically fitted to the histogram of spectral angles from the hyperspectral image of an SLN	199
B.12	Spectral angle classified version of the hyperspectral image of an SLN . .	200
B.13	Average spectra of the classes found by classification with spectral angle .	200
B.14	2D histogram of spectral angle against the dominant plane of the SPM of a hyperspectral image of an SLN	201

B.15	Partially classified image of the hyperspectral image of an SLN, classified using the spectral angle and dominant plane of the SPM	201
B.16	Average spectra of the classes found by partial classification by spectral angle and dominant plane of the SPM	202
B.17	Pseudo colour image of a hyperspectral image of an SLN	203
B.18	Histogram of spectral angles from the hyperspectral image of an SLN . .	204
B.19	Gaussian peaks automatically fitted to the histogram of spectral angles from the hyperspectral image of an SLN	204
B.20	Spectral angle classified version of the hyperspectral image of an SLN . .	205
B.21	Average spectra of the classes found by classification with spectral angle .	206
B.22	2D histogram of spectral angle against the dominant plane of the SPM of a hyperspectral image of an SLN	206
B.23	Partially classified image of the hyperspectral image of an SLN, classified using the spectral angle and dominant plane of the SPM	207
B.24	Average spectra of the classes found by partial classification by spectral angle and dominant plane of the SPM	208
B.25	Pseudo colour image of a hyperspectral image of an SLN	209
B.26	Histogram of spectral angles from the hyperspectral image of an SLN . .	210
B.27	Gaussian peaks automatically fitted to the histogram of spectral angles from the hyperspectral image of an SLN	210
B.28	Spectral angle classified version of the hyperspectral image of an SLN . .	211
B.29	Average spectra of the classes found by classification with spectral angle .	211
B.30	2D histogram of spectral angle against the dominant plane of the SPM of a hyperspectral image of an SLN	212
B.31	Partially classified image of the hyperspectral image of an SLN, classified using the spectral angle and dominant plane of the SPM	213
B.32	Average spectra of the classes found by partial classification by spectral angle and dominant plane of the SPM	214

List of Tables

2.1	UV laser specification table	35
3.1	A table of vectors and measures compared to a reference	44
3.2	A table of vectors and their measures compared to a reference	45
3.3	A table of vectors	47
4.1	A table showing the class and values of spectral position for various areas on figure 4.19	75
5.1	Clinical descriptions of the stages of breast cancer at diagnosis	91
5.2	Summary of several spectroscopic studies of cancerous tissue	100
5.3	As-built specifications of the Mark I imaging system	108
5.4	As-built specifications of the Mark II imaging system	109
5.5	Initial specification of Specim hyperspectral imager	110
5.6	Final specification of Specim hyperspectral imager	111
C.1	Layout of the HDI file header	216
C.2	Layout of the HDI table of contents	216
C.3	Layout of the HDI image data block	217
C.4	Layout of the HDI metric data block	217
C.5	Layout of the HDI class data block	218
C.6	Layout of the HDI image mask data block	218
C.7	Layout of the HDI channel data block	218
C.8	Layout of the HDIx file header	219
C.9	Layout of the HDIx table of contents	219
C.10	Layout of the HDIx image data block	220
C.11	Layout of the HDIx metric data block	220
C.12	Layout of the HDIx class data block	221
C.13	Layout of the HDIx image mask data block	221
C.14	Layout of the HDIx channel data block	221

Academic Thesis: Declaration of Authorship

I, Jack D. O'Sullivan, declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

Imaging Through A Scanner, Darkly: Spectral Imaging for Sentinel Lymph Node Biopsies

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at the University of Southampton.
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
3. Where I have consulted the published work of others, this is always clearly attributed.
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
5. I have acknowledged all main sources of help.
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.
7. Either none of this work has been published before submission, or parts of this work have been published as:

O'Sullivan, J.D., Hoy, P.R., Rutt, H.N., An extended spectral angle map for hyperspectral and multispectral imaging. In *CLEO:2011 - Laser Applications to Photonic Applications*, OSA Technical Digest(CD), Optical Society of America, 2011;

O'Sullivan, J.D., Hoy, P.R., Rutt, H.N., Spectral imaging as a potential tool for optical sentinel lymph node biopsies. In Nirmala Ramanujam and Jürgen Popp, editors, *Clinical and Biomedical Spectroscopy and Imaging II*, volume 8087, SPIE, 2011.

Signed:

Dated:

Acknowledgements

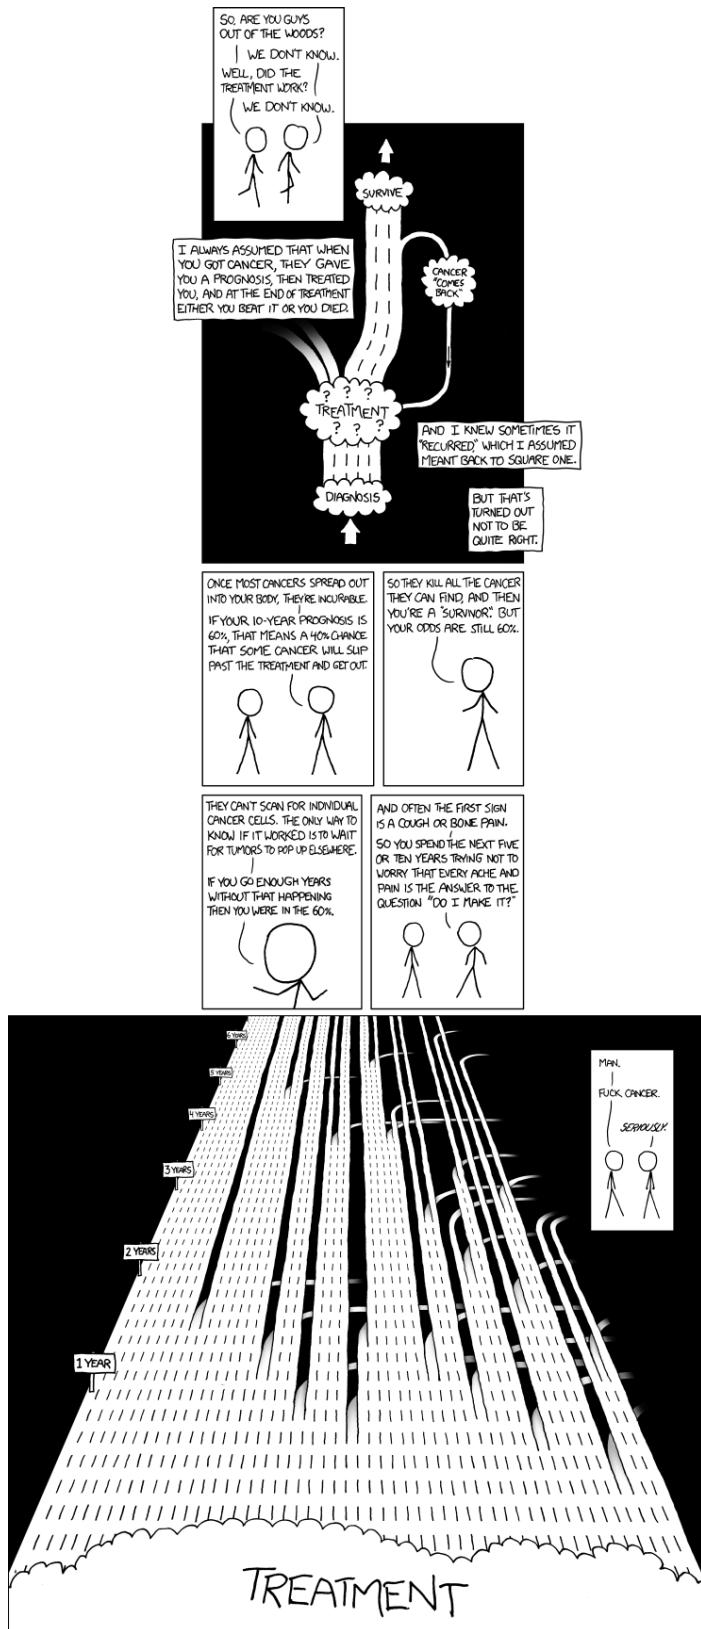
There are three people without whom this project simply would not have been possible. Firstly, my supervisor Prof. Harvey Rutt, his enthusiasm and insight have always been evident and despite the many calls on his time, he has always been available. Secondly, my main medical collaborator, Mr. Richard Sainsbury who provided the application for this work and who has been generous with his time and access. And finally, Dr. Paul Hoy, his assistance, knowledge and patience have been invaluable to me. A sincere and humble thank you to all three.

I would like to thank the theatre staff at St Mary's Hospital on the Isle of Wight who have often gone out of their way to accommodate me and provide any every assistance possible. A big thank you also to the staff at Red Funnel, who always allowed me to stretch the definition of "personal baggage" to absurd lengths. Thank you also to Profs. Eugene Healy and Brian Diffey for sharing their time and expertise.

I would like to acknowledge my funding, from the School, EPSRC and from various sources of Harvey's, into which I felt it best not to enquire...

For the sunshine, moonlight, good times, boogies, and just about everything else imaginable, I blame it all on Sadie, thank you so much, LUB. A big thank you is also owed to everybody else who has made my time in Southampton so enjoyable; for the innumerable lunches, dinners, coffees, distracting chats, games of Scrabble, crosswords, parties and sessions, a big hug to Grace, and firm handshakes to Amoré, Tom, Chris, Little Dude, and everyone else, the list is long. Extra special thanks of course to Tana, whose love and support has kept my sanity in check these last two years.

Finally, a big thank you to my family for all of their love, help and support through this project and always.



"Each quarter of the lanes from left to right correspond loosely to breast cancer stages one through four (at diagnosis)".

*Image and caption reproduced under GNU Creative Commons License by kind permission of XKCD
(<http://www.xkcd.com/931>)*

Abbreviations and selected terms

-ectomy	General suffix to denote removal of tissue.
<i>a-posteriori</i>	Posterior to, <i>with probability</i> : probability based on knowledge of the outcome of an event upon which it is conditional.
<i>a-priori</i>	Prior to, <i>with information</i> : previously known information. <i>With probability</i> : probability based on no prior knowledge of the system.
<i>in situ</i>	In position, where something naturally lies.
<i>in vitro</i>	Within the glass, tissue removed from its normal biological context.
<i>in vivo</i>	Within the living, using whole, living tissue.
ND	<i>N</i> -dimensional.
ALND	Axillary Lymph Node Dissection.
ANN	Advanced Neural Networks.
AOTF	Acousto-Optical Tunable Filter.
Axillary	Of, or pertaining to, the armpit.
BCC	Basal Cell Carcinoma.
BCE	Bhattacharyya Co-Efficient.
BMP	BitMaP.
CCD	Charge Coupled Device, an image capture device consisting of an array of photosensitive capacitors and a shift register to transfer accumulated charge to a storage area where it can be converted into a digital signal, a passive pixel sensor (c.f. CMOS).
CMOS	Complementary Metal Oxide Semiconductor, an image capture device consisting of an array of photodetectors and active amplifiers, due to the active nature of the pixels, processing can be incorporated on chip (c.f. CCD).

CMY	Cyan, Magenta, Yellow; a system of representing colours as subtractive quantities of pure cyan, magenta and yellow pigments, typically used for printing purposes.
CSV	Comma Separated Variable, a generic spreadsheet format where values in adjacent columns are separated by commas and each row is separated by a line break.
CW	Continuous Wave.
DCIS	Ductal Carcinoma <i>in-situ</i> .
DNA	De-oxy riboNucleic Acid.
DTOF	Diffuse Time Of Flight.
ESS	Elastic Scattering Spectroscopy.
fps	Frames Per Second, measure of the speed of video.
FTIR	Fourier Transform InfraRed.
FWHM	Full Width at Half Maximum.
H&E	Hematoxylin and Eosin, a common combination of staining agents, broadly speaking hematoxylin stains cell nuclei and acidic material blue, and eosin stains other structures and basic material pink.
Hartley Window	An region of low absorption of radiation by ozone in the Hartley bands, which are bands of high absorption between 200nm and 300nm and centred at 255nm. The Hartley window extends from around 280-300nm.
HCA	Heirarchical Cluster Analysis.
HDI	High Dimensional Image.
HDR	High Dynamic Range.
HSL	Hue, Saturation, Lightness; a system of representing colours in a polar co-ordinate space.
HSV	Hue, Saturation, Value; a system of representing colours in a polar co-ordinate space.
IDC	Invasive Ductal Carcinoma.
Interstitial Fluid	A solution surrounding cells and tissue.
Intra-operative	During the course of an operation.

IR	InfraRed, the part of the electromagnetic spectrum with a wavelength greater than 700nm.
LCTF	Liquid Crystal Tunable Filter.
LDA	Linear Discriminant Analysis.
Lymphatic System	A system of vessels and nodes that drains interstitial fluid from tissues and returns it to the bloodstream.
MED	Minimum Erythema Dose, the energy density required to create a burn on human skin.
Metastasis	The transfer of diseased tissue, especially cells of malignant tumours, from one part of the body to another.
PCA	Principal Components Analysis.
PCC	Pearson Correlation Co-efficient, also known as PMCC, Product Moment Correlation Co-efficient.
Pixel	Picture element, single indivisible point of an image, the smallest addressable unit of an image, the base unit of image size. A spectral image pixel in this thesis refers to a single spatial point, a single vector of spectral data.
Resection	The removal of part of the body, or tissue from the body.
RGB	Red, Green, Blue; a system of representing colours as additive intensities of pure red, green and blue light, typically used for display/projection purposes.
RNA	RiboNucleic Acid.
SAM	Spectral Angle Map.
SCC	Squamous Cell Carcinoma.
SDK	Software Development Kit.
Sensitivity	A measure of the performance of a binary classification, the fraction of actual positives that are classified positive. $\frac{\text{TruePos.}}{\text{TruePos.} + \text{FalseNeg.}}$. Sometimes called recall rate.
Sentinel Lymph Node	The first node in a cluster to which fluids from the local area will drain.
SFG	Sum Frequency Generation.
SHG	Second Harmonic Generation.

SID	Spectral Information Divergence.
SLNB	Sentinel Lymph Node Biopsy.
SLOC	Source Lines of Code.
SNR	Signal to Noise Ratio.
SOIR	Spectroscopic Oblique Incidence Reflectometry.
Specificity	A measure of the performance of a binary classification, the fraction of actual negatives that are classified negative. $\frac{TrueNeg.}{TrueNeg.+FalsePos.}$.
SPM	Spectral Position Map.
SPV	Spectral Position Vector.
SRS	Stimulated Raman Scattering.
Thermal relaxation time	The time taken for 50% of heat absorbed from a pulse of laser radiation to dissipate.
TP	Touch Print or Touch Prep(aration).
TPI	Terahertz Pulsed Imaging.
UV	UltraViolent, the part of the electromagnetic spectrum with a wavelength less than 400nm, typically split in to UVA (315-400nm), UVB(280-315nm) and UVC(100-280nm).
VNIR	Visible and Near InfraRed, the wavelength range in which silicon CCDs are effective, \sim 400-1000nm.
Voxel	Volume element, single indivisible point of a volume, the smallest addressable unit of a volume.

Chapter 1

Introduction

1.1 Brief Project Overview & Summary of Prior Work

This project is an extension of work begun in 2005 by Dr. Paul Hoy for his own doctoral thesis (Hoy, 2009). It began as an investigation into the possibility of developing a low cost device for the remote detection of water. This quickly developed into a study to assess the viability of multispectral imaging to guide brain tumour resection surgery. On the basis of some promising early results, other biomedical collaborations were sought, this thesis describes the work undertaken in two such areas. Some work on spectral imaging of skin pathologies was performed, although the work quickly fell out of the scope of the project as a whole as described in chapter 2. The main work described in this thesis investigates the possibility of using spectral imaging for quantitative optical biopsies, with particular focus on biopsies of the lymph nodes of breast cancer patients.

The central hypothesis is that if different types of visually similar biological tissue can be distinguished by their spectra, then a spectral imaging system could provide benefits to patients and treatments centres alike that are far in excess of the cost of the components required to build it. Further, if this system can be made to run unsupervised, or at least supervised by a person with little training or expertise, then the cost of operating the system can be kept very low. The purpose of this project is thus to test the assertion that different types of tissue are indeed distinguishable by their spectra, and that the separation of different types of tissue can be performed with minimal user intervention or direction.

Medical applications of imaging typically fall into two categories, diagnostic and intra-operative. X-rays, mammograms and MRI scans are examples of the former, whereas endoscopic imaging for key-hole surgery is an example of the latter. The work on lymph node biopsies has aspects of both diagnostic and intra-operative imaging.

As the pathway for metastasis, the health of the lymphatic system is an important prognostic indicator in cancer patients. Biopsies of axillary lymph nodes are routinely taken during lumpectomy and mastectomy surgeries to enable the extent of the spread of cancer to be determined; in that sense any imaging used to perform this biopsy could be considered intra-operative. The aim of such imaging is to determine the health of the lymph nodes, and so in that sense this imaging is akin to traditional diagnostic imaging.

Work on this project has concentrated on the visible and near infrared (VNIR) section of electromagnetic spectrum. This enables the use of standard digital cameras based on silicon Charge Coupled Devices (CCDs). These have an effective spectral responsivity in the wavelength range from about 400-1000nm (the response curve for one of the CCDs used in this project, a Sony Ex-View model, is shown in figure 1.1). Both the CCDs and the requisite optics (lenses, filters, beam-splitters etc.) are widely and commercially available at relatively low cost, making this region ideal for preliminary investigation.

The traditional “finger-print” region for biological substances is further into the infrared (IR) spectrum (around $5-20\mu\text{m}$) and it is anticipated that spectral imaging in this region would be very useful for biomedical applications. Work is ongoing by Dr. Hoy to develop a camera with sensitivity in this region as part of his brain tumour detection work, however it has not yet been possible to take any spectral images with this.

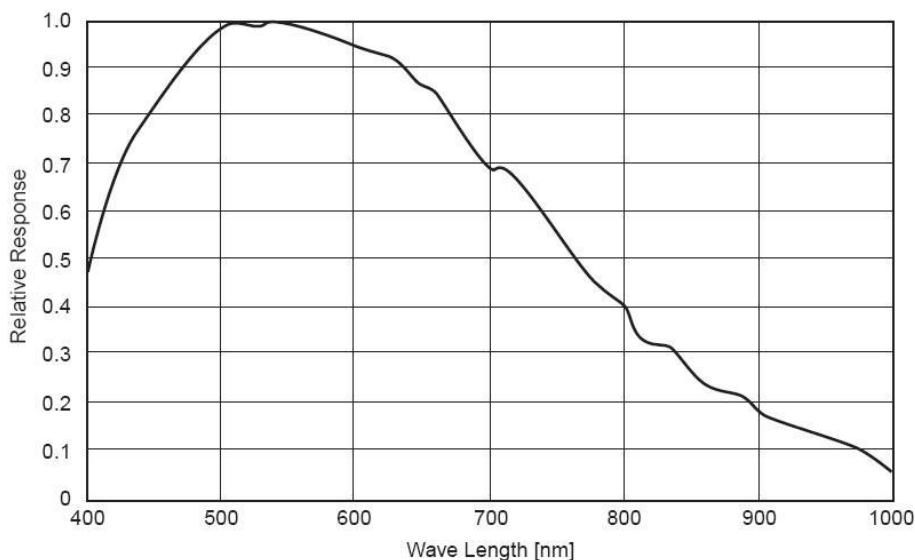


FIGURE 1.1: The spectral responsivity curve for a Sony Ex-View CCD (Sony Electronics, 2003).

1.2 Thesis Structure

The remainder of this chapter will describe some of the fundamental interactions of light and matter that are important to biological imaging and spectroscopy. Further, the key chromophores in the VNIR region are presented and discussed. The concept of spectral imaging is then introduced along with a discussion of the technologies, applications and issues involved. Chapter 2 presents the research carried out on skin damage imaging and treatment.

Chapter 3 discusses how images are assessed in order to classify which areas are similar to and which are different from each other. There follows a discussion of spectral similarity metrics and in particular, the Spectral Angle, which leads into an examination of a proposed extension to this metric that has been dubbed the Spectral Position.

Chapter 4 presents tests of the spectral position metric on a series of artificially generated colour and spectral images, before applying the metric to some real datasets from various disciplines.

Chapter 5 presents the work carried out on the development of an optical sentinel lymph node biopsy.

Finally, a discussion of the future direction of this work, as well of the conclusions drawn from existing work is presented in chapter 6.

1.3 Photon Interactions with Biological Tissue

Light can interact with matter in a number of ways and whilst there are differences in the details of the mechanisms, these can essentially be divided into scattering and absorption processes. With regards to the specific case of visible and IR radiation interacting with biological tissue, the key interactions are absorption by chromophores in the visible region, scattering in the near IR, and molecular absorption in the mid to far IR. This section first presents a brief overview of how these interactions are measured and then outlines some of the key absorbing and scattering components of tissue relevant to this project.

1.3.1 Absorption

Absorption occurs when incident radiation causes an energy level transition from some stable or metastable low energy state, to some excited high energy state; this transition can be a change in the energy level of electrons (more important for visible

light) or the vibrations of a molecule (more important for infrared radiation).

Absorption causes either a decrease in the intensity of light transmitted or a complete attenuation of transmission through the material.

In a purely absorbing material of thickness l , each layer dl will absorb the same proportion of the incident intensity, I . This relationship is known as the Lambert-Bouguer law and is formally expressed as:

$$\frac{dI}{I} = \mu_a dl, \quad (1.1)$$

where μ_a is called the absorption coefficient. This relationship was subsequently modified by August Beer who noted that when a compound is diluted in a non-absorbing material, the absorption coefficient is linearly related to the concentration (see equation 1.2 below where α is the specific absorption coefficient and c is the concentration). The subsequent relationship is known as the Beer-Lambert law and is important in tissue where the absorbing material is usually dissolved or immersed in a non-absorbing material:

$$\mu_a = \alpha c. \quad (1.2)$$

Where multiple absorbing materials are present in the same sample, as is the case in biological tissue, the total absorbance is the sum of the absorbance due to each. The absorbance is defined as the logarithm of the ratio of the incident radiation and the transmitted/reflected (measured) radiation. According to equation 1.1, the transmitted radiation is given by:

$$I = I_0 e^{-\mu_a l}, \quad (1.3)$$

where I_0 is the intensity of the incident radiation. This shows that the absorption coefficient can be thought of as a probability that a photon will be absorbed in a given length of material. Therefore the absorbance of the i th material, A_i , is given by:

$$A_i = \ln \frac{I_0}{I} = \alpha_i c_i l, \quad (1.4)$$

and thus the total absorbance is:

$$A = \left(\sum_i \alpha_i c_i \right) l. \quad (1.5)$$

1.3.2 Scattering

Photon scattering occurs when light incident on a material has its propagation direction altered due to an absorption and re-emission interaction with the medium. Specifically, scattering occurs where the medium is not composed of atoms with a high degree of spatial order; where the material does have a high degree of order, i.e. a crystalline structure, the light is considered as having been diffracted rather than scattered, leading to sharp, well defined patterns of light intensity. Propagation direction can also be altered by light passing through a material in which it has a different speed, when this occurs the change in direction is called refraction, accordingly the speed of light in a given material is determined by the refractive index.

Scattering can be categorised as occurring from a single scattering centre or from multiple centres. In single scattering it is assumed that a photon interacts with only one molecule or particle; deflection in this case is highly dependent on the size of the particle. With a diffuse light source, for small particles, ($\lesssim \frac{\lambda}{10}$) the scattering is described as Rayleigh scattering. Deflection in this case is usually random, the precise location and orientation of the scattering centre and hence the incoming trajectory of the light before the interaction is uncertain. Scattering from larger particles is described as Mie Scattering. Mie theory is only analytically described for spherical and elliptical particles and so there is less particle orientation dependence. Scattering in Mie theory is much more intense in the same direction as the initial propagation (see figure 1.2).

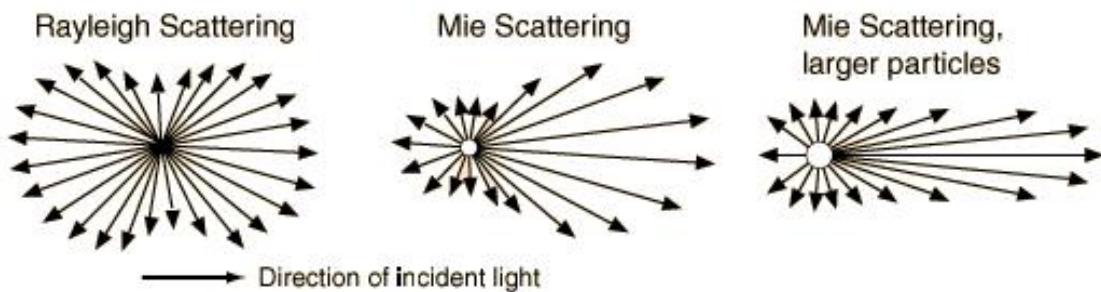


FIGURE 1.2: A diagram showing the relative intensities of scattered light in various directions (relative to the direction of incident light) for Rayleigh scattering, Mie scattering and Mie scattering from large particles. (Image courtesy of <http://hyperphysics.phy-astr.gsu.edu/hbase/atmos/blusky.html>)

Multiple scattering can be considered as a sequence of single scattering events; the random deflection angles of each single scattering can be considered statistically, tending to average out to a deterministic distribution. The effect of multiple scattering events is analogous to a diffusion of the light.

In biological tissue the large number of scattering centres (particles/molecules) and their close proximity to each other means that all scattering should be considered as

multiple centre scattering. A scattering coefficient can be defined in a similar manner to the absorption coefficient described in the previous section by considering the probability that a photon will be scattered in some given length:

$$I = I_0 e^{-\mu_s l}. \quad (1.6)$$

I and I_0 again represent the measured and incident intensity of radiation respectively, l is the length of material traversed and μ_s is the scattering coefficient (in a given direction). Generally the scattering coefficient is angle-dependent and the total scattered intensity is a solid angle integral over these values. This angular dependency is particularly strong for materials with a high degree of spatial order.

By combining equation 1.3 and equation 1.6 it is possible to determine the intensity of radiation at a given distance in a material that is both scattering and absorbing the light present:

$$I = I_0 e^{-(\mu_a + \mu_s)l} = I_0 e^{-\mu_t l}, \quad (1.7)$$

where μ_t is known as the total attenuation coefficient.

It is now clear that for a photon entering an absorbing and scattering material there are four potential outcomes:

1. the photon is scattered randomly a number of times within the material, traversing completely through the material and emerging with a new propagation direction;
2. the photon is scattered randomly a number of times through the material before being absorbed;
3. the photon is scattered randomly a number of times within the material and re-emerges with a new propagation direction;
4. the photon passes completely through the material without being absorbed or scattered.

These outcomes are shown in figure 1.3.

In the first outcome, the light has been diffused by passing through the material; in the second, the light has simply been absorbed by the material. Clearly, in the third case, light re-emerging from the incident surface can be considered as reflectance of the incident light and is indeed known as diffuse reflectance. In contrast to specular reflectance (mirror like), the light that is diffusely reflected does not form an image of the incident light source, but it is this light that is recorded when imaging a non-backlit scene.

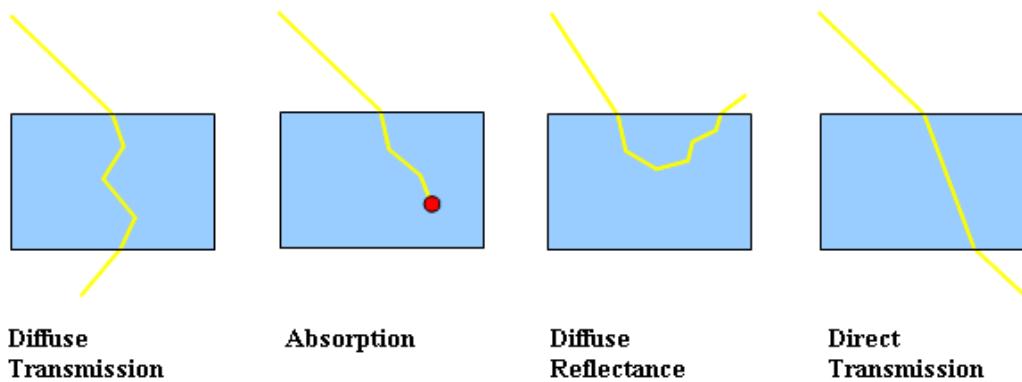


FIGURE 1.3: A diagram showing photon paths through a medium that is both scattering and absorbing.

There is another way to categorise scattering and that is whether it is elastic or inelastic. Elastic scattering conserves the energy of the incident light, inelastic scattering does not. In this case scattering is considered as an absorption/emission event; a photon is absorbed by a molecule promoting it into an excited virtual energy level. When the molecule relaxes back to the lower energy state a photon is emitted in a random propagation direction. This is shown in figure 1.4. Where the excitation and relaxation energies are the same, the emitted photon will have the same wavelength as the incident photon and it can be considered as the same photon having undergone elastic scattering (where the particle is small compared to the wavelength of the light, this is Rayleigh scattering, as described above).

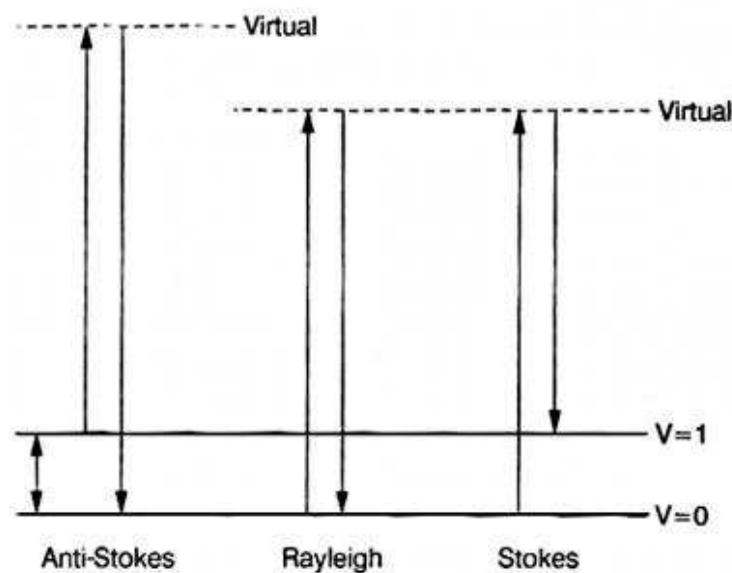


FIGURE 1.4: An energy level diagram illustrating Rayleigh, Stokes and anti-Stokes scattering events (Hendra, 1998).

It is possible that a molecule will be promoted to an excited state and relax into a non-excited state of slightly different energy from its original state. The emitted photon now has a slightly different wavelength and hence a different energy and the light is said to have been inelastically scattered. Where the scattering is caused by molecular vibrations the scattering is known as Raman, if it caused by an interaction with an acoustic phonon it is known as Brillouin. The energy changes associated with Brillouin scattering are typically very small. This discussion is concerned only with Raman scattering. When the scattered light is of a lower energy it is said to have undergone Stokes Raman scattering, if it has a higher energy than the incident light it has undergone anti-Stokes Raman scattering (see figure 1.4).

Measuring the spectra of Raman scattered light from a sample can provide useful information about the sample that is not provided by normal infrared spectroscopy as different processes give rise to the spectra. However, Raman scattering is a weak process, only around 0.00001% of photons (1 in 10^7) will produce a Raman shift.

1.3.3 Key Chromophores

This section will consider some of the key biological chromophores in the VNIR region.

Haemoglobin is an iron based chemical found in red blood cells, it binds to molecular oxygen and is responsible for transporting oxygen throughout the body. It is referred to as oxy-haemoglobin (HbO_2) whilst bound to oxygen and deoxy-haemoglobin (Hb) when it is free of oxygen. The absorption spectra for HbO_2 and Hb from the near ultraviolet (UV) to the near IR are shown in figure 1.5.

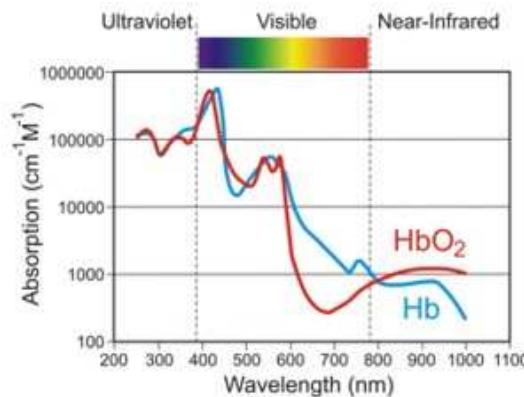


FIGURE 1.5: Absorption spectra for oxy- and deoxy-haemoglobin from the near UV to the near IR (UCL Biomedical Optics Research Laboratory, 2006).

The absorptions are sufficiently different for the two to be distinguishable and certain fingerprints are immediately noticeable. Both show a strong absorption peak in the blue end of the visible spectrum ($\sim 450\text{nm}$) although the peak is shifted to the red end for Hb compared to HbO_2 . HbO_2 has two further absorption peaks at $\sim 540\text{nm}$ and

~576nm whilst Hb has a broad absorption centred at ~554nm (Zijlstra et al., 1991). Absorption of light from around 600-800nm is significantly lower in HbO₂ than Hb, this trend reverses above 800nm.

Water accounts for around 60% of a human's body mass and so unsurprisingly it is an important chromophore when considering biological tissue. The absorption of light by water is fairly low throughout most of the VNIR region, as is intuitive from the fact that water appears highly transparent and colourless. In fact, the absorption coefficient for water does rise slightly from 400nm through to around 900nm, at this point it rises sharply to a peak at 970nm before it starts to fall off again (see figure 1.6 below). This peak is the first harmonic of the symmetric molecular stretching mode combined with the fundamental anti-symmetric stretching mode (Chung et al., 2010).

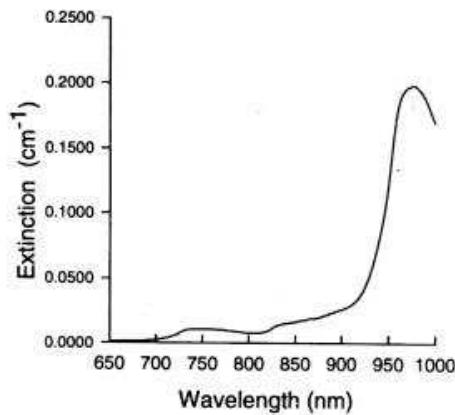


FIGURE 1.6: Absorption spectrum of water from 650-1000nm (Elwell and Hebden, 2008).

Most of the water content in the human body is not pure water but water that has hydrogen-bonded to other biomolecules (particularly proteins). The existence of these hydrogen bonds can lead to the exhibition of new peaks and the altering of existing peaks in the spectra of corporeal water compared to pure water. Whilst this is an important consideration in the mid and far infrared it has little effect in the VNIR region.

Lipids are fatty compounds found in biological tissue, in the human body they are most commonly found as triglycerides (fats), cholesterol and phospholipids (a constituent of all cell membranes). Figure 1.7 below shows the absorption spectrum of pure pork fat from 800nm to 1080nm. This fat is composed mainly of lipids and will be very similar to a spectrum taken of lipids from the human body. The spectrum is similar to that of water in the same region, however the peak is centred rather lower at around 930nm.

Melanin is the substance responsible for the colouration of human skin. It is a pigment produced in the epidermal (top-most) layer of the skin. It has an absorption peak in the near UV (~330nm) which serves to protect the skin from photodamage. The

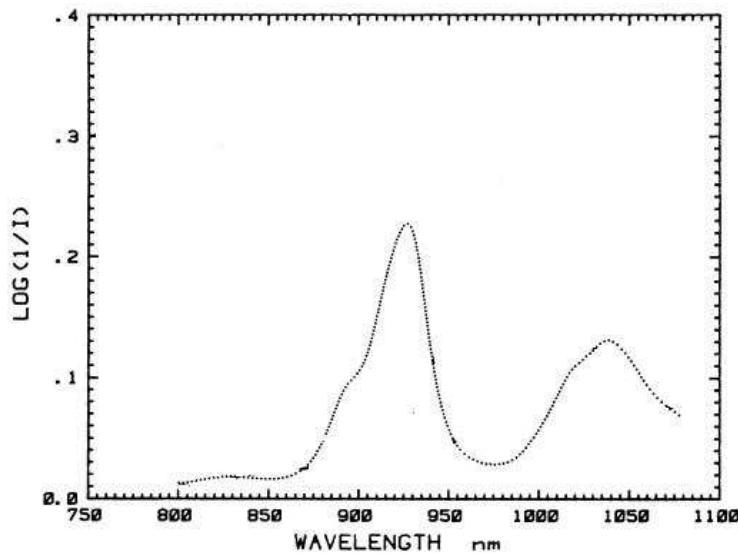


FIGURE 1.7: Absorption spectrum for pure pork fat, the vertical axis is a relative measure of absorption coefficient (image modified to remove unnecessary information from the original) (Conway et al., 1984).

absorption gradually decreases across the visible range, as can be seen in figure 1.8 below. This gradual decrease continues beyond the visible and into the near IR (Matas et al., 2002).

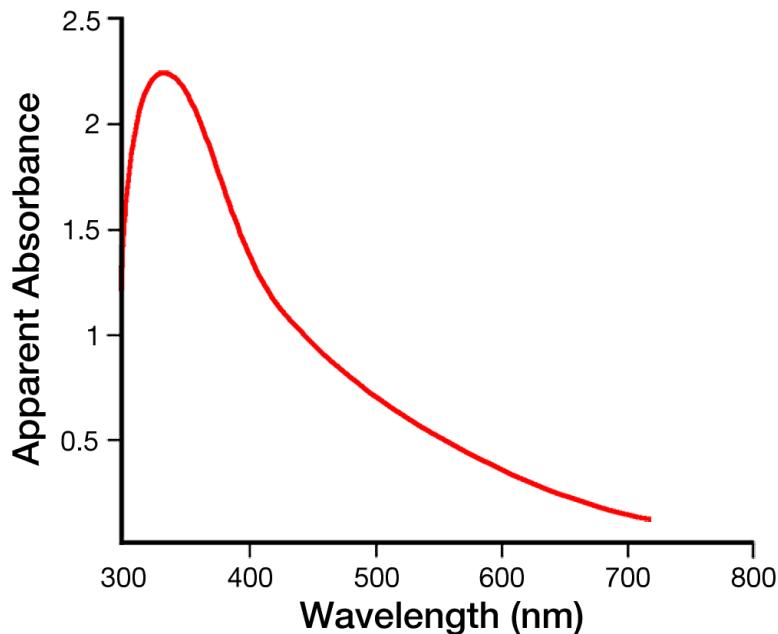


FIGURE 1.8: Absorption spectrum of melanin from 300nm to \sim 700nm (Clinuvel Pharmaceuticals, 2009).

1.4 Spectral Imaging

The word ‘image’ has many definitions, and even within the scope of science and technology it is not uniquely defined. It is important then to set out how it is understood within the context of this work. An image is a spatial map where each discrete element is associated with a vector of values, each of which is a proxy measure for some other variable. These variables are typically all intensities of radiation, however other metrics or derived statistics can be mapped in the same space to form an image. This implicitly only considers “digital” images, and whilst qualitative assessment of “analogue” images is still practised in the medical field, e.g. many dentists still use film x-ray examinations to assess patients, to be used in the development of quantitative techniques, images will need to be in a digital format.

The rest of this section explains spectral imaging as a specific imaging modality within the above definition before moving on to discuss briefly the technology used to record these images, the ways in which they are used and the problems and limitations faced in using them.

1.4.1 What Is Spectral Imaging?

From the above definition, one of the simplest forms of image would be a black and white, or monochrome photograph. Each picture element (pixel) has a single value that is proportional to the total intensity of light reflected by an object onto a photosensitive detector. To record the same object in a colour image, each pixel requires a vector of three values.

The human eye detects light incident on the retina using two types of cells, rods and cones. Rods detect only absolute intensity and as such convey no colour information, they do however work in lower lighting conditions than cones. There are three types of cone cells known as blue, green and red, each is sensitive to only a limited section of the visible spectrum. Blue cones have a peak sensitivity around 420-440nm, green cones peak around 530-550nm and red cones at around 560-580nm. The brain perceives colour by comparing the response in each of these cells.

In a similar way, digital cameras that record colour images separate each pixel into sub-pixels, these are small enough to be considered co-located. Each of these is behind a broadband filter, designed to transmit either blue, green or red light; there are four sub-pixels in a standard camera, two of which are behind green filters, and one each behind a red and a blue filter in a Bayer pattern (see figure 1.9), designed to reflect the fact that human vision has an overall peak sensitivity to yellow-green light. Each colour can be thought of as a point in a three dimensional (3D) space, whose axes are intensities of red, green and blue light. This is known as the RGB Colour Space,

however, as with any space, the axes can be rotated, or indeed different axes altogether can be defined, and various other spaces, such as the CMY (Cyan-Magenta-Yellow), HSV (Hue, Saturation and Value) and HSL (Hue, Saturation and Lightness) are also commonly used.

The CMY space is widely used in colour printing as it is a “subtractive” space, that is to say that each colour is composed of combinations of cyan, magenta and yellow pigments, each of which absorb (subtract) colour from white light (red, green and blue respectively). This is as opposed to the RGB space, which is an “additive” space mainly used in connection with displays; red, green and blue emitting elements are combined (added) to create different colours. The CMY space can be described as a rotation of the RGB space, although in practice there is no universal conversion as the gamut of the spaces (the complete subset of colours that can be reproduced) is highly device dependent.

The HSV and HSL are related spaces that are cylindrical polar co-ordinate models of the RGB space. The hue is the angle around the central axis and measures the basic “colour” (from red at 0° through green and blue at 120° and 240° respectively, back through to red at 360°). The saturation measures the distance from the central axis and moves from grey at $S=0$ through to the fully saturated colours at a normalised $S=1$. The lightness and value attributes both measure the distance along the central axis, from black at $L/V=0$, through the shades of grey to white at $L/V=1$. The difference between the two spaces is subtle and is best explained by considering the case of $S=1$ for changing L/V . For both spaces all points at $L/V=0$ are black, as V increases the colours move from black to fully saturated colours, as L increases they pass through fully saturated colours at $L=0.5$ to white at $L=1$.

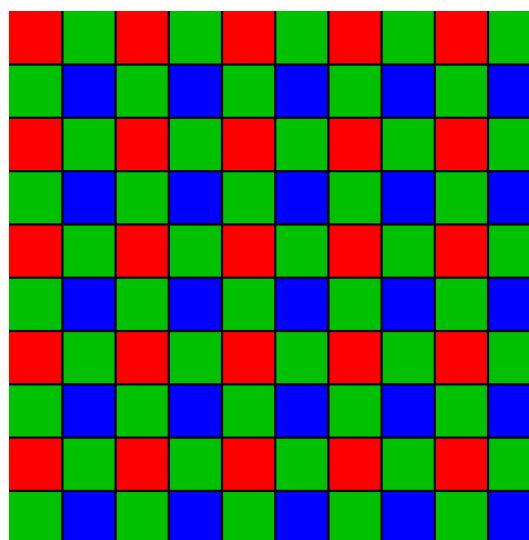


FIGURE 1.9: The Bayer pattern.

Extending this idea, it is clear that a different number of filters with different bandwidths could be used in place of the red, green and blue broadband filters described, the resultant images are spectral images. The potentially high number of dimensions of information for each pixel means that the image is often referred to as a datacube; it is also common to see each datum in the cube referred to as a volume-element, or voxel.

Within this definition, there are two types of spectral imaging that are generally referred to, multispectral and hyperspectral. The distinction between multi- and hyper- spectral images is not well defined. In general multispectral is used to refer to images where each vector contains information from up to around 20 dimensions; these are typically, but not necessarily, discrete, non-contiguous spectral bands, and can cover situations where some bands are in the visible range, some in the near IR and some further into the IR (or indeed in the UV). Hyperspectral refers to images where each vector contains information from more than 20 bands (typically up to hundreds). These bands are generally very narrowband and contiguous. The nature of hyperspectral imaging makes it a good approximation to spectroscopy, and indeed it is sometimes referred to as “imaging spectroscopy”. It is possible then to say that in a given spectral range, like the VNIR covered by a silicon detector, a hyperspectral image will generate a full spectrum at each pixel whereas a multispectral image will just sample the spectrum at discrete locations. This difference is demonstrated in figure 1.10 below.

Throughout this report, spectral imaging will be used to refer both hyperspectral and multispectral imaging; the multi- and hyper- prefixes will be used where a technique or technology is only applicable to one modality. Imaging spectroscopy will be used to describe situations where point source spectroscopy is used to form an image by scanning the point in two dimensions across the scene.

1.4.2 Technologies and Applications

Another potential way of distinguishing multi- and hyper- spectral imaging would be to consider the acquisition of the image. In their textbook on hyperspectral image analysis, Grahn and Geladi outline two acquisition methods for spectral imaging, “Line Scanning” and “Focal Plane Scanning” (Grahn and Geladi, 2007). A third method, “Point Scanning” is also mentioned, this is not considered here as it is a fully spectroscopic method and is thus more properly imaging spectroscopy than spectral imaging.

Focal plane scanning simply involves taking consecutive images of an object through different filters, changing the filter after each image acquisition. In the VNIR range a standard silicon CCD camera together with standard optics can be used. The filter

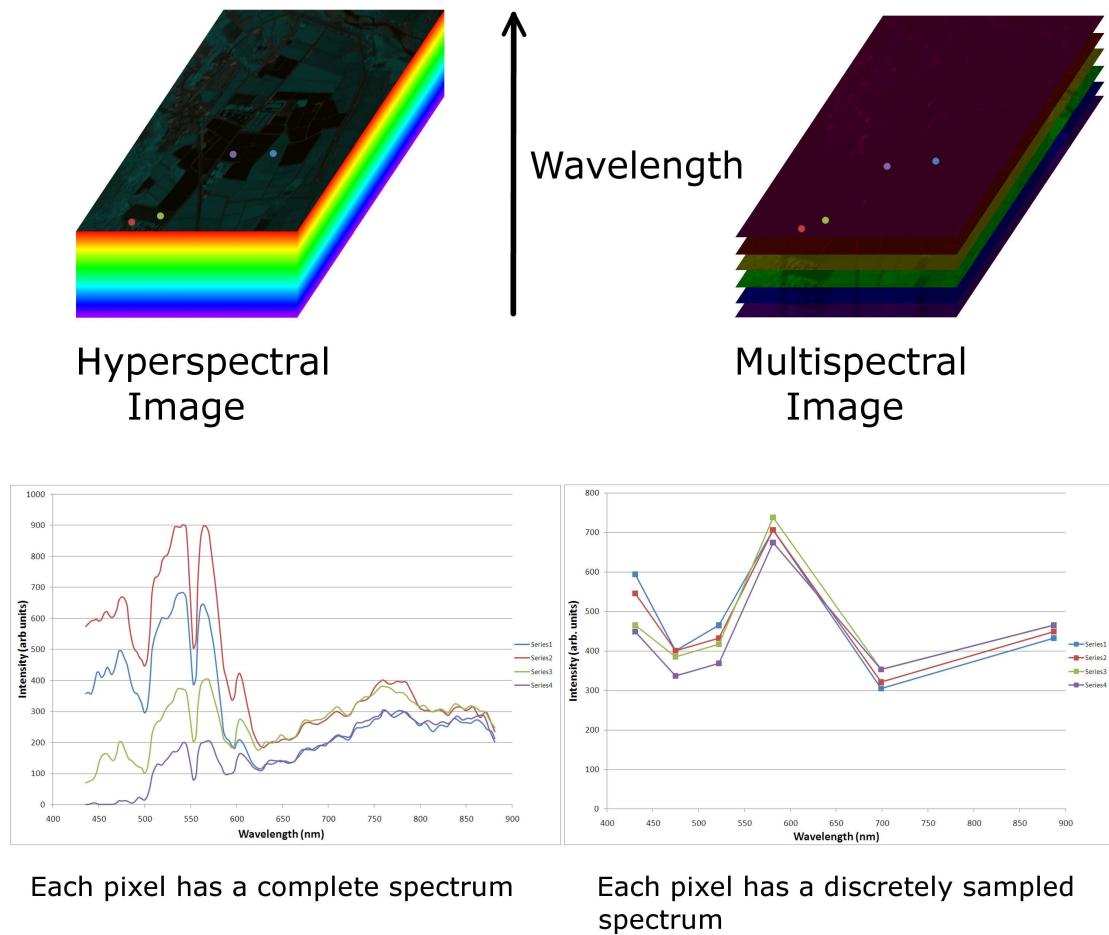


FIGURE 1.10: An illustration of a multispectral and hyperspectral image of the same scene, shown are the two datacubes and the spectral output from each method.

change is often achieved in practice by placing the filters in a wheel, or series of wheels, and rotating each into place in turn. Alternatives such as Liquid Crystal Tunable Filters and Acousto-Optical Tunable Filters can be used to replace the potentially large number of filters with a single optical component (as used for example by Zuzak et al. (2002) and Tuan (2004) respectively). A schematic of the set-up for focal planes scanning is shown in figure 1.11.

Line scanning places the detector behind an optical set-up that spectrally disperses the light reflected by the object. This is typically achieved using a thin slit to deliver a single line of the object scene to a diffraction grating or prism, which disperses the light towards the detector. The 2D image formed on the detector is then comprised of one spatial and one spectral dimension. The object has to be scanned line by line in order to get full spatial coverage. Again, for the VNIR range a standard silicon CCD camera can be used for detection. A schematic of the set-up for line scanning is shown in figure 1.12.

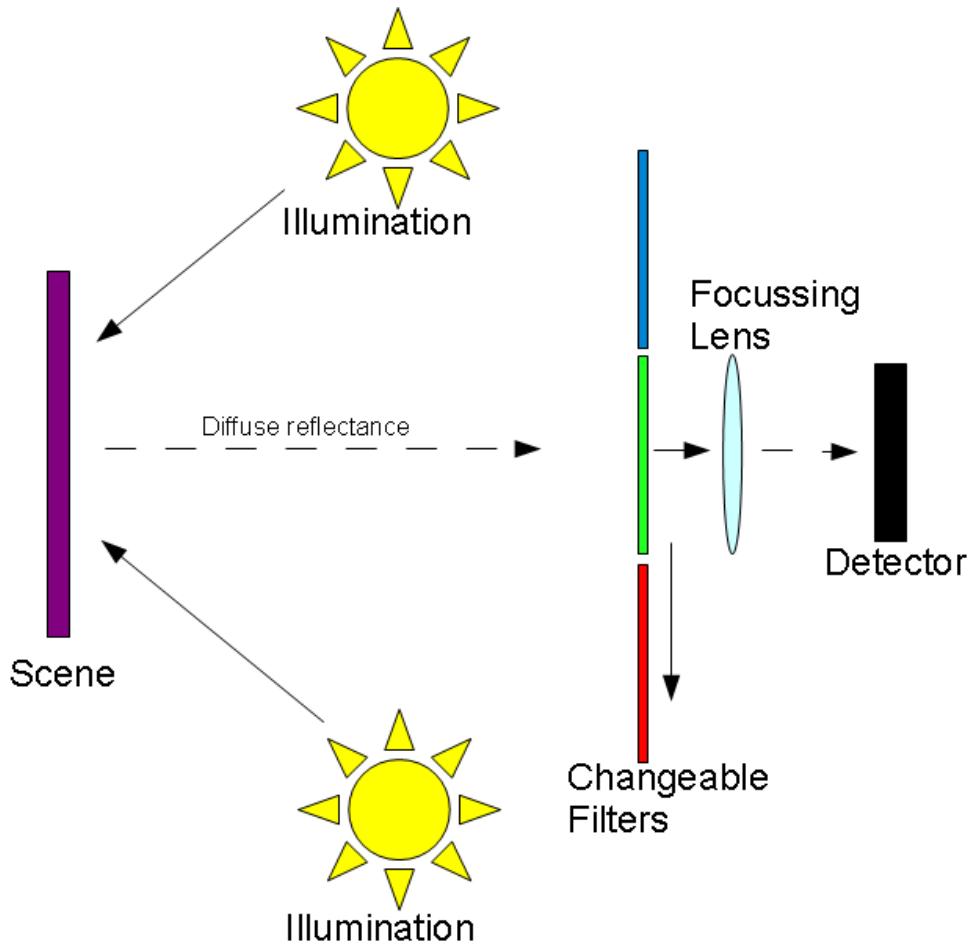


FIGURE 1.11: A schematic of an experimental set up for Focal Plane Scanning acquisition.

Both focal plane and line scanning build up the datacube plane by plane using a number of consecutive image acquisitions, each of which is a very rapid process. In line scanning the requirement to scan the scene line by line implies a trade-off between overall acquisition time and spatial resolution; the spectral dimension and one spatial dimension will be resolution limited by the optical set-up or the detector resolution, the other spatial dimension will be limited by the scan. With focal plane scanning this trade-off still exists, albeit in a slightly different form. Here, both spatial dimensions are resolution limited by the optics or detector and the spectral resolution is limited by the number and bandwidth of the filters used, or by the resolution of the tunable filter. Just as repositioning the scanner (or the object) in between acquisitions is the overall rate-limiting step in line scanning, changing filters, or retuning the filter, is the rate limiting step in focal plane imaging.

Line scanning, by its very nature, captures hyperspectral data, the continuous spectrum of light passing through the slit is dispersed onto the detector such that one dimension of the detector measures the complete spectrum in a contiguous manner. This process captures, in parallel, an entire spectrum for each point along the line of

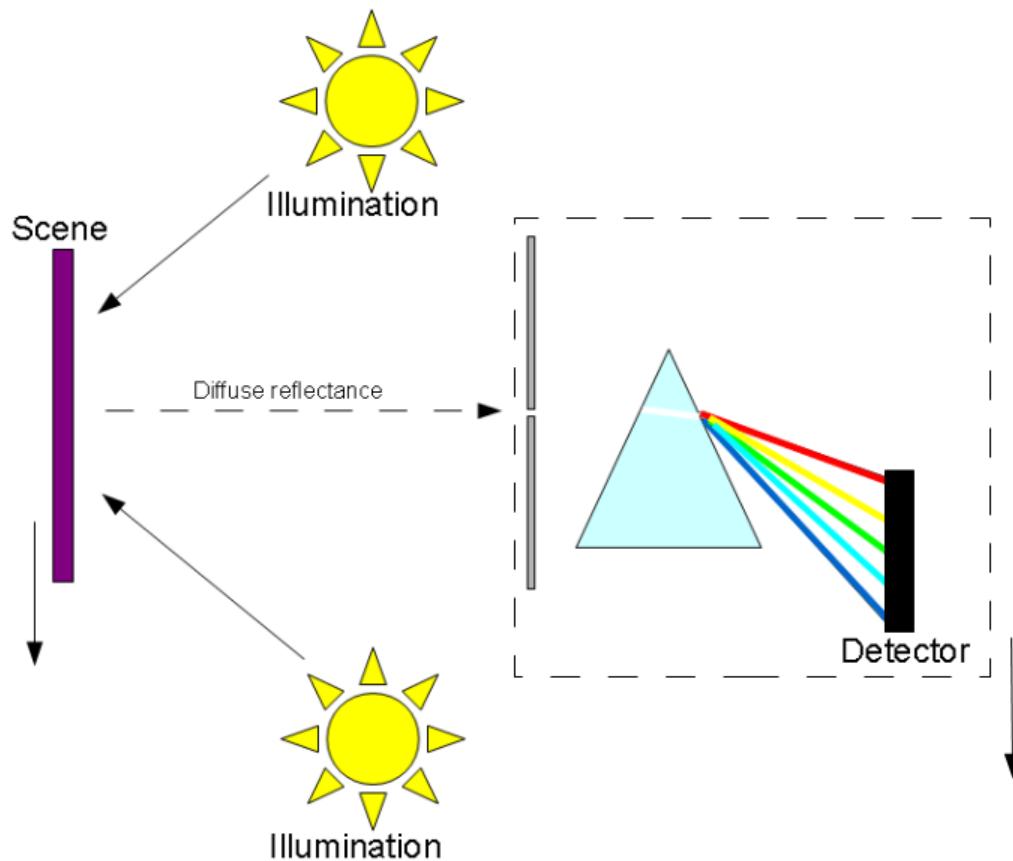


FIGURE 1.12: A schematic of an experimental set up for Line Scanning acquisition. In order to acquire a full 2D scan of the scene, the scene can be moved (if it is moveable) or alternatively the set-up within the dashed-line box on the right hand side can be moved.

the scan, making it many times faster than imaging spectroscopy where the entire spatial plane would be raster scanned. On the other hand, in order for focal plane scanning to be truly hyperspectral, many filters would be needed to cover the entire spectral range, significantly slowing down the process. A smaller number of broadband filters could be used but this would be at the expense of spectral resolution. As such, it is often more practical to limit focal plane scanning to the acquisition of multispectral imaging.

Spectral imaging has traditionally been used primarily for earth remote sensing. Spectral imagers are mounted onto aircraft or satellites and images are then used for assessing land usage or coverage of particular species of plant. NASA have two major remote sensing projects, Landsat and AVIRIS (Airborne Visible/Infrared Imaging Spectrometer). Landsat is a satellite based remote sensing project with 7 satellites providing multispectral images of the earth's surface, whereas AVIRIS is an aircraft mounted hyperspectral imager with a working spectral range of 400-2500nm. It is not intended to present here a thorough review of the literature, suffice it to say that there is a large body of work, and indeed entire research groups dedicated to remote sensing;

a list of recent Landsat-based publications is available at
http://landsat.gsfc.nasa.gov/references/recent_pub.html.

More recently other applications of spectral imaging have emerged such as food contamination detection (Park et al., 2007a), solar cell defect detection (Li et al., 2010), detection of illicit substances (Ng et al., 2009), and disaster management (Jusoff, 2010).

1.4.3 The Curse of Dimensionality

Large data volumes have long been an issue in image processing due to the computational time required to perform calculations on large numbers of pixels. This problem becomes increasingly acute in spectral imaging paradigms when each pixel is represented by values in tens or hundreds of dimensions. Compensating for this increase in complexity and computation time is the expectation that increasing the number of spectral channels or dimensions in which data is collected will increase the ability of a suitably trained system to discriminate between different signals. The implicit assumption in this expectation is that each dimension added contains contrast between different signals; often however, the additional dimensions show little difference between signals, as such they act like noise, actually decreasing the contrast between signals. The fact that increasing dimensionality increases the complexity of data handling without necessarily improving discrimination is referred to as “the curse of dimensionality”.

The matter of high dimensional space is further complicated by a number of properties that are not intuitive from extrapolation from a three dimensional space. One of the simplest to demonstrate mathematically is the relationship of the volume of hyperspheres and hypercubes in multiple dimensions.

The derivation given here follows the argument presented by Landgrebe (2003). The volume of an n D hypercube with sides of length $2r$ is given by the equation:

$$V_c(r) = (2r)^n. \quad (1.8)$$

The volume of the hypersphere of radius r in the same space is given by:

$$V_s(r) = \frac{\pi^{\frac{n}{2}}}{\Gamma\left(\frac{n}{2} + 1\right)} r^n, \quad (1.9)$$

where $\Gamma(x)$ is the Gamma function, given by:

$$\Gamma(x) = (x - 1)!, \quad (1.10)$$

or, for half-integral arguments:

$$\Gamma\left(\frac{1}{2} + x\right) = \frac{(2x)!}{4^x x!} \sqrt{\pi} = \frac{(2x-1)!!}{2^x} \sqrt{\pi}, \quad (1.11)$$

Setting $\frac{n}{2} + 1 = x + \frac{1}{2}$ or $x = \frac{n+1}{2}$ in equation 1.11, the value $\Gamma\left(\frac{n}{2} + 1\right)$ can be reduced to:

$$\Gamma\left(\frac{n}{2} + 1\right) = \frac{\sqrt{\pi} (n+1)!}{2^{n+1} \left(\frac{n+1}{2}\right)!}. \quad (1.12)$$

Equation 1.12 is required to calculate the volume where n is odd, where n is even the volume can be calculated directly using the Gamma function in equation 1.10.

The fraction of the space of a hypercube occupied by a concentric hypersphere of the same major dimension can be calculated as:

$$f_o = \frac{V_s}{V_c} = \frac{\pi^{\frac{n}{2}}}{2^n \Gamma\left(\frac{n}{2} + 1\right)}. \quad (1.13)$$

For even n , this reduces to:

$$f_o = \frac{\pi^{\frac{n}{2}}}{2^n \left(\frac{n}{2}\right)!}, \quad (1.14)$$

and for odd n :

$$f_o = \frac{2\pi^{\frac{n-1}{2}} \left(\frac{n+1}{2}\right)!}{(n+1)!}. \quad (1.15)$$

The factorial in the denominator for both equations 1.14 and 1.15 is the most rapidly increasing term for increasing n , thus the fraction of space in the hypercube that is occupied by the hypersphere diminishes rapidly (see figure 1.13), tending to zero as $n \rightarrow \infty$. This shows that as dimensionality increases the volume of a hypercube concentrates into the corners.

The implication of these findings is that high-dimensional space becomes, in some sense, increasingly empty. The space is concentrated away from the centre of the hypercube, in the corners; similar arguments can be followed to demonstrate that the volumes of a hypersphere and a hyperellipsoid concentrate increasingly in an outer shell (see Chapter 5 of Landgrebe (2003)), i.e. the volume of the space is concentrated away from the origin of the space.

From the point of view of statistics in high dimensional spaces, it can be shown that the increase in the volume of space in a shell at a distance r for increasing n is more rapid than the decrease of the normal probability density for the same r and n (again, see Landgrebe (2003)). As a result, the peak of the “probability mass” moves away from the mean of the normal distribution as n increases, which means that for

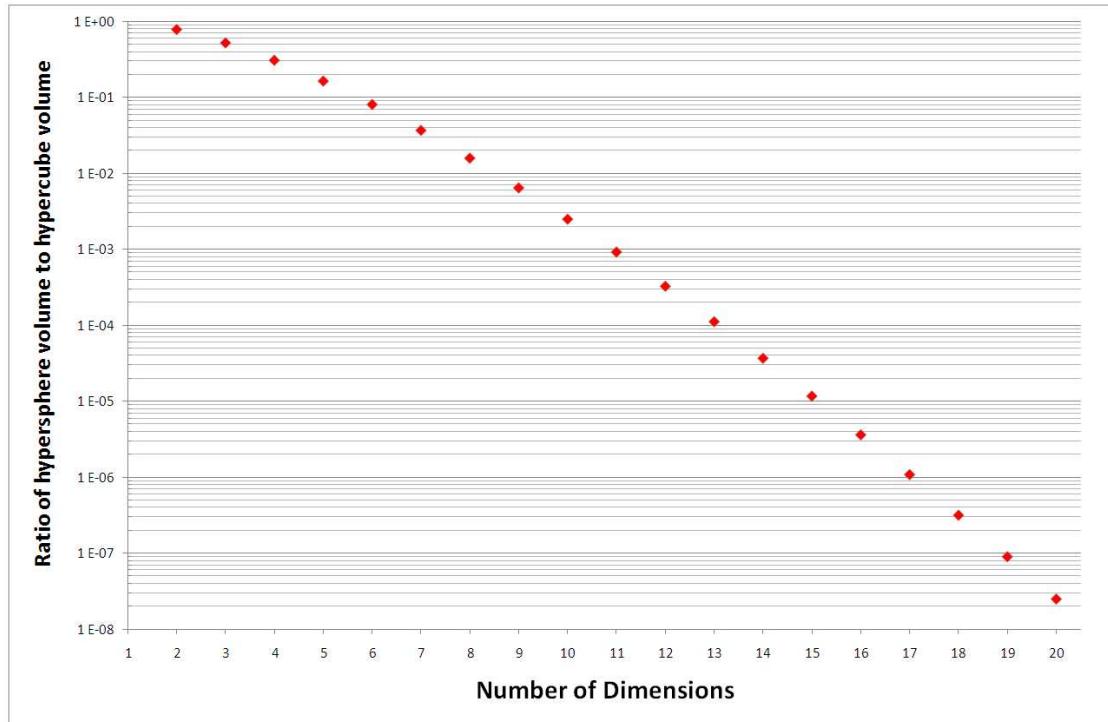


FIGURE 1.13: A graph showing the ratio of the volume of an n D hypersphere of radius r to an n D hypercube of side $2r$.

normally distributed data, increasing the number of dimensions has the effect of placing increasing amounts of data into the tails of the distribution.

These considerations of high dimensional space suggest that in most real-world situations, high-dimensional data, which are likely to be approximately normally distributed, can be projected into a lower dimensional subspace without significant detriment to the information content. There is thus a trade off in deciding how many dimensions in which to work in order to create an efficient and effective system and the working maxim must be data≠information.

Chapter 2

Skin Damage Imaging and Treatment

As the most visible and accessible organ of the human body, skin is an obvious starting point for consideration of medical applications of spectral imaging. Discussions on the matter with Professor Eugene Healy of the School of Medicine, University of Southampton, and Professor Brian Diffey, Emeritus Professor of Biomedicine at Newcastle University, revealed that one area in which dermatologists could be assisted by spectral imaging would be in the detection of psoriasis. These discussions resulted in the formulation of a project whose ultimate aim is the construction of an unsupervised, or minimally supervised, phototherapy treatment system for psoriasis.

This chapter details the work done in developing this idea. It will begin by providing a background to psoriasis to explain the need for an improved treatment system and explaining what role a detection system would play in that. Following this is a review of the medical literature as to the detection and monitoring of psoriatic plaques, and the nature of the phototherapy currently on offer. A review of the viability of producing a laser source for such phototherapy is then presented, followed by a specification of the system requirements.

This part of the project is no longer under active development, the reasons for this, and the conclusions of the research conducted, are presented in a discussion at the end of this chapter.

2.1 Background

The background to this work is split into three sections, a general introduction to psoriasis and the problems encountered by current treatment regimes, a review of the medical literature relevant to the detection and phototherapy treatment of psoriasis,



FIGURE 2.1: A psoriatic plaque on the skin of a *psoriasis vulgaris* sufferer(Kane et al., 2002).

and finally a review of the laser science relevant to the development of a new light source for phototherapy.

2.1.1 A Brief Introduction to Psoriasis

Psoriasis is a chronic skin condition that causes acceleration of the replacement of skin cells. This excessive regeneration causes a build up of dead, often scaly, skin on the surface, which is known as a psoriatic plaque (see figure 2.1). These plaques can affect any area of the surface of the body from small, localised spots through to almost complete body coverage; they are often itchy or painful. It is a condition that affects an estimated 2-3% of the population of the UK. Around 80% of sufferers are affected by common plaque psoriasis (*psoriasis vulgaris*), which tends to manifest around the knees, elbows, scalp and lower back¹.

Diagnosis is made following a visual examination of the affected area by a GP or dermatologist. For mild cases, a range of treatments consisting chiefly of ointments and lotions to be applied to the affected area can be prescribed by a GP. More severe cases are typically referred to a specialist for further treatment options. There are a number of medications available for treatment and management of severe cases of psoriasis, including biologics; these are generally designed to slow down the production of cells. Many cases (an estimated 60,000 each year in the UK (Diffey, 2008) however are treated by UV phototherapy.

¹Figures based on statistics published on The Psoriasis Association website, <http://www.psoriasis-association.org.uk/what-is.html>

It has long been established that sufferers of psoriasis can benefit from exposure to sunlight and it is now known that it is the UV part of the spectrum, specifically UVB (280-315nm), that confers benefit by triggering chemical reactions that retard the rate of cell production. Phototherapy involves irradiating the skin with UVB light, traditionally from a broadband source such as a mercury-xenon arc lamp. More recently narrowband and monochromatic sources have been developed to minimise unnecessary exposure.

UVB exposure is a known risk factor for carcinogenesis and both UVA and to a lesser extent UVB are known to cause erythema (sunburn) after sufficient exposure; it is also known that psoriatic plaques are more tolerant of UV radiation than normal skin. Consequently, dosimetry is an important factor in the extent of applicability of phototherapy; in individual sessions in order to minimise erythema risk, and across the treatment period, and indeed lifetime, of a patient to minimise the cancer risk. To this end, “sun bed” style phototherapy, where the patient is placed in a light-cabinet which exposes the whole body to UV light, is being replaced with more targeted treatments, which are discussed in more depth in section 2.1.2 below.

At present these treatments consist of narrowband light sources that are directed to the skin by optical fibre arrangements, targeted by hand-held probes. The treatment sessions are thus supervised by healthcare professionals. This increases the costs of this form of therapy as opposed to the light-cabinet form, which can be largely unsupervised. It is anticipated that by combining the unsupervised nature of the light-cabinet style treatment with the benefits conferred by targeted treatment, cost savings could be made by treatment centres and time savings could be made for patients. The development and requirements of such a system are discussed in section 2.2 below.

2.1.2 Medical Review

This section seeks to outline the current state of the art in two areas; firstly the detection and monitoring of psoriasis, and secondly, the phototherapy treatments available. It should be noted that it is not intended to present evidence here in support of the use of phototherapy as a treatment as this would be beyond the scope of the work undertaken; however, the specific use of narrowband UVB irradiation is considered.

There is a relative paucity of research into spectroscopic methods for identifying and demarcating psoriasis and psoriatic plaques. Clearly this is because visual identification of such plaques is comparatively straightforward. The plaques exhibit on the surface of the body and are typically red and/or scaly and will often be identified by the patients themselves.

Although a plaque itself is visually distinct from the skin surrounding it, visual inspections are not sufficient to determine the likely direction of development of the plaque if left untreated. This 'advancing' or 'active edge' has been studied using Laser Doppler Flowmetry (Davison et al., 2001) and by Laser Doppler Imaging (Hull et al., 1989; Murray et al., 2005). Hull et al. (1989) report that increased blood flow was measured in advancing edges before immunohistological examinations could identify any changes.

Murray et al. (2005) use illumination at two laser wavelengths (633nm and 532nm), these penetrate to different depths in the body due to differences in the absorption coefficients at each wavelength. The paper quotes a 20 fold increase in the absorption coefficient of haemoglobin at 532nm compared to 633nm (Gush and King, 1991). This makes absorption of 633nm light by blood in small capillaries, which are closer to the surface of the skin, much less likely than light at 532nm, thus larger blood vessels, which are deeper into the body, were more likely to be imaged by the red wavelength and smaller vessels by the green. It was found that blood flow was increased in all vessels within the plaque and that flow was lower in uninvolved skin away from the plaque than skin adjacent to it. This effect was seen to be more marked for the red wavelength and hence the larger vessels.

Bissonnette et al. (1998) report finding an auto-fluorescence peak at 635nm in psoriatic plaques illuminated with blue (442nm) light. This is attributed to the presence of the biological molecule protoporphyrin IX. Although this fluorescence was only detected in 45% of the plaques studied it was not found at all in clinically normal skin or diseased skin of other dermatoses.

Delgado Gomez et al. (2007) present an imaging system designed to track changes in psoriatic plaques across a number of weeks. This system acquires and registers digital images of the plaques and automatically separates lesional from normal skin in the image. Changes are detected using a number of methods, the merits of which are compared. It is found that a principal component analysis is useful for displaying the changes across a series of images.

The human erythema action spectrum is shown in Figure 2.2 below. The relative effectiveness on the vertical axis is a measure of how effective energy at a given wavelength is in causing skin to burn after some exposure period, as such it represents a relative risk of burning due to exposure to radiation at that wavelength. In this case, it can be seen that for wavelengths greater than 300nm the risk of burning after exposure to radiation decreases dramatically.

A number of narrowband sources have been developed with emissions in the 305nm - 315nm range; for example Solarc Systems Inc.² manufacture systems with lamps

²http://www.solarcsystems.com/us_selection_guide.html

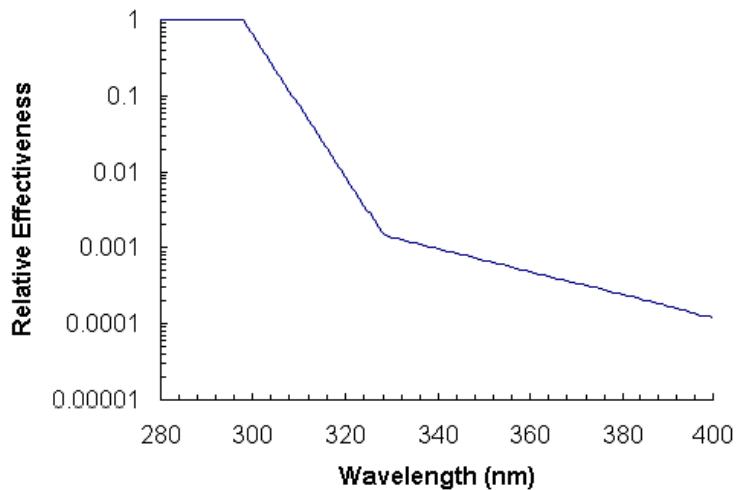


FIGURE 2.2: The human erythema action spectrum for UV light (Wong and Parisi, 1999) (image modified to remove unnecessary information from the original).

emitting at 311nm and PhotoMedex³ manufacture the VTRAC Excimer Lamp which emits at 308nm. The choice to operate at these wavelengths is informed by a ubiquitous reference to Parrish and Jaenicke (1981).

This study presents an action spectrum for phototherapy of psoriasis based on measurements of the effectiveness of radiation doses at discrete wavelengths between 254 and 313nm. It concludes that radiation at 254, 280 and 290nm is erythemagenic without being therapeutic. The action spectrum generated shows a maximum at 300nm with a sharp drop off either side of this maximum. The paper makes note of the rapid decline in therapeutic effectiveness below around 290nm, similar levels of effectiveness are visually observable above 310nm although these are not specifically mentioned by the paper.

Many studies exist to demonstrate the efficacy of treatments in the 305-315nm regime, for example Köllner et al. (2005) and Asawanonda et al. (2000).

2.1.3 UV Laser Generation

An effective targeted treatment will require a suitable light source. The combination of the erythema and phototherapy action spectrum discussed in section 2.1.2 above restrict the light source to a narrow spectral range of around 300nm to 310nm (hereafter referred to as the therapeutic window). Below this the risk of erythema becomes significant and above this the therapeutic effectiveness decreases sharply. In this section the availability of lasers for the UV regime shall be discussed. This narrow therapeutic window means that a monochromatic light source would be more

³<http://www.photomedex.com/vtrac/specifications.htm>

appropriate than a broadband source. A laser source would minimise the time required to deliver a therapeutic dose and thus would be the ideal monochromatic source for this application. This has the further benefit that laser beams can generally be restricted to small spot sizes, thus delivering very concentrated and localised doses of radiation, and minimising the risk of delivering a harmful dose to the healthy tissue surrounding the lesion.

As mentioned in section 2.1.2 above, PhotoMedex manufacture a treatment system based on an excimer lamp, which is a monochromatic light source at 308nm. Lasers are available commercially at this wavelength based on a xenon chloride (XeCl) excimer.

Excimers (a contraction of excited dimers) are short lifetime molecules formed when one monomer is in an excited electronic state. Their short lifetime is due to the fact that when the monomer relaxes back to the ground state the chemical bond is unsustainable and the molecule breaks apart. In the case of XeCl the xenon atom is promoted into an excited energy state leaving it with a free outer-shell electron. This configuration enables it to form a bond with a chlorine atom, which has one incomplete outer-shell orbital. At relaxation the excimer dissociates very rapidly. The atomic separation of the monomers will be in a local energy minimum whilst they are bound, as the excited monomer relaxes this same separation will be at a lower energy level on a repulsive curve, the difference between these energy levels is emitted as a photon.

Other wavelengths are achievable using this same principle by changing the inert gas or the halogen used. These lasers are typically also UV, for example krypton fluoride (KrF) excimers will lase at 248nm and xenon fluoride (XeF) at 351nm. XeCl is the only excimer laser in the therapeutic window.

Although commercially available and in use for laser-eye surgery, excimer lasers are not without their drawbacks. The combination of high voltage and corrosive gases in the lasing medium and resonator means that early generation cavities had very short lifetimes. Extensive engineering has mitigated the potential for this combination to damage the laser however regular servicing is still required to keep the laser operable. A further drawback is the gas lifetime, this is of the order of 10^7 pulses (for XeCl), however operating at near kilohertz repetition rates as anticipated (see section 2.2.2 below) this gives only a few hours operation between gas refills.

There are no other standard commercial lasers in the therapeutic window and as such an alternative to excimer lasers would need to be designed specifically for this application.

A common way to achieve new laser wavelengths with existing gain media is to use a non-linear combination technique called sum-frequency generation (SFG). Two photons incident on a non-linear crystal (with frequencies f_1 and f_2) emerge as a single photon (with frequency f_3). By energy conservation we see that:

$$hf_3 = hf_1 + hf_2, \quad (2.1)$$

where h is Planck's constant, or, since wavelength is inversely proportional to frequency:

$$\lambda_3 = \frac{\lambda_1 \lambda_2}{\lambda_1 + \lambda_2}. \quad (2.2)$$

If the two incident photons have the same frequency, the frequency of the resulting photon will be twice that of the incident photons, this process is known specifically as second harmonic generation (SHG) or frequency doubling. Green lasers at 532nm are a common example of this, the laser light is the frequency doubled 1064nm line of an Nd:YAG (neodymium-doped yttrium aluminium garnet) laser. Strictly speaking this radiation is caused by non-linear processes and not stimulated emission so it is technically incorrect to term such sources lasers; however this is the prevailing nomenclature in the field as the two types of radiation share many characteristics such as coherence and monochromaticity, as such the term laser will be used in this report for such a source.

Generation of UV lasers by SFG is long established, Blit et al. (1978) report a tunable continuous wave (CW) UV laser as early as 1978. This system uses a krypton laser and an argon pumped dye laser; it was tunable from 285nm-400nm with a peak power output of $750\mu\text{W}$ at 313nm. In order to achieve this power output the krypton laser was running at 5W and the argon laser was pumping the dye laser at 5W. The output from the dye laser was 1.2W. Such systems are typically very large and require water cooling.

Efficiency of conversion is a major factor in the use of SFG, in order to obtain efficient output care must be taken to phase-match the input photons by an appropriate choice of non-linear crystal and crystal orientation. Further, if input lasers are pulsed rather than CW, then care must be taken to match pulse lengths and/or timings as SFG will only occur for the overlapping parts of each pulse.

Some of the other practical issues associated with efficient SFG are described by Mimoun et al. (2008) in their paper demonstrating efficient generation of a 589nm laser using intra-cavity SHG. The discussion in this paper is focussed on the reflectivities of the cavity mirrors that will minimise power loss in the cavity.

Even with SFG, commercially exploitable lasers to achieve a wavelength in the therapeutic window are rare, however employing this in combination with a technique called Stimulated Raman Scattering (SRS) could potentially allow the development of a suitable laser.

SRS is a process based on the Raman scattering discussed in section 2.1 above. Laser radiation is passed through a material, most commonly a gas cell containing hydrogen (H_2), deuterium (D_2), methane (CH_4), or some combination of these, where it is absorbed and stimulated into being emitted in a process analogous to lasing. The emerging radiation is Raman shifted by a frequency characteristic of the material. Again, if the shift is due to a loss of energy the resulting frequency is referred to as a Stokes line, if due to a gain it is an anti-Stokes line. Typically the strongest line is the first Stokes line as emission at higher Stokes lines requires an interaction of this line with the scattering medium.

The characteristic wavenumber shift (k_S) of the first Stokes line is measured in wavenumbers, for H_2 it is 4155cm^{-1} , for D_2 it is 2987cm^{-1} and for CH_4 it is 2916cm^{-1} . The resulting wavelength, λ_o , is thus related to the original, λ_i by the equation:

$$\lambda_o = \frac{\lambda_i}{1 - \lambda_i k_S}. \quad (2.3)$$

The energy “lost” for a Stokes shift is actually absorbed by the shift medium, and is proportional to k_S , causing heating of the medium. For an anti-Stokes shift, the energy “gained” is drawn from the shift medium, causing a cooling.

As with SFG, SRS is a long established process in the UV. Loree et al. (1979) demonstrated the generation of new UV lines by broadband non-resonant scattering of various excimer lasers in various media (including the aforementioned gases and liquid nitrogen). Using KrF (at 248nm) scattering in H_2 they demonstrate the generation of a line at 276nm with a power of around 25% of that of the pump, demonstrating that high efficiency is possible for SRS.

More recently, Stoffels et al. (1997) demonstrate the generation of 224nm radiation from an ArF laser (193nm) by performing SRS in a H_2/D_2 mixture. Here the 224nm radiation is generated due to the shifting of the first H_2 Stokes line (at 210nm) in the D_2 combined with the shifting of the first D_2 Stokes line (at 205nm) in the H_2 . The result of both multiple shifts is a total shift to 224nm hence the efficiency of the generation of this line is much higher than would normally be expected for a second order shift.

SRS pumped by solid state lasers has also been demonstrated. Uesugi et al. (2000) demonstrate SRS pumped by the second harmonic of a tunable Ti:sapphire laser (spectral range 375-425nm, up to $200\mu\text{J}$ at 1kHz repetition rate), producing an output range of 421-657nm using the first and second Stokes lines. Conversion efficiencies of > 10% for H_2 and > 20% for CH_4 are reported for a high pressure (up to 60atm) gas cell.

Xu et al. (2003) used a third harmonic Nd:YAG (355nm) to pump a gas cell filled with D₂ and mixtures of D₂ with Ar and He. This was performed at low repetition rate (10Hz) at 85mJ. Again good conversion efficiencies are found as pressure increases, up to around 25% for the first Stokes line in pure D₂ at \sim 2600kPa (\sim 25atm). The report notes that efficiency can be increased to around 33% in the D₂/Ar mix however this occurs at the expense of the stability of the resulting pulses (as measured by the energy drift of the resulting radiation).

Ermolenkov et al. (2005) use a combination of frequency doubling followed by SRS followed by a further frequency doubling to produce a UV laser at 281nm from the 1064nm Nd:YAG line (1064nm \longrightarrow 532nm \longrightarrow 563nm \longrightarrow 281nm). Again, this was performed at a low rep rate (10Hz) with the output UV wavelength having a pulse energy of 4.2mJ. A beam divergence of 1.3mrad was measured for this UV radiation.

This laser was developed for diagnosing tropospheric ozone by differential absorption lidar (DIAL). This area of research has provided several other SRS UV laser systems as measurements need to be made in the Hartley window, 280-300nm. Haner and McDermid (1990) discuss Raman shifting of a fourth harmonic Nd:YAG laser (266nm) in H₂, resulting in a laser wavelength of 299nm. Again a 10Hz pulsed laser is used with a pulse energy of around 100mJ. The conversion efficiency is studied as a function of pump energy at 400lbf/in² (27atm). Experiments in D₂ determined that conversion efficiency increased with gas pressure up to 27atm but no significant change was seen as the gas pressure increased from there to around 50atm. Efficiency was found to decrease with increasing pump energy from around 70% at 10mJ to around 40% at 70mJ. Above 20mJ the total of the residual pump plus first Stokes radiation began to decrease as energy was lost to higher order processes.

Raman shifting of a 266nm laser is also investigated by de Schoulepnikoff et al. (1997). Again a 10Hz laser is used and single gas scattering in H₂, D₂ and CH₄ are tested along with various mixtures of these. In H₂ conversion efficiency for the first Stokes line is found to maximise at low gas pressure (>50% at \sim 2atm), decreasing to around 15% at 30atm. The effect of adding buffer gases (Ar, He and Ne) was studied and a conversion efficiency of 61% was measured in H₂ at a partial pressure of 10atm and Ar at a partial pressure of 33atm.

In investigating other gas combinations, it was found that a mixture of H₂, D₂ and Ar could produce close to 50% conversion efficiency to 299nm radiation at an Ar pressure of 16atm, a H₂ pressure of 7atm and a D₂ pressure of 19atm. The beam quality was also measured for these partial active gas pressures and an Ar pressure of 33atm and a full angle divergence of 0.42 ± 0.02 mrad was found for the 299nm line.

Further studies by Simeonov et al. (1998) and Milton et al. (1998) also demonstrate similar conversion efficiencies in H₂ to generate 299nm light. The details of these studies differ slightly but broadly they add little new fundamental insight.

2.2 System Requirements

The ultimate aim of this project is, as stated above, the development of an unsupervised, or minimally supervised, treatment system that is able to detect psoriatic plaques and safely guide UV treatment. As envisaged, this system would consist of two components, an imaging system and a treatment beam. The initial technical requirements of each of these components can be derived from what is already known about human skin in general, and psoriasis in particular. This section addresses the key specifications of each component in turn.

2.2.1 Imaging System

An imaging system that merely demarcates a plaque on a static image, whilst an important first step, will ultimately be of limited use to this project and has been demonstrated elsewhere (Delgado Gomez et al., 2007).

To be of value to an integrated treatment system the imaging system will need to be capable of real-time or near real-time tracking of any movement of a plaque. Real-time video is typically thought of as a frame rate of 20-30 frames per second (fps); standard-definition television broadcast and cinema have frame rates of 24/25 fps (depending on the precise format), a typical digital camera will record video at around 30fps, and high-definition television broadcast is around 50-60fps (again, precise format dependent). For this system, the minimum acceptable frame rate will ultimately be determined by the specification of the laser system. Movements that place uninvolved skin in the path of the treating radiation will need to be measurable in sufficient time for the laser to be redirected or shut off before it poses an erythema risk.

Further, any imaging system will need to be able to provide information on plaques that follow the natural contours of the body. It is entirely conceivable for example that a single plaque could extend around a large proportion of the circumference of a limb and could not be presented in its entirety to a single imaging plane.

The first issue is largely a software problem. Raw imaging systems exist that enable real time image capture, indeed the Prosilica camera used in this project and detailed in section 5.2.1 below, is capable of video capture at 20fps and a high speed counterpart is available with 30fps at full resolution. This camera connects to a computer using gigabit ethernet, which allows data transfer at 10^9 bits per second or ~ 120 megabytes (MB) per second. Each frame is around 5MB, which is compressible to around 2.65MB by the removal of zero-padding. Clearly then, neither the CCD nor the data connection will limit the ability to image in real time.

The rate limiting step would then seem to be processing the image to determine the border of the plaque. It is anticipated however that the image processing used to

demarcate the plaque will be performed along similar lines to the processing discussed in section 2.3 above. A process diagram is shown in Figure 2.3. Clearly in order for this to work the plaque would need to be imaged at two (or possibly more) wavelengths, each frame displayed (or used to instruct the laser) then becomes at least two exposures and the rate at which filters can be changed becomes a further constraint to the maximum frame rate. It is likely that this change of filter can be achieved at a sufficient rate for it not to be limiting, which again leaves the image processing as the determining factor.

Developing an automatic processing code and optimising this will thus need to be the focus of work in the initial development. The choice of filters in this case can be informed with some confidence from the literature. This would suggest that focussing on changes in blood content should be sufficient to distinguish plaque tissue from surrounding healthy tissue. Although any choice would need to be experimentally verified, a filter around 540nm or 580nm (resonant with the HbO₂ absorption peak) and one in the infrared at around 800-900nm (off-resonant) would seem a sensible starting point.

Mapping plaques that do not present in their entirety to a single image plane presents an interesting choice of approaches that will need to be considered early in the design stage of an integrated treatment system.

The obvious approach would be to mount the laser and the imager co-linearly and constrain the movement of the treatment beam to within the image boundaries (see Figure 2.4a). The imager and beam could then rotate around the body on a rail until full coverage has been achieved. This approach is implicit in Figure 2.3. An alternative approach would be to have two cameras mounted at some angle to each other in order to form stereo images to create a 3D image; the treatment beam would then be able to move between the cameras whilst still being able to see the region of interest (see Figure 2.4b). In this situation the whole set-up would still need to rotate around the object, although the approach could be extended to a number of cameras placed all around the object.

The first approach (single imager) has the advantage of being easier to implement. Implementing a two imager approach will require an additional image processing step to create and interpret the 3D image. Since processing as the rate limiting step has been discussed, it may well be that this extra burden will make the whole process too slow to be of use.

If the second approach could be made to proceed at sufficient speed to make it useful it has the advantage of being able to “see” where the body is curving away from the image plane. Looking back at Figure 2.4a, it is clear that not only is the distance of the object from the image plane greater at the edges than in the centre, but the angle of incidence of the treatment beam is also different. Such a difference will affect the

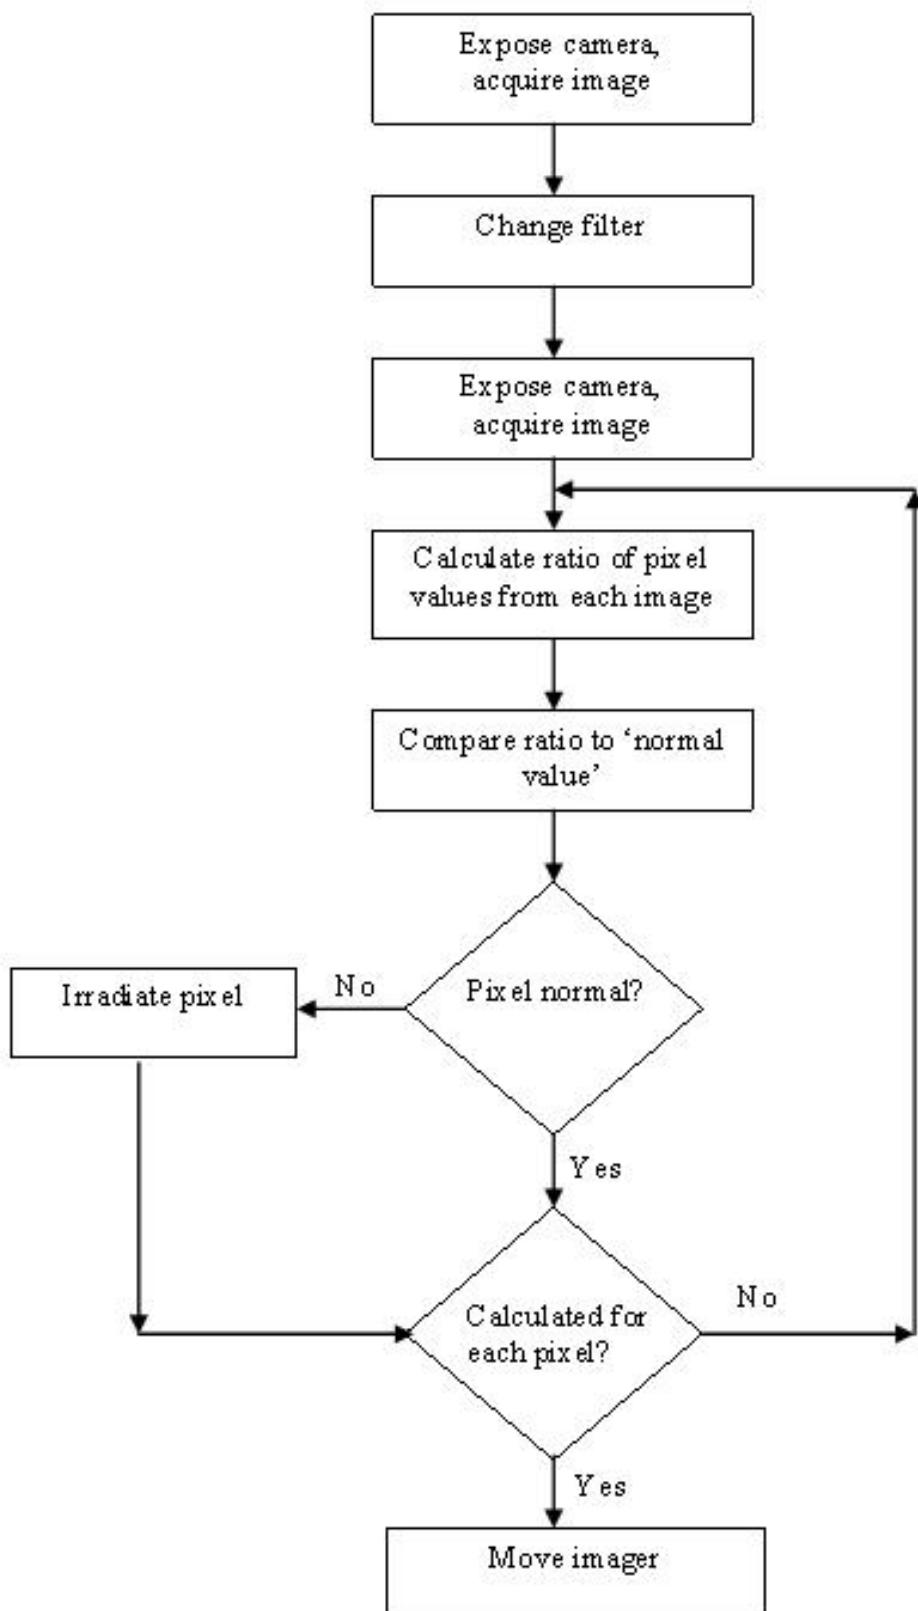


FIGURE 2.3: Process flow diagram showing simplified anticipated imaging sequence.

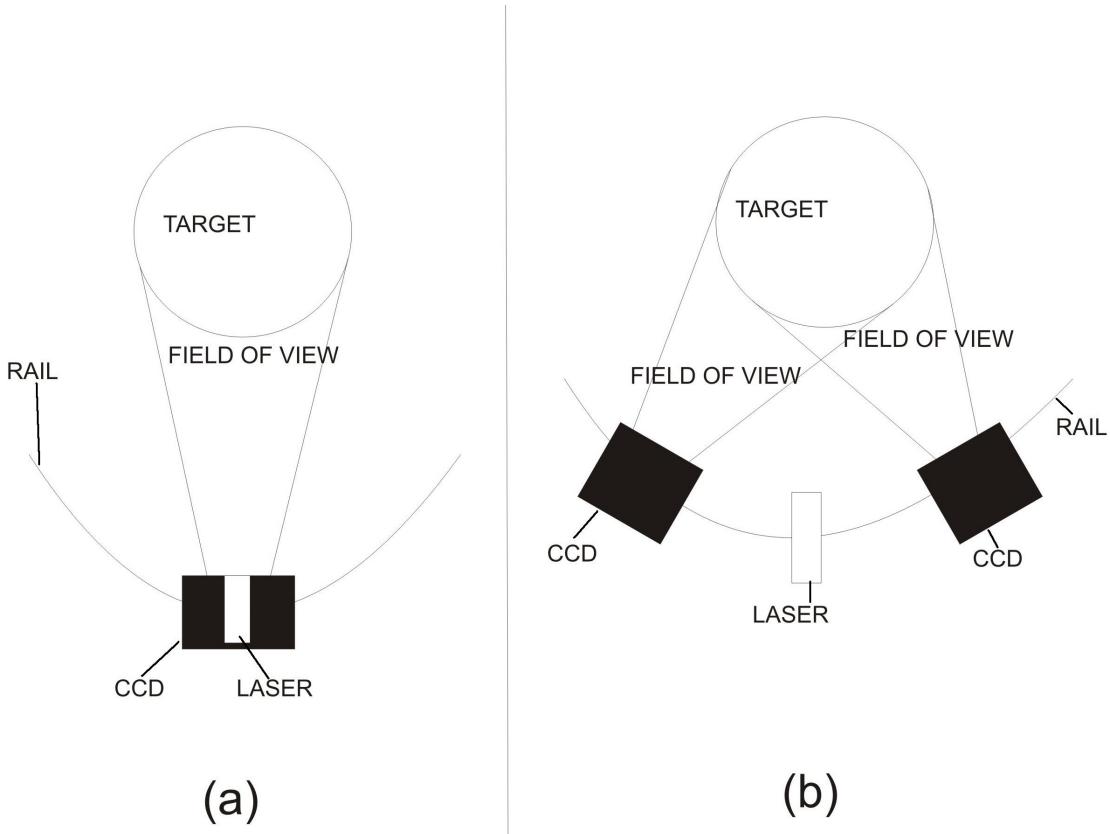


FIGURE 2.4: Alternative approaches for the layout of the imaging/treatment system.

received dose of radiation for different areas of the object. Although this illustration is likely to prove an exaggeration of most real situations in terms of the curvature, it is important to keep these issues in mind from an early stage so as to avoid designing in an inherent flaw.

A compromise solution may be to use a single imager and restrict the beam movement to a central region of the image where the object can be regarded as being flat and parallel to the image plane. Doing this would combine the ease of implementation with the advantage of consistent dosimetry to all points. This approach could also bring an extra benefit of reducing the effective size of the image that requires processing, allowing the speed of this step to be increased.

2.2.2 Laser System

These specifications for a laser system to provide a treatment beam have been drafted according to current treatment regimes.

As discussed in section 2.1.3 above, to provide a safe and effective treatment the laser should have a wavelength in the therapeutic window ($\sim 300\text{-}310\text{nm}$), also discussed is the lack of availability of off-the-shelf solutions in this area. It was proposed that a

bespoke laser could be developed by taking the fourth harmonic of the 1064nm Nd:YAG laser line (which is a common laser) and subjecting this to SRS in H₂. By equation 2.3, the resulting radiation would be at a wavelength of 299nm, which whilst just outside the window should be usable (Diffey, 2008).

From discussions with Profs. Healy and Diffey, an initial radiation dose of 80mJ/cm² at 299nm has been proposed as broadly representative of existing treatments. It has further been suggested that a treatment time in the region of 20 minutes for a patient with approximately 35% body coverage would provide a suitable target.

The body surface area of a typical adult⁴ is usually given as around 1.7m², thus 35% coverage equates to $\sim 0.6\text{m}^2$. For a radiation dose of 80mJ/cm² in 20 minutes, the patient would need to be exposed to 480J in 1200s, this gives an average power rating for the treatment beam of 400mW. 266nm lasers are available that are consistent with the conversion efficiencies achievable in SRS, for example, the AVIA 266-3⁵, which has a power output of 3W at a 30kHz rep. rate. For this particular laser a conversion efficiency of around 13% would be required.

The power absorbed by the gas cell in generating a 400mW, 299nm beam from a 266nm beam can be estimated to first order by considering the wavenumber of the output beam ($k_o = 33438.98\text{cm}^{-1}$) compared to the wavenumber of the shift ($k_S = 4155\text{cm}^{-1}$). Since energy, $E \propto k$, the absorbed power, P_a can be calculated from the output power, P_o using the equation:

$$P_a = \frac{P_o k_S}{k_o}, \quad (2.4)$$

using $P_o = 400\text{mW}$, P_a is found to be $\sim 50\text{mW}$. The precise absorbed power will be exacerbated by any higher order Stokes shifts, and mitigated by any anti-Stokes shifts. This will cause a localised heating of the H₂ gas, which could in turn create shock-waves through the entire medium due to the pressure fluctuations resulting from the heat. This is similar to the problems faced by a gas based laser system, and can be controlled by having a steady flow of gas through the cell rather than a static volume. Such a flow, or such fluctuations caused by heating will affect the beam quality, but as will be shown below, the beam quality is not likely to be critical in this system.

A resolution of the order of a millimetre is expected to be sufficient, implying that a beam size of 1mm at an anticipated source-target separation of 1m will be required. This implies a beam divergence of around 1mrad. To compare this to the diffraction limited spot size in this case an aperture of 1cm is assumed, both for mathematical convenience, and as this would allow for low cost and lightweight optics. At a focal

⁴<http://www.emedicinehealth.com/script/main/art.asp?articlekey=39851>

⁵http://www.coherent.com/downloads/AVIA_266-3_DS_final.pdf

length of 1m, this implies an F -number of 100. The diffraction limited spot size, d , is given by:

$$d = 1.22\lambda F, \quad (2.5)$$

which, for $\lambda = 299\text{nm}$, gives $d \approx 36\mu\text{m}$. This system would thus only need to work at $\sim 30\times$ the diffraction limit.

The minimum erythema dose (MED) is taken to be 27mJ/cm^2 (Giacomoni, 2007) at 300nm (0.27mJ/mm^2). For a pulsed laser at 400mW , this means that a rep. rate of at least $400\text{mW}/0.27\text{mJ} = 1.48\text{kHz}$ will be needed to keep the pulse energy below the MED. The higher the rep. rate the safer the system will be, ideally a CW laser would be used. The thermal relaxation time for skin is around 1ms (Bernhardt, 1998), which, with a rep. rate $> 1.5\text{kHz}$, means that the pulse length of the laser would always be lower than the relaxation time, thus the entire energy of the pulse should be considered to be deposited simultaneously. Although each pulse is delivering less than the MED, the effect of multiple pulses being incident on the same area of tissue would need to be considered to ensure that no thermal damage is done.

The specification is summarised Table 2.1.

Property	Value
Wavelength	300-310nm ideal, 299nm acceptable
Power	400mW
Rep. Rate	CW Ideal, $>1.5\text{kHz}$ minimum
Beam divergence	$\sim 1\text{mrad}$

TABLE 2.1: Initial specification for UV laser

2.3 Discussion

The aim of this project was to consider the potential application of spectral imaging to dermatology. From the initial discussions with Prof. Healy, it is clear that there is interest in the technology. The specific application of the treatment of psoriasis was identified as an area in which spectral imaging could be put to use, and a well defined problem was formulated. Could a spectral imager form the basis of a detection, monitoring and targeting component of an unsupervised, or minimally supervised treatment system?

A review of the literature indicated that the basic principle of detecting and demarcating psoriatic plaques on an image had been established Delgado Gomez et al. (2007), which is an encouraging start. In order to detect plaques *in-situ*, it was realised

that there were two realistic options for the set-up of the imaging system (as shown in figure 2.4). The first of these, co-locating the laser to a single CCD, was considered to be the most feasible. In the time since then however, some moves towards low cost, real-time 3D imaging have been made using the Microsoft Kinect system as a base (Kreylos, 2010). This work would seem to be somewhat premature for this project, coming as it did before the release on an official Software Development Kit for non-commercial users.

Concurrent with this research, work was undertaken to establish the viability of the second component of the system, that of an effective, targeted treatment beam.

Research on phototherapy established that a laser with well defined properties was required to achieve this. There seems to be no indication in the literature that the particular combination of requirements has been demonstrated in a laser before.

Various steps required had been demonstrated on their own; for example the Raman shifting of a 266nm laser to produce a 299nm source, the requisite conversion efficiency, and beam divergence have all been achieved. However, these have always been in low rep. rate lasers (10Hz), and for beam powers much lower than would be needed here. Further, there does not seem to be any investigation as to the effect of rep. rate on conversion efficiencies or beam quality. This provides a good starting point for further research on the development of such a laser.

It was decided that further research on this treatment system was actually out of the scope of the research upon which this thesis is based. The imaging system described by Delgado Gomez et al. (2007) is not based on spectral imaging but rather on standard RGB imaging; further, as the principle has been demonstrated, the imaging component is more properly the subject of development than research. There is scope for spectral imaging research into the advancing edge of psoriatic plaques, and if it could demonstrated it would make for a useful add-on to the treatment system, however, this research was deemed to be outside of the scope of the project.

By contrast, the laser system would indeed be the subject of research rather than development, however, it is more properly within the remit of a UV laser group.

Chapter 3

Pixel Classification

From a scientific viewpoint, an image is only as useful as the information that can be gleaned from it. With that in mind, this chapter discusses what kind of information can be extracted from a spectral image; some basic concepts with roots in more conventional image processing are discussed before moving on to methods that are more spectral in nature. A novel extension to one of these methods, the Spectral Angle Map (SAM), is also discussed in this chapter.

For conventional imaging, the aim is often to identify an object (or subject) to compare to some target or to track across a time series of images. The ultimate aim for most spectral imaging is to classify pixels according to the material they represent and this is certainly true for the biomedical applications which are the subject of this thesis. The overlap of function here, i.e. matching a candidate pixel, or region to a known target, is diminished by the fact that for conventional imaging this matching is normally a spatial matching, whereas for spectral imaging the spatial configuration of pixels is less important, and the matching is spectral-based. This distinction makes it unsurprising to find that the techniques used in spectral imaging are quite specific to this particular modality. The first section of this chapter however discusses how spectral images can be analysed using techniques that would not be uncommon in conventional imaging.

3.1 Image Processing Concepts

3.1.1 Difference Mapping

Difference mapping is a conceptually simple method of comparing the content of two images. A composite image is created by subtracting the value of a pixel in one image from the value of the corresponding pixel in the second image, and assigning that difference to the corresponding pixel in the composite.

The obvious use of such a method is to look for changes in images of the same scene taken at different times, in a spectral imaging system however it can be used to compare the same scene at different wavelengths. As described thus far, this method is highly prone to the spectral characteristics of both the illumination system and the image capture system. The graph shown in figure 1.1 demonstrates that a typical silicon CCD is around twice as sensitive to light at 500nm than it is to light at 750nm. For a scene illuminated by spectrally uniform illumination, captured at the same exposure length in all bands, we would therefore still expect to see a noticeable difference in the composite. A similar argument can be made in the case that the illumination of the scene is not constant across all wavelengths.

Shadows within an image, even when they appear at the same locations on both images, will also cause a problem. In areas of shadow the absolute pixel intensity values are much lower than in those areas in bright illumination, as such small absolute changes in pixel intensity represent much larger relative changes in areas of shadow. By extension, the signal to noise ratio (SNR) will be lower in these areas than in brighter areas.

Taking into consideration these effects, when deployed in practice, the calculated difference is usually normalised in some way, often by the average intensity of the pixel in the two images. If a pixel at position (x, y) in image 1 is denoted as $I_1(x, y)$, and the corresponding pixel in image 2 is $I_2(x, y)$, then the pixel in the composite image is given by:

$$C(x, y) = 2G \frac{I_1(x, y) - I_2(x, y)}{I_1(x, y) + I_2(x, y)}, \quad (3.1)$$

where G is a gain factor used to stretch the contrast.

For display purposes, the matrix of composite pixel values is converted to a suitable format by scaling all values to integers with a suitable bit depth, and shifting the zero level to half the maximum possible value, this ensures that all pixels in the composite image have non-negative values. It is also useful to map this composite into a colour space, where the hue represents the magnitude of the difference, as human vision is better at distinguishing colours than it is at distinguishing grey-levels.

3.1.2 Ratio Mapping

Ratio mapping is a method very much related to difference mapping as described above. The pixel values for the composite image in this case are calculated by dividing the value of the pixel in one image with the the value of the corresponding pixel in the second image. Using the nomenclature established in the previous section:

$$C(x, y) = \frac{I_1(x, y)}{I_2(x, y)}. \quad (3.2)$$

Unlike the difference map, this method is inherently robust to local illumination changes, provided that they are consistent between images. The problem of having negative values in the composite is also avoided, in areas where $I_1(x, y) < I_2(x, y)$, the composite value will be in the range [0,1] and $I_1(x, y) > I_2(x, y)$, the composite will be in the range (1, maxvalue]. The disadvantage of this is apparent, the composite output is very non-linear, and the composite values will need to be processed so that the ranges [0,1] and (1, maxvalue] both inhabit the same size integral space. This extra processing step means that in practice this method is very similar to a difference map as normalised in equation 3.1.

Both of these methods are most useful when there is a particular spectral feature in a single waveband that discriminates between a target and the background. From figure 1.6, it is clear that the presence of water is a particularly useful feature for this method. Having a strong absorption peak at 970nm means that selecting the image at this on-resonance wavelength for one of the images, and an image at an off-resonance wavelength (say 900nm) as the second, will show any areas with water present as highly contrasting with the background. A two-band spectral image would be very quick to capture and process in this way and would be ideal in a finished spectral imaging system (such systems are in use in the food sorting industry, known as “bi-chromatic cameras”). However, in early stage research, having such a clear cut spectral feature to use as a discriminant is rare, and this has been the case in this project.

The next section describes some analysis methods that use the full spectral range of an image, and which do not require *a-priori* knowledge of discriminant spectral features.

3.2 Measuring Spectral Similarity

One of the key challenges in spectral imaging is presenting a large volume of data in a meaningful and comprehensible way. This can be achieved by reducing the dataset to a smaller, representative subset, or by calculating a metric or statistic from the data. This is typically calculated at the level of individual pixels, rather than from the entire dataset, and for presentation, these metrics are often mapped to a false colour image. Whereas the methods outlined in section 3.1 are based on the premise of comparing entire image planes, or on a pixel-scale, comparing different dimensions of the same pixel, the methods outlined in this section consider the entire pixel at once, comparing it to other pixels or known reference pixels.

Five measures are briefly introduced before a brief comparison is made on some simple example data.

At this point it is useful to distinguish between a measure and a metric; the former is a generic term for anything that can be measured whereas the latter has a precise

mathematical definition. A metric is a function, d , that defines the distance between two elements, $x, y \in S$, where S is a set, according to the following axioms, $\forall x, y, z \in S$:

1. $d(x, y) \geq 0$, with $d = 0$ iff $x = y$, i.e. distance must be non-negative and only coincident elements have a zero distance;
2. $d(x, y) = d(y, x)$, i.e. the distance must be the same when measured in either direction;
3. $d(x, z) \leq d(x, y) + d(y, z)$, this is the triangle inequality, the distance between two points must not be greater than the sum of the distances from each point to an intermediate point.

3.2.1 Distance Measures

The Pearson Correlation Co-Efficient (PCC) or the Product Moment Correlation Co-efficient (PMCC) is used in statistics to measure the linear dependence of two variables. It is defined as the covariance of the two variables divided by the product of their standard deviations (see equation 3.3). It varies in value from -1 for two variables that negatively linearly correlated, to +1 for positive linear correlation (as such it violates the first metric axiom above). A value of zero means that there is no linear correlation, the variables may still be correlated in some other way.

$$r = \frac{Cov(x, y)}{\sigma_x \sigma_y} = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2 \sum_i (y_i - \bar{y})^2}}. \quad (3.3)$$

In a spectral imaging paradigm, the variables x and y are vectors of data assigned to a pixel. For two similar spectra, one would expect the values in each dimension to be similar and thus the two spectra will be strongly linearly correlated. A negative correlation is less straightforward in this application, the implication is that as values in x increase, the equivalent values in y are decreasing and vice-versa. This would be the case where a peak in x corresponds to a trough in y , but to be true across the entire spectral range it would imply that one of the vectors represented an absorption spectrum and the other a transmission spectrum.

The Bhattacharyya Measure (Bhattacharyya, 1943) is used chiefly in statistics to measure the similarity of two probability distributions. For two probability distributions a and b the measure, or Bhattacharyya Coefficient (BCE) is calculated as:

$$\rho(a, b) = \sum_i \sqrt{a(i)b(i)}. \quad (3.4)$$

As a and b are probability densities, we know that $\sum_i a(i) = 1$ and $\sum_i b(i) = 1$, we can also see that if a and b are identical then:

$$\rho(a, b) = \sum_i \sqrt{a(i)b(i)} = \sum_i \sqrt{a(i)a(i)} = \sum_i a(i) = 1. \quad (3.5)$$

The BCE thus cannot be termed a metric. It can be used to define the Bhattacharyya Distance:

$$D_B = -\ln \rho. \quad (3.6)$$

The distance defined by equation 3.6 does not obey the triangle inequality axiom, and as such is also not technically a distance metric, however it does provide a measure ranging from $0 \rightarrow \infty$. A true distance metric can be defined from the co-efficient however:

$$d = \sqrt{1 - \rho}. \quad (3.7)$$

Spectral Information Divergence (SID) is a measure derived from information theory (Chang, 1999). It relies on the definition of self-information, which is the information associated with observing a particular outcome of a random variable. This information should increase as the outcome becomes increasingly unlikely, and the information gained from observing a number of independent events should be cumulative. Further, no information is gained from observing an outcome that is certain (i.e. has a probability of 1). From these considerations, the self-information for a random event with probability p is:

$$I(p) = -\log_b p. \quad (3.8)$$

In principle, b can be any base, however as information theory is linked to computation, using $b = 2$ is common, and the information is thus measured in bits. From this definition, Kullback and Leibler (1951) introduced a “directed divergence” measure, known as the Kullback-Liebler (KL) Distance or Relative Entropy, to calculate the difference between two probability distributions, x and y :

$$D(x||y) = \sum_i x(i) \log \left(\frac{x(i)}{y(i)} \right). \quad (3.9)$$

This is a non-symmetrical measure, and as such does not represent a true metric as it violates axiom 2 above. Chang (1999) defines the SID as a symmetric divergence measure based on the KL Distance:

$$SID(x, y) = D(x||y) + D(y||x). \quad (3.10)$$

In contrast to the statistical approaches discussed thus far, the Euclidean distance requires a geometric interpretation of the spectral image pixel. Each dimension of the spectral image can be considered as a dimension of a high-dimensional space, each pixel is then a vector from the origin to a point in that space.

In two or three dimensions, it is intuitive that the straight line distance between two points measures the difference between them, and this holds as the dimensionality of the space increases. For two vectors, x and y in an n D space, the distance, d , is defined as:

$$d = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}. \quad (3.11)$$

This is a very straightforward metric, easily understood and inexpensive to calculate. However, it is of very limited use. One major problem is that it is heavily scale dependent. This is easiest to understand in two dimensions. If $x = (1, 1)$, $y = (2, 1)$ and $z = (4, 2)$, then the distance d_{xy} will be 1 unit and the distance d_{xz} will be $\sqrt{10}$ units. In both imaging and spectroscopic interpretations y and z are both representative of the same measurement, the scale factor being explicable by a local change in intensity of illumination, or a local change in the amount of absorbing, emitting or fluorescing material. If what is being assessed is not concentrations of material or illumination conditions, but simply whether two areas of the image are intrinsically similar, the Euclidean distance can be very misleading. This can be countered by normalising each vector and working with the unit vectors.

Like the Euclidean distance, the spectral angle is obvious measure when considering a spectrum as a high-dimensional vector (first introduced as a spectral imaging concept by Kruse et al. (1993)). Any two vectors from a coincident origin are co-planar, forming two sides of a triangle. There is thus an angle between the two vectors, which can be calculated from the scalar product. So for two vectors x and y , the spectral angle, θ , between them is given by:

$$\theta = \arccos \left(\frac{\sum_i x_i y_i}{\sqrt{\sum_i x_i^2} \sqrt{\sum_i y_i^2}} \right). \quad (3.12)$$

Again, much like the Euclidean distance this is a metric that is easily understood and easy to calculate, however, unlike the Euclidean distance it is also inherently scale independent. A scalar multiple on either x or y , or indeed both, would be carried through in both the numerator and the denominator of the arccos term in equation 3.12, cancelling and thus leaving the same angle. This property makes the the spectral angle robust to changes in illumination across an image. It should further be noted

that the spectral angle is calculated using the \arccos function, thus, whilst in principle negative angles could be measured, in practice the domain of \arccos , $[-1, 1]$, limits the calculated value to the range $[0, \pi]$. Further, in spectral imaging the values that are being considered are typically all positive (CCDs record positive values), this restriction further limits the range of spectral angles to $[0, \frac{\pi}{2}]$.

3.2.2 Comparing Spectral Similarity Measures

In this section, some relationships between the measures introduced in section 3.2.1 are considered and their performance on some simple, illustrative data is explored.

For two unit vectors separated by an angle θ , the length of the arc that is swept between them is also given by θ , for small angles this arc length closely approximates the Euclidean distance between the vectors, and thus for small angles the Euclidean distance and the spectral angle are equivalent measures. This intuitive result can be shown mathematically by considering the geometry in figure 3.1.

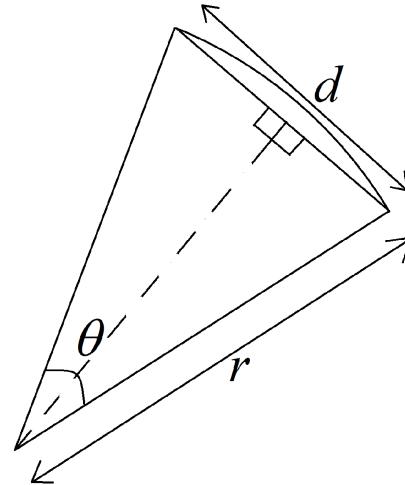


FIGURE 3.1: The geometry of an arc and the Euclidean distance between unit vectors.

Clearly:

$$\frac{d}{2} = r \sin\left(\frac{\theta}{2}\right). \quad (3.13)$$

Since $r = 1$,

$$d = 2 \sin\left(\frac{\theta}{2}\right), \quad (3.14)$$

and since, for small angles, $\sin \alpha \approx \alpha$,

$$d \approx 2 \left(\frac{\theta}{2} \right) \approx \theta. \quad (3.15)$$

This explicit correlation suggests only one of these metrics need be considered any further.

Perhaps more surprising than the relationship between two explicitly spatial measurements, is the connection of the spectral angle to the BCE. As a and b from equation 3.4 are probability distributions, we can also consider them to be the unit vectors $\alpha = (\sqrt{a_1}, \dots, \sqrt{a_i})$ and $\beta = (\sqrt{b_1}, \dots, \sqrt{b_i})$, the spectral angle between these two vectors is then:

$$\theta = \arccos \left(\frac{\sum_i \sqrt{a_i} \sqrt{b_i}}{\sqrt{\sum_i \sqrt{a_i}^2} \sqrt{\sum_i \sqrt{b_i}^2}} \right) = \arccos \left(\sum_i \sqrt{a_i b_i} \right) = \arccos(\rho(a, b)). \quad (3.16)$$

Thus, the BCE is actually the cosine of the spectral angle between the the unit vectors α and β .

The PMCC is also related to the spectral angle, comparing equations 3.3 and 3.12, it is clear that the PMCC is the cosine of the spectral angle of the vectors $x_i - \bar{x}$ and $y_i - \bar{y}$. In probability distribution terms, this would be shifting the distribution to centre it about zero.

To illustrate the different behaviour of these measures in different conditions their performance on some simple data is considered, the first two datasets are taken from a paper by De Carvalho and Meneses (2000), which itself was comparing the performance of the PMCC to the spectral angle.

Spectrum	Spectral Angle	PMCC	BCE	SID
Reference [0.7, 0.6, 0.5, 0.6, 0.7]	-	-	-	-
A [0.5, 0.6, 0.7, 0.6, 0.5]	0.2484	-1	0.9922	0.0629
B [1, 1.2, 1.4, 1.2, 1]	0.2484	-1	0.9922	0.0629
C [0.005, 0.006, 0.007, 0.006, 0.005]	0.2484	-1	0.9922	0.0629

TABLE 3.1: A table of vectors and measures compared to a reference

In the examples in table 3.1, it is clear that none of the measures finds a difference between the test vectors, this is both encouraging and unsurprising as the vectors A, B

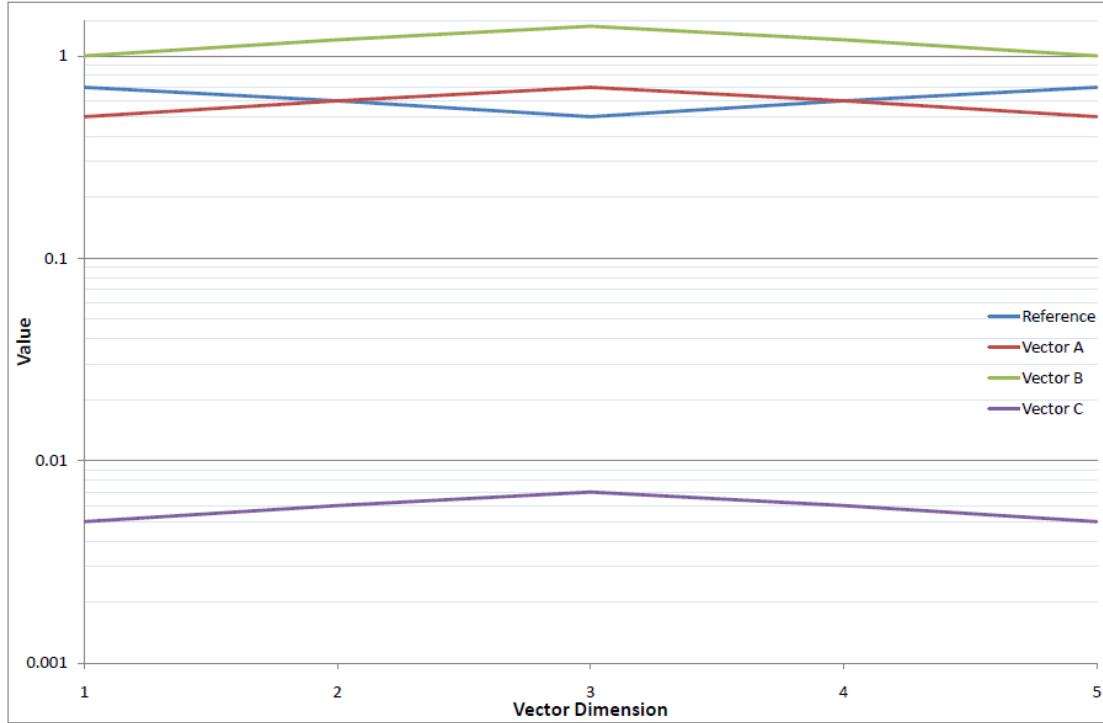


FIGURE 3.2: A graphic representation of the vectors used in table 3.1.

and C are clearly all scalar multiples of the same vector, in spectral imaging terms these are all pixels of the same material subjected to different levels of shading or illumination. The PMCC value of -1 confirms that the reference is negatively correlated to the test vectors. From the graphic in figure 3.2 it is clear that this confirms that a negative correlation occurs when a peak in one spectra is compared to a trough in another.

The low values for spectral angle and SID, and the value close to 1 for BCE all suggest that the test vectors are similar to the reference vector. The implication here is that all three are measuring a similarity of shape and say nothing about the direction or sense of the vector.

Spectrum	Spectral Angle	PMCC	BCE	SID
Reference	-	-	-	-
[15, 10, 5, 10, 15]				
A				
[15, 12, 9, 12, 15]	0.1515	1	0.9958	0.0337
B				
[15, 13, 11, 13, 15]	0.2166	1	0.9920	0.0642
C				
[15, 14, 13, 14, 15]	0.2752	1	0.9878	0.0984

TABLE 3.2: A table of vectors and their measures compared to a reference

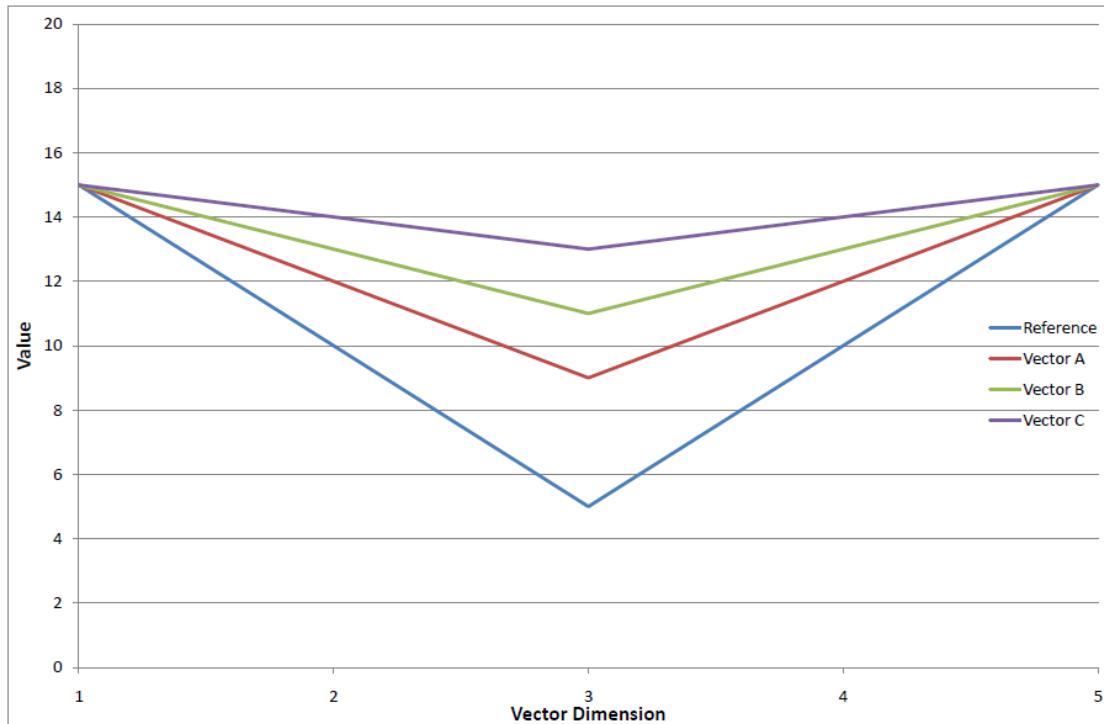


FIGURE 3.3: A graphic representation of the vectors used in table 3.2.

The examples in table 3.2 show that once again the PMCC finds no difference between the test vectors, or between the test vectors and the reference. However, each of the other measures shows the vectors A-C becoming increasingly different from the reference. De Carvalho and Meneses (2000) cite this as a limitation of the spectral angle, attributing the difference between vectors A, B and C as being due to a shading effect, to which the PMCC is “indifferent”. Careful consideration demonstrates however that this is a misinterpretation of the data.

The difference between the vectors should not be regarded as a shading effect, this effect was illustrated in the comparison in table 3.1. The values in the highest and lowest dimension are constant across the four vectors, this suggests that any shading is occurring preferentially in the three central dimensions. This behaviour is what would be expected in the presence of a material which absorbs in these dimensions, or in the central dimension, and which is present in differing volumes or densities in the measurements of each vector. The difference should thus be regarded an absorption trough, which is getting shallower in each subsequent spectrum. The change in the values of the spectral angle, BCE and SID therefore show that each is potentially sensitive to the amount of absorbing/reflecting/radiating material present in the spectral image.

Having established that the spectral angle, BCE and SID are changing in the above example, the natural question to ask is do they all change in the same way and at the same rate. This is not a question that is clearly answered by the above example and so

in the following example the test vector is systematically altered, and the effect of this change on the various measures is studied.

Spectrum	
Reference	[20, 20, 20, 20, 20]
Test(i)	[20, 20 + 2 <i>i</i> , 20, 20 - <i>i</i> , 20]

TABLE 3.3: A table of vectors

The first thing to note from this set of reference and test vectors is that the PMCC is undefined for all test vectors here; since all values in the reference vector are equal, the denominator in equation 3.3 will be zero. Graphs plotting the spectral angle, BC and SID against changing i are shown in figure 3.4.

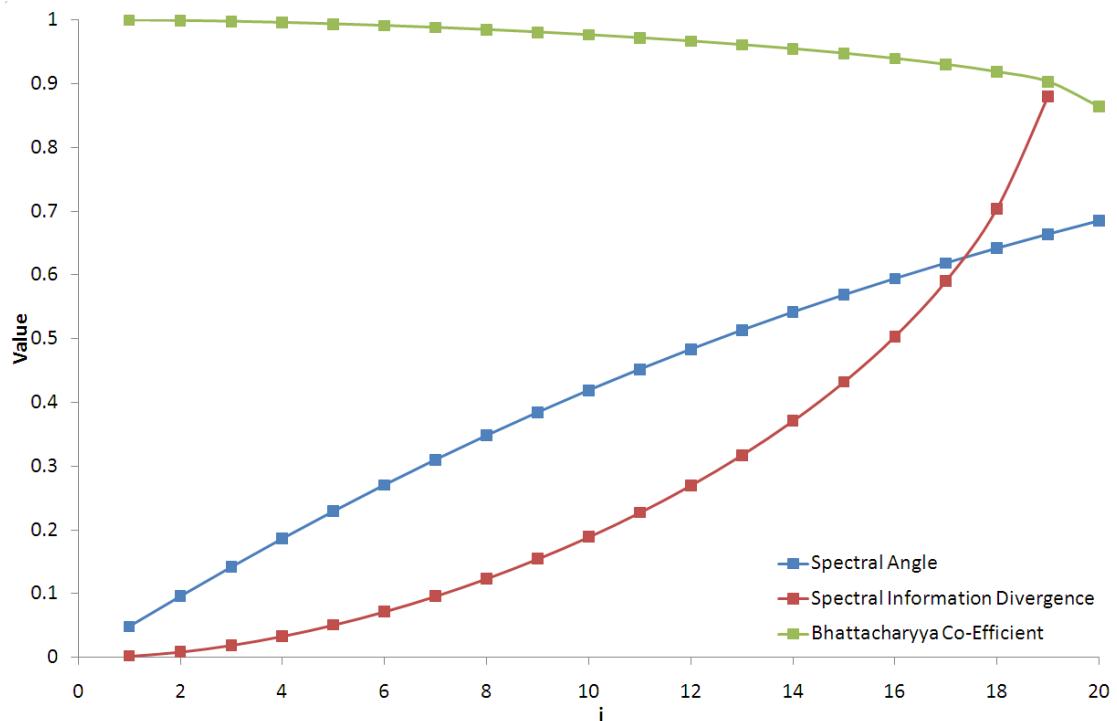


FIGURE 3.4: A graph showing the change in spectral angle, SID and BC with increasing i creating test vectors from table 3.3 that are increasingly different from the reference.

From this figure it is clear that the three measures change in very different ways as the test vector is systematically distorted. It is also apparent that the SID is not shown for $i = 20$; the fourth dimension at $i = 20$ has a value of zero, this causes a problem as the SID explicitly requires the logarithm of the data in each dimension and the logarithm of zero is undefined. This demonstrates a minor weakness of the SID, that any zero values in the spectral image will cause an undefined SID at that pixel.

There is a similar, but less pronounced problem with the spectral angle and BC whereby vectors whose values are all zero will cause an undefined spectral angle. This is again explicit in the calculation of the spectral angle as there will be a zero in the

denominator of equation 3.12. The problem is more subtle in the BC, which would appear to just generate a value of zero for a zero vector, however it must be remembered that as the BC is technically defined for probability densities, the vectors used must have been normalised to sum to 1 before being used in equation 3.4, such normalisation not being possible for a zero vector. This highlights a general problem with such “hidden” normalisations whereby errors can be introduced which are hard to trace.

Figure 3.5 shows the first derivative of each plot from figure 3.4. Considering these two sets of plots in tandem, several conclusions can be drawn. The BC and SID both behave in an exponential type manner over this range, as evidenced by the similarity of the shapes of the original plots and their derivatives. These measures both change slowly at lower deformations and much faster at higher deformations. The spectral angle behaves the opposite way, the changes here being more marked at lower deformations and gradually tapering off at higher deformations. This demonstrates an increased sensitivity to change at lower deformations in the spectral angle than in either of the other measures.

The derivative of the spectral angle in this case appears to be linear, careful analysis shows this not to be true. The derivative actually follows a backwards, stretched ‘S’ shape, however in this particular case a linear decrease is a very good approximation.

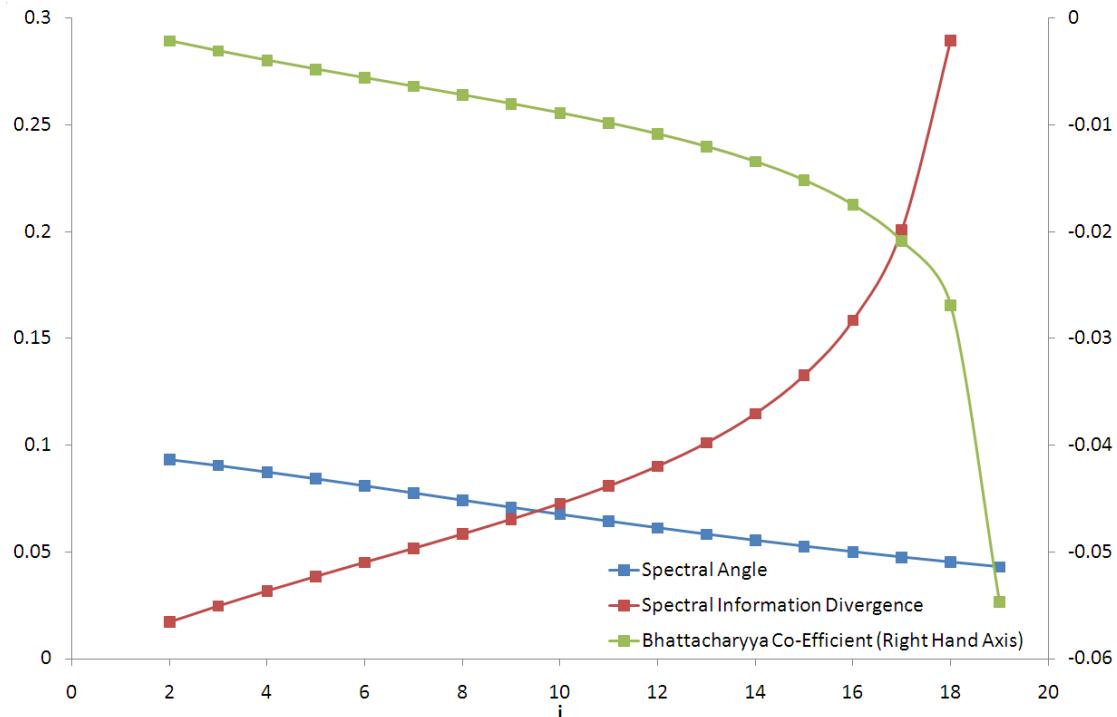


FIGURE 3.5: A graph showing the derivatives of the plots in figure 3.4, the spectral angle and SID are plotted on the same vertical axis (left), the BC is plotted against the right hand vertical axis.

Within the context of spectral imaging it is more common to be working in a regime where the requirement is to detect small changes between spectra that are both similar to a reference than it is to be detecting the same changes in spectra that are very different from a reference. Further, within the context of this project, this is the case. Therefore, whilst there are situations where the SID or the BC would be entirely appropriate measures, the rest of this thesis shall be concerned with the spectral angle. This has two further advantages, firstly, that the spectral angle requires no hidden normalisations, and secondly, this geometric interpretation of spectra leads naturally to an extension of the spectral angle, which will be further discussed in the next section.

3.3 Extending the Spectral Angle

The spectral angle is widely used in spectral image analysis, this is largely due to reasons mentioned in section 3.2.1, it is a straightforward concept to comprehend, it is simple and computationally inexpensive to calculate and it is inherently scale invariant. The spectral angle can be used to compare a pixel of unknown spectral type to various known target spectra in order to classify it; alternatively, several unknown pixels can be compared to a standard target, thus giving a relative measure of how similar each is to the others. In either modality, it is common to create a spatial map of how this metric varies across the image plane, the resultant image is known as the Spectral Angle Map (SAM).

The benefits offered by the spectral angle are further demonstrated by Sohn et al. (1999), who use spectral imaging to map deforestation in North-Central Yucatan. This paper tests the performance of the spectral angle against the Euclidean distance and a transformed Euclidean distance called the Mahalanobis distance, in terms of separating sagebrush spectra from soil. The spectral angle is shown to provide consistently smaller intra-class values and consistently larger inter-class values than either distance measure, thus making it a good candidate for a classification metric.

The same author further explores the suitability of the spectral angle as a classification metric in a follow up paper (Sohn and Rebello, 2002). This confirms the robustness of the spectral angle as a classifier and further expands on its particular suitability as being a measure of the spectral shape rather than a statistical distribution pattern. In that sense, the spectral angle is very much a direct measure and one that is entirely appropriate to spectral imaging. This gives it a further advantage over most of the measures considered in the previous section, which rely on approximating the spectral response as a probabilistic distribution.

As a single number being used to represent a vector of data that was originally perhaps tens or even hundreds of dimensions, the spectral angle is necessarily a lossy

data reduction and suffers from a number of inherent limits. It is these limits, and the attempt to mitigate them, which form the basis of the rest of this section.

3.3.1 Limitation of the Spectral Angle

The most fundamental limit of the spectral angle, is that for a given reference vector, it does not specify a unique test vector.

Considering the 2D case for example, relative to the reference vector $(1, 1)$, both vectors $(\sqrt{3}, 1)$ and $(1, \sqrt{3})$ subtend an angle of 15° and appear identical to a spectral angle metric. This is a point noted by De Carvalho and Meneses (2000), and it is in answer to this that they introduce the idea of using the PMCC as a spectral similarity measure, as discussed in section 3.2 above.

In two dimensions, it is clear that an angle θ subtended from a reference can describe any length multiple of one of two possible vectors. In three dimensions this is extended to a cone of vectors centred on the reference, which is, in principle, an infinite number of vectors. Taking a resolution limit into account it is clear that degeneracy is still a problem, and that indeed the problem gets worse as the number of dimensions increases (discussed in section 3.3.2 below), although beyond three dimensions it becomes difficult to picture what is happening.

With the exception of the above mentioned De Carvalho paper, this degeneracy is a point that seems to have been neglected in the literature; however, evidence does exist to suggest that the spectral angle has practical limitations. In a study investigating aerial mapping of previously unmined regions that compared the utility of the SAM and another spectral analysis algorithm known as Tricorder, Crósta et al. (1998) note the SAM misclassifying some areas of interest.

3.3.2 Numerical Degeneracy - The Other Curse of Dimensionality

In two and three dimensions, the issue of which spectra subtend the same angle from a reference is easily pictured. In higher dimensions this is much more troublesome. Extrapolating and remembering the relationship between the spectral angle and the Euclidean distance, it may be imagined that the points of the vectors subtending the same angle from a reference in four dimensions form a sphere around the reference point. Extrapolating further, it can be imagined that these vectors form hyperspheres around the reference in even higher dimensions. In this scenario, it seems that as the number of dimensions increases so does the number of degenerate vectors at any spectral angle. The situation seems particularly bleak when it is remembered that even in three dimensions, the number of degenerate vectors is, in principle, infinite.

The infinity of degenerate vectors is a result of considering the geometric space in which the calculation takes place to be continuous. In any practical situation, this is not the case. The space is actually discrete, each dimension measures values only as integers, from zero to an upper limit determined by the precision of the image. A standard RGB bitmap image is described as being 24 bit colour (or 16 million colours), it provides a 3D space in which each dimension takes a 8 bit integer, i.e. only values between 0 and 255. The space therefore consists of 16,777,216 unique points, and the number of vectors that can subtend any given spectral angle is therefore bounded.

Clearly, improving the precision of the image will increase the size of the space, for example, the Sony Ex-View camera whose spectral response is shown in figure 1.1 captures 12bit images. An RGB image generated from this camera will form a 3D space in which each dimension takes only values between 0 and 4095, generating over 68 billion unique points.

There are thus two factors determining the size of the discrete space, the number of dimensions and the precision of the image. Increasing either will increase the number of degenerate vectors for a given spectral angle. Adding more dimensions, irrespective of whether they provide extra discriminatory information, will increase the number of points in the space at a given spectral angle from a reference. This necessitates a decrease in the confidence that two vectors measured to be at the same spectral angle from the reference can be considered as being the same; this decreasing confidence from increasing the number of dimensions can be considered a second “curse of dimensionality”.

3.3.3 Azimuthal Angles in Three Dimensions

As mentioned in section 3.3.1, in a 3D space, the set of vectors, \mathbf{T} , at an angle θ from some reference vector, \mathbf{R} , form a cone around it (see figure 3.6).

By rotating the co-ordinate frame such that the reference vector is used as an ordinate axis, i.e. $\mathbf{R} = Rz'$, and looking down this axis, the situation is as pictured in figure 3.7.

From figure 3.7, it is clear that projecting each member of \mathbf{T} onto the $x'y'$ plane will form a circle centred on z' . Each vector can now be uniquely identified by specifying the angle, φ , between this projection and the x' axis, i.e. by taking the spectral angle of this projection with the axis. The choice of axis in this case is entirely arbitrary as there is a constant relationship between the angle from x' and the angle from y' .

Thus by considering this pair, (θ, φ) , one can uniquely distinguish between distinct spectra, even where they are at coincident spectral angles from a reference.

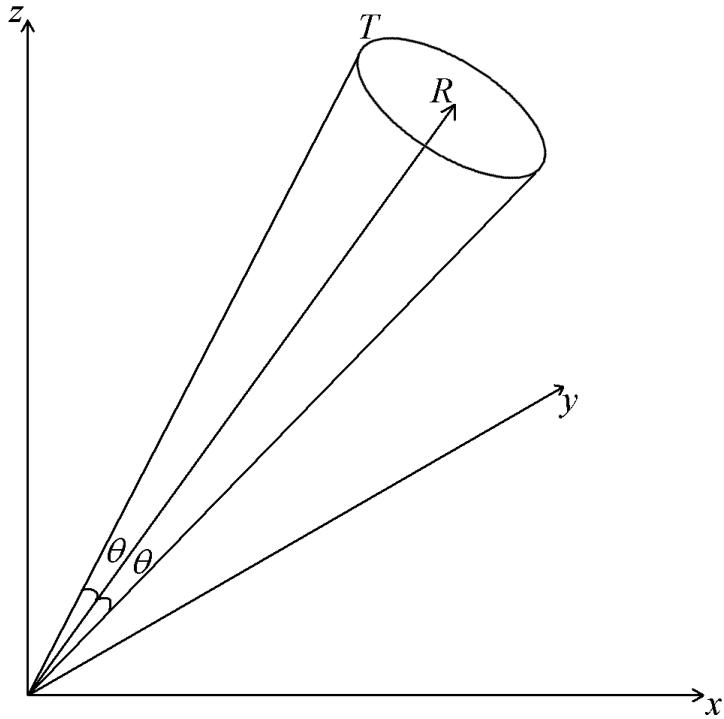


FIGURE 3.6: A vector \mathbf{R} in three dimensions with a cone of vectors \mathbf{T} subtending an angle θ around it.

3.3.4 The Spectral Position Map in Many Dimensions

In $N + 1$ dimensions, the set of vectors \mathbf{T} is described mathematically as:

$$\mathbf{T}_i = \sum_{j=1}^{N+1} T_{ij} x_j. \quad (3.17)$$

For convenience, the co-ordinate space is rotated such that $\mathbf{R} = R\mathbf{x}'_1$. The projection considered is that of \mathbf{T} onto the ND subspace orthogonal to x'_1 . This leaves the set of vectors τ :

$$\tau_i = \sum_{j=2}^{N+1} T'_{ij} x'_j. \quad (3.18)$$

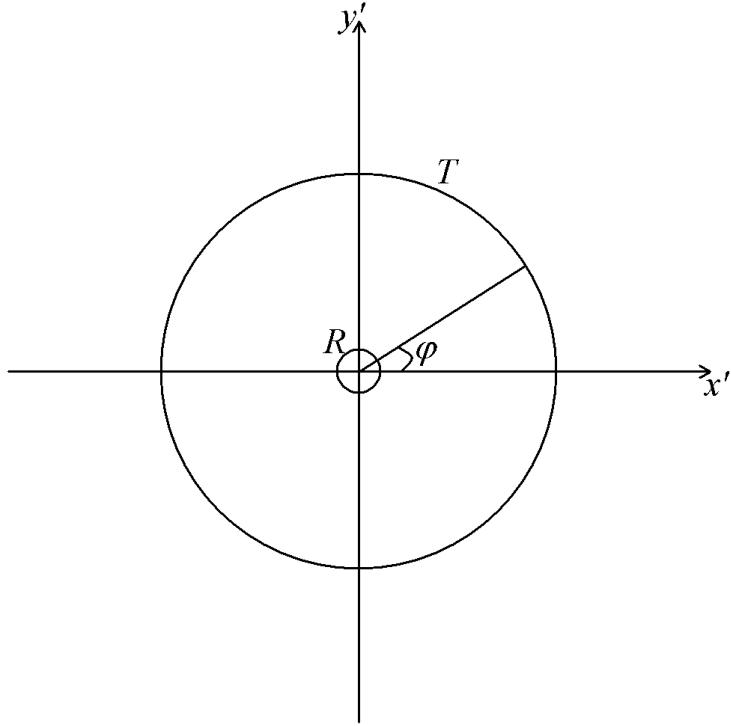


FIGURE 3.7: The scene from figure 3.6 after the co-ordinate space has been rotated such that the vector \mathbf{R} is now coincident with the ordinate axis out of the plane of the paper. Each member of the set of vectors \mathbf{T} now subtends a unique angle φ from \mathbf{x}' .

A set of N angles, φ_k , can now be calculated for each τ_i , describing the angle between the vector and the relevant axis x'_k :

$$\varphi_{ik} = \arccos \left(\frac{T'_{ik}}{\sqrt{\sum_{j=2}^{N+1} T'_{ij}^2}} \right). \quad (3.19)$$

It should be noted that this applies for an $N > 1$; in the special case of $N = 1$, it is sufficient simply to consider the projection of T'_2 . $N = 2$ is also a special case in that there is a constant and well defined relationship between φ_2 and φ_3 , making the two metrics degenerate.

Each set of φ calculated is referred to from this point as a spectral position vector (SPV). Much like the spectral angle can be spatially mapped to form an SAM, the SPVs can be mapped spatially to form a Spectral Position Map (SPM).

For an $(N + 1)$ D image, the full SPV has N dimensions, if the spectral angle is included then the full spectral position map has the same dimensionality as the

original image. This means that the combination of SAM and SPM represent an intensity normalised re-encoding of the data.

3.4 Established Alternatives - K-Means Classifiers and Principal Components Analysis

The problem of unsupervised classification is not unique to this project, and as such other approaches have been studied. Further, strategies for dealing with the challenges of working with high dimensional data have also been studied. In this section, one such classifier, or more properly one family of classifiers, is described alongside a method for managing high dimensional data.

3.4.1 The K-Means Classifier

The K-Means classifier is an algorithm which seeks to sort a set of data into subsets, or classes, such that variations within a class are minimised and variations between classes are maximised.

The algorithm is seeded with a number of classes, k , into which to sort the data and the first step is to randomly assign each data point to one of the classes. The remainder of the algorithm is an iterative resorting following the steps outlined below.

1. For each class, calculate the average value, \bar{x} , of the n data points x_i , in the class, where, $\bar{x} = \frac{\sum_{i=1}^n x_i}{n}$, these are the “class centres”;
2. for each data point, calculate a distance from it to each of the class centres, $d_{ij}(x_i, \bar{x}_j)$;
3. reassign the data point to the class to whose centre it is closest;
4. calculate the new set of class centres;
5. repeat steps 2-4 until some condition, ϵ , is satisfied.

There are two parts of this algorithm that can be modified, thus creating a family of classifiers; the first is the condition ϵ , which is often along the lines of “the sum of the distances of each data point from its class centre is less than a given value”, but can be as simple as a maximum number of times to iterate over. The second is the distance measure, d_{ij} used, which can in principle be any function of the two points, but is most commonly a distance measure like those outlined in section 3.2.1.

3.4.2 Principal Components Analysis

Principal Components Analysis (PCA) is a common data transformation and reduction method that seeks to rewrite data in terms of uncorrelated variables. It can be thought of as a rotation of the axes in multidimensional space such that the new (orthogonal) axes represent the directions of greatest variance in the original data. These new axes are the principal components (PCs). Once this rotation has been achieved, data reduction can be performed by projecting the data into a lower dimensional space defined by the higher order PCs. As these PCs have been calculated to maximise the amount of variance retained, the information loss from this reduction is minimal.

PCA is calculated on zero-centred data, so the first step in calculating the PCs is to calculate the mean data value in each dimension and subtract this from each datum.

$$x'_{i,k} = x_{i,k} - \bar{x}_k \quad (3.20)$$

Having done this, the co-variance matrix C_{kl} is calculated where:

$$C_{kl} = \frac{\sum_{i=1}^N x'_{i,k} x'_{i,l}}{N - 1} \quad (3.21)$$

This generates an $n \times n$ diagonal matrix measuring the co-variance of each dimension with every other dimension. The eigenvectors of this matrix form the PCs, the first PC being the eigenvector with the greatest corresponding eigenvalue.

The eigenvalues are proportional to the variance of the data accounted for by each eigenvector (and hence the total variance of the data is proportional to the sum of the eigenvalues). It is worth noting that eigenvectors are not unique solutions to the matrix equation $M\mathbf{x} = \lambda\mathbf{x}$, in that if \mathbf{a} is a solution then $\kappa\mathbf{a}$ is also solution (where κ is some constant). This implies that eigenvectors at least provide unique directions since scaling a vector does not change the direction in which it points; however since κ could be negative, this direction is not unique since the opposite direction is an equally valid solution.

Having calculated the PCs it is possible to perform a data reduction by rotating the data into PC-space. This is a vector operation where both the data and the PCs can be treated as vectors in n dimensions. The component of each data point in each PC is then simply the projection of the original data vector onto the PC vector (i.e. the dot product of the two vectors).

$$y_{i,q} = \sum_{k=1}^n x'_{i,k} p_{q,k} \quad (3.22)$$

where p_q is the vector corresponding to the q th PC.

As the eigenvalues are proportional to the variance they capture, it is possible to reduce the dimensionality to the first n eigenvectors, retaining a known proportion of the variance, and typically a high proportion of the variance is captured in a small number of PCs.

Rotating all data into a PC space where all PCs are retained represents a re-encoding of the data, not dissimilar to that provided by using the SAM and SPM. However, PCA does not normalise the data with respect to intensity, unlike the SAM and as such is a reversible (lossless) re-encoding. The axes of the new PC space are labelled in such a way as to consider them to be weighted, the first PC and hence the first axis, is known to account for the largest proportion of the variation in the data and so on down the list. This follows from PCA being a statistical consideration of the data. In contrast, with the SAM and SPM rotation, which is based on the geometry of the data in the high dimensional space, this weighting is not apparent, and it cannot be declared without further measurement, which of the axes in the newly rotated SPM space might contain the greatest variance.

3.5 Conclusions

In this chapter a number of metrics of spectral similarity have been discussed and compared and it was concluded that for this project the spectral angle is a sensible metric to use for analysis of spectral images.

The original motive for introducing a spectral similarity metric was to represent high volumes of data in a meaningful way. Reducing a vector of $N + 1$ dimensions to a single meaningful value is by necessity a many to one reduction and degeneracy will be introduced into the system. There is no requirement however for a meaningful representation to be limited to a single dimensional measure; for example, as has been intimated throughout this report, a colour image is a meaningful representation of information in three dimensions. In this chapter the possibility of a multidimensional representation of spectral information has been explored, and a measure called the Spectral Position Vector (SPV) was introduced.

The motivation for the SPV was to find a measure that would complement the spectral angle in an attempt to reduce the inherent degeneracy. The full SPV achieves a complete elimination of this degeneracy, but only at the cost of not actually reducing the data. As a compromise it is envisioned that the spectral angle along with a small

number of the SPV dimensions will ultimately be used to generate a meaningful representation of the data, this is discussed further in chapter 4.

Chapter 4

Testing and Validating the Spectral Position Map

If an image is only as useful as the information that can be gleaned from it, then the same must be true of a similarity measure. It is, after all, through these measures that the information is obtained. Having been introduced in chapter 3, the SPV will, in this chapter, be tested on a range of images to verify that it does actually provide additional information to the spectral angle, and that that information is actually useful.

4.1 A Demonstration Using Specially Generated Images

The test spectral images in this section were generated using code kindly written by Dr. Paul Hoy. The analysis was performed by both Dr. Hoy and the author.

Each “test” image in this section was generated in the same manner. A uniform grey reference spectrum, i.e. a spectrum in which the intensity in each dimension is the same, is assumed. The pixels are set such that all pixels in the same column are at the same spectral angle from this reference; the angle increases from zero at the left hand side of the image, to a value that varies with the number of dimensions on the right hand side. In each column, each successive row is the result of the vector in the top row being incrementally rotated around the reference vector, from zero to 2π radians.

The 3D case, i.e. a test image with three planes, is the obvious starting point for verifying that the spectral position behaves as would be expected; the properties of this case are straightforward both to comprehend and to calculate. Figure 4.1 shows the same view of the 3D co-ordinate space as figure 3.7 from the previous chapter, and shows how points in the space relate to pixels in the test image. Clearly, all points lying on the same circle subtend the same spectral angle from the reference, and thus

lie in the same column. Points on the same radial lines are at the same spectral position and thus lie in the same row.

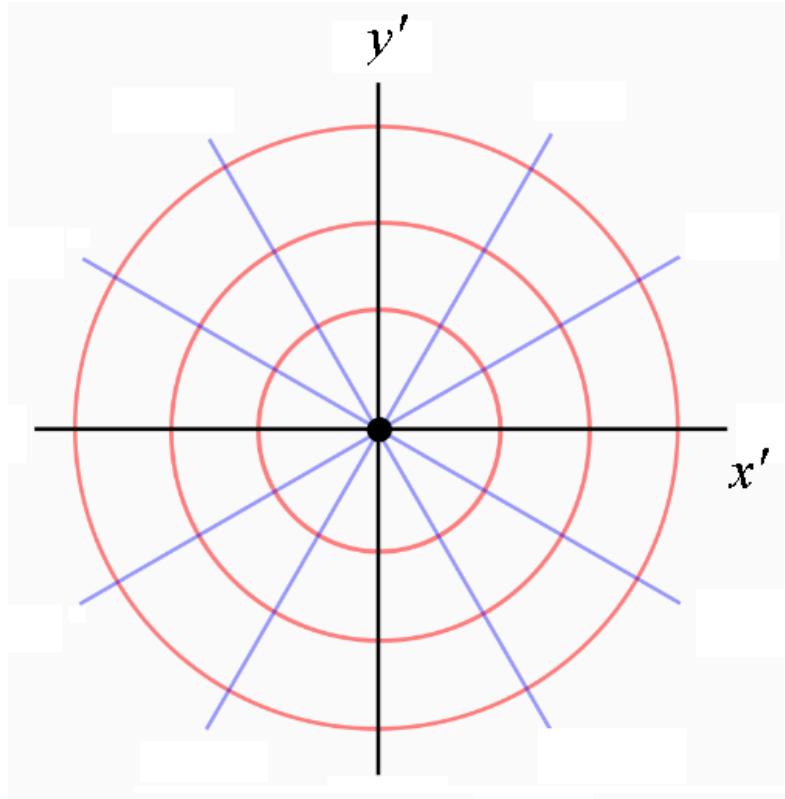


FIGURE 4.1: An image¹ showing the 3D co-ordinate space in the same configuration as in figure 3.7, and how the points in the space correspond to the spatial pixels of the test images generated. Points on the same radial line lie on the same row of the test image, points on the same concentric circle lie on the same column.

There is the further benefit from working in three dimensions, i.e. that the test image can be thought of, and displayed, as an RGB image. Figure 4.2 shows the RGB display of the 3D test image.

It is clear from figure 4.2 that the colours in the image change both horizontally and vertically. On the extreme left, there is a band which appears to be all grey, this should not be surprising as the reference vector was defined to be grey and pixels at the left hand side have a small spectral angle from the reference. It is also clear that this area should vary only slightly through the vertical, as for smaller spectral angles, the circle defined by a rotation of spectral position is much smaller than for larger angles (see figure 4.1 to verify this), and thus the colour values will be little changed. On the right hand side of the image however, the spectral angle is much larger and changes in spectral position cause much more rapid changes in colour. A full rotation about the reference vector rotates the test vector through 360° of hue, which can be

¹Image modified from an original at http://en.wikipedia.org/wiki/File:Polar_graph_paper.svg under a GNU Creative Commons License by kind permission of Mets501

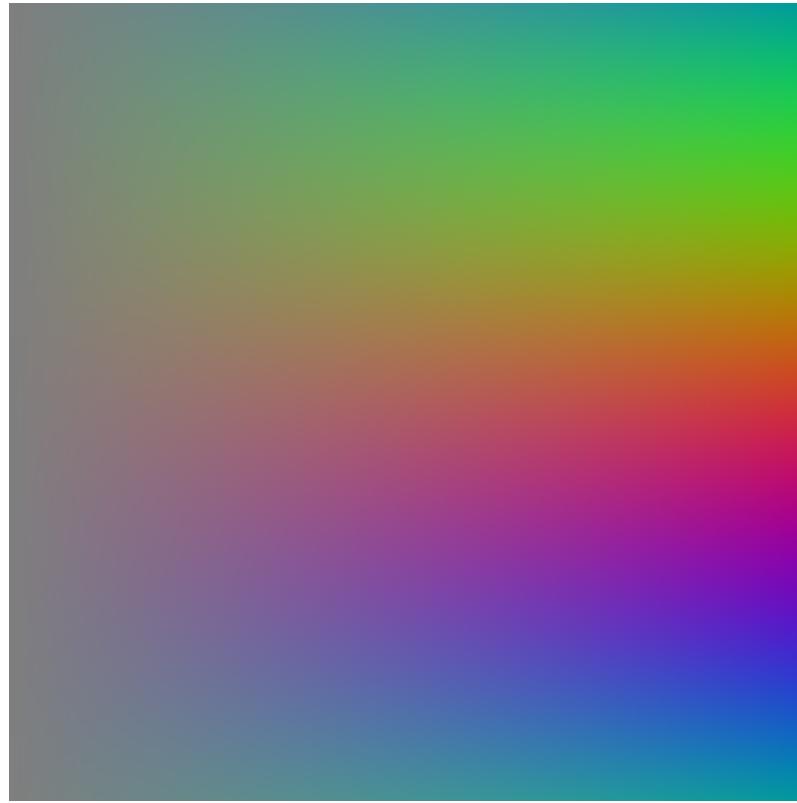


FIGURE 4.2: The RGB view of the 3D test image.

seen as the colour moving from a greenish blue at the top through a complete spectrum of colours to the same greenish blue at the bottom.

Figure 4.3 shows each pixel from the test image as a point in a 3D space, the RGB colour space. This suggests that only a small portion of the available colour space is being represented, however it should be remembered that the intensity independence of the spectral angle means that each point is in fact representative of the entire straight line, from the origin to the edge of the space, on which it lies. The “satellite-dish” like shape is formed by the points because the the image is generated such that each pixel is of a similar intensity, and hence the points are all at a similar distance from the origin. Figure 4.4 shows the same space rotated such that the reference vector is pointing out of the plane of the page, the points lie on concentric circles, and along radial lines, further confirming the image from figure 4.1.

The SAM of this image was calculated, yielding the expected result as shown in figure 4.5. In this image the spectral angle is mapped to an intensity in grey-scale, an angle of zero corresponds to a completely black pixel (zero intensity) and the maximum angle present (in this case $\sim \frac{\pi}{5}$) corresponds to a completely white pixel (100% intensity), with all intervening values scaling linearly. The resulting image is thus a horizontal graduation from black on the left to white on the right, constant through the vertical dimension.

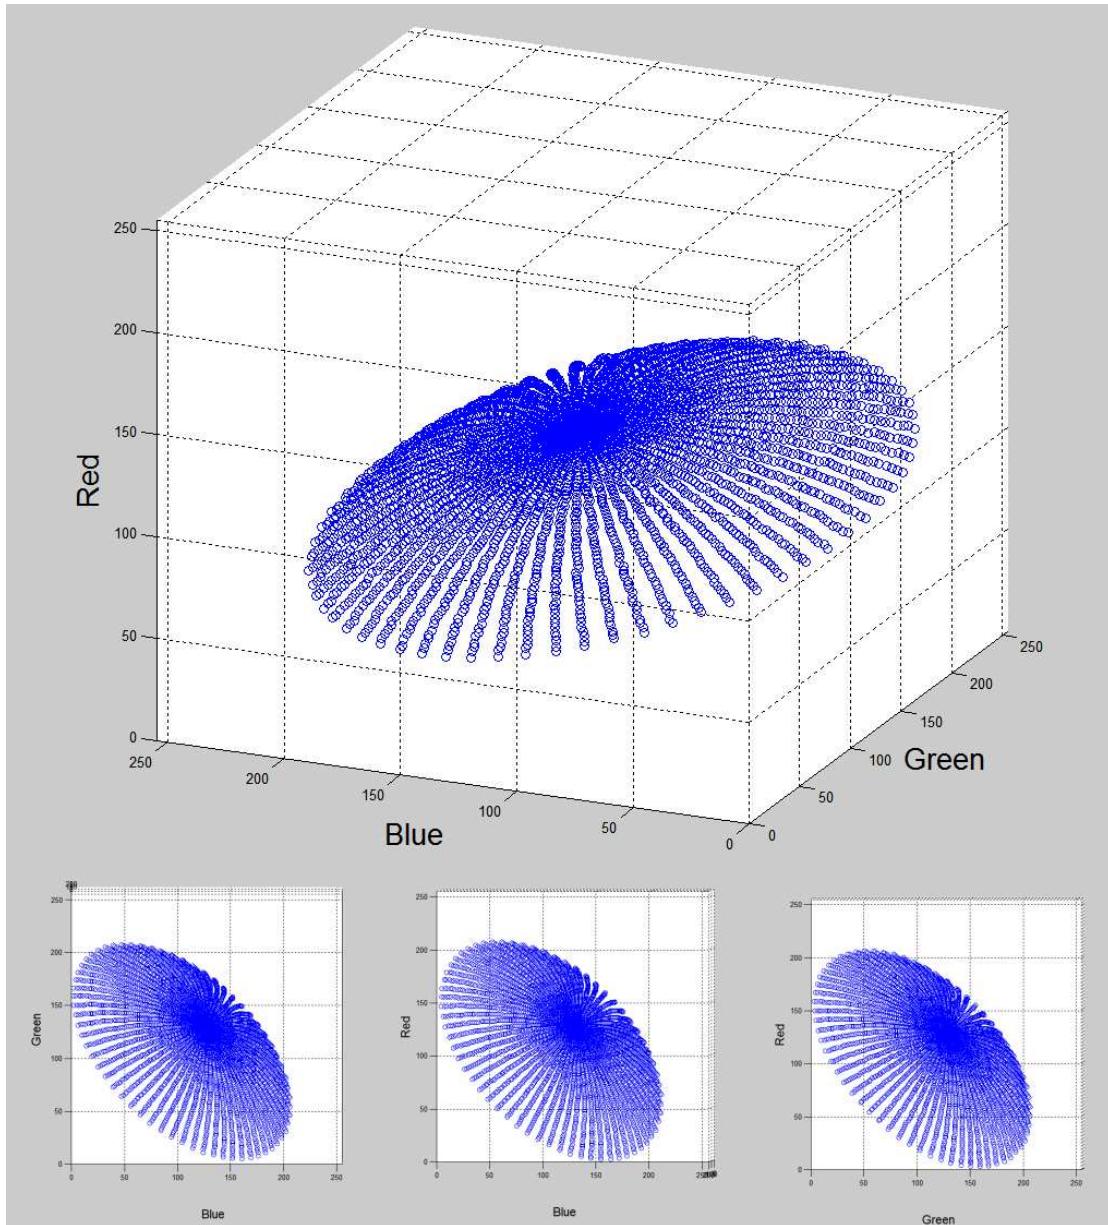


FIGURE 4.3: An image showing the pixels of the test images generated as points in a 3D space, with axes rotated to demonstrate the shape of the distribution of the points. The points form a “satellite-dish” shape centred on the line $R=G=B$, with a circular distribution. At the bottom the projection of the distribution onto the BG, BR and RG planes are shown.

The combination of figures 4.2 and 4.5 demonstrate that the spectral angle alone is unable to discriminate between very different spectra.

The SPM was then calculated for the image. The SPV in this case is a 2D vector at each pixel, and thus there are two image planes in the SPM; these are both shown in figure 4.6. Again, angles of zero map to black pixels, with increasing angles linearly mapping to increasingly light grey values, white pixels correspond to an angle of π radians. The scale bar for all of the SPMs in this thesis is shown in figure 4.7.

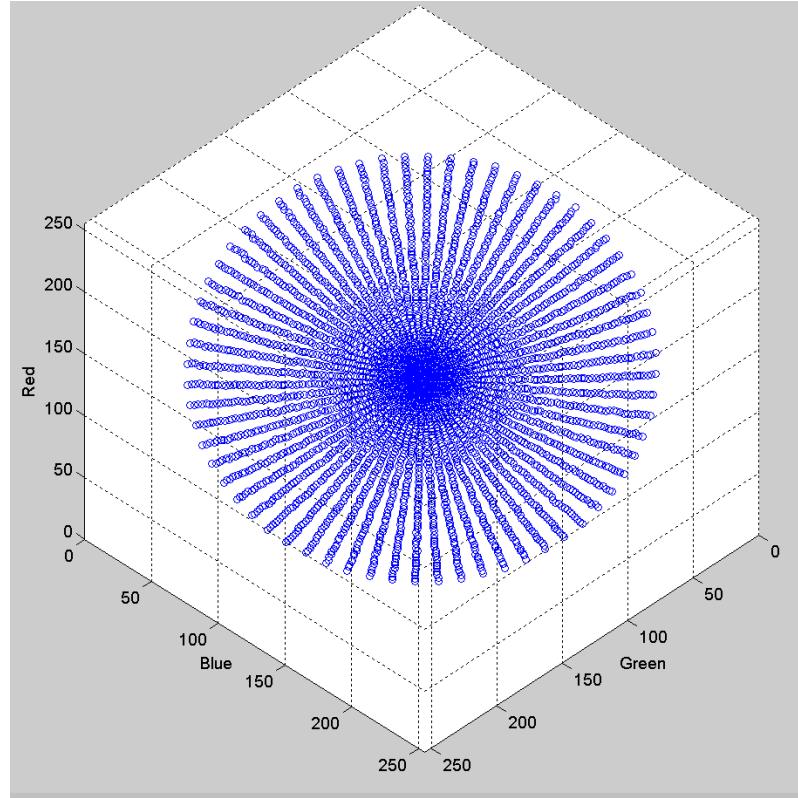


FIGURE 4.4: An image showing the spatial pixels in the test image as points in a 3D space with axes rotated such that the reference vector is orthogonal to the plane of the page, the general shape here confirms figure 4.1.

Moving vertically down the φ_1 image (left hand side of figure 4.6), the angle increases from zero at the top through π radians in the middle and back to zero again at the bottom. The φ_2 image on the right hand side can be seen as a vertical translation of the φ_1 image by one quarter of the image height (with the bottom rows looping back up to the top). This is intuitive since in three dimensions, the axes of the orthogonal plane in which these angles are calculated translate to each other by a rotation of $\frac{\pi}{2}$ radians (a quarter of the full rotation). The decision as to which axis to label as the φ_1 axis is entirely arbitrary, and hence the angles should actually be degenerate.

The fact that the φ_1 image moves from black at the top, through white in the middle and back to black at the bottom, or alternatively, that π radians rather than 2π radians maps to white, can be explained by noting that whilst the image is generated by rotating the test vector by 2π radians, when the SPV is calculated, the angles are determined using the arccosine function; this function has a principal domain of 0 to π radians. The cosine function is symmetrical about π , i.e. $\cos(\pi + \alpha) = \cos(\pi - \alpha)$, so for angles greater than π , as the angle increases the arccosine decreases. In order to determine whether the angle is $\pi + \alpha$ or $\pi - \alpha$ from a given cosine value, the value of the pixel in the φ_2 plane must also be known.

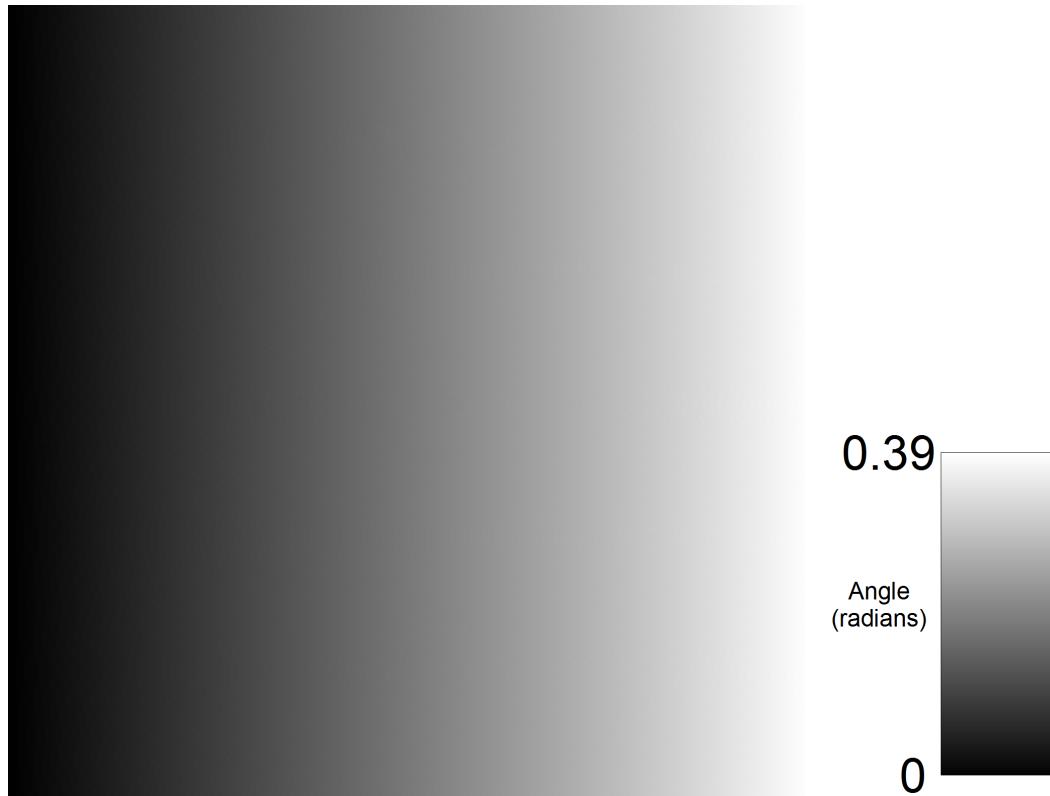


FIGURE 4.5: The SAM of the 3D test image.

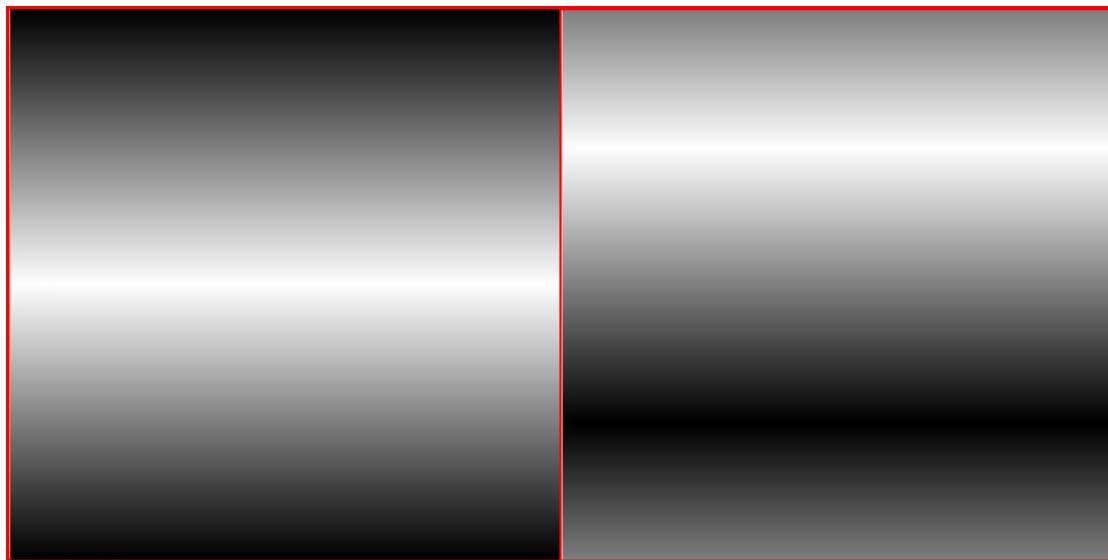


FIGURE 4.6: The SPM of the 3D test image, the first plane is shown on the left hand side and the second plane on the right hand side.

This indicates that whilst in principle the dimensions of the SPV are degenerate, in practice the limited principal domain of the arccosine function means that both dimensions are actually required to uniquely specify a vector.

In any more than 3 dimensions, there is no useful or insightful visualisation of the test

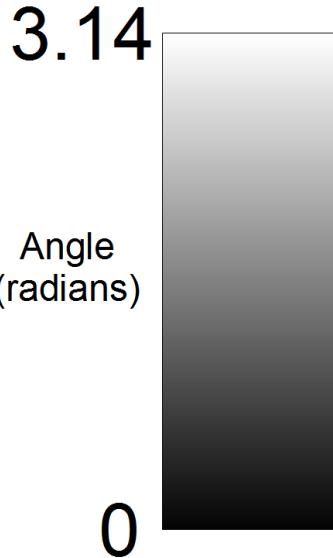


FIGURE 4.7: The scale bar for all SPM images used in this thesis.

image, however, there continues to be interesting information in the SPMs. In general, for a spectral image each plane represents a specific wavelength or waveband and hence mapping a selection of these to RGB channels can produce a colour or pseudo-colour image that reveals something about the underlying scene, however for these test images each plane is merely a dimension in an abstract space. Another standard visualisation is to approximate the total intensity at each pixel by summing or averaging the values in each plane, again this is not relevant for the test images since, as mentioned above, they are generated in such a way that the total intensity of each pixel is constant (or approximately constant) across the entire image plane. An intensity visualisation would thus be a uniform (or near-uniform) grey square, which reveals nothing of interest.

The SAM for all higher dimensional test images is the same as that shown in figure 4.5; the maximum angle present actually decreases as the number of dimensions increases, however the SAMs are generated such that the maximum angle present always maps to a white pixel, this is done to maximise the contrast in each image. Further, the first dimension of the SPV is also the same for all test images (φ_1 image from figure 4.6); this merely confirms that the images were generated by the method described above.

The full set of SPV planes for the 5D and 6D cases are shown in figures 4.8 and 4.9 below. These results are typical for test images where the number of dimensions are odd and even respectively.

Figures 4.10 and 4.11 below show the intensity profile through a column of each plane of the SPV in the 5D and 6D cases respectively. The images for planes 2 and 3 in figures 4.8 and 4.9 appear to be the same, this is confirmed by comparing their profiles in figures 4.10 and 4.11.

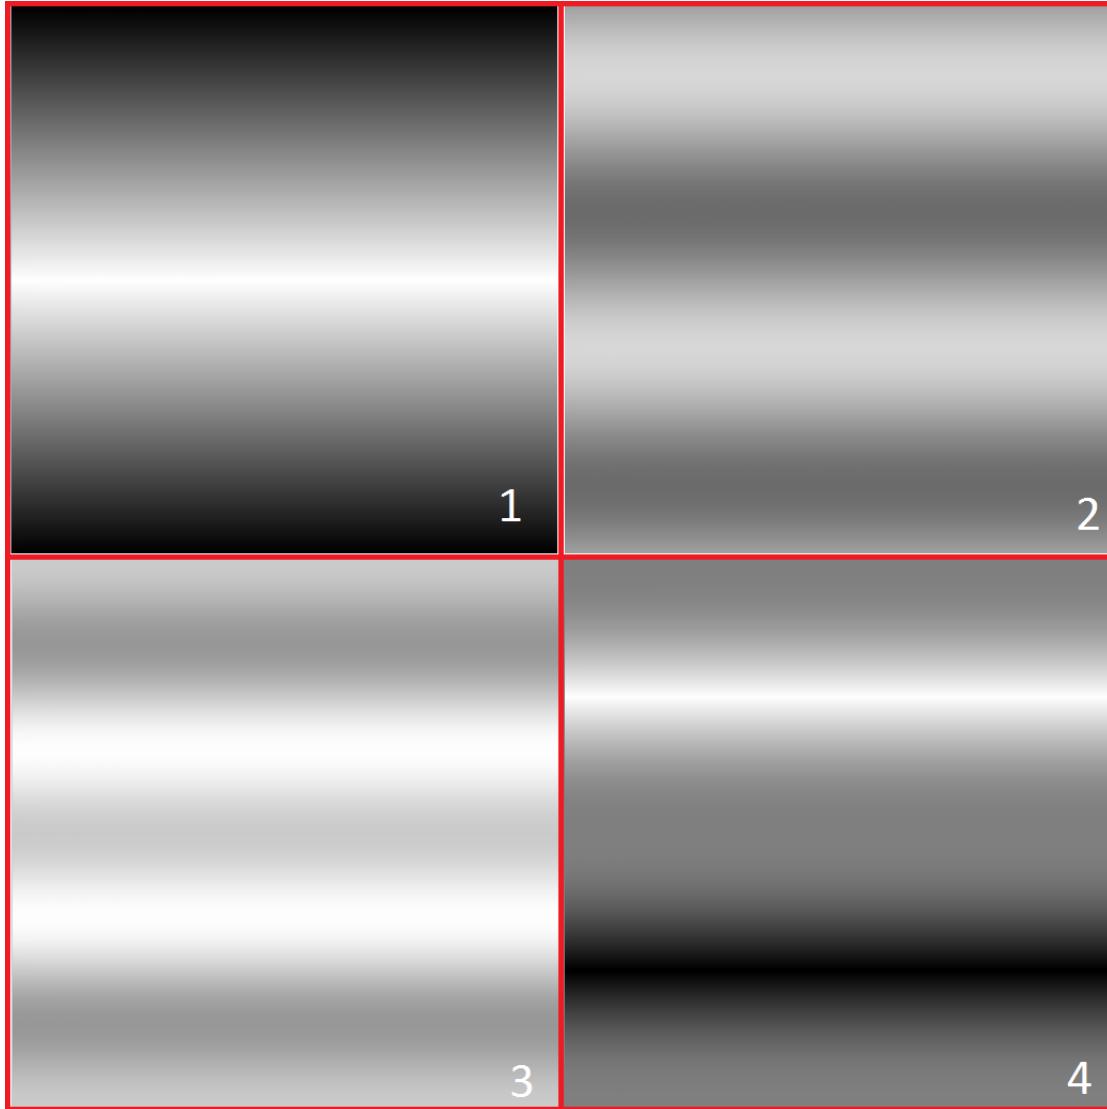


FIGURE 4.8: The SPM for the the 5D test image, the plane number is indicated in the bottom right hand corner of each plane.

The profiles of the last plane for each share similarities in their shape, although for the 5D case it is an odd function about pixel 256, whereas for the 6D case it is even about pixel 256. This is a pattern followed as the dimensionality increases, the last plane has a maximum and a minimum for odd numbers of dimensions and two minima for even numbers. As dimensionality increases the peaks become sharper and increasing numbers of pixels are at 50% intensity.

This is part of a wider trend in which the profile for each plane has a symmetry about pixel 256 that is odd for odd numbered planes and even for even numbered planes. For an ND case, the $(N - 2)$ th plane introduces a new profile, which appears in the same plane all subsequent cases $> N$.

The purpose of the set of test images presented here was to demonstrate that the SPV

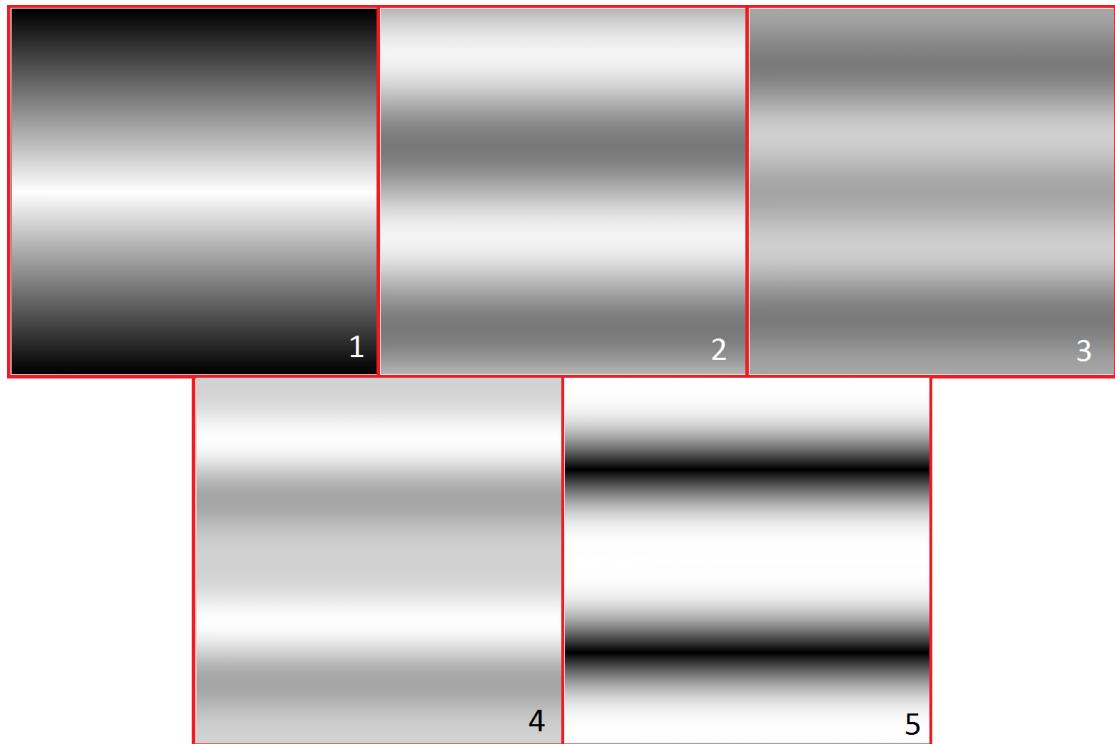


FIGURE 4.9: The SPM for the the 6D test image, the plane number is indicated in the bottom right hand corner of each plane.

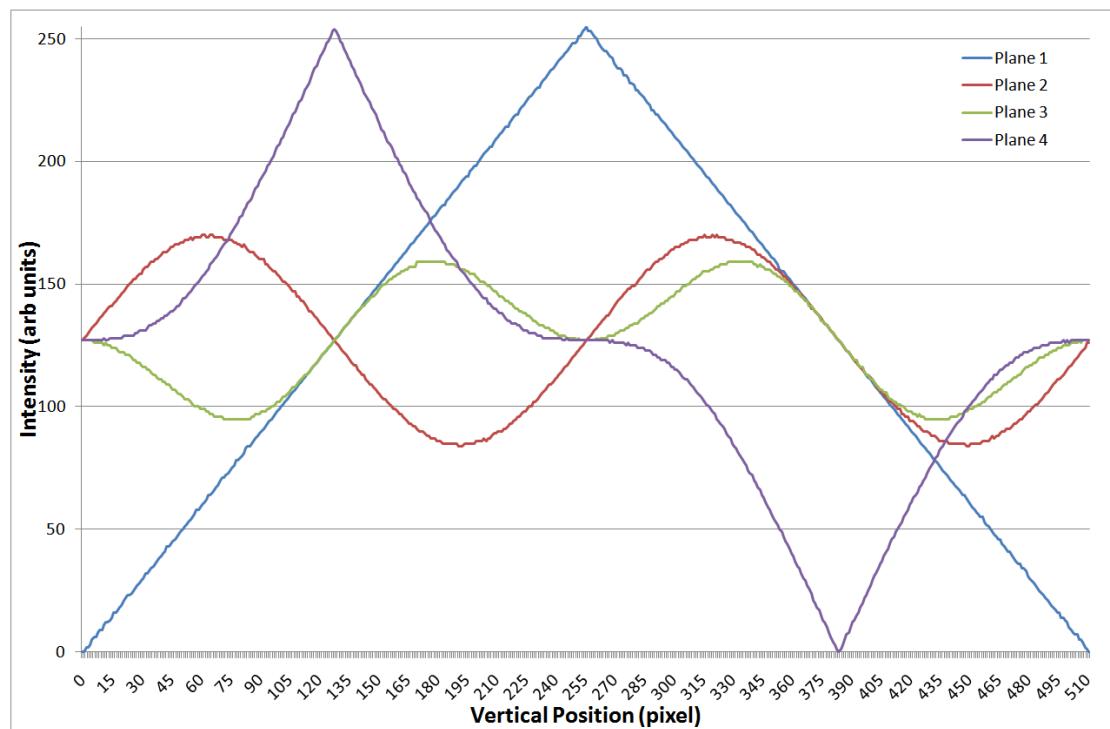


FIGURE 4.10: The intensity profiles for a vertical line through each of the planes of the SPM for the 5D test image.

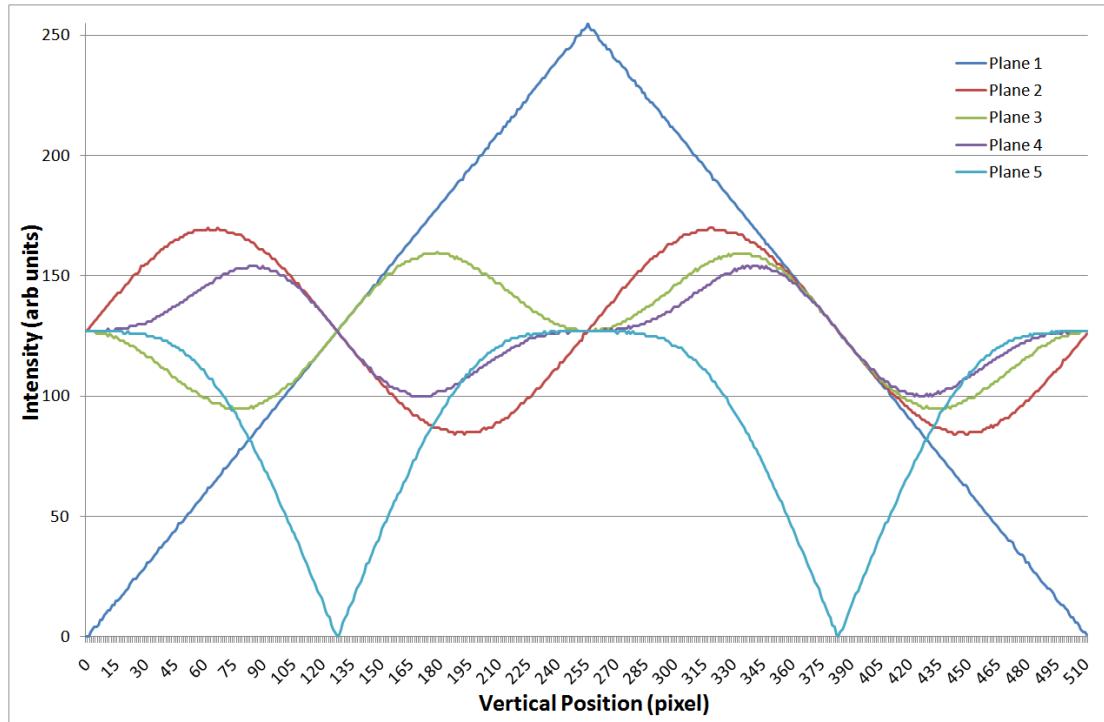


FIGURE 4.11: The intensity profiles for a vertical line through each of the planes of the SPM for the 6D test image.

contained useful information that was independent from, and in some way complementary to, the spectral angle. The discussion above, whilst not an exhaustive interrogation of the possible set of images, demonstrates that this is the case, and further that there is information in each plane of the spectral position, none of which is dependent on the spectral angle.

4.2 A Demonstration in Three Dimensions - RGB Images

Having established that it is possible to extract information independent from the spectral angle, the SPM will now be applied to a number of RGB images to demonstrate how this information can be used. This will start with an image that is a staple target in image processing, known simply as “Lena” or “Lenna”.²

Figure 4.12 shows the original image. This is dominated by a small number of tones and colours; much of the image is a soft skin tone, the colours of Lena’s skin being barely distinguishable from most of the background. Her hair colour is very similar to the frame of the mirror visible in the image, and there is a yellow reflection in the mirror which also provides some contrast. The most striking contrast however is clearly the blue feather in the hat.

²The story of this image can be found at <http://www.lenna.org>



FIGURE 4.12: Lena - A staple image processing target.

The SAM for this image, again computed against a grey reference, is shown in figure 4.13 below.



FIGURE 4.13: The SAM for the Lena image.

It is noticeable that the darker areas on the SAM include the top right hand side of the hat, the brim of the hat, the right hand edge of the shoulder and the feather on the hat. If this was the only information we might conclude that these areas looked the

same, or similar on the original image; this is somewhat true of the hat and shoulder areas, the feather however, is markedly different. The hair and mirror frame do have similar spectral angles, as one would expect, but these are far from unique in the map. For example, from right to left across the hat the spectral angle increases, passing through a value where it has the same spectral angle as the hair. The yellow reflection in the mirror is clearly visible in the SAM, although this is largely because the borders are lighter than the internal area, which is of a similar spectral angle to the rest of the mirror face area.

The previous paragraph provides only a vague and qualitative analysis of the SAM. A more quantitative analysis can be performed by examining the distribution of image pixels in the metric space, i.e. by plotting a histogram of the spectral angle. The histogram for the Lena image is shown below in figure 4.14.

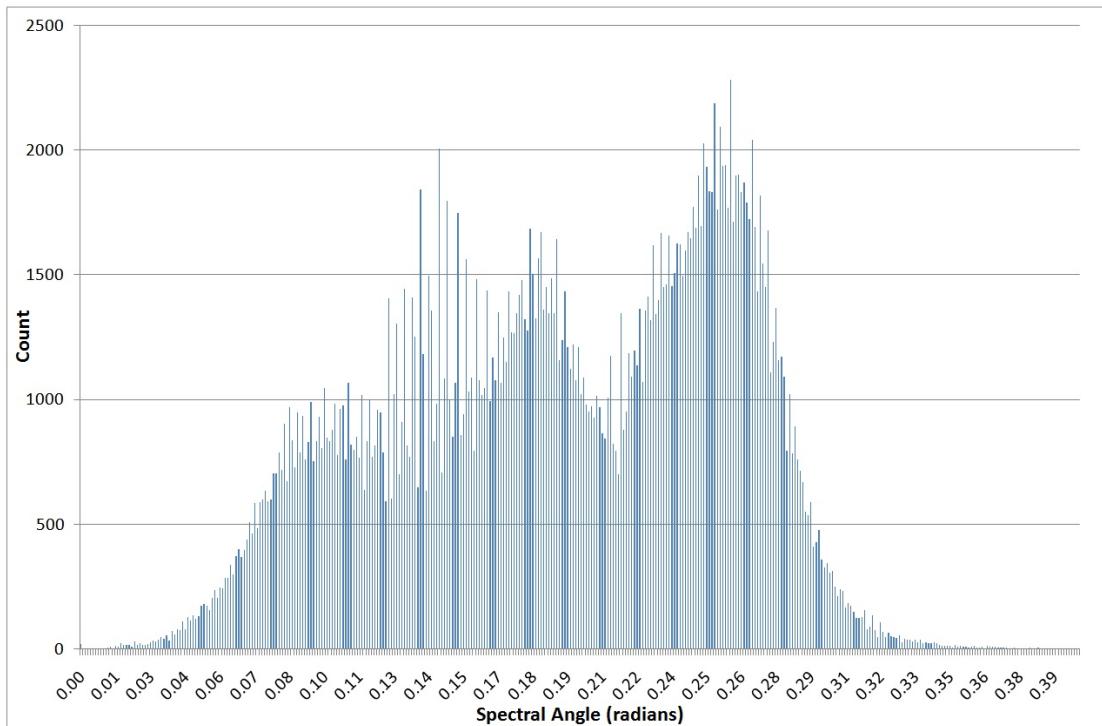


FIGURE 4.14: The histogram of spectral angles present in the SAM of the Lena image.

A number of peaks can be seen within the histogram, which is what would be expected. In an ideal case, each material in the image scene would produce a unique spectral response, or in this case, colour, which would stand at a unique spectral angle; the idealised histogram would then be a series of lines. In reality, the signals received at each pixel are rarely pure, stochastic noise from the detection process, non-uniformities in the materials themselves and the effect of “mixed” pixels, i.e. where more than one material is present in a single pixel area, all contribute to slight variations of spectral angle. According to the central limits theorem, these would be

expected to broaden the lines into Gaussian curves. Such Gaussian curves can be fitted to the histogram as shown below in figure 4.15.

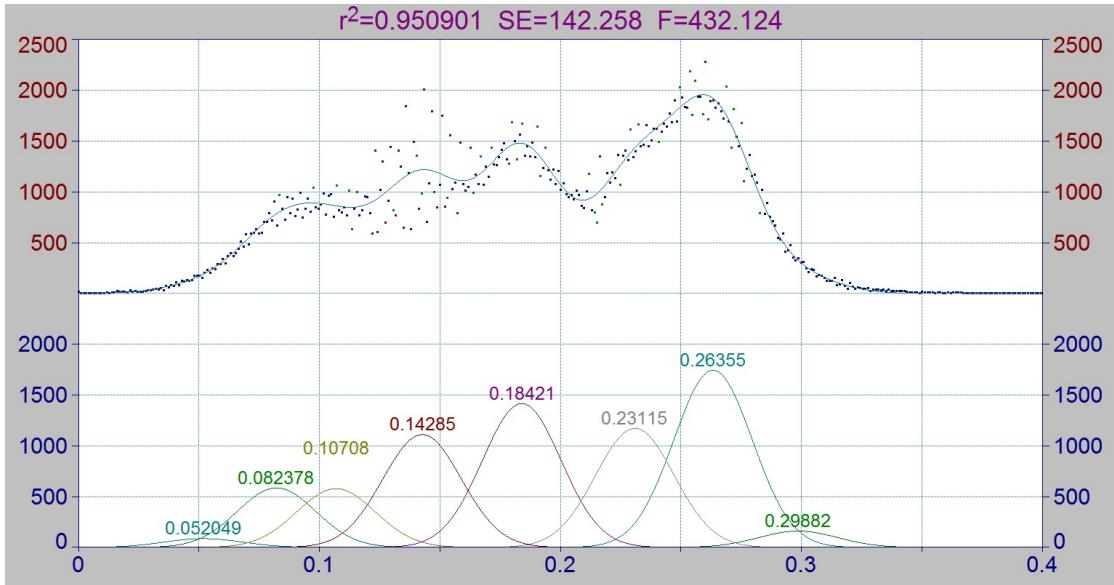


FIGURE 4.15: Gaussian peaks that have been automatically fitted to the histogram from figure 4.14.

These curves can be used to define classes for an unsupervised classification; the peak of each curve can be considered to be the centre of a class, each pixel can then be assigned to a class based on its proximity, in metric space, to the class centres. Performing this with the Lena image gives the classifications shown in figure 4.16 below. The classes are coloured in order of spectral angle, from lowest to highest, according to the scale in figure 4.17. This scale was selected to ensure contrast between adjacent groups.

From figure 4.16, it can be seen that the qualitative observations made above have been largely confirmed in this quantitative approach. The soft skin tones are represented by the consecutive classes 6, 7, 8 and 9. One feature that is shown up is the contrast between the reflection from the mirror and the rest of the scene, this was not so marked in figure 4.13. The darker areas of the image are represented by the consecutive classes 3, 4 and 5. This confirms that the mirror frame and the hair are indeed very similar, however as discussed above, the feather is indistinguishable from the front of the hat.

In three dimensions, it is straightforward to examine both planes of the SPM, and these are shown in figure 4.18 below. It can be seen from this figure that both planes provide good contrast between the feather and the rest of the image. It is clear then that using the SPM in conjunction with the SAM allows the separation of areas of the image that appear the same to the SAM alone.

Figure 4.19 shows a scene of some snooker balls on a snooker table. The same analysis is performed for this image, and the classified scene is shown in figure 4.20.



FIGURE 4.16: Lena image with pixels classified by spectral angle, using the Gaussian peaks from figure 4.15 as class centres.

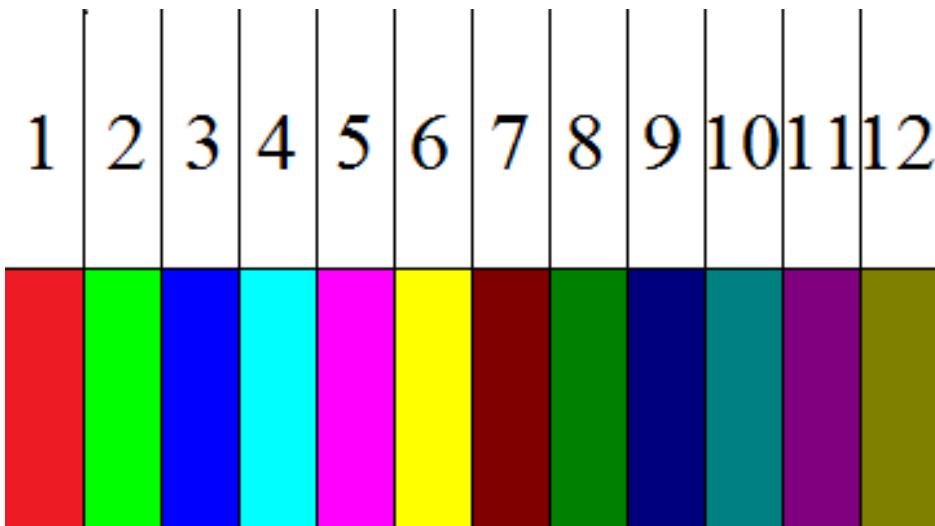


FIGURE 4.17: Scale bar showing class numbers with the colours assigned to them.

It is clear from this image that the white, black, blue, brown and pink balls are entirely indistinguishable by spectral angle alone, all being classified into class 2. The baize of the table consists of three consecutive classes, 3, 4, and 5. These classes also make up the majority of the pixels of the green and yellow balls, which all but disappear from the image; they remain distinguishable only because of the specular reflections of the light source and the shading around the edges of the balls which renders their shapes recognisable, statistically they are indistinct from the background.

The fact that there are four classes making up the red balls (6, 7, 8 and 9) would

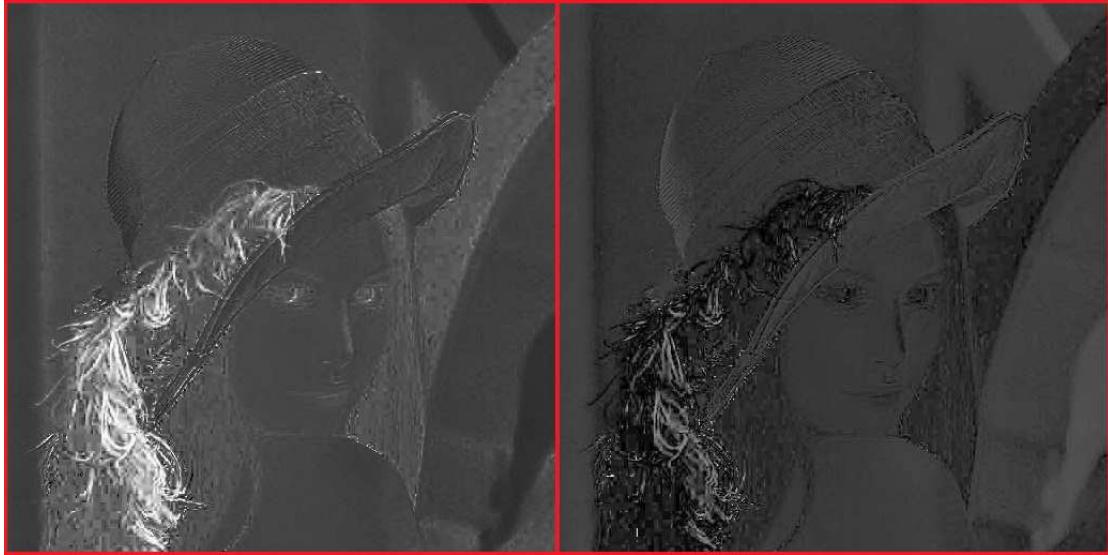


FIGURE 4.18: The planes of the SPM of the Lena image, plane 1 on the right hand side and plane 2 on the left hand side.



FIGURE 4.19: An RGB image of some snooker balls on a table.

suggest that more peaks have been fitted than may actually have been present, however, as the distribution on each ball is similar, it might also suggest that this is a function of the lighting. The SAM is, as explained above, intensity independent, however it is possible that the colour response of the camera is not. It is possible, and indeed figure 4.20 suggests that it is even likely, that the vectors recorded by the camera do not scale linearly with changing light intensity.

Again, the SPM was calculated for this image and is shown in figure 4.21 below.

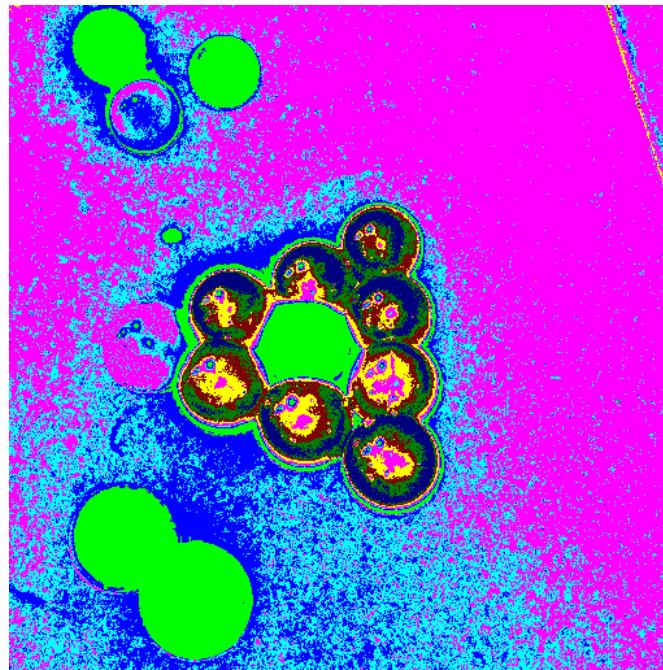


FIGURE 4.20: Snooker ball image with pixels classified by spectral angle, using the Gaussian peaks fitted to the histogram of spectral angles.

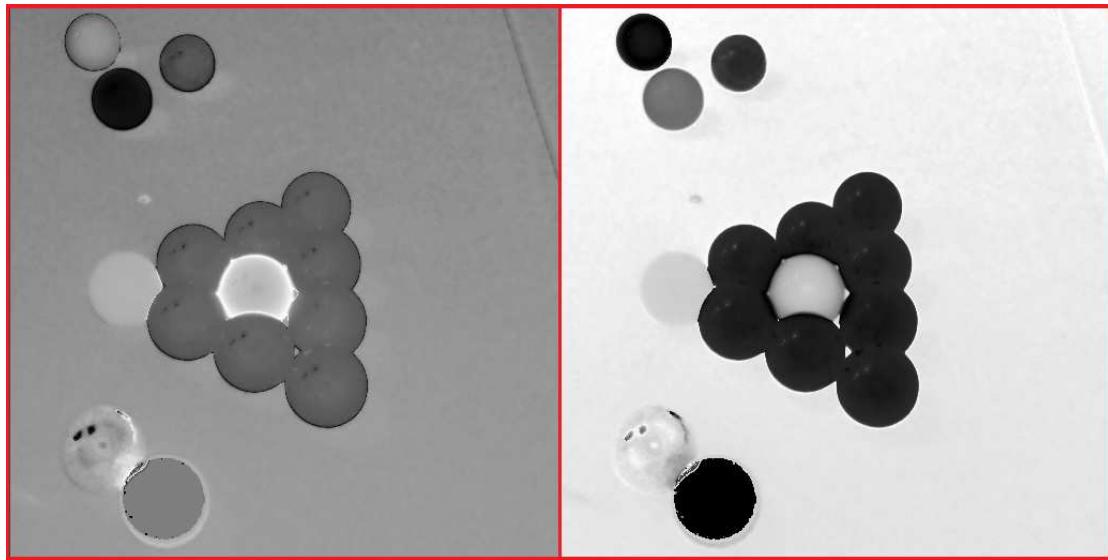


FIGURE 4.21: The planes of the SPM of the snooker balls image, plane 1 on the right hand side and plane 2 on the left hand side.

It is immediately obvious from the SPM that the yellow ball can be clearly distinguished from the baize and the green balls; the green ball is still similar in appearance to the baize, which is not unexpected given the two share a similar colour, however it is significantly lighter than the baize in the φ_1 (left) image and significantly darker in the φ_2 (right) image. In order to assess how well the SPM separates out the classes assigned by the spectral angle, the average intensity values of each component of the image are presented in table 4.1 below. Each component is represented by a

5×5 pixel area, representative of the area as a whole.

Region	SAM Class	φ_1 Intensity	φ_2 Intensity
Background/Baize	3,4,5	56%	93%
Red Balls	6,7,8,9	35%	14%
Blue Ball	2	80%	70%
Green Ball	3,4,5	68%	81%
White Ball	2	50%	0%
Black Ball	2	57%	91%
Pink Ball	2	58%	7%
Yellow Ball	3,4,5	5%	43%
Brown Ball	2	32%	17%

TABLE 4.1: A table showing the class and values of spectral position for various areas on figure 4.19

The combination of the spectral angle and φ_1 metric separates all but the black and pink balls, which are similar in both maps; the introduction of the φ_2 map removes this last remaining inseparability. However, using the spectral angle and the φ_2 map separates all areas without the need for the third metric at all.

The analysis of both of these images demonstrates that in seeking to represent the multidimensional data by a single metric, important discriminatory information is lost. The SPV provides an additional source of information that, whilst in some way related to the spectral angle, is independent from it. By its use, some of the lost information can be regained, allowing improved discrimination whilst still using only a subset of the original data.

Whilst they are illuminating and easily understood, RGB images are not typical of spectral images as a whole. They are easily visualised and can be processed without recourse to metrics like the spectral angle. Further, the SPV associated with them is a 2D object and so it is feasible to examine the entire SPM for a given image; as the number of dimensions increases, the assumption that the entire SPM can be visually studied becomes increasingly invalid. The next section will introduce new ways of using the SPV data to perform a similar enhancement of discriminatory power for real spectral images.

4.3 A Demonstration with Spectral Images

The images in this section are real spectral images made publicly available through The University of Eastern Finland (University of Joensuu Color Group, 2006), they have varying numbers of dimensions and have been collected with various filter settings. The colour image representations have been generated by simply selecting a single plane to represent each of the red (nearest filter to $\sim 650\text{nm}$), green ($\sim 510\text{nm}$)

and blue ($\sim 425\text{nm}$) planes of an RGB image, as such they are better described as pseudo-colour images, but give a sense of the visual scene being imaged.

The first image is a landscape scene of a tree in the foreground with fields and bushes making up the background. It is comprised of 7 spectral dimensions; these correspond to wide-band filters (full width at half maximum (FWHM) of around 40nm) evenly spaced between 400nm and 700nm (approximately centred at 400nm, 455nm, 505nm, 555nm, 605nm, 655nm and 700nm). The pseudo-colour image is shown in figure 4.22.



FIGURE 4.22: Pseudo-colour image of a countryside landscape spectral image.

Performing the same unsupervised classification on this image as described in section 4.2 above results in the image shown in figure 4.23. From this it can be seen that the track in lower part of the image is assigned to the consecutive classes 1, 2 and 3; these classes also pick out the top left corner of the image, which appears to be distant bushes in the pseudo-colour image. The bushes in the middle distance are largely classified to class 4 although there is some crossover into class 5, which largely covers the field in the middle of the image. Both the tree and the short grass in the foreground of the image are put into the same class. The major misclassification in this example would appear to be that the tree and the short grass were not separated.

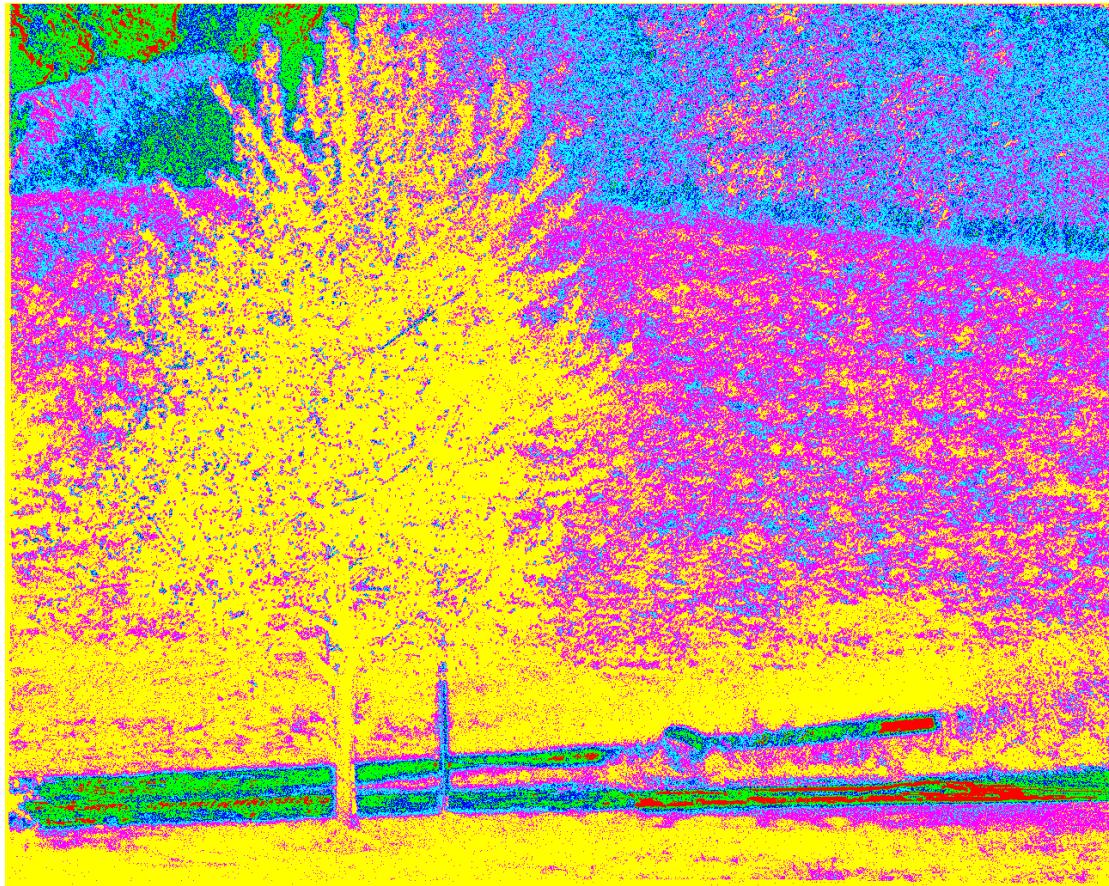


FIGURE 4.23: Landscape image with pixels classified by spectral angle, using the Gaussian peaks fitted to the histogram of spectral angles.

With this small number of dimensions it is still possible to look at all of the SPM planes in order to assess where useful extra information can be obtained, however even at 7 planes, this approach is reaching its limits of feasibility. An alternative approach is to consider a 2D histogram of spectral angle against each of the SPM planes in turn. The plots for this example are shown below in figure 4.24 below. In each plot, the spectral angle runs across the horizontal and the spectral position runs along the vertical. Each pixel is shaded according to how many image pixels have the combination of spectral angle and spectral position values represented at that point, the more image pixels, the darker the histogram pixel.

Instead of looking for contrast across the image plane, it is possible to look briefly at each histogram to decide where the interesting information might be. Identifying the spectral angle range in which misclassification is occurring, it is straightforward to assess which plane shows the greatest variance in spectral position or the greatest number of separate peaks for this range. Having established which plot to use, specific “structures” on the plot can be isolated and the pixels contributing to them on the spatial image can be identified. In this example, the first spectral dimension is selected and two regions of interest identified. Figure 4.25 below shows the two regions of

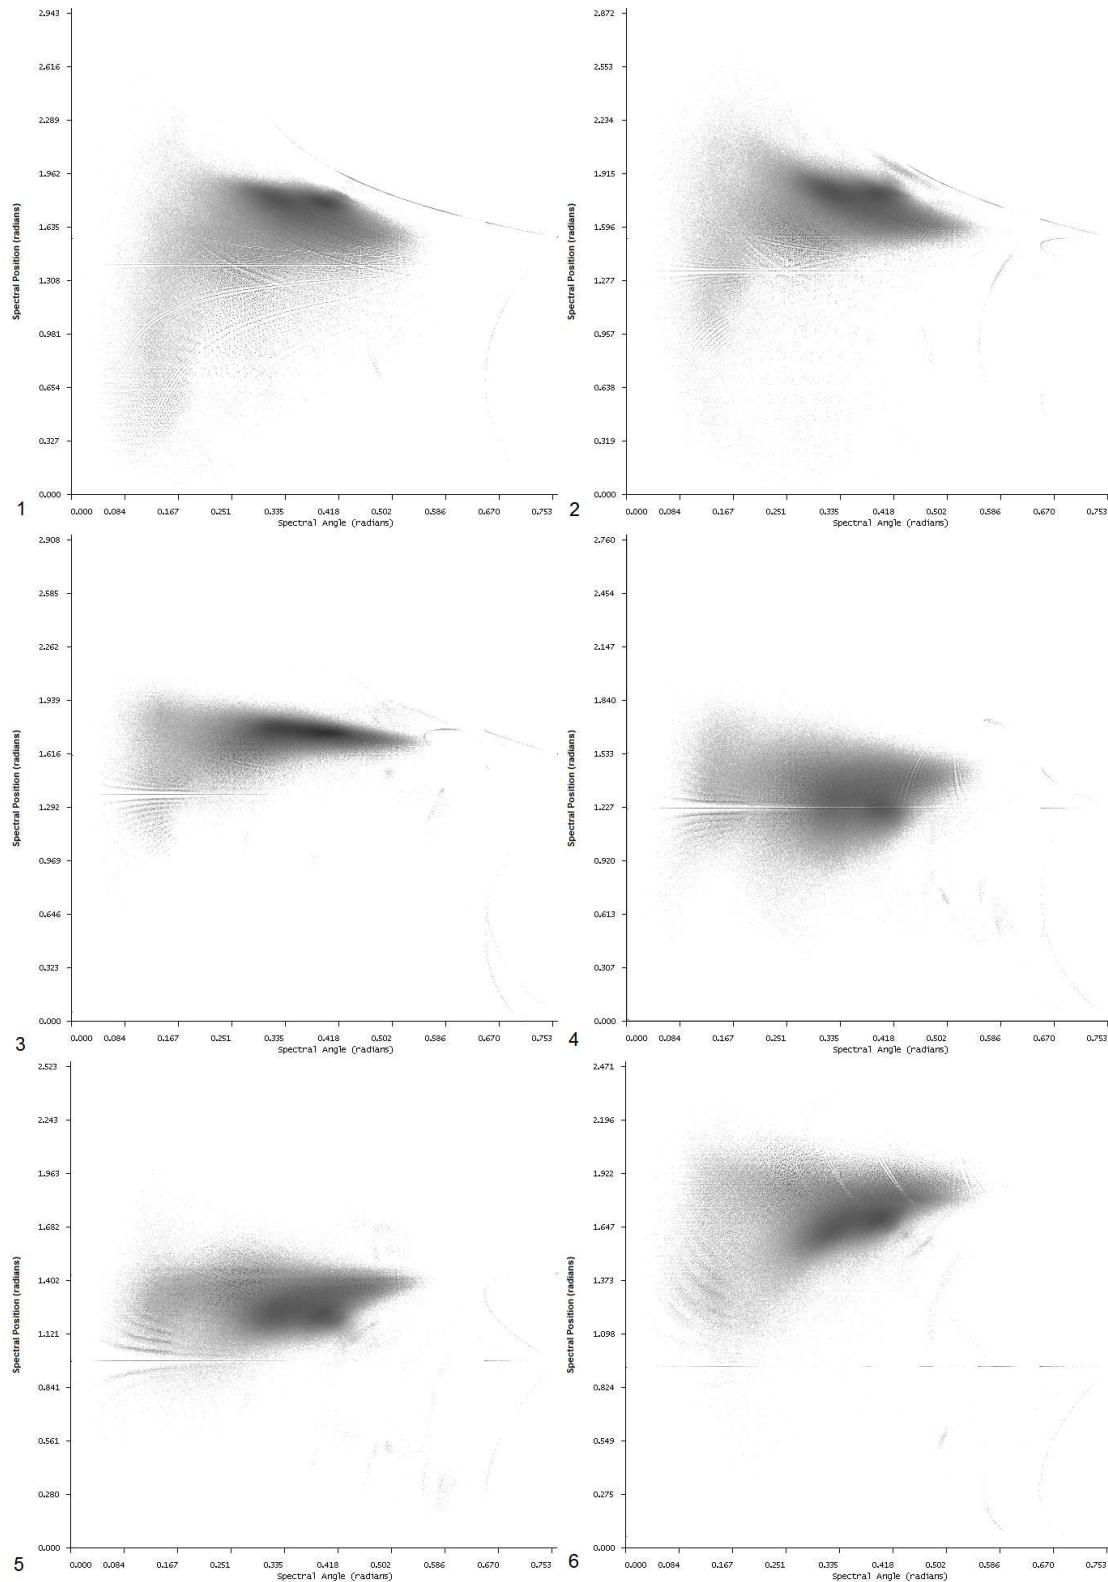


FIGURE 4.24: 2D histograms of spectral angle (horizontal axis) against planes of spectral position (vertical) for the landscape image. The plane of the SPM used in each plot is marked in the bottom corner of the plot. The darker pixels on the plot represent more counts in the histogram.

interest selected, the areas within the red and green boxes. The boxes cover the same spectral angle range, covered by group 6 from figure 4.23.

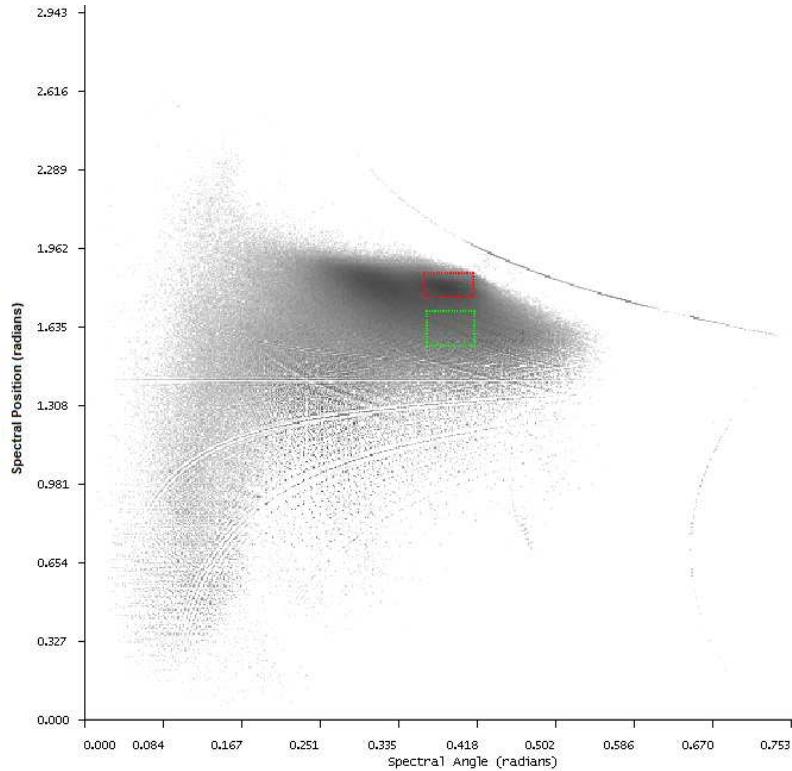


FIGURE 4.25: 2D histogram of spectral angle against the first plane of the SPM for the landscape image. Two regions of interest have been enclosed in a red and a green rectangle.

The pixels in the spatial image that contribute to these boxes are highlighted on the spatial image itself in figure 4.26. Those in the red box are coloured red, those from the green box are coloured green, and those in neither are shown in grey-scale.

It can be seen from figure 4.26, that using information from the first SPM plane, the tree and the short grass can be successfully separated.

The second image is a standard colour wheel target; this image has 16 spectral dimensions representing broad-band filters (FWHM of $\sim 25\text{nm}$), evenly spaced ($\sim 20\text{nm}$ apart) between 400nm and 700nm. The pseudo-colour image is shown in figure 4.27.

Again, classification was performed as described above using only the spectral angle information; the resulting image is shown in figure 4.28. In a number of ways this classification performs as would be expected; the background is a single class and is matched by the white/grey areas in the centre of the two colour wheels. In the large wheel, there is reasonably good separation of areas of different colour; however, there are also some unexpected class shapes, in particular, group 8 representing the yellow area of the original image extends beyond the yellow and through areas that are

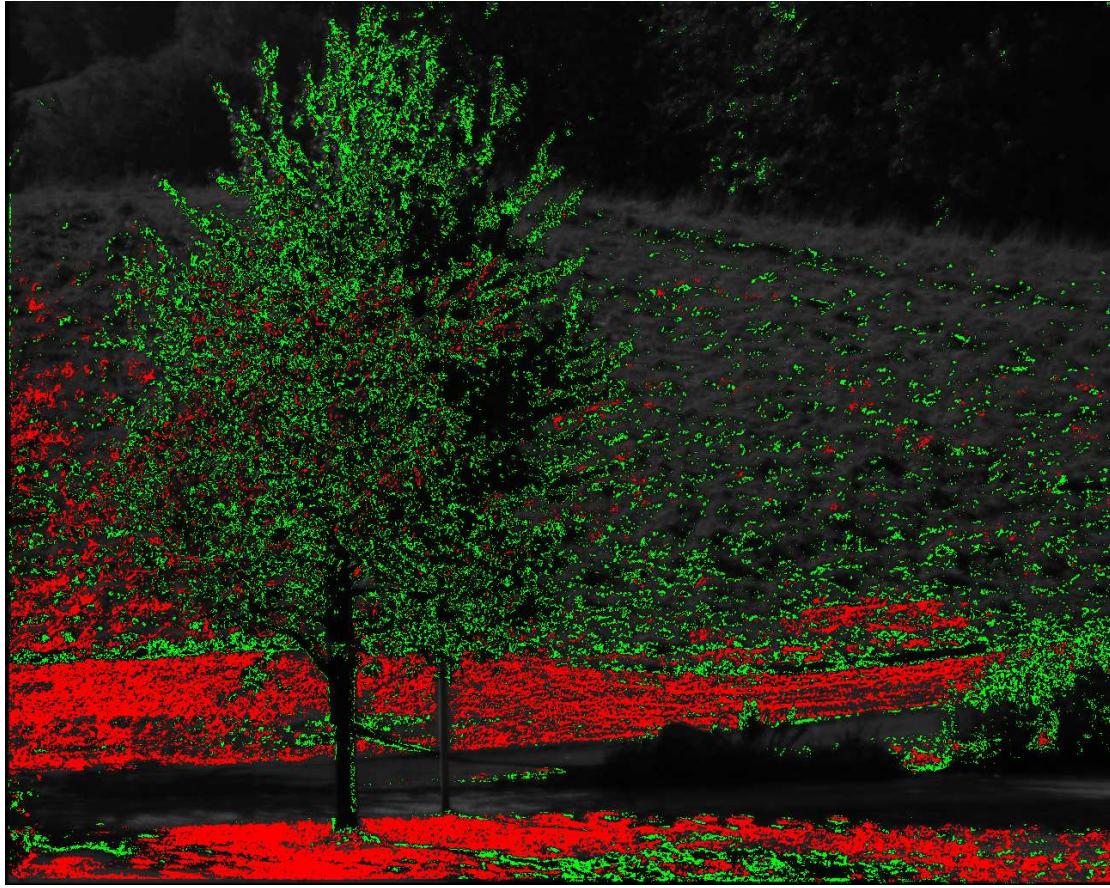


FIGURE 4.26: The landscape image with those pixels from the red and green boxes from figure 4.25 shown in red and green respectively.

orange and pink. There is a further area of this class at the top left of the wheel, in the green coloured area.

With 15 SPM planes here to choose from, visual assessments are starting to be impractical. One way of selecting a plane in which to look for extra information is to take a quantifiable measure of the contrast existing in a given plane. In this context, contrast is a measure of the size of the spread of values present, and so measuring the variance or standard deviations of the values of the SPM in each plane is a sensible starting point. However, these are absolute measures of spread, and so should be normalised against a representative measure of magnitude, for example, the mean value in each plane. This way, the plane showing the most significant spread of values can be calculated. The contrast metric thus used is defined for each plane as:

$$C = \sqrt{\frac{\sum_i (x_i - \bar{x})^2}{\frac{N}{\bar{x}}}}, \quad (4.1)$$

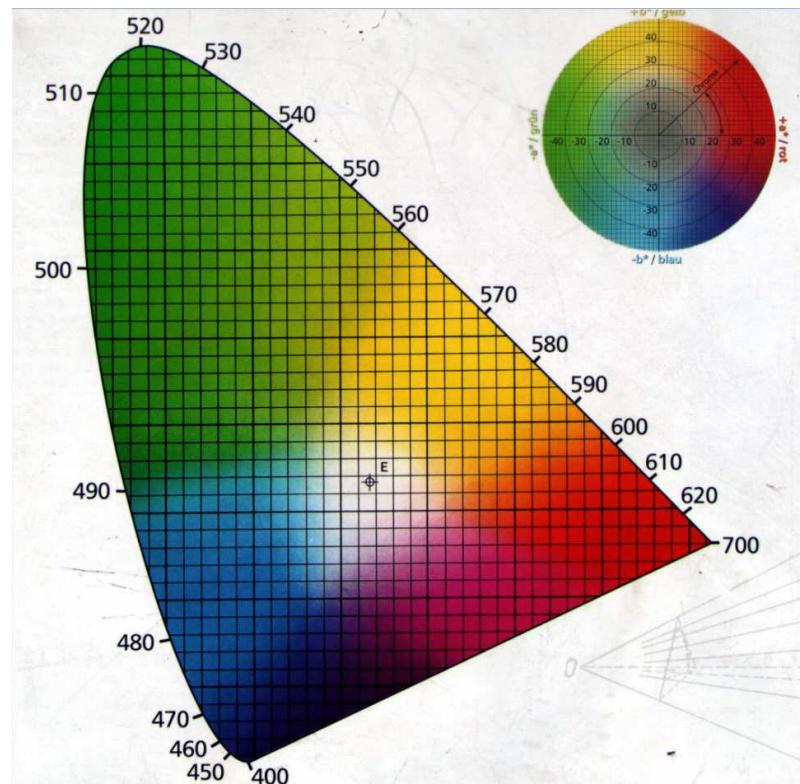


FIGURE 4.27: Pseudo-colour image of a colour wheel target spectral image.

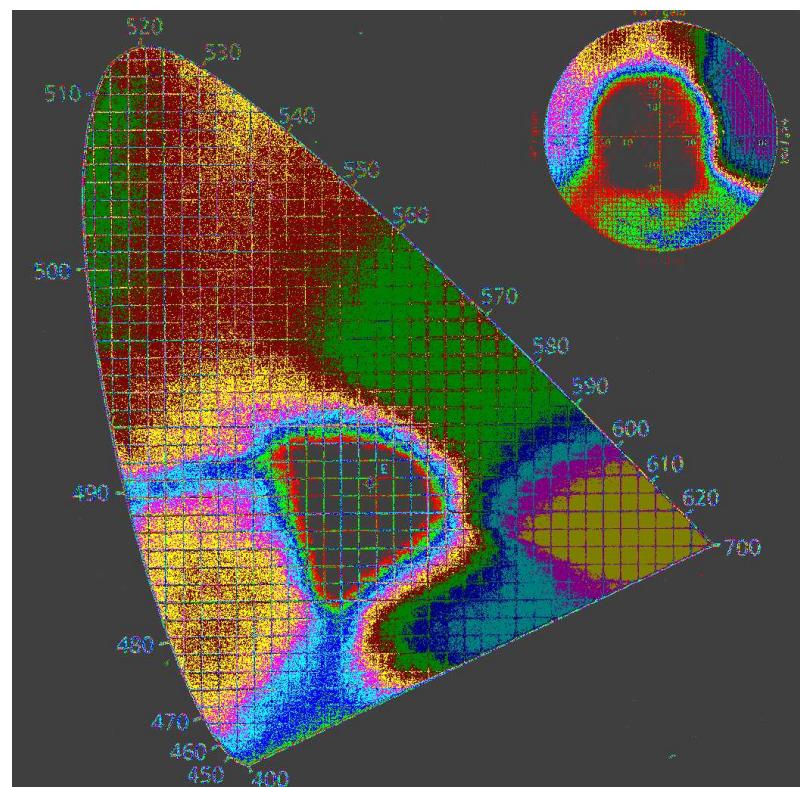


FIGURE 4.28: Colour wheel target image with pixels classified by spectral angle, using the Gaussian peaks fitted to the histogram of spectral angles.

where N is the number of spatial pixels in the image, x_i is value of the SPM plane at pixel i , and

$$\bar{x} = \frac{\sum_i x_i}{N}. \quad (4.2)$$

Calculating C for each plane of the SPM in this example reveals that the 10th plane has the most contrast, and so this is used for the 2D histogram in figure 4.29.

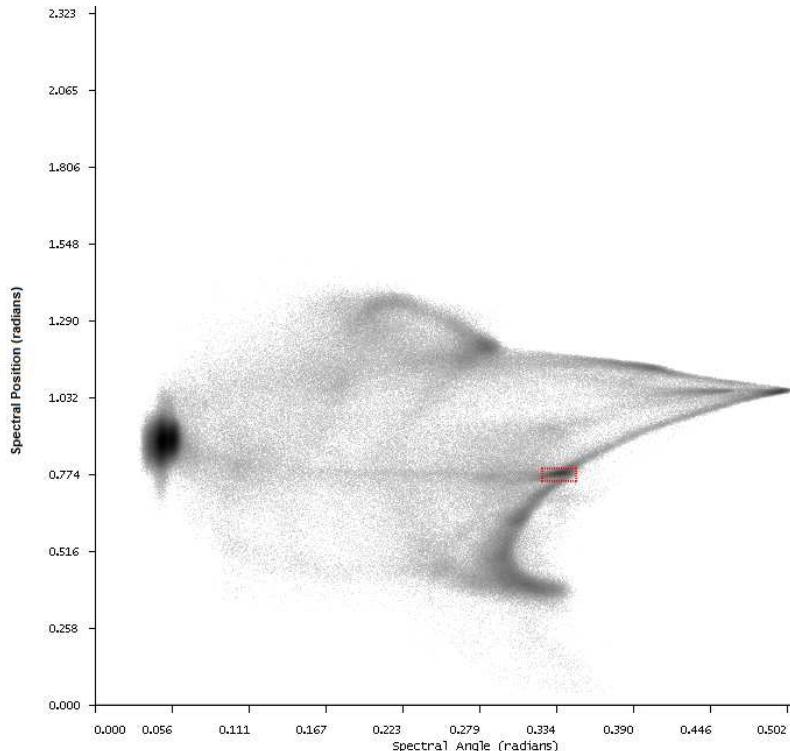


FIGURE 4.29: 2D histogram of spectral angle against the first plane of the SPM for the colour wheel. A region of interest has been enclosed in a red rectangle.

The red box in figure 4.29 highlights a single peak lying in the spectral angle region coinciding with class 11 in figure 4.28. It is clear that this represents just a subset of the pixels in this spectral angle range, but also that it is the most significant structure entirely in this range. Mapping this region back into the spatial domain gives the image shown in figure 4.30. From this it can clearly be seen that the peak represents the yellow area from the original image, separating it entirely from the rest of class 11.

These two spectral images have demonstrated that classification information that is demonstrably present in the original images and subsequently lost with the reduction to spectral angles, can be recovered by use of additional information provided by the SPM. The spectral position is a natural extension to the spectral angle and one which provides extra, useful information, thus helping to improve unsupervised classification.

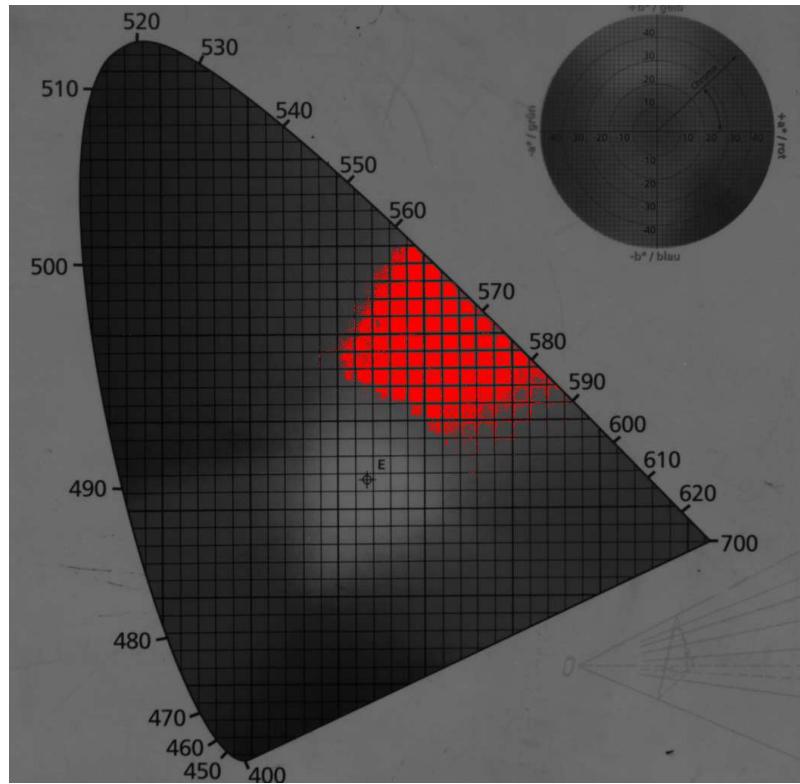


FIGURE 4.30: The colour wheel image with those pixels from the red box from figure 4.29 shown in red.

4.4 A Comparison to Established Alternatives

In this section the spectral image from figure 4.22 is examined with a standard unsupervised classification using the commercially available ENVI software³, one of the most common tools in use for spectral image analysis.

ENVI's k-means classifier was set up to classify data into six classes (to match the number found in figure 4.23), the result of which is shown in figure 4.31.

From this it is immediately apparent that a number of differences are exhibited in the classes found in figure 4.23 and those in figure 4.31. The path that was so apparent in figure 4.23 is completely indistinguishable in the k-means image, and whilst the tree trunk and the left hand side of the tree foliage are separated from the middle-ground field and the foreground, they are not separated from the shrubs either in the lower left hand side of the image, or along the top of the image. Further, under the k-means classifier, the field in the middle of the image appears to transition, from being dominated by magenta pixels towards the top to more green, yellow and blue pixels towards the bottom.

³<http://www.exelisvis.com/language/en-us/productsservices/envi.aspx>

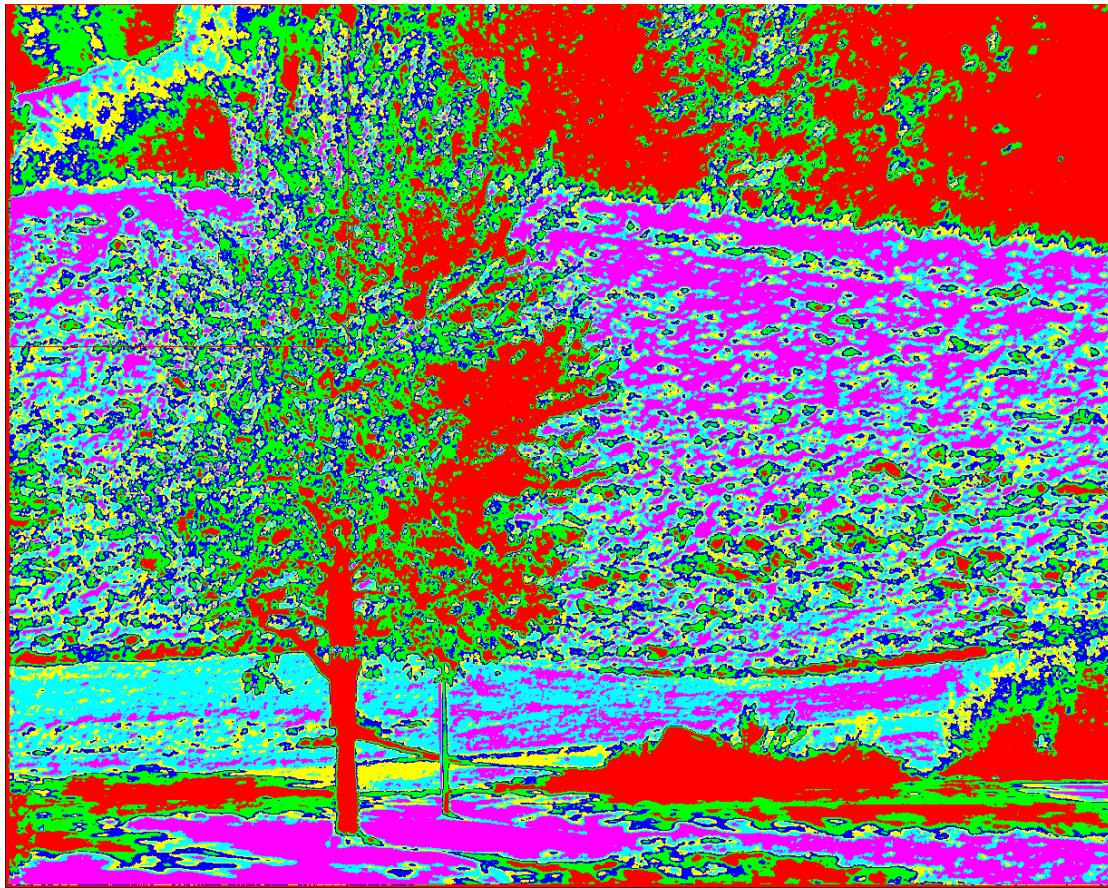


FIGURE 4.31: Landscape image with pixels classified using ENVI's k-means classifier.

Clearly, although they use the same basic premise of measuring distances in a metric space, the two classifications produce different results. It is difficult to definitively say that one performs inherently better than the other as the tree trunk is very much clearer in the k-means, the path is much better defined in the image in figure 4.23. Both methods also fail to separate areas of the image which would be expected to be different, although the details of how they do this differ.

The image was also analysed using PCA. The image in figure 4.32 shows a gray scale image where the value of each pixel is the value in the first principal component. This would suggest that the first principal component is very related to the total intensity of illumination. Since this is not very interesting information, and as the spectral angle is not affected by total intensity, this first principal component is disregarded for the rest of this discussion.

The image in figure 4.33 shows the second, third and fourth principal components mapped into the red, green and blue channels of a bitmap respectively.

In figure 4.33 the path very distinguishable, much like in figure 4.23, and, as in figure 4.31, the tree trunk is distinguishable from the areas around it, yet appears very similar to bushes on the lower and upper left areas of the image.



FIGURE 4.32: The first principal component of the landscape image mapped spatially.

As PCA can be thought of as a re-encoding of the data, k-means classification can be performed on the resulting dataset. This was performed (again, not including the first PC) and the result is shown in figure 4.34.

Again, the path here is clearly visible, however much like the image in figure 4.23, the tree is not very clearly distinguished, particularly from the grassy areas in the foreground and to the front of the middle-ground.

4.5 Conclusions

The analysis of the test images in this chapter has confirmed that in a situation where the properties of the spectral position vector can be predicted, it does indeed behave as expected. Further, this has suggested that there are some systematic properties of the spectral position that are not yet fully explained and may provide interesting future research.

The analysis of the spectral images has demonstrated that classification information that is demonstrably present in the original images and subsequently lost with the



FIGURE 4.33: The second, third and fourth principal component of the landscape image mapped spatially.

reduction to spectral angles, can be recovered by use of additional information provided by the SPM. Further, it is clear from the analysis using standard alternative classification techniques, that the spectral position provides extra, useful information that can help with unsupervised classification.

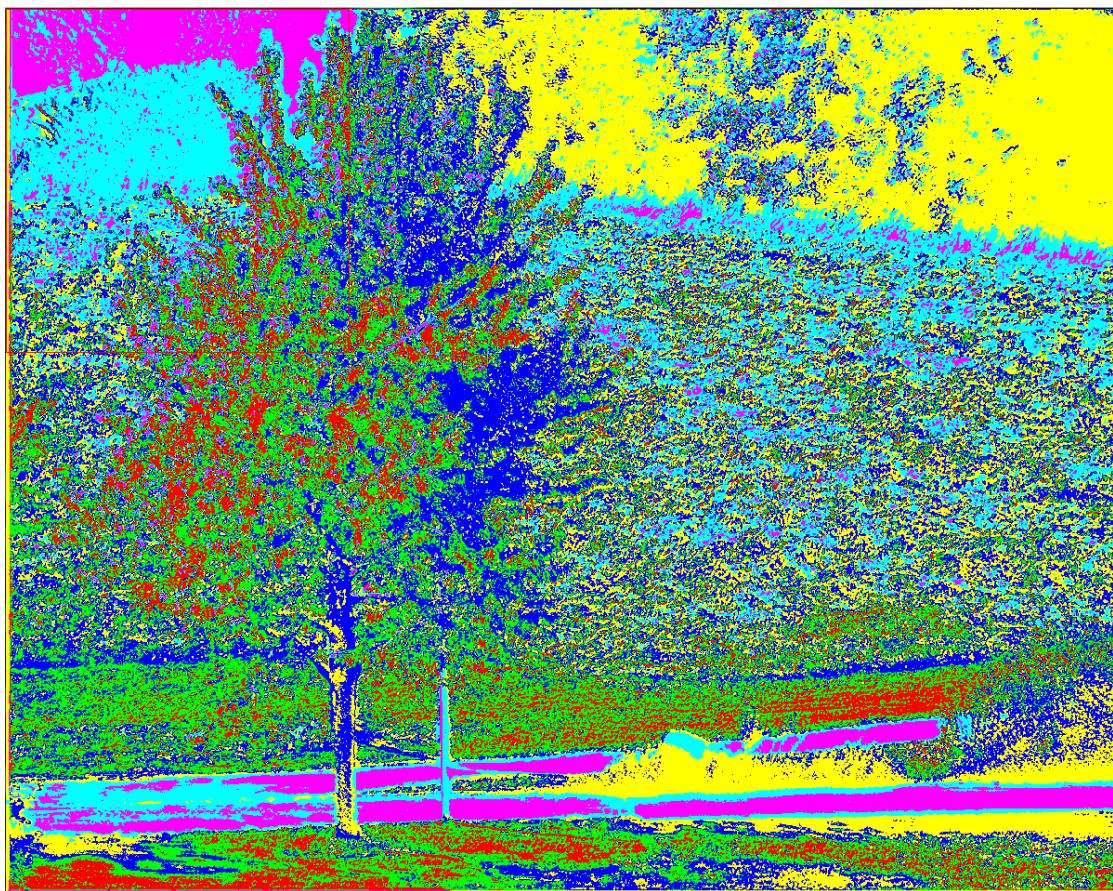


FIGURE 4.34: Landscape image with pixels rotated to principal component space before being classified using ENVI's k-means classifier.

Chapter 5

Towards an Optical Sentinel Lymph Node Biopsy

The spectral imaging techniques described in previous sections are generic techniques that are applicable to any form of spectral imaging in any discipline. However, their development has arisen from the study of the potential application of spectral imaging to medical applications; specifically, to the use of spectral imaging as an optical technique to perform tissue biopsies. The idea of using spectral imaging for this purpose arose from discussions with Mr. Richard Sainsbury, a consultant oncologist at the University of Southampton. Mr. Sainsbury and his surgical team at St Mary's Hospital, Newport, Isle of Wight, have kindly collaborated in this work since those initial discussions. This work was performed under local ethics clearance from St Mary's Hospital, which was kindly organised by Mr. Sainsbury.

This chapter describes the work undertaken in this regard; section 5.1 explains the clinical setting and provides a review of the literature which motivates the work, section 5.2 describes the work performed in setting up the imaging system and the data processing suite necessary for this work, and section 5.3 presents and discusses some preliminary results obtained.

At this stage, it should be noted that non-melanoma skin cancers are conventionally omitted from the quoted cancer statistics; this is due to uncertainties and inconsistencies in their registration, and because they are, on the whole, highly treatable and even curable. The statistics quoted in 5.1.1, unless otherwise referenced, are obtained from Cancer Research UK (2011) and all follow this convention.

5.1 Background

5.1.1 Breast Cancer and the Sentinel Lymph Node Biopsy

Breast cancer is an extremely prevalent form of cancer worldwide, and is the oldest recorded cancer. It was first described in the Edwin Smith Papyrus, part of an Egyptian textbook on surgery dating to around 3000BC, it was noted that there was no treatment and cauterisation was recommended as a palliative care.

Today, for women in the UK, the lifetime risk of developing breast cancer is estimated to be 1 in 8. Despite its very low incidence in men, breast cancer is the single most common form of cancer in the UK, accounting for around 16% of all cancers. This figure is around 31% of cancers in women, nearly three times higher than the second leading form (lung cancer, accounting for around 11%). Less than 1% of all breast cancer patients are men. In real terms, there were 45,822 new cases of breast cancer in the UK in 2006, of which 45,508 were in women. Breast cancer occurs at a rate of 75.6 cases per 100,000 people (147.3 cases per 100,000 women) per year.

Whilst accounting for some 16% of cases, breast cancer only accounts for around 8% of all cancer deaths in the UK. Indeed, the mortality rate (deaths per 100,000 people in a year) has been in slow decline since around 1990, even as incidence has been increasing. The decrease in mortality rates can be attributed to improved treatment and earlier diagnoses.

Survival rates are measured as a percentage of those diagnosed who are still alive at yearly intervals after initial diagnosis. These have been increasing in the UK since 1971 but are still very heavily linked to how early the cancer is diagnosed. The overall five year survival rate for cases diagnosed between 2001 and 2003 is 80%, for those diagnosed between 1971 and 1975 this rate was just 52%. It is predicted that the 20 year survival rate for the 2001-2003 cases will be over 60%, which is a clear indication of the improvement in survivability.

The development of cancer is split into four clinical stages, details for which are given in table 5.1. This table pre-empts a discussion on lymph nodes and the lymphatic system, which is presented later on in this section, however it is necessary to introduce these stages here.

A study of breast cancer cases diagnosed in the West Midlands between 1990 and 1994 showed that for patients diagnosed with stage I cancers the five-year survival rate was 92%; this dropped to 73% for stage II cancers, 50% for stage III and just 13% for stage IV. Ten year survival rates were 87% for stage I, 63% for stage II, 40% for stage III and 8% for stage IV. Early detection is therefore vital to effective treatment, and being able to accurately assess the stage of the cancer is important to being able to plan a treatment regime and to provide a reliable prognosis for the patient.

Stage	Clinical Description
I	Tumour size < 2cm No lymph nodes affected No evidence of spread beyond the breast
II	Tumour size 2-5cm Axillary lymph nodes affected No evidence of spread beyond axilla
III	Tumour size > 5cm Axillary lymph nodes affected No evidence of spread beyond axilla
IV	Tumour of any size Non-axillary lymph nodes affected Cancer has metastasised to another part of body

TABLE 5.1: Clinical descriptions of the stages of breast cancer at diagnosis

Diagnosis of breast cancer is typically made by way of x-ray mammography; the size and location of the tumour can be determined from this mammogram. In many cases, a patient will have discovered a palpable lump, however routine screening programmes have been implemented to improve detection rates and these also aid in early detection. Treatment is typically a surgical resection of the tumour, for particularly large tumours a complete mastectomy may be required. Sometimes adjuvant treatments such as radiotherapy and chemotherapy may be performed to reduce the risk of recurrence, or where micro-metastasis is suspected.

The lymphatic system is a circulatory system that is responsible for draining interstitial fluid from tissues and returning it to the bloodstream. It consists of a network of small capillary-like vessels called “initial lymphatics”, filter-like structures called “lymph nodes”, and larger lymph vessels; the fluid transported is referred to simply as lymph. Fluid filters into the initial lymphatic through a series of valve-like structures on the vessel walls before being “pumped” by muscular contractions of the vessel walls into the larger lymph vessels. At various points in the system there are clusters of lymph nodes; these nodes produce and process white blood cells, which play a key role in the immune system by destroying harmful foreign particles. The nodes thus filter the lymph and also release these white blood cells for eventual deposition into the blood stream. The lymphatic vessels return lymph to the main circulatory system near its entrance to the heart at the right atrium.

As a consequence of being the system by which excess interstitial fluids are transported around the body, the lymphatic system is also a significant pathway for metastasis of cancer. Cancerous cells from the primary tumour site are transported in the lymph and can accumulate in a secondary tumour site in another organ. The first indication that this may have started to happen will be an accumulation of cancerous cells in the lymph nodes nearest to the primary tumour site. As the process progresses, other node

clusters may become involved before a secondary tumour forms. In this way, the extent of involvement of the lymph nodes is an indicator of the stage of the cancer. Assessing this involvement is therefore a key task to reliably determine the stage of the cancer.

The first cluster of nodes that fluids from the breast area will encounter are located in the axilla (armpit), and within this cluster, there is single node that is the first to which fluids drain, this is the “Sentinel Lymph Node” (SLN).

The historical approach to assessing the lymphatic involvement was to completely remove the axillary cluster of nodes during the course of tumour resection surgery and perform histological examination of the whole cluster; this procedure is referred to as an Axillary Lymph Node Dissection (ALND). Due to their role in the immune system, such a removal of nodes poses a long term health risk. Further, it can lead to blockages in the lymph vessels that stop fluid from being effectively removed from tissue; this can lead to severe swelling due to a build up of excess fluid, particularly common in the limbs, known as edema or oedema.

More recently, the technique at the heart of this thesis, the Sentinel Lymph Node Biopsy (SLNB), has become increasingly standard as a way to preserve healthy lymph nodes whilst assessing lymphatic involvement. The patient is injected with a tracer, or combination of tracers, at the tumour site before surgery. The common tracers are a colloidal solution of a radioactive material, metastable technetium 99 (^{99m}Tc), and a blue food dye called Patent Blue V (E131). These tracers will drain rapidly to the sentinel node, which can then be identified during surgery by its blue colour and/or its high radioactivity. This sentinel node is then resected and examined for the presence of cancerous cells.

The node is sliced open through a longitudinal axis and opened or “butterflied”. A visual examination of the cut surface is sufficient only in the most clear cut cases when the cancer has spread completely through the node. Figures 5.1 and 5.2 show the cut surfaces of two lymph nodes that have been resected, demonstrating that colour images are also not reliable, even as indicators of the extent of the node compared to the surrounding tissue. The cut surface is pressed against a microscope slide, which is then examined cytologically, a technique known as TP Cytology (standing for either Touch Print or Touch Prep.). The node itself is then fixed in formalin for subsequent slicing, staining and histological examination. The TP cytology is a rapid process (compared to traditional histology, which has a long preparation time), and an indication of whether cancerous cells are present or not can be obtained in around an hour. Evidence suggests that if the sentinel node is unaffected by cancerous cells, then the probability that other nodes are involved is very low (Howard et al., 1998); the rest of the node cluster can thus be left intact. If cancerous cells are detected then a full dissection will still be required to assess the extent of the spread.

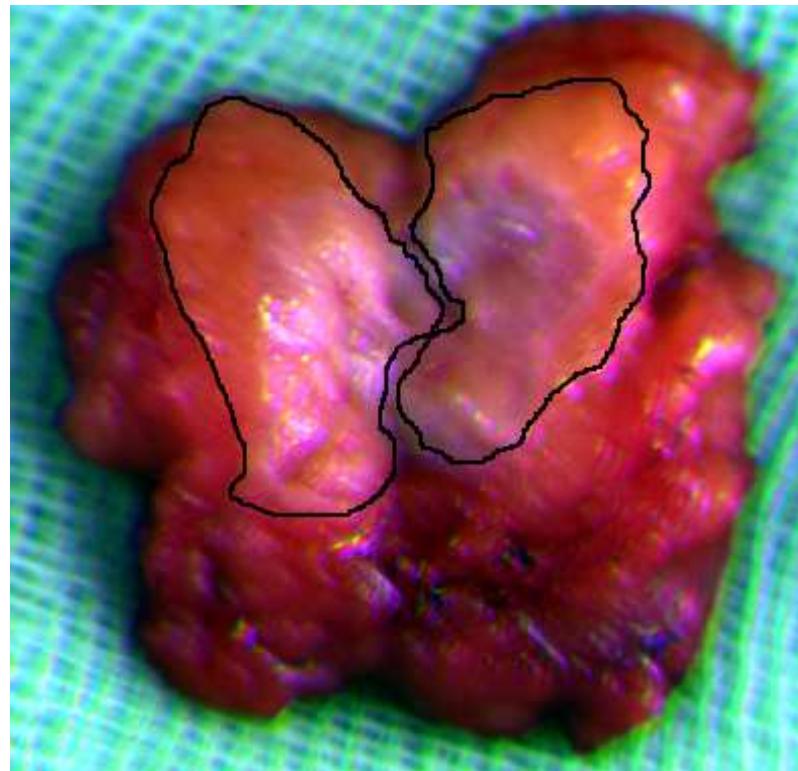


FIGURE 5.1: The cut surface of a butterflyed lymph node, set in some fatty tissue. The approximate extent of the node tissue is marked in black, and although the central part is distinguishable by a bluish appearance, the exact border is not clear from this image.



FIGURE 5.2: The cut surface of a butterflyed lymph node, set in some fatty tissue. The approximate extent of the node tissue is marked in white, and as with the above example, some of this is distinguishable from subtle changes in the colour, but the border is not at all clear.

Although a comparatively rapid process, the TP Cytology is still often the longest part of the surgery. The time taken for a result to come through is time that the patient is anaesthetised and “open” in the operating theatre, and often time in which nothing

else is happening, the tumour having already been resected. This is a cost not only to the patient in terms of increased risk of infection or complications from anaesthetic, but also to the treatment centre in terms of the turnaround time of the operation; reducing this waiting time would allow more surgeries to be performed in the same period.

The situation is actually potentially worse, if the facilities for performing the TP Cytology do not exist on-site at the treatment centre, then a SLNB can only be performed if the node is sent to external laboratories for assessment. In the case that cancerous cells are detected, the patient will require further surgery to perform the ALND, potentially at a delay of several weeks. Again, this presents a cost to the patient, not only in terms of the risks that are inherent with any surgery, but also the anxiety caused by the lack of certainty over the need for further surgery; again, there is also a cost to the treatment centre in terms of reducing the throughput of patients.

As a result of these considerations, the SLNB has been identified as an area that would greatly benefit from a technique that will reduce the time taken to produce a result. In order to increase the utility of such a technique, it will also have to be low cost and ideally automated, or at least a technique that requires no specialised operators. These two factors would allow it to be readily available at all treatment centres.

Following discussions with Mr. Sainsbury, and on the basis of a review of existing spectroscopic work on lymph nodes detailed in section 5.1.2, spectral imaging in the VNIR region has been identified as potential technology for performing such a biopsy; the technologies involved are readily commercially available at low cost, and current image processing work suggests that the processing could be automated and rapid.

5.1.2 Medical Review

Whilst the SLNB has been presented here as a routine procedure, like much in the field of medicine, it continues to be an active field of research. Whilst a full discussion of the medical literature would be beyond the scope of this report, it is illuminating to know in broad terms how the procedure is being studied by experts in the field, not least to ensure that this project is starting on a firm clinical footing. Some background on the medical literature surrounding the SLNB is thus provided in section 5.1.2.1 before sections 5.1.2.2 and 5.1.2.3 present a review of the literature in relation to spectroscopic studies of cancer and lymph nodes.

5.1.2.1 The Efficacy of the SLNB

This section provides a brief a overview of the ongoing research as to how surgical intervention can be best applied in cancer patients. As with any new surgical

procedure, two main concerns from the outset have been whether the technique is safe, and whether it provides demonstrable benefit.

A 2004 study by Veronesi et al. (2003) compared two groups of patients, one group who had SLNBs followed by a full ALND (the ALND group), and a second who had SLNBs followed by ALND only if the sentinel node showed metastases (the SLNB group). This study followed up with the patients for 46 months post-surgery and found no significant difference in the outcomes between the two groups. Further, it was found that the proportion of those in the ALND group who had metastatic axillary lymph nodes was almost the same as the proportion of those in the SLNB group whose SLNBs showed metastases. In the ALND group the SLNB had a false negative rate of around 8.8%. The SLNB group reported fewer post-surgical problems with pain, arm mobility, paresthesia and swelling.

A more recent review of the safety of avoiding an ALND on the basis of SLNB results has been conducted by Pepels et al. (2011). This review includes the above mentioned Veronesi study as well as a later follow up from the same group. Studies from the pre-SLNB era show that where patients did not have ALNDs, 12-18% of patients developed an axillary recurrence of cancer in 5-10 years and there was a 5% decrease in the overall survival rate. In studies where patients had negative SLNBs and no ALND, the recurrence rate was 0.6%. In studies such as the Veronesi study, where two groups existed, the recurrence rate for those with negative SLNBs was 0.3% at 47 months for those who had ALND and 0.4% at 31 months for those who did not. In studies where patients had a positive SLNB but no ALND the recurrence rate was 1.7% at 30 months. The review concludes that it is safe to avoid an ALND following a negative SLNB.

Another recent review and meta-analysis of existing reports covering has been published by van der Ploeg et al. (2008). This review covered nearly 15,000 patients with negative sentinel node biopsies who were followed up for a median of 34 months post surgery. The review reports a tumour recurrence rate of just 0.3% and a sensitivity of the biopsy as high as 100%. It is conceded that this sensitivity may be skewed by the study selection process, those papers that included the information required to contribute to the sensitivity happened to report few if any recurrences. It would seem highly improbable that the process is as foolproof in reality as this reported figure would indicate.

A further point of interest to note from this review is that of the 48 studies included only one is reported as not using a radioactive tracer, whereas six studies report not using blue-dye tracers (this detail is omitted from a further four studies). This suggests that the method of detecting the sentinel node is not universally agreed. Indeed, even the precise radiopharmaceutical employed varies between treatment centres (mainly between ^{99m}Tc nanocolloid and ^{99m}Tc -sulphur colloid). Whilst this

does lead to a significant difference in recurrence rates (0.14% for sulphur colloids, 0.53% for nanocolloids, $P < 0.01$) it is far from clear whether this is due to the tracer itself or the fact that a number of specialist centres use the sulphur colloid.

Goyal and Mansel (2008) have also recently published a review of the advances in sentinel node biopsy for breast cancer. This review touches on the choice of tracer used and concludes that a combination technique of both blue dye and radiopharmaceutical optimises the detection of the sentinel node. They further note that data from the Memorial Sloan-Kettering Cancer Center (a leading specialist centre) suggests that blue dye provides only a marginal benefit in detection with increasing surgical experience. This review also refers to the typical sizes of metastases detectable in lymph nodes, noting that an in-vitro analysis test called the GeneSearch BLN Assay (produced by Veridex) can detect metastases as small as $200\mu\text{m}$ with a turn-around time of around 30 minutes. This size of metastasis is considered by the US Food and Drug Administration (FDA) as clinically relevant and actionable (US Food and Drug Administration, 2006), however this same document accepts that there is a lack of evidence that metastases smaller than 2mm have a significant impact on long term prognosis.

A further review by Cheng et al. (2011) discusses some of the remaining controversies over the SLNB. These are technical and procedural debates rather than a fundamental assessment of safety and effectiveness, which now appear to be well studied and widely accepted. A further discussion of the choice of tracers and how best to detect these tracers is presented alongside questions such as whether more than just the sentinel node should be tested; these issues are tangential to this report however, and so a more full discussion is not rehearsed here.

5.1.2.2 Spectroscopy and Cancer

There are many studies in the literature into the feasibility of distinguishing abnormal from normal tissue by optical means. These studies cover a range of different cancers and even incorporate benign as well as malignant tumours. Some relevant and representative publications are discussed here.

Many of these studies are performed in the IR rather than the visible or the VNIR. Spectroscopy in the IR is dominated by molecular absorptions and vibrations, thus providing often very detailed information about the specific molecules, and sometimes the specific configurations of molecules, present in a substance; for this reason it is considered the fingerprint region for biological tissue. The units of spectroscopy are wavenumbers, measured in cm^{-1} (wavelengths per cm), rather than wavelengths and the typical IR region is between 500 and 4000cm^{-1} , or 2.5 to $20\mu\text{m}$. The units quoted in this review will be cm^{-1} for the IR region and nm for the VNIR.

A group at the Robert Koch Institute in Berlin investigated using IR microspectroscopy to examine breast tissue. Initially, the focus was on benign tumour sections Fabian et al. (2003); a spectral image was formed using a 64×64 detector array at an approximate spatial resolution of $4\mu\text{m}$. A comparison of this image with standard stained tissue examination was made, which showed that similar tissue structures were visible with both methods. A follow up study (Fabian et al., 2006) presented a comparison of benign and malignant tissue. This work showed that differences between benign and malignant tissue in the regions of the spectrum associated with lipid ($2800\text{-}3000\text{cm}^{-1}$) and collagen ($1000\text{-}1450\text{cm}^{-1}$) content.

A study by Rehman et al. (2007) presents Raman spectroscopic differences between normal tissue and tissue from two types of breast cancer, invasive ductal carcinomas (IDC) and ductal carcinomas *in situ* (DCIS). The study used archived tissue samples, which had previously been fixed in formalin and embedded in paraffin wax. The samples were de-waxed and cleaned in alcohol before examination. Differences were found between the three groups of tissues that were again associated with lipid and protein/collagen content.

In addition to finding differences between the types of tissue, this study also found differences between grades of tumour for both IDC and DCIS. These were in the same spectral regions as above and again indicate that concentrations of lipids and proteins change as tissue progresses from healthy to cancerous and then on to increasing malignancy.

Two notable studies of breast cancer in the VNIR range provide evidence of spectral differences, or optical properties that can be measured spectrally. Both are based on diffuse reflectance spectroscopy, with Grosenick et al. (2005) using a distributed time of flight (DTOF) measure to infer optical properties and Tromberg et al. (2005) taking direct spectral response measurements. Grosenick et al. (2005) conclude that total haemoglobin concentrations were higher in tumourous tissue than normal but blood oxygen saturation was not significantly different. Scatter power (the dependence of the reduced scattering co-efficient on wavelength) was also found to be higher in tumourous tissue, however only marginally and the 95% confidence intervals of normal and tumourous tissue overlapped.

Tromberg et al. (2005) calculated oxy- and deoxy-haemoglobin concentrations, lipid and water contents, tissue haemoglobin concentration, tissue oxygen saturation, scatter power, and tissue optical indices for normal and malignant tissue. With the exception of scatter power and oxygen saturation, these were measured to be significantly different in malignant tissue compared to normal. This study also found that tissue optical index (defined in the paper to be a parameter derived from haemoglobin, water and lipid concentrations) was much higher in cancerous tissue; for patients undergoing

adjuvant treatment such as chemotherapy, the optical index for the tumour decreased over the timespan of the treatment.

A group led by Christoph Krafft in Dresden has used IR spectroscopic imaging to study malignant gliomas (brain tumours). In one study on a single patient (Krafft et al., 2007), images were obtained using a spectroscope with a 64×64 detector array at a spatial resolutions of around $63\mu\text{m}$. Tissue from inside the tumour, at the tumour margins and outside the tumour were identified based on a supervised linear discriminant analysis (LDA) classifier. The spectral features used to train the classifier correspond to lipid, haemoglobin and protein contents. In a separate study using LDA (Krafft et al., 2006), images were generated by raster-scanning a microscope coupled to a spectroscope across the tissue with a spatial resolution of $180\mu\text{m}$. Again good separation between normal and malignant, and between grades of malignancy was achieved using a classifier trained with spectral features corresponding to protein to lipid ratios, and further features to represent haemorrhage.

Fujioka et al. (2004) performed Fourier Transform IR Spectroscopy (FTIR) on samples of tumours resected from gastric cancer patients. Differences between normal and cancerous tissue spectra were determined and then used to train an LDA classifier, which was used to discriminate between 23 cancerous and 12 normal samples. All but one of the cancerous samples and 9 of the 12 normal samples were correctly classified. The 10 spectral features that were used to train the classifier correspond to absorption bands in protein, collagen, phosphates, and some amino and nucleic acids.

Gastric cancers were also the focus of an IR spectroscopy study by Park et al. (2007b). Spectra from samples of cancerous, normal and adenoma (a benign tumour) tissues were compared using Principal Components Analysis (PCA). Score plots of the first and second principal components were generated and whilst these showed fairly good discrimination between cancerous and normal, there was significant overlap between both adenoma and normal, and adenoma and cancer tissues. This study presented a “whole spectrum” approach, i.e. the sources of the differences were not identified, only their existence.

FTIR was used by Andrus and Strickland (1998) in their study of non-Hodgkin's Lymphomas. It was found that the ratio of the shoulders of the nucleic acid peak (1084cm^{-1}) at 1121cm^{-1} and 1020cm^{-1} was higher in cancerous tissue than in normal tissue and that it increased with increasing clinical grade. The 1121cm^{-1} shoulder is characteristic of absorption by RNA whereas the 1020cm^{-1} is characteristic of DNA. This suggests that the ratio of RNA to DNA increases with increasing lymphoma grade.

A second ratio, between the peaks at 1240cm^{-1} and 1084cm^{-1} is also observed to increase with increasing grade; this is a measure of the ratio of collagen to nucleic acid.

It is noted that this trend is opposite of what is seen in a number of other cancers, although it is the same as has been observed in breast cancer by Jackson et al. (1997).

There are many similar studies for skin cancer, as would be expected for a common cancer on such an accessible organ, some of which are summarised in table 5.2 below.

Reference	Tumour/Lesion Type	Method	Molecules Attributed
Wong et al. (1993)	BCC	IR Spectroscopy	Phosphates ($1190\text{-}1270\text{cm}^{-1}$) C-O ($1140\text{-}1185\text{cm}^{-1}$) C-H ($2800\text{-}3050\text{cm}^{-1}$)
McIntosh et al. (1999)	BCC, SCC, Melanocytic lesions	IR Spectroscopy	Lipids ($2780\text{-}3020\text{cm}^{-1}$) Collagen ($900\text{-}1500\text{cm}^{-1}$)
McIntosh et al. (2001)	Various lesions	VNIR Spectroscopy	Haemoglobin (618-698nm) Water (918-998nm)
Gniadecka et al. (2004)	BCC	Raman Spectroscopy	Amides ($1500\text{-}1800\text{cm}^{-1}$, 1270cm^{-1}) Proteins (940cm^{-1}) Lipids (1450cm^{-1})
Garcia-Uribe et al. (2004)	Various lesions	SOIR	Haemoglobin

TABLE 5.2: Key: BCC - Basal Cell Carcinoma, SCC - Squamous Cell Carcinoma, SOIR - Spectroscopic Oblique Incidence Reflectometry.

A group in Michigan led by Mary-Ann Mycek is investigating the use of VNIR spectroscopy for detection of pancreatic cancers. The methodology used is discussed by Chandra et al. (2009), which notes that both reflectance and autofluorescence (see section 5.1.2.3 below) spectra are measured using a fibre-optic probe assembly. One paper from this group is of particular interest to this project, Wilson et al. (2011). This describes the use of a Photon Tissue Interaction (PTI) model developed by the group (developed in Wilson et al. (2009) and Wilson et al. (2010)) to fit reflectance spectra to parameters related to tissue morphology and biochemistry. One such parameter is referred to as the “nuclear enlargement parameter”, which relates to the nuclear size of the cells involved. It was found that by increasing this parameter systematically, the reflectance in the 450nm-500nm region of the spectrum also increased systematically, specifically, a peak around 475nm starts to emerge. In their measurements, the spectra generated with larger nuclear size corresponds well to spectra from adenocarcinoma, whereas the small nuclear size corresponds to normal tissue. A similar difference in this spectral region is seen between metastatic and normal lymph node tissue by Bigio et al. (2000) (see discussion in section 5.1.2.3 below).

Water content is mentioned by Tromberg et al. (2005) and is also the likely source of the spectral differences arising in the 918-998nm band in McIntosh et al. (2001). The use of water as a distinguishing feature is further supported by a paper from Gniadecka et al. (2003), although it is suggested here that the structure of the water in carcinoma cells is the source of the difference.

Gniadecka et al. (2003) studied absorption at a band around 180cm^{-1} ($\sim 55\mu\text{m}/5\text{THz}$) representing what is described as a tetrahedral arrangement of water molecules. In this arrangement water molecules are bound to each other and not to other macro- (or bio-) molecules. They found that photo-aged skin and malignant tumours exhibited an increased amount of this tetrahedral water. They postulate that this could be due to alterations in protein structure and decreased water-protein interactions.

The use of water gains further credence from studies, particularly in skin cancers, in the terahertz (THz or T-ray) region of the electromagnetic spectrum. This region is dominated by water absorption bands and when studying biological systems this is likely to be the most significant factor. Teraview Ltd, a commercial operation specialising in THz applications, together with groups at Cambridge University and Addenbrooke Hospital, led by Woodward are the most significant contributors to this field. They have published a number of studies (see Woodward et al. (2002), Woodward et al. (2003) and Wallace et al. (2004)) which claim some success in using Terahertz Pulsed Imaging (TPI) to differentiate normal tissue from diseased tissue. This technology is now being commercially developed by Teraview.

Although the above represent a range of different cancers, it is clear that similar diagnostic factors are present in most cases. The evidence in the literature would

suggest that protein and lipid concentrations are significantly different in normal tissue than they are cancerous and that haemoglobin levels are higher in cancerous tissue than normal.

The discrepancy in haemoglobin concentration is readily explained by the standard description of tumour growth. As the tumour starts to develop it requires the creation of new blood vessels to keep it supplied; this process is known as angiogenesis and can lead to hypervasculature around a tumour site. The fact that oxygen saturation does not seem to be commensurately greater at tumour sites may be explained by parts of the tumour becoming hypoxic. Parts of the tumour can still be inadequately supplied with oxygen as the tumour growth outstrips the increase in blood supply.

5.1.2.3 Spectroscopy and the Lymph Node

Using spectroscopic techniques to detect cancer in lymph nodes is by its nature, a niche subject. There are three main groups however who are working or have worked in this area; the first is a group at l’Istituto Nazionale per la Fisica della Materia (INFM) in Florence, the second is led by Max Diem at Northeastern University, Boston, and the final group is jointly led by Irving Bigio at Boston University and Stephen Bown at University College, London.

The Florentine group have two major publications in this area based on multispectral imaging of autofluorescence in lymph nodes. Fluorescence is a phenomenon occurring in or close to the visible spectrum that is analogous to Raman scattering as described in section 2.1, however in the case of fluorescence the molecule is promoted to a real rather than a virtual energy state; light at one frequency is absorbed by the molecule and light of a different frequency is emitted. This process is often exploited in biomedical imaging by injecting a fluorophore that is known to bind certain types of biological molecule and then imaging the fluorescence as a proxy measure for the biological molecule. Autofluorescence implies that no exogenous fluorophore is used but instead that a molecule naturally present within the tissue is the cause of the fluorescence. Again, like Raman scattering, this is a low intensity event, although typically more intense than Raman scattering.

The first paper (Rigacci et al., 2000) presents findings of a study in which frozen sections of lymph nodes from patients suffering from adenopathy (swelling of the lymph nodes) were illuminated with 365nm radiation in order to excite autofluorescence from extracellular collagen and elastin fibres. Images were taken behind 50nm bandwidth filters at 450nm, 550nm and 658nm, and were combined into an RGB image. Most of the autofluorescence signal was noted to occur in the 450nm channel. Changes in the morphology were observed between hyperplastic nodes and those with Hodgkins lymphoma. Further, autofluorescence spectra were measured

using a spectrometer. These spectra showed that the fluorescence peak from lymphoma tissue was shifted slightly to the red when compared with the peak from hyperplastic tissue, from 433nm to 440nm. A difference spectrum of the two signals showed a maximum at 408nm and a minimum at 460nm.

A follow up paper (Pantalone et al., 2007) presents a study using the same methods to examine lymph node sections from gastric and colo-rectal cancer patients. The red filter was 650nm for this study rather than 658nm as for the previous. Here it is found that the autofluorescence peak from metastatic lymph nodes was shifted by more than 20nm to the red from the peak from the healthy lymph nodes (459nm from 437nm). The metastatic peak was much broader than the normal and in some cases a shoulder or second peak at 602nm was observed. Again, it was found that significant morphological changes could be seen between the healthy nodes and the metastatic. It was also noted that the metastatic multispectral images contained a greater contribution from the red channel, which is consistent with the broadening of the peak and potential contribution from a second peak around 600nm, that was measured. The red-shifting of the peak is not explicitly explained in the paper, the discussion talks more generally about the differences in the autofluorescence peaks being due to differences in the concentrations and distributions of endogenous fluorophores between normal and metastatic tissue.

Diem's group have published a number of papers on IR spectroscopy of individual cells, and more recently have moved on to looking at infrared spectral imaging. It is these spectral imaging papers that shall be discussed in more detail here.

In a paper by Romeo and Diem (2005), tissue sections sliced from paraffin fixed samples were de-paraffinized and imaged by scanning a spectrometer with a 16 detector array across the tissue. The images had a spatial resolution of $25\mu\text{m}$ and a spectral resolution of 4cm^{-1} in the range from 800 to 1800cm^{-1} . All pixels in an image were classified using an unsupervised hierarchical cluster analysis (HCA) algorithm. The resulting images were compared to standard H&E stained slides. It was found that restricting the spectral range to the low frequency range (below $\sim 1500\text{cm}^{-1}$) gave better discrimination between different cell types, and the morphology observable in the H&E stained images could be replicated in the spectral images. The restriction of the spectral range prevented a shift in the amide band at $\sim 1650\text{cm}^{-1}$ from dominating the classification algorithm; this shift was attributed to "reflective components in the spectra" rather than chemical components of the cells.

A further paper by Bird et al. (2008) presents a study where IR spectra were obtained from frozen or de-paraffinized sections of routinely excised lymph nodes. The same spectrometer was used and again it was scanned across the tissue in order to generate a spectral image. An additional spectral region between 2800 and 3100cm^{-1} was added in order to improve classification. The images were classified using an unsupervised

HCA algorithm and also by an advanced neural network (ANN) classifier. Again here the HCA processed images showed information comparable to that obtained by H&E staining; the ANN classified images, which produced just two classes, normal tissue and cancerous, also performed well and correlated with the findings from traditional histology.

A similar methodology is employed in a later paper (Bird et al., 2009). Here the detector array is 8×2 as opposed to 16×1 from the previous papers, but again sections of de-paraffinized lymph node tissue is scanned to produce an image. As in the first paper, the spectral range here is from $800\text{-}1800\text{cm}^{-1}$, and again unsupervised HCA analysis is performed. Again, good agreement was found with the histological images, and metastatic areas of the tissue were successfully classified. The major advance in this paper is the detection of “micrometastases”, collections of cancerous cells measuring $100\text{-}150\mu\text{m}$. This paper also includes some discussion on the type of material responsible for responses in given spectral areas. The $900\text{-}1350\text{cm}^{-1}$ range (similar to the reduced region from the first paper) is described as the “phosphate region” and is identified as being indicative of DNA and RNA content. The $1450\text{-}1700\text{cm}^{-1}$ region is described as the “protein region” and changes here are attributed to changes in protein composition (both quantity changes and structural changes).

The final group to be discussed in this section is perhaps the most relevant as their work concerns reflectance spectroscopy of axillary lymph nodes in the visible spectrum. An early paper (Bigio et al., 2000) describes their “elastic scattering spectroscopy” (ESS) method. A broadband light source (pulsed xenon arc lamp) is directed to the tissue via fibre optic cable, an adjacent fibre returns the reflected signal to a spectrometer which records a spectrum in the range 330nm-750nm.

This paper describes not only measurements taken from the resected lymph nodes, but also from the breast tumours themselves. Much of the discussion concerns the tumour measurements rather than the lymph nodes, however representative spectra for normal and metastatic sentinel nodes are shown. This shows that overall the metastatic tissue reflects more light than normal tissue below $\sim 520\text{nm}$ but less above this. The other significant difference in the spectra is the presence of a large shoulder at $\sim 470\text{nm}$ in the metastatic tissue that is not apparent in the normal tissue. Unsupervised HCA and ANN analysis were performed whereby metastatic sentinel nodes were detected with a sensitivity of 91% for HCA and 58% for ANN, and specificities of 76.5% and 93% respectively.

A follow up paper by Johnson et al. (2004) focusses on the sentinel lymph node; the spectra in this case were reduced by PCA before being subject to an LDA classifier. Two separate analyses were carried out, one compared only normal nodes and those that had become completely metastatic (referred to a “per-spectrum” analysis), and

the second added nodes which showed only partial metastasis (“per-node” analysis). The per-spectrum analysis classified metastatic nodes with a sensitivity of 84% and a specificity of 91%, whereas the per-node analysis figures were 75% and 89% for sensitivity and specificity respectively.

A more recent paper describes movement towards making this an imaging system (Keshtgar et al., 2010). The probe is scanned across the tissue at 0.5mm intervals, over an area of 1cm², creating a 20×20 pixel spectral image. The spectra were analysed in the same manner as before, the images are then colour-coded using the scores from the LDA. A single image is presented in the report, which shows an area of metastases, however no other image of the tissue is given for comparison. Within the subset of metastatic nodes, macrometastases (>2mm) were detected with a sensitivity of 76% and micrometastases (0.2-2mm) with a sensitivity of 29%. There were three nodes with submicrometastases, but none of these were detected. In the subset of nodes that also underwent TP cytology, ESS managed a sensitivity of only 69% as compared to 88% for macrometastases; however, it detected micrometastases with a sensitivity of 40% where TP failed to detect any. Two false positives were recorded with ESS, whereas TP recorded none.

Another paper from the same year presents more images on the basis of the work above (Austwick et al., 2010). Here, the colour-coded spectral images are shown next to photographic images of the tissue, allowing for some comparison. Here, the metastases detected do appear to match up with the histology, however the metastases are not visually obvious on the tissue and so it is still not certain that the classification is accurate spatially.

The work reviewed in this section provides encouraging evidence that spectral differences do exist between normal lymph node tissue and metastatic deposits. In particular, the work from Bigio’s group suggests that even in the visible region significant differences exist, which may be exploitable by imaging based rather than spectroscopy based spectral imagers.

5.2 System Development

Any spectral imaging system consists of two components, which can be loosely classified as hardware and software; the hardware component is the image capture system, which by necessity involves some software to control image acquisition, the software component deals with the processing of the data. Each component is detailed in turn in this section, outlining the specifications of the hardware and the algorithms used to derive information from the data.

5.2.1 Image Capture Systems

Two imagers have been developed for this project. The first is a multispectral imager working on the “Focal Plane Array” modality as discussed in section 1.4.2; the second is a hyperspectral imager working on the “Line Scanning” modality (also discussed in section 1.4.2).

The imaging systems discussed in this section were both initially commissioned by Dr. Hoy for his own doctoral work, and the details of their first stage of development are detailed in his thesis (Hoy, 2009).

The main specification requirements to which he was working can be summarised in four points:

1. the imaging system should be non-contact so as not to disturb or damage biological tissue;
2. the system should provide real-time imaging of the surgical field at a resolution equal to or better than the precision of the surgeons cut (estimated to be 0.25mm by Dr. Hoy’s surgical collaborator);
3. the system should have minimal impact on the surgical environment; and
4. the system should be low cost and have low infrastructure overheads.

Of these, items 1 and 4 are very much true of imaging system for this project. As the resected nodes are subject to full histopathological examination after the surgery, it is important that any biopsy technique is non-destructive. Further, in order to realise the aim of making this technology viably available at all treatment centres, it must be low cost, and not have onerous requirements for the installation of specialised equipment.

With regards to item 2, Dr. Hoy’s work was aiming to provide true *in-vivo* intra-operative imaging, which is why a real-time display was considered necessary. In this project the aim is to provide a rapid biopsy; as a result, whilst it is important for image acquisition and processing times to be minimised, it is not necessary to provide “real-time” imaging. If the entire process could be performed within 10-15 minutes, this would still provide a tangible benefit over traditional biopsy or spectroscopy.

Resolution for Dr. Hoy’s system was given an upper bound of the surgical precision on the basis that the aim of the project was to provide accurate tumour demarcation for excision. A neurosurgeon will know approximately where the tumour lies in the brain, even if it is not visually distinct from the healthy tissue around it; it is the precise margins that are unknown. The upper bound for this work would sensibly be set such that the system be capable of detecting clusters of metastasis that are clinically

relevant and on the same scale as can be detected visually by a pathologist. By coincidence, the 0.25mm resolution Dr. Hoy was working to is also a reasonable starting point for this work. As was demonstrated by Johnson et al. (2004), TP cytology is capable of detecting metastases greater than 0.2mm with good reliability, but is not reliably able to detect micrometastases. As noted above, the FDA consider metastases of 0.2mm to be clinically relevant and actionable (US Food and Drug Administration, 2006).

Item 3 is again a product of the intra-operative nature of the equipment being developed. In order to be useful, the imaging system had to be in theatre with a clear line of sight to the surgical field. This project is examining resected tissue and as such can be located outside of the theatre.

Aside from being low cost and having minimal installation overheads, in order to make this technology appealing, it should be capable of being operated by a non-specialist operator. This means that the process should be as automated as possible and where user intervention is required, it should be for routine tasks, such as “inputting” the node.

From these considerations, a new specification requirement list can be made:

1. the imaging system should be non-contact;
2. the imaging system should have a resolution not worse than $\sim 0.25\text{mm}$;
3. the imaging system should produce results within a time frame of minutes;
4. the imaging system should be operable by a minimally trained, non-specialist operator; and
5. the imaging system should be low cost and capable of being installed in the near vicinity of the operating theatre.

5.2.1.1 Multispectral Imager

A summary of the as-built specifications of Dr. Hoy’s mark I multispectral device is given in table 5.3 below.

In addition to the camera, two filter wheels were placed in front of the lens. The first contained a clear window and two linear polarising filters set orthogonally to each other; the second contained one neutral density filter transmitting $\sim 1\%$ of incident radiation, and five 10nm bandwidth (FWHM) interference filters centred at 470nm, 580nm, 610nm, 900nm and 970nm. The transmission of the neutral density filter is similar to the total transmission of each of the interference filters when integrated

Camera model	Apogee U260
Sensor	1 inch format Kodak KAF-0261E CCD
Resolution	512 × 512 pixels ($\sim 0.25\text{MP}$)
Pixel size	$20 \times 20\mu\text{m}$
Dynamic range	16 bits (65,536 levels)
Lens	Pentax TV lens, F1.4
Lens focal length	75mm
Field of view	$\sim 10 \times 10\text{cm}$
Minimum working distance	$\sim 50\text{cm}$

TABLE 5.3: As-built specifications of the Mark I imaging system

across the spectral range of the CCD, thus allowing a “clear” image to be exposed with similar integration time to the filtered images.

It soon became apparent that this set-up had been optimised for the intra-operative brain tumour project and that this was not the ideal configuration for the sentinel node work. A particular concern was the large field of view. Lymph nodes vary greatly in size, however it is rare to find one much bigger than 2cm (4cm total cut “diameter” after it has been butterflied) and typically they are around 0.5-1.5cm. This means that they have a total cut surface area in the order of 10cm^2 , which is just 10% of the total field of view of the mark I imager. A second iteration was thus designed with the priority of increasing the percentage of the total field of view that would be imaging tissue.

The key issue in this effort to decrease the size of the field of view was the lens that was used. In an initial experiment, a 10mm spacer was added between the lens and the camera body. This increases the distance between the lens and the focal plane thus decreases the minimum working distance (see equation 5.1 below, where u is the distance from the lens to the focal plane, v is the distance from the lens to the object plane and f is the constant focal length of the lens) and increases the magnification (see equation 5.2 where M is the magnification); this comes at the cost of reducing the ability of the lens to focus at long distances (specifically, it will be unable to focus to infinity). A further problem that can be encountered when using spacers is “vignetting”, this means that a dark ring forms around the edge of the image. This effect initially becomes apparent at the corners of the image.

$$\frac{1}{f} = \frac{1}{u} + \frac{1}{v}. \quad (5.1)$$

$$|M| = \frac{u}{v}. \quad (5.2)$$

In this application, the working distance was small, and so the long range focussing did not present a problem, and with a 10mm spacing the vignetting was not occurring in the area covered by the detector. However, whilst the field of view was improved, the improvement was not substantial, decreasing from $\sim 10 \times 10\text{cm}$, to $\sim 8 \times 8\text{cm}$.

A new zoom lens from Tamron with adjustable focal length between 70mm and 300mm was thus incorporated. This was specified to have a minimum working distance of 1.5m in normal mode and 0.95m in macro mode. The lens was used in macro mode, and with the addition of a conversion mounting (the Alta camera accepts only c-mount lenses, the lens was designed for mounting to Nikon cameras) effectively acting as a spacer ring, the actual minimum working distance was in the region of 0.7m, with a minimum field of view of $\sim 2 \times 2\text{cm}$ on the Alta camera's $10.24 \times 10.24\text{mm}$ image plane.

With a 2cm spatial field and 512 pixels, the minimum theoretical resolution of the system is around $40\mu\text{m}$. A standard calibration is to image line pairs at increasing spatial “densities”, and thus decreasing feature size, in order to determine the real resolution of the camera. Doing this, line pairs at 50cm^{-1} were readily separable, at higher densities they were not. This equates to a minimum visible feature size of $\sim 100\mu\text{m}$ ($\sim 0.1\text{mm}$). This brings the specification of the multispectral imager in line with the requirements set out above.

Table 5.4 summarises the key specifications.

Camera model	Apogee U260
Sensor	1 inch format Kodak KAF-0261E CCD
Resolution	512×512 pixels ($\sim 0.25\text{MP}$)
Pixel size	$20 \times 20\mu\text{m}$
Dynamic range	16 bits (65,536 levels)
Lens	Tamron AF70-300mm, F4-5.6, Macro
Lens focal length	70-300mm
Minimum Field of view	$\sim 2 \times 2\text{cm}$
Minimum working distance	$\sim 690\text{mm}$

TABLE 5.4: As-built specifications of the Mark II imaging system

A description of the full schematic set-up of this “mark II” system is given in section 5.2.1.3 below, following a discussion of the hyperspectral imaging system.

5.2.1.2 Hyperspectral Imager

As mentioned above, the hyperspectral imager was also initially purchased for Dr. Hoy's doctoral research project; however no work was actually undertaken with it as it failed to perform as specified. The initial specifications are given in table 5.5 for reference.

Model	Specim V10E Spectral Camera
Sensor	CMOS
Resolution	1280×1024 pixels ($\sim 1.3\text{MP}$)
Dynamic Range	12 bits (4,096 levels)
Spectral range	400-1000nm
Spectral resolution	2.8nm

TABLE 5.5: Initial specification of Specim hyperspectral imager

The Specim V10E is a diffraction grating based spectral imager. Light from a single horizontal line across the image scene passes through a $30\mu\text{m}$ slit, from where the grating disperses the light spectrally across the vertical dimension of the sensor. In this way, each image column represents the spectrum from a single point in the spatial scene, and each image row represents the spatial variation of reflected light at a single wavelength. As described in section 1.4.2, this “line-scanning” technique requires the line to be swept across the entire image scene in order to get a full spatial measurement.

Since the light reaching the detector is spread spectrally across the entire sensor height, the total intensity of light reaching any given pixel is low, and thus long integration times and a low noise sensor is required to capture useful images. The supplied CMOS proved to have too low a signal to noise ratio (SNR) to provide usable images. A more appropriate CCD imager was thus acquired as a replacement.

The replacement imager is a Prosilica GC1380 camera, using a CCD sensor. This has a resolution of 1360×1024 pixels; this high vertical resolution means that the CCD is actually over-sampling the spectral components. The performance of this camera far exceeded that of the supplied CMOS imager and it was able to generate usable spectral images.

With a spectral range of 600nm and a resolution of 2.8nm, the Specim produces only ~ 200 resolvable spectral elements, thus the CCD is taking around 5 readings for each element. The vertical dimension was binned such that each image pixel measured the output of two consecutive sensor pixels, reducing the amount of over-sampling. It was decided not to correct this entirely at the acquisition stage, this allowed the image to be processed by smoothing the spectra in software without compromising spectral resolution. This binning effectively doubles the collection area of the sensor pixel thus reducing the exposure time required for the same overall intensity of light.

The dynamic range of this imager was expanded by taking a number of exposures and averaging them. The signal is proportional to the number of exposures, N whereas the random noise is proportional to \sqrt{N} . Thus N additional exposures produce a \sqrt{N} improvement in SNR. It was decided to average 16 images, producing a 4 fold increase

in SNR, equating to 2 bits extra dynamic range. To gain a 3rd extra bit would require averaging 64 images, which represents a significant time penalty.

In order to allow scanning, the sample was mounted on a translatable stage. At this point in development, the stage is manually controlled, with each rotation of the control representing 1mm of travel. Half-rotations can be reliably controlled, giving a spatial resolution of 0.5mm in the scan direction. The field of view in the spatial dimension of the imager itself is around 40mm, this gives a spatial resolution of $\sim 30\mu\text{m}$, at a working distance of 340mm.

The relevant specifications for the second iteration of this imager are set out in table 5.6 below.

Model	Specim V10E Spectral Imager
Camera Model	Prosilica GC1380
Sensor	Sony Ex-View ICX285AL CCD
Native Resolution	1360 \times 1024 pixels ($\sim 1.3\text{MP}$)
Pixel Size	$6.5 \times 6.5\mu\text{m}$
Operating Resolution	1360 \times 512 pixels (spatial \times spectral)
Native Dynamic Range	12 bits (4,096 levels)
Operating Dynamic Range	14 bits (16,384 levels)
Spectral range	400-1000nm
Spectral resolution	2.8nm
Spatial resolution	$30 \times 500\mu\text{m}$
Field of View	$\sim 40\text{mm}$ (single dimension)
Working distance	340mm

TABLE 5.6: Final specification of Specim hyperspectral imager

5.2.1.3 System Set-Up

In order to make the system a viable prospect for Dr. Hoy's application of imaging of brain tissue during surgery it was necessary to make it compact and portable. For this reason the cameras, filter wheels and control electronics were all mounted in a metal caddy, which itself was mounted to a tripod so that it could be moved easily around the operating theatre in order to get a clear line of sight to the surgical field. A photograph of this is shown below in figure 5.3.

A schematic diagram detailing the layout of this is shown below in figure 5.4.

This layout was based on the original Pentax lens working distance and the requirement to minimise the overall footprint of the device. By switching to the Tamron lens, the requirements for the optical path lengths from the target scene to the imagers was altered. Further, the change in the nature of the work meant that the whole device could be made into a bench top rather than a portable application, and

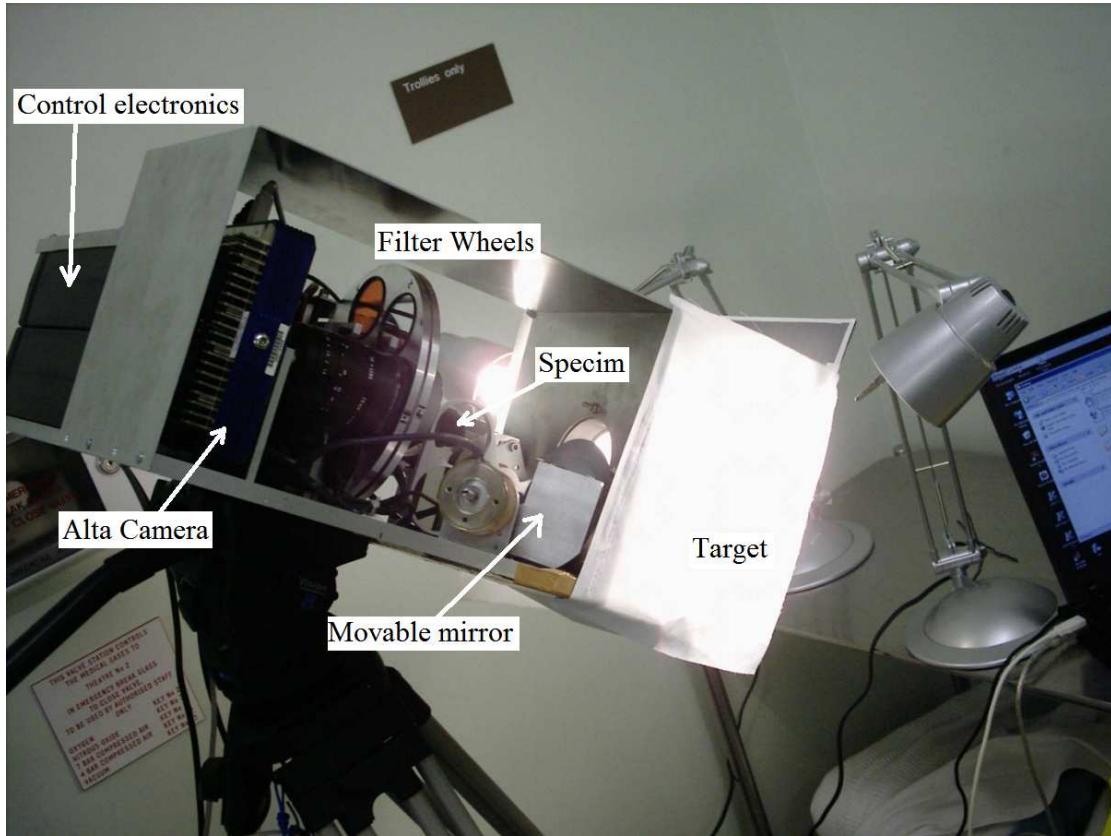


FIGURE 5.3: A photograph of the mark I spectral imaging system.

the overall footprint was less critical. The position of the tissue also became much more controllable, and it was possible to incorporate into the system a fixed translation stage upon which the tissue could be placed. The other significant change was to replace the mirror from the mark I device with a beam-splitter. The original mirror was heavy and was moved into place when required by a large solenoid; replacing it with a static beam-splitter reduced the number of moving parts and removed a high power consuming piece of electronics.

A schematic diagram showing the revised layout of the device is shown below in figure 5.5. In order to achieve the required optical path lengths whilst still keeping the overall footprint of the device small, a mirror is used to raise the optical level to a height of $\sim 200\text{mm}$; this can be seen in the photograph of the layout in figure 5.6. The sample area was illuminated from either side with light being diffusing through spectrally neutral diffusers.

Software was also developed (in C and C++) in order to control image acquisition and storage. Much of the work on this was in place as a result of Dr. Hoy's work, however significant additions were required in order for the Prosilica camera to be controlled. Software size is generally measured in source lines of code (SLOC), for comparison, in C/C++ the most basic "Hello World!" program would require around 10 SLOC.

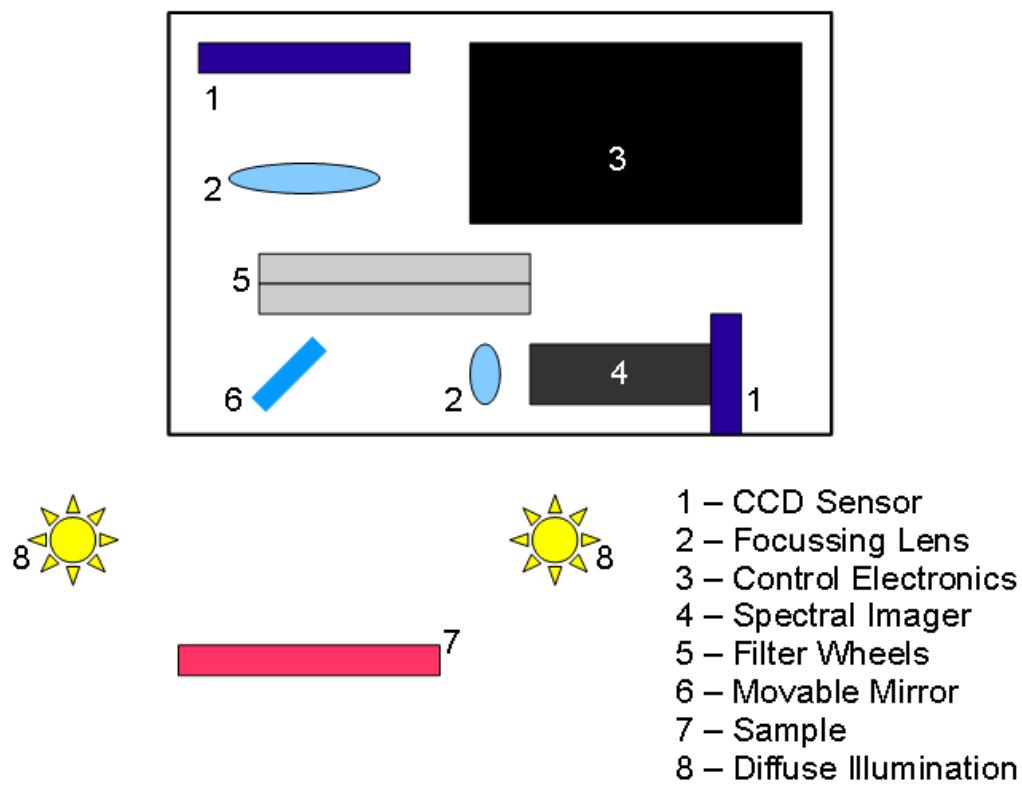


FIGURE 5.4: A schematic diagram of the mark I spectral imaging system. All equipment within the black outline is mounted to a metal base. The sample stage is mounted orthogonally to the plane of the base, because the whole system is mounted on a TV camera dolly the base can be angled down so that the sample is not vertical to the ground.

The main control program, which is a procedural based piece of code, is around 2250 SLOC, of which around 500 were a direct contribution from the author; the remaining code was the existing base developed by Dr. Hoy, although it was heavily edited and re-ordered by the author. This main program also relied on a number of libraries, including an extensive image processing library, written by Dr. Hoy ($\sim 12,000$ SLOC). A flow chart showing the main tasks of this program is shown in figure 5.7.

The Prosilica camera was controlled by use of a library building on the manufacturer issued SDK (Software Development Kit). SDKs provide descriptors and functions for accessing the low-level functionality of a piece of hardware, allowing developers to create high-level abstractions and functions. The library was developed in an object-oriented paradigm and consists of around 3500 SLOC, of which the author directly contributed around 1000. A separate program for testing new library functions and controlling the camera independently of the main control program was written by the author, running to around 1100 SLOC.

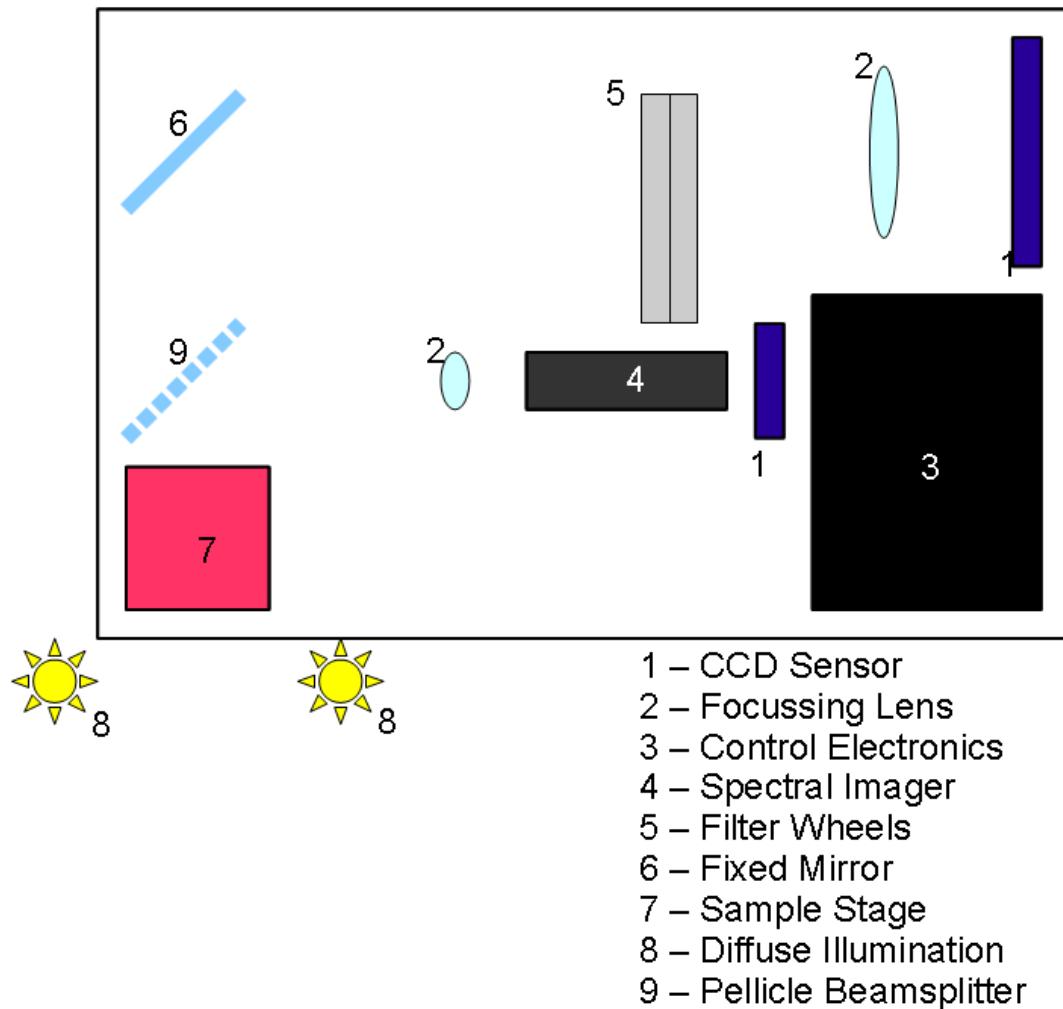


FIGURE 5.5: A schematic diagram of the mark II spectral imaging system. All equipment within the black outline is mounted to a metal base. The sample stage is mounted in the plane of the base and a mirror is used to raise the optical height and achieve the required working distance.

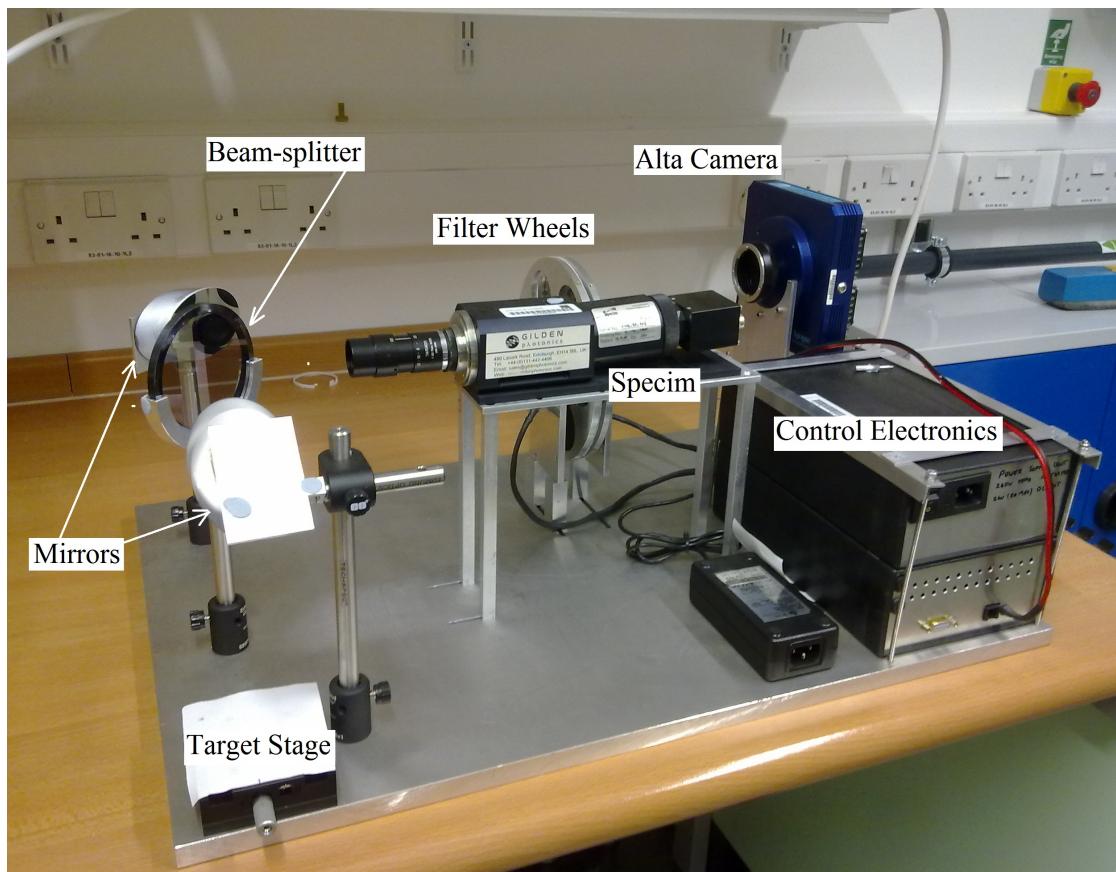


FIGURE 5.6: A photograph of the mark II spectral imaging system.

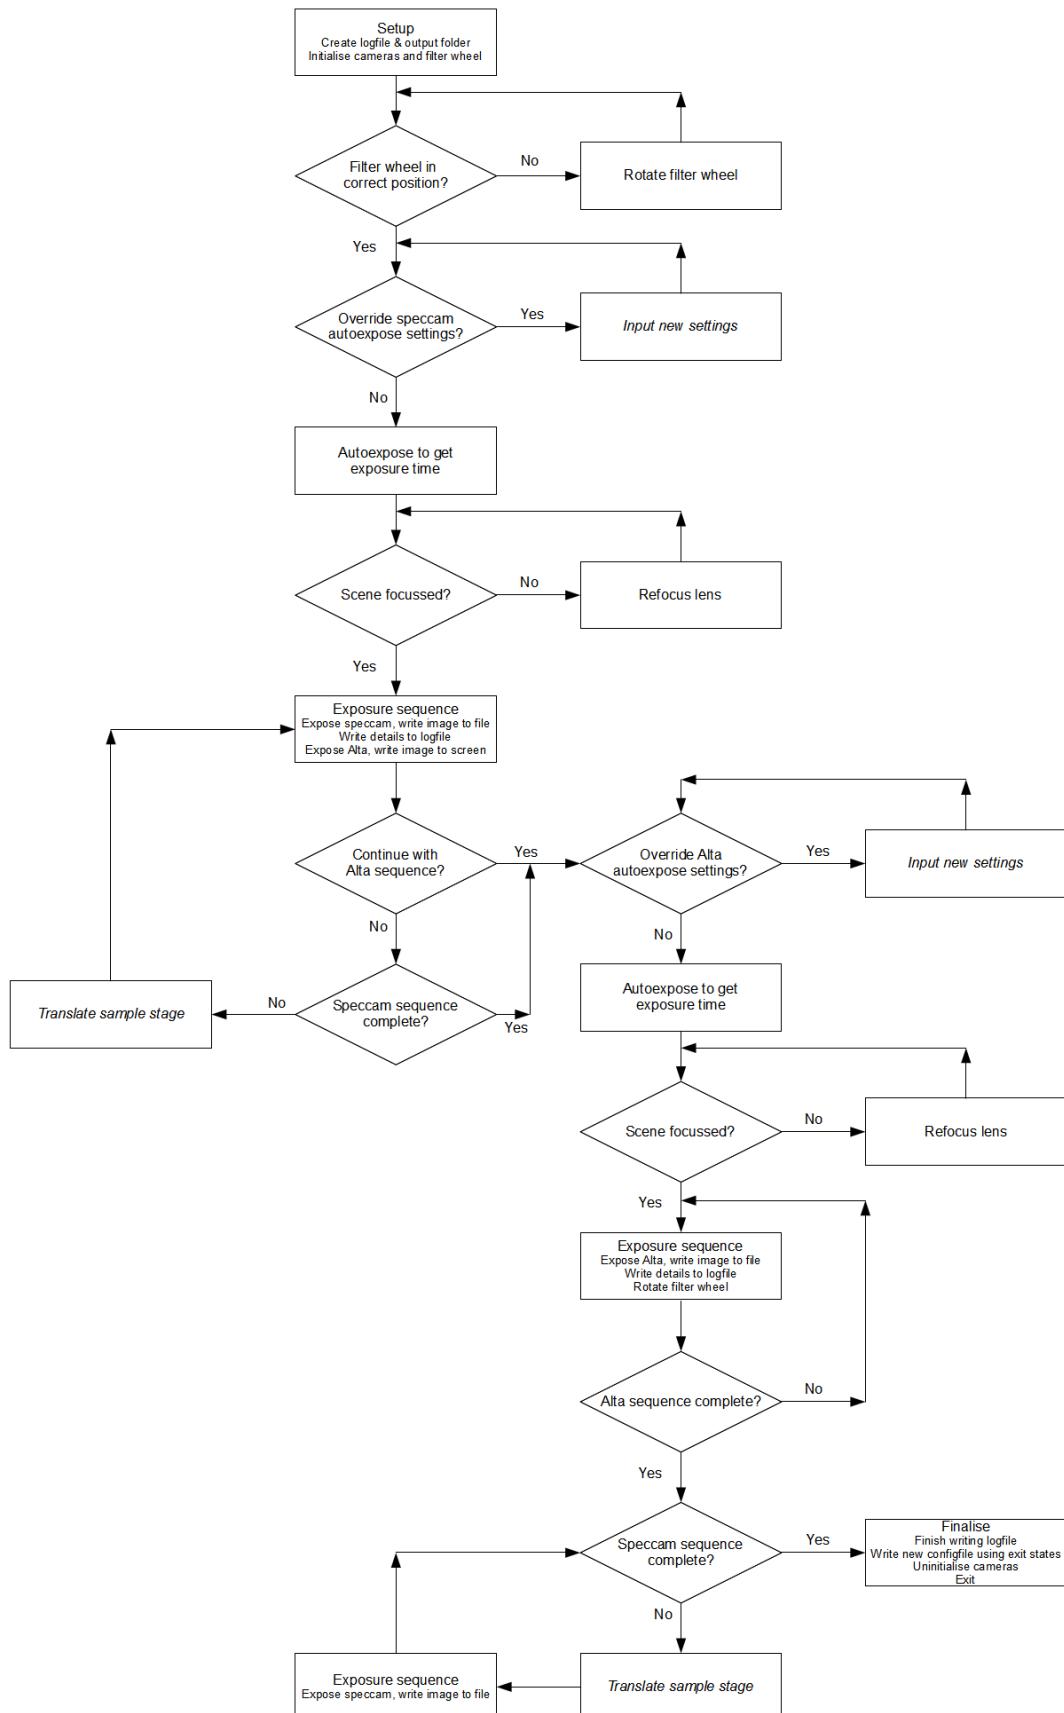


FIGURE 5.7: A flow chart showing the operation of the image acquisition system.

5.2.2 Data Processing Suite

All of the spectral images in this project were captured in single plane (monochrome) images stored in a proprietary High Dynamic Range (HDR) image format. A large image processing library (around 8750 SLOC) was developed by Dr. Hoy to manipulate these images, however, the real power of spectral imaging lies in its spectral nature. In order to more fully exploit this, processing suites that treat the images as either true multispectral or true hyperspectral images have been developed. There is much overlap in how these two classes of image should be treated and so the multi- and hyper- spectral projects are both written as offshoots of a base spectral image suite. Nearly 20,000 SLOC have been written for these suites.

In this section some of the key routines and algorithms will be discussed, separated into “pre-processing” and “real” processing. Pre-processing is a general description to cover how the data is taken from its native HDR format and converted into a spectral image that is ready for analysis. Real processing is a general description of those routines that analyse and interrogate the data to extract useful information.

5.2.2.1 Pre-Processing Routines

Having acquired a number of single plane images, the primary task is to convert them into a single spectral image. For multispectral images the key steps are ensuring that the images are correctly registered (i.e. that the images align and overlap properly) and then combining them into a single image.

For hyperspectral images this step requires, first “hotspot” filtering the images (explained below) and then combining and rearranging them such that the primary projection of the image is spatial-spatial rather than spatial-spectral. None of the images taken for the hyperspectral imager overlap spatially and as such registration is not an issue.

The Prosilica camera has a number of known “hot” pixels, i.e. pixels which consistently have high values in comparison to their neighbours. These were isolated by taking a dark exposure and searching for pixels with a raw value greater than 410 (10% of the maximum pixel value of 4,095), creating a hotspot map. It should be noted that only 10 such pixels were identified on the CCD, and whilst they are prominently visible, they represent a very low error rate on the detector (approximately 1 in 10^6). These pixels represent a characteristic form of noise known as “salt and pepper” noise, described by Nixon and Aguado (2002) as *“isolated black and white points... within an image”*. The most appropriate way to deal with such noise is to apply a median filter.

In image processing, a median filter is typically applied across an entire image. Each pixel in turn is considered, and a list of it and its neighbouring pixels (most often the 8

so-called “nearest neighbours”, which along with the subject pixel form a 3×3 grid centred on the subject pixel) is constructed and then ordered. The subject pixel value is then replaced with the median value from the ordered list; this approach is ideal for removing very light and very dark spots. Some care has to be taken with this filter however as real information, in particular features smaller than the filter size (3×3 for the nearest neighbour case) can be filtered out.

Features at the single pixel size are not likely to be significant in this project and so the filter can be safely used. In this application the filter is only applied to pixels at known hotspots as salt and pepper noise has not been observed more generally. Hotspots have not been observed on the Alta CCD. This filter is an irreversible process and once modified the original data can not be recovered.

Once the images have been arranged into a single spectral image datacube, they are written to a new file type. The HDR image format allows for an arbitrary number of planes in a single image and as such is extendible to spectral imaging, however, there is no ability within the format to increase the metadata available. Thus it is not possible to store derived data, such as a spectral angle map, or information such as labels of spectral channels, within the image file. A secondary, more extensible format, based on the HDR was defined as the default storage format for spectral images. This is the High Dimensional Image (HDI) format, which is documented in appendix C. From this point in the processing, all actions performed during a single session of the spectral imaging software is recorded in a log file, making it possible to trace all modifications made to the image files.

A list of spectral channels is added to each image. For multispectral channels this is merely a list of the central wavelengths of the various bandpass filters used, for hyperspectral channels, this list is based on calibration against an Argon lamp. This calibration identified 18 peaks for comparison between 695nm and 970nm. The calibration was confirmed by comparison to halophosphate fluorescent lamps (standard “office” strip lights) which provided 4 peaks between 450nm and 790nm. Figure 5.8 shows the output spectrum of a spectral image taken of a white target under illumination from both the fluorescent lamps and the argon lamps.

The hyperspectral image goes through some further pre-processing at this point. The raw spectra recorded in the hyperspectral images are observed to have a high frequency noise component. This is partly due to the fact that the sensor is over-sampling the signal, and partly due to the inherent electronic noise on the CCD. By studying a dark exposure it can be determined that the average pixel value is around 15, in an ideal detector this would obviously be zero; this CCD noise level is still however, very low (about 0.3% of maximum signal).

This high frequency noise is dealt with by applying a low pass filter, in this case a Gaussian filter, along the spectral dimension. The width (σ) of the Gaussian curve

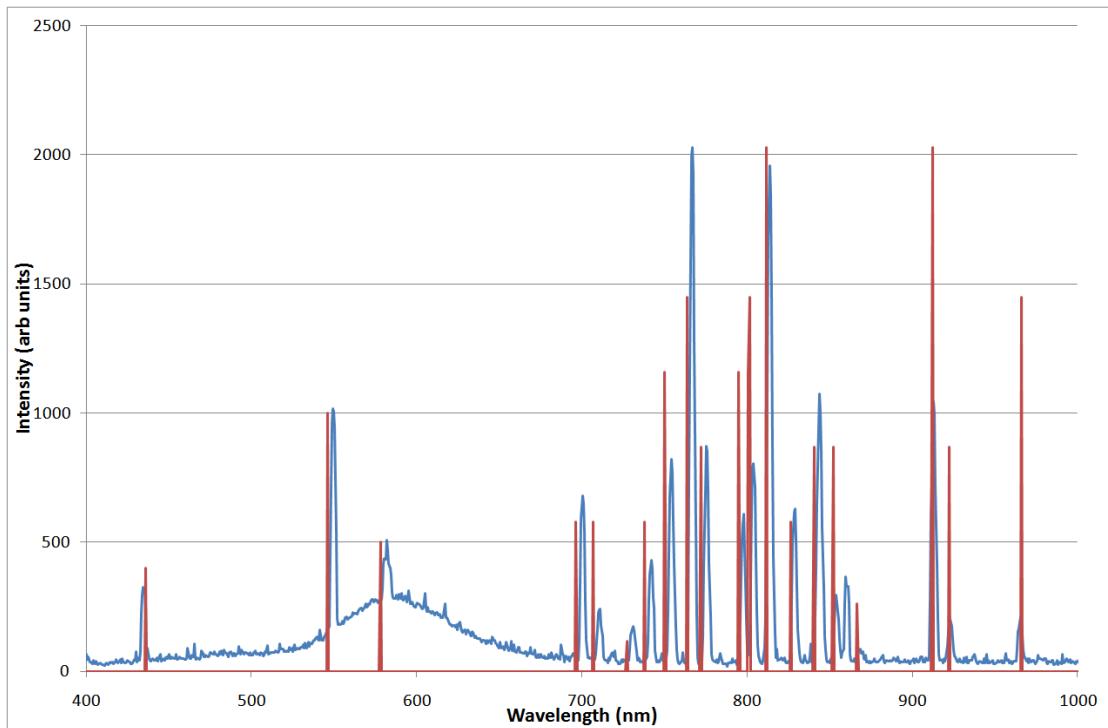


FIGURE 5.8: Output spectrum from a spectral image of a blank target illuminated by fluorescent strip lights and an argon lamp. The output spectrum is the blue line, the red lines indicate the expected peaks, the height of the red lines shows the relative sizes that these peaks are expected to be. There is generally close agreement with all peaks, no linear scale could be fit to ensure all the peaks aligned exactly, however with the understanding that the imager is only specified to 3nm, and that no exact spectroscopy was being performed, it was deemed that this calibration was adequate.

applied as a filter can be altered, as can the size of the filter (how many pixels either side of the subject pixel are considered in the calculation). The size of the filter is described by its 'radius', and whilst a greater radius achieves better smoothing, increasing the radius has two drawbacks. Firstly, for a radius of n , the first and last n pixels cannot be properly considered by the filter, and thus $2n$ pixels are lost; secondly, the Gaussian filter smooths out any features that are higher frequency than (smaller than the size of) the filter; whilst this includes the high frequency noise it could also include real features and the larger the filter, the more likely it is that real features will be lost. For this application a radius of 3 was determined empirically to be a satisfactory compromise. A sigma of 1.7 was used, based on an empirical observation that $\sigma \sim 1 + \frac{(2n+1)}{10}$ achieves appropriate attenuation of the curve at the limits of the filter size. The resulting convolution matrix is thus [0.011, 0.135, 0.606, 1, 0.606, 0.135, 0.011]. Each pixel in turn is convolved with this matrix and the pixel value is replaced with the result of the convolution. This filter is an irreversible process.

The CCD does not have a uniform response across the spectral region (see figure 1.1). In addition, the illumination used in this project is from two broadband halogen lamps, approximating black-body emitters with a peak wavelength of around 800nm;

as such, illumination is not constant across the spectral range. The total spectral response is dependent on the combined effect of the emitted radiation and the detectors responsitivity.

For some analysis purposes this is not an important factor as the spectral response is the same for each pixel and is constant across images. However, in order to interpret each pixel as a real spectrum and not just an abstract vector, this spectral response should be corrected for. In order to do this and in order to ensure that no systematic effects exist across the spatial axis of the CCD, each row of the spectral image is normalised against the spectra from a row of background pixels. In effect this normalises the images on a pixel by pixel basis, removing any fixed pattern noise that may be present in the CCD. The background in this project is a blank piece of white card and so its spectrum is representative of the system's spectral response.

The final stage of pre-processing for both sets of images is to crop the image spatially such that the amount of background is minimised. For this, the software prompts the user to select the extent of tissue on the image before discarding all data from outside this region of interest. This crop is an irreversible process.

5.2.2.2 Real Processing Routines

This section details the way in which concepts such as the SAM and SPM described in chapter 3 have been implemented for this project.

The first real processing performed on the spectral images is to calculate an SAM. The spectral angle can be measured between any two arbitrary spectra and often a reference spectrum of a known material is used alongside a test spectrum from an unknown material; in this case the spectral angle becomes a measure of how well the spectra match and often a threshold, of say 5° , is set below which the test material is classified as being the known material. In this project there is no library of spectra of relevant materials and so the same reference is used for all pixels in all images.

The choice of reference is somewhat arbitrary and the software allows for a number of options to be used. These include using a background pixel's spectrum, the mean of all image spectra, a spectrum composed of the median values across the image for each channel, and a uniform "grey" spectrum, i.e. one in which the value of the spectrum is the same in all channels. At the current stage of research, the spectral images are be considered largely in isolation, in this situation these are all equally valid choices and arguments could be forwarded for each. However, only the grey spectrum will provide spectral angles that can be compared across images at a later stage and so this is the reference used in all images.

At each pixel in the image, the spectrum from that point is fed as an input, along with the spectrum from the grey reference, to a function that calculates the angle between the two vectors according to equation 3.12. At this point the result is normalised by dividing it by $\frac{\pi}{2}$, thus the returned spectral angle value is a dimensionless value between 0 and 1. Once constructed in this manner, the SAM is then saved in the spectral image file, and a flag denoting that it has been calculated is set.

The SPM is then calculated for each pixel. In order for the SPM to carry relevant meaning, the same reference as was used for the SAM should be used again. The first step in calculating the SPM is to calculate the vector of rotation angles that will rotate the reference vector onto the primary ordinate axis. An alternative interpretation of this step is to rotate the space so that the ordinate axis is coincident with the reference vector (this is referred to as the reference axis); this creates a new space that will be referred to as the reference space. For an ND reference vector, there will be $N - 1$ rotations, which can be calculated using a recursive algorithm, outlined below.

1. Copy reference vector to temporary array, $temp[N]$, and set counting integer $i = N - 2$.
2. Calculate i th rotation angle, $rotate[i] = \arctan\left(\frac{-temp[i+1]}{temp[i]}\right)$.
3. Set $temp[i] = \sqrt{(temp[i]^2 + temp[i+1]^2)}$.
4. Decrement i .
5. Repeat steps 2-4 for all i down to and including $i = 0$.

Using this, each test vector in turn can be rotated into the reference space. In two dimensions, a rotation through an angle α can be computed using the matrix

$$\begin{pmatrix} \cos(\alpha) & \sin(\alpha) \\ -\sin(\alpha) & \cos(\alpha) \end{pmatrix}. \quad (5.3)$$

The dimensions for the test vector in the reference space can be calculated by a series of such 2D rotations through the angles from the vector of rotations calculated above; this is an inverse process of the algorithm detailed above.

1. Copy test vector to a temporary array, $temp[N - 1]$, and set counting integer $i = N - 2$.
2. Calculate i th and $i + 1$ th reference space dimensions of test

$$T_i = temp[i] * \cos(rotate[i]) - temp[i + 1] * \sin(rotate[i])$$

$$T_{i+1} = temp[i + 1] * \cos(rotate[i]) + temp[i] * \sin(rotate[i]).$$

3. Copy T_i and T_{i+1} into $temp[i]$ and $temp[i + 1]$ respectively.
4. Decrement i .
5. Repeat steps 2-4 for all i down to and including $i = 0$.

The SPM is now the vector of angles between the projection of the test vector onto the subspace orthogonal to the reference axis, and the axes of this subspace. These can be calculated using the inner product. Since each axis is a vector with only one non-zero dimension, and since the non-zero dimension can be of arbitrary size, for convenience, each axis is represented in the inner product calculation by a unit vector. The angle, φ_i , is thus

$$\varphi_i = \arccos \left(\frac{temp[i]}{\sqrt{\sum_{j=0}^{N-1} temp[j]^2}} \right). \quad (5.4)$$

Once this has been constructed for each pixel, the entire SPM is saved in the spectral image file and a flag denoting that it has been calculated is set. The SPM is saved in the same data block as the SAM, thus creating an array with the same dimensionality as the original image. Throughout the software the SAM is treated as the zeroth dimension of the SPM.

There are a number of ways to classify datasets but the principle behind all of them is the same; the data should be divided into subsets such that datapoints within a subset are similar to each other and different from datapoints in other subsets. This can be achieved by examining the position of the datapoints in some, generally abstract, space; in this way datapoints that are close to each other form subsets with the aim that there are large inter-subset distances. The spectral angle is a simple representation of a datapoint and exists in a simple 1D metric space. A simple classification strategy based on spectral angle therefore, is to assign classes based on a map of the datapoints in the metric space, such a map being a histogram of spectral angles.

In order to construct a histogram, it is necessary to determine how coarsely (or finely) to bin the measurements. Using too many bins will distribute the measurements too much, real features may be lost by spreading them too thinly; using too few bins may also lose features by combining them into larger groups. A number of strategies have been investigated for optimising the number of bins to use, and whilst it is beyond the scope of this report to present a review of these, two of the simplest choices are to use the square root of the number of datapoints (n) and Sturges' formula (Sturges, 1926):

$$No. of Bins = \log_2 n + 1. \quad (5.5)$$

For a typical 512×512 image, the square root approach would advocate 512 bins, whereas Sturges' formula would suggest only 19 bins be used. It was determined that Sturges' formula was not providing sufficient detail for the histogram to be useful. Typically only two groups were seen, a set of background pixels and a set of biological tissue; the large numbers of datapoints in each bin also suggested that the histogram could be binned more finely. The square root approach gave a much better spread and further structure emerged.

Once the histogram of spectral angles has been calculated, it is written to a Comma Separated Variable file (CSV). This spreadsheet is imported into commercial software called PeakFit (published by Systat Software), wherein Gaussian peaks can be fitted to the histogram. The centres and widths of each peak were exported into a text file which could be read by the spectral imaging software.

The PeakFit software provides a graphical user interface that makes the fitting process highly configurable. Any parameter of the fitting process can be manipulated, although defaults are provided. Perhaps the two most important parameters for this project are the number of peaks to fit and the width of the peaks. The key decision in determining the peak widths to use is whether all peaks in a single fit should be the same width, or be allowed to vary. It was decided that since all peaks are formed by the same distortions of pure signals, and there is no reason to suspect that the noise causing these distortions should be different at different spectral angles, all peaks in a single fit should be same width. Having set that constraint, the actual peaks widths used were determined entirely by the PeakFit software.

For all the work presented in this thesis, the PeakFit software was set automatically to detect the number of peaks present in a given histogram. This decision was made on the basis of some experimentation with changing the default number of peaks (not documented) and a qualitative assessment of whether or not adding or subtracting peaks improved the fit. It was determined that any benefit in fit improvement was outweighed by the cost in time taken to manually improve on the automatic fit. As a result, the number of peaks was only manually changed if it was evident that the software had detected coincident peaks (as happened occasionally). With further time for development, this parameter could be systematically altered to determine the optimal fit, which would ultimately be measured by the r^2 of the fit.

The peak centres are used to define the centres of the classes present within the image. For each pixel in the image, the spectral angle is used to compute the probability that the datapoint belongs to each of the given peaks. The Gaussian curve, suitably normalised, represents a probability distribution function; so for a spectral angle of θ ,

the *a priori* probability that it belongs to a peak with centre (or mean) μ , and width (standard deviation) σ is

$$p(\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \frac{-(\theta - \mu)^2}{2\sigma^2}. \quad (5.6)$$

These probabilities can then be used to classify the pixel in one of two ways; either by assigning it to the class centred on the peak to which it has the highest *a priori* probability of being connected, or by assigning it probabilistically to one of the classes. The second method involves summing the probabilities of the pixel belonging to each class (this need not equal one as they are independent probabilities, and for two closely located peaks, a pixel falling in between them could have a >50% of belonging to both) and normalising each by this sum, giving the *a posteriori* probabilities of belonging to each class (i.e. given there is a measurement at θ and given the relevant classes, what is the probability that the measurement belongs to each class). If there is then a 30% chance it belongs to peak A, then there is a 30% chance it will be assigned to peak A. Both ways have been tested and both have advantages and drawbacks.

Assigning purely on the basis of which has the highest probability means that measurements from the tails of each peak are not assigned to the correct class; it is not possible to determine which are the misclassified points, however it is known that they exist. The misclassifications in this case are most likely to be in favour of adjacent groups and so the “margin” of misclassification is small. However, it is arguable that misclassification is a binary state where a pixel is either classified correctly or incorrectly, in which case this margin is meaningless. Assigning probabilistically means that although each class will contain approximately the appropriate number of pixels, there are still likely to be misclassifications; here it is not possible to identify the misclassified points, or indeed to positively know that they even exist. Further, it is likely that a small number of pixels will be misclassified by a large margin, as with large numbers of pixels, events with low *a posteriori* probabilities are increasingly likely to occur.

Having run classifications in both ways, it was decided to assign pixels to the peak with which they had the highest *a priori* probability of being connected. It was observed that class boundaries were more clearly defined spatially by this method and that the classified image appeared to suffer less from what appeared visually to be statistical noise. Further, where it could be determined that adjacent classes were actually over-representing one type of response, these could be merged with confidence that the misclassification rate would decrease. Both methods remain available in the software.

Having been classified, the average spectrum for each group is calculated along with standard deviations in each spectral channel. These are output as a CSV file as a record of the group centres and 68% confidence limits. The classification creates a map

of group assignations, which is stored in the spectral image file and again, a flag is set to indicate that classes have been calculated.

The SPM extends the metric space in which the datapoints can be clustered, and so in order to start a classification using the SPM a 2D histogram is generated. The function for creating this histogram will work with any two SPM dimensions, however for the purposes of this project only the zeroth dimension, i.e. the SAM, is used as the horizontal axis.

At this stage of development, the further classification is supervised in that groups are not yet automatically mapped across the entire space. Instead, at this point the software user is prompted to highlight features of interest on the histogram. This is currently performed by drawing a coloured rectangle around the feature, this is a crude method and boundaries can not be drawn tight to the feature. Once all interesting features (as determined by the user) have been selected, each pixel is considered in turn and if it maps into ones of the boxes in the metric space, it will be displayed as the same colour as the corresponding box. If the pixel does not map into any of the boxes its average value across the spectral range is displayed in grey-scale. Both the histogram with the coloured boxes, and the corresponding spatial image are saved as separate bitmap image files (BMPs). The average spectrum for each box, and standard deviations for each spectral channel, are calculated and output as a CSV as for the classes above.

An alternative functionality exists whereby instead of the second dimension being a dimension of the SPM, it measures which dimension of the SPM has the largest value for each pixel. This “dominant plane” measure is akin to considering the direction in which the original test vector is pointing, or its heading. The same method of creating a histogram and allowing the user to examine features is then used.

For this project, this is the extent of the processing. Possible extensions and developments of these methods are discussed in chapter 6, and the source code for the full suite of spectral image processing software, which includes functionality developed but not discussed in this thesis, and all other software developed for this project, is included on a CD-ROM.

5.3 Results

5.3.1 Test Hyperspectral Results

In order to test the viability of the SAM with VNIR spectral images, a dummy image using visible spectra of healthy and metastatic sentinel nodes of breast cancer sufferers reported by Bigio et al. (2000) was constructed.

A hyperspectral image was created in which each pixel was either assigned the normal spectrum or the metastatic spectrum according to a binary mask image. Further hyperspectral images were created such that random deviations from the spectrum were added to simulate noise in the measurements. Each spectral channel of each pixel was adjusted by adding or subtracting a number randomly selected between zero and a given percentage of the maximum signal value, creating systematic distortions of the “pure” signals. This results in noise with a uniform probability distribution function. This is not representative of noise in real-world measurements, which would have a Gaussian probability distribution function. The purpose of this exercise was not however to model real-world conditions, but rather to assess the performance of the spectral imaging techniques in non-ideal conditions and to ensure their robustness. Such uniform noise was used in this section firstly for convenience of writing the algorithm to generate it, and secondly because it is not anticipated that the distribution of the noise would have an impact on the resulting effects or the validity of the conclusions drawn from these.

In this section, the term “noise level” will be used to refer to the maximum possible perturbation generated in a give image; i.e. a “noise level” of 10% means that the number to be added or subtracted to the original signal value varied between zero and 10% of the peak signal value. The noise level was increased from 1% to 20% in integral increments, each with a new hyperspectral image.

For each image an SAM was then created, figure 5.9 shows the SAMs for two of these hyperspectral images; the image on the left shows the original (all pixels are either filled with the normal, or metastatic spectrum as reported), the image on the right shows the hyperspectral image with the addition of a 20% noise level. It is evident from figure 5.9 that the background (normal) pixels are significantly less well separated from the foreground (metastatic) pixels in the noisy image than they are in the original binary image. In particular the finer spatial details in the foreground are less well defined.

Figures 5.10-5.12 show histograms of spectral angle from three of these images, the original image, a 10% noise level image and a 20% noise level. On the histogram of the original image there are just two lines at the values of the spectral angles of the normal and metastatic spectra. As we increase the noise, these lines are spread out into peaks, which gradually merge into each other. This is confirmation of what would be expected to be the case. In the presence of noise on the signal, the normal pixels exhibit a range of closely matched spectra, and the metastatic pixels also exhibit this range. It also confirms what can be observed visually from the SAMs, that the normal and metastatic pixels become less separable in the metric space as the noise increases.



FIGURE 5.9: SAMs for dummy spectral images with noise levels of 1% (left) and 20% (right).

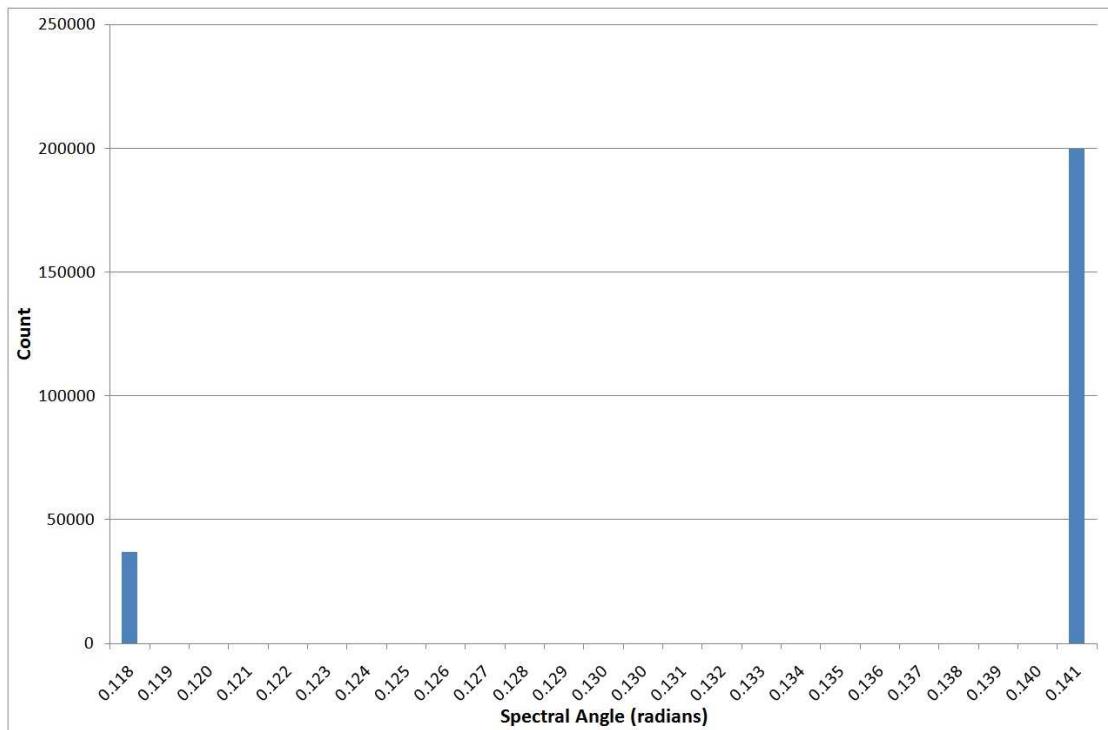


FIGURE 5.10: Histogram of spectral angles for the original dummy image. As only two types of pixel are present on the image, there are two lines, one at the spectral angle of each type.

Using these Gaussian profiles, each pixel was classified as normal or metastatic using the algorithm described in the previous section. Figure 5.13 below shows the results of classification for a low noise level (5%) and the 20% noise level image.

The 5% noise level image shows virtually no misclassification whereas much can be seen on the 20% noise level image; however, the foreground is still distinguishable from the background, and the classification seems to be correct for the majority of pixels. Using the binary mask, the number of misclassified pixels was counted and this is

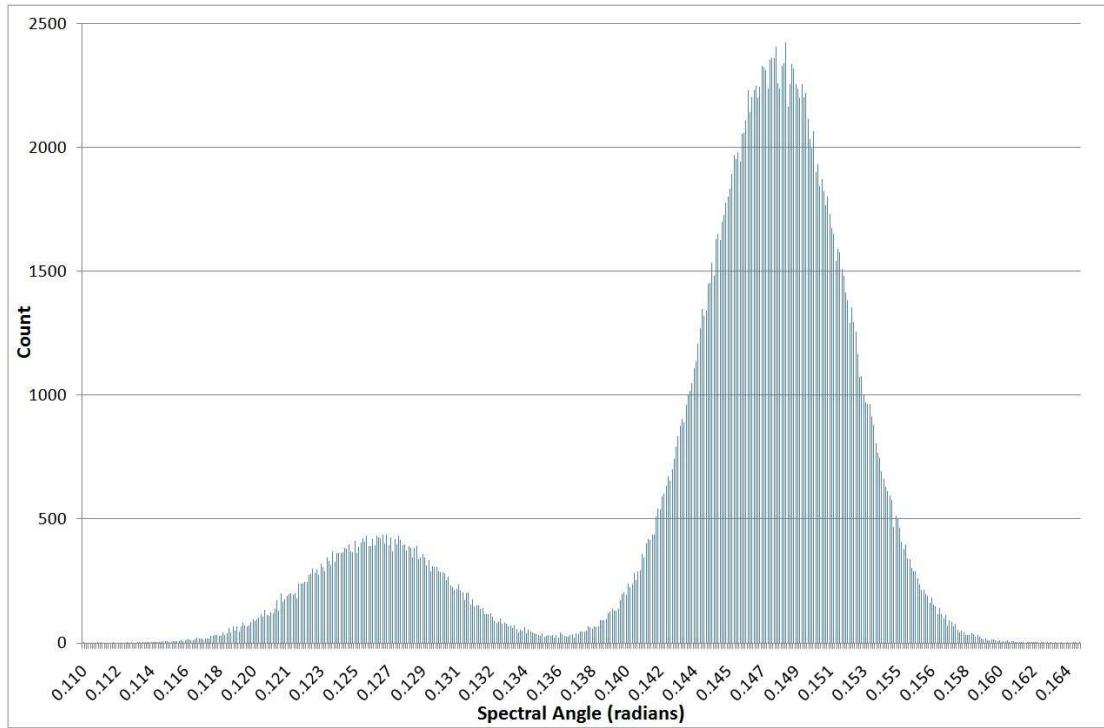


FIGURE 5.11: Histogram of spectral angles for the dummy image with 10% noise level.
The two lines from figure 5.10 are now two Gaussian shaped curves.

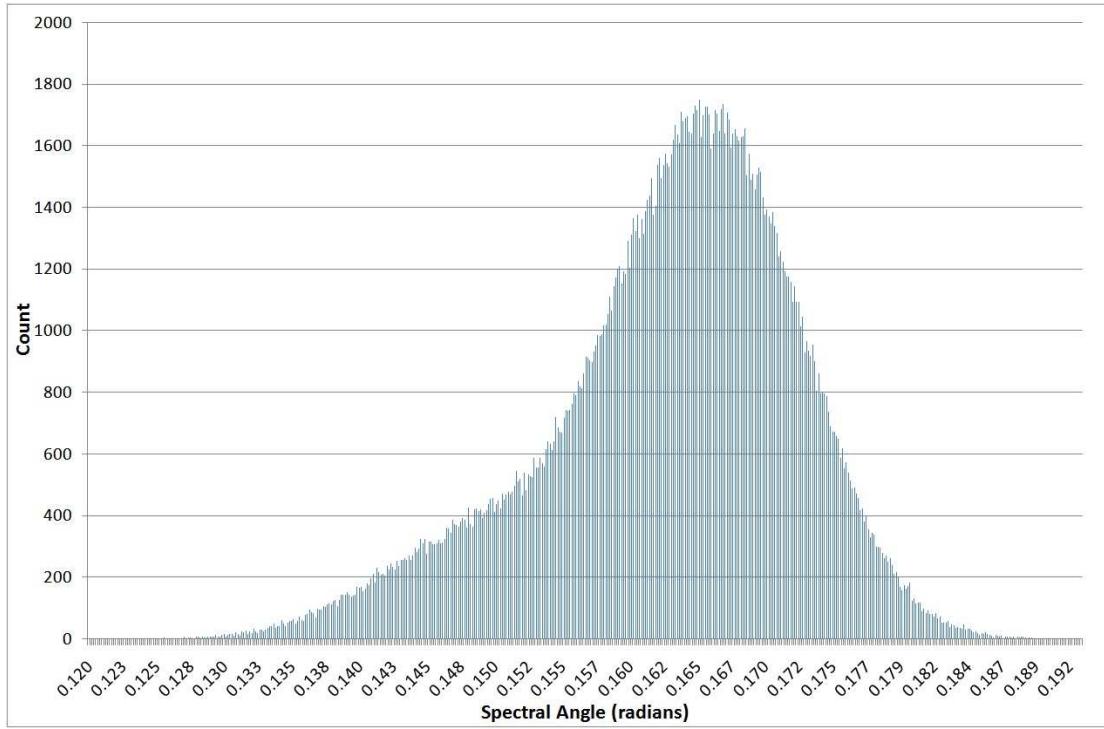


FIGURE 5.12: Histogram of spectral angles for the dummy image with 20% noise level.
The curves from figure 5.11 have largely merged into each other.

shown below for increasing noise levels in figure 5.14. This shows that the number of misclassifications was steady and very low up until noise levels of around 9-10% before



FIGURE 5.13: Dummy images with 5% noise level (left) and 20% noise level (right) with pixels classified by spectral angle using Gaussian peaks fitted to the histograms of spectral angle.

rising steadily. It is around the 9-10% level at which the peaks on the histograms begin to merge significantly.

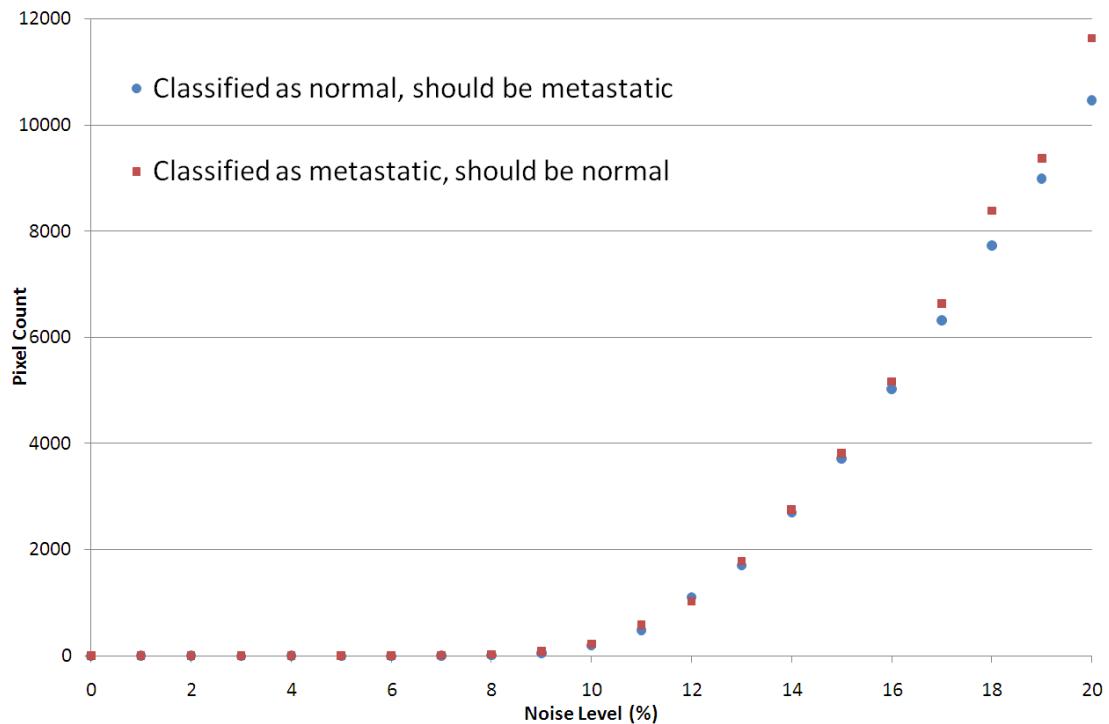


FIGURE 5.14: A graph showing the number of pixels misclassified by the spectral angle classification as a function of the noise level of the dummy image.

The average spectra of the normal and metastatic classes were calculated and these are shown for the 5% and 20% noise level images in figure 5.15. Although still clearly distinguishable, the differences between the spectra begin to diminish at higher noise levels. In particular, the peak in the metastatic spectrum at around 460nm is significantly reduced, and in the area between that and the start of the decline at

around 510nm the spectrum becomes much flatter. The separation between the spectra at higher wavelengths is also markedly reduced. The normal spectrum is much less changed than the metastatic.

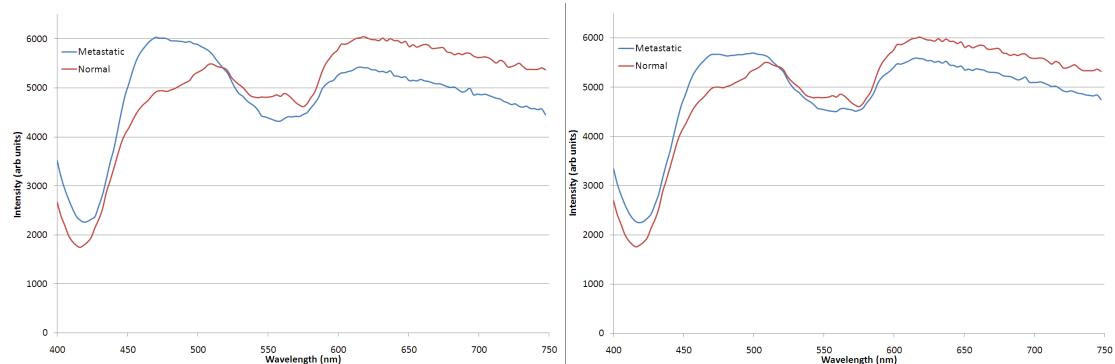


FIGURE 5.15: Graphs showing the average spectra of the normal and metastatic classes of the classified dummy images with 5% (right) and 20% (left) noise levels. The base spectra used were based on pure spectra reported by Bigio et al. (2000).

There are approximately 5 normal pixels for every metastatic pixel in the original image; it is anticipated that in a real world situation there would be less cancerous tissue in a node than normal tissue, although it is likely that the ratio would be higher in favour of normal pixels. This means that every misclassified pixel affects the metastatic spectrum a lot more than the normal spectrum, so whilst the differences between them are changing, most of this can be attributed to the metastatic spectrum becoming more like the normal.

These average spectra represent the class centres and as they change so too do their spectral angles. These class centres are represented on the histograms from figures 5.10-5.12 by the peak centres, and so as the spectra change, so do the positions of the peak centres. These changes are shown by the graph in figure 5.16, where each Gaussian curve is plotted as a point representing the centre and an error bar representing the peak width.

Although based on dummy data, these results indicate that the differences exhibited by the spectra of normal and metastatic sentinel lymph node tissue should suffice to separate the tissue types in a metric space. However, as a caveat it should be noted that this data is heavily idealised in a number of respects. Firstly, the number of spectral types present in each image is known, indeed defined, to be two; in real images there can be less confidence that this will be the case. Secondly, the distribution of these spectral types is known, meaning that misclassifications can be positively identified; in real images such misclassifications will be a fundamental limit to the sensitivity of the system. Finally, each pixel is made up of either one spectral type or the other; in a real image it is possible that a pixel could contain more than one type tissue and the spectral response at that point would be a mixture of the constituent spectra.

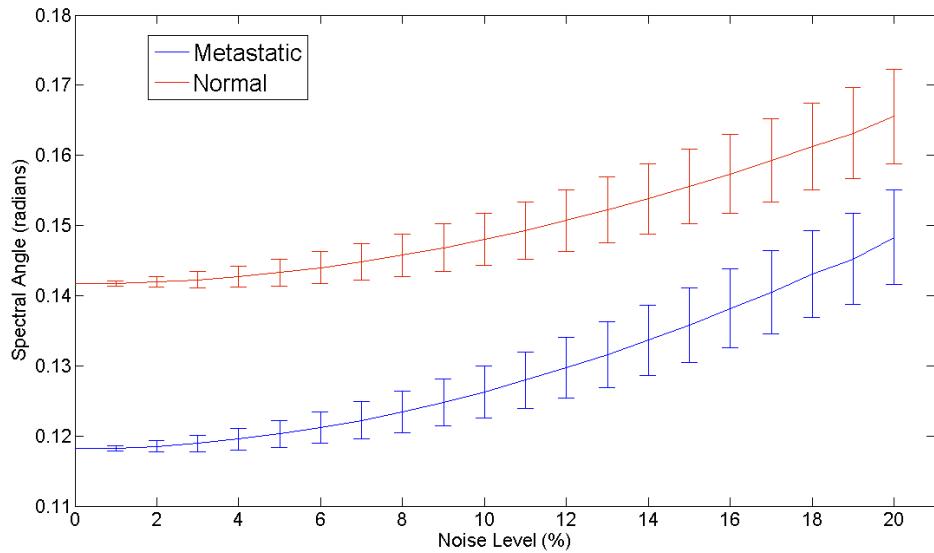


FIGURE 5.16: A graph showing the positions of the “normal” and “metastatic” Gaussian peaks as a function of noise level of the dummy image. The FWHM of each curve is shown as an error bar.

The changing of the peak position with noise level also demonstrates that looking for a fixed response to identify cancerous tissue is unlikely to be successful and so efforts at this stage should focus on identifying contrast within an image.

The SPMs for these images were also calculated, and a 2D histogram showing spectral angle against the dominant plane of the SPM was generated. The histogram for the 20% noise level image is shown in figure 5.17. It is apparent that there are two distinct groups and these are highlighted by the red and green boxes. The groups map back into the spatial domain as shown in figure 5.18.

When compared with figure 5.13, it is clear that whilst there are still a number of misclassified pixels, the separation achieved using the SPM is far superior to that achieved with SAM alone. This is confirmed by counting the number of misclassifications in each image, the graph for this is shown in figure 5.19. By comparison with the graph in figure 5.14, it is apparent that even at high noise levels the total number of misclassifications is far smaller (~ 180 c.f. $\sim 12,000$), and also that the level of noise at which misclassifications start to occur is higher (15% c.f. 10%).

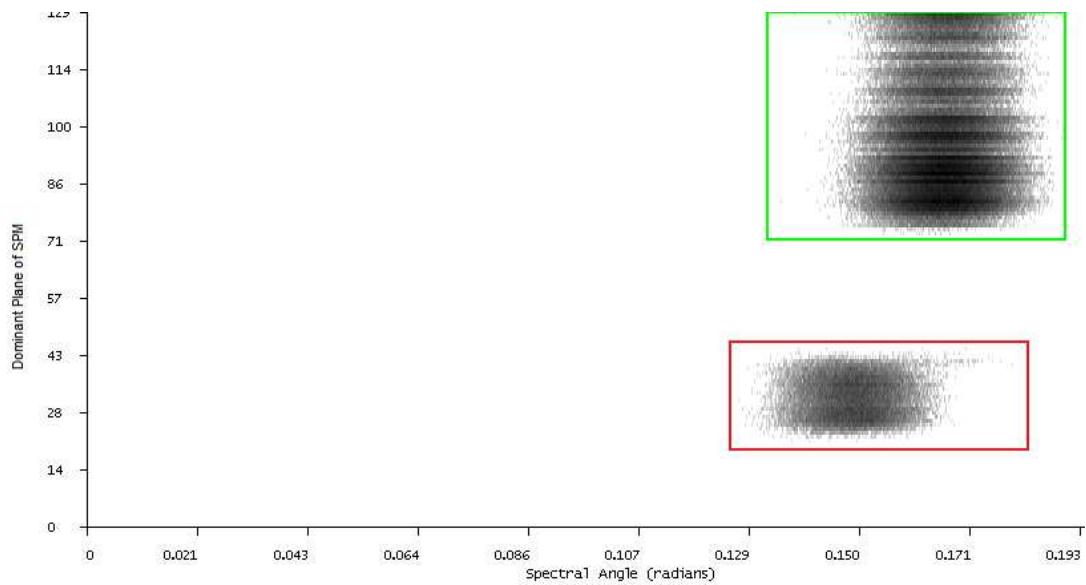


FIGURE 5.17: 2D histogram of spectral angle against the dominant plane of the SPM for the dummy image with 20% noise level. Two distinct groups are highlighted with red and green boxes.



FIGURE 5.18: Dummy image with 20% noise level with those pixels in the red and green boxes from figure 5.17 shown in red and green respectively.

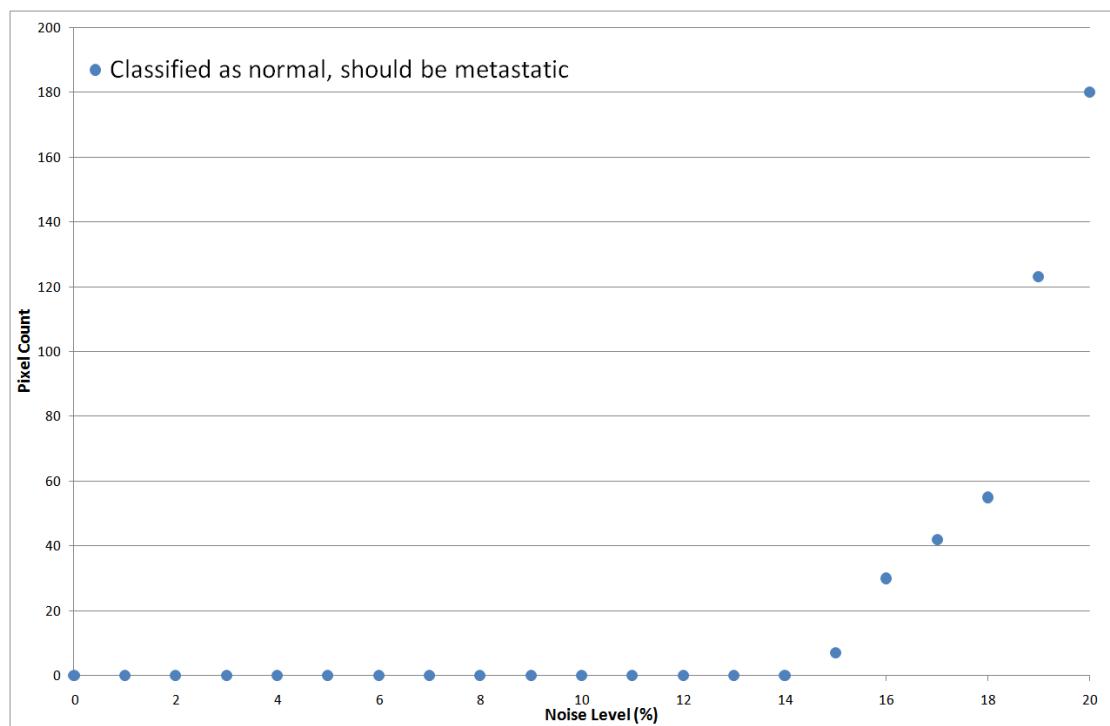


FIGURE 5.19: A graph showing the number of pixels misclassified by the spectral angle and SPM classification as a function of the noise level of the dummy image.

5.3.2 Real Hyperspectral Results

The lymph nodes in this project have been identified as the sentinel nodes of breast cancer patients and excised during the normal course of surgery. They were imaged after having been butterflied and pressed against a microscope slide for TP Cytology and before being prepared for histopathology. The images are thus of the cut surfaces and the fatty tissue surrounding them. The nodes were identified by a combination of blue dye and radioactive tracer and show varying levels of blue staining. Identification and labelling of the excised tissue as node or fatty tissue was initially provided by Mr. Sainsbury or his clinical staff, and with increasing exposure to excised samples, by the author himself. The ease with which this process of identification can be performed is highly variable, and as a first step, an automatic detection would be a useful tool.

In order to image the nodes they were placed on a microscope slide, which was in turn placed on the translation stage. The stage was translated in 0.5mm increments, at each position a line across the tissue was imaged through the $30\mu\text{m}$ slit at the face of the Specim imager. When the tissue had moved into the field of view of the Alta camera, the multispectral image was taken, before resuming the process of scanning to complete the hyperspectral image (see figure 5.7). The hyperspectral images thus acquired have high spatial resolution in the horizontal direction but much more coarse spatial resolution in the vertical direction ($\sim 30 \times 500\mu\text{m}$, see section 5.2.1.2 and in particular table 5.6).

The images presented in this thesis have been stretched in the vertical direction to make it easier to see this dimension and to make the aspect ratio a closer reflection of the real image dimensions. This makes the images appear very quantised in the vertical dimension. It should also be remembered that whilst these horizontal lines are shown a continuous progression, they in fact represent discrete lines separated by around 0.5mm.

A good example of a hyperspectral image of a sentinel lymph node is shown in figure 5.20. This is a pseudo-colour image where the red, green and blue planes of a bitmap image are comprised of the individual image planes at approximately 650nm, 510nm and 475nm, and as such it gives an approximation of how the tissue appears to the human eye. The image is 725 pixels wide, 35 pixels high and comprised of 183 spectral dimensions between 453nm and 1000nm. It should be noted that this spectral range has been used for all hyperspectral images in this project; this is because the low responsivity of the camera, coupled with the low output of illumination of the halogen lights at wavelengths lower than 450nm means that the signal at these low wavelengths is too low to provide reliable information. The approximate extent of the node itself is outlined in black, showing two, approximately elliptical cut surfaces of the bi-valved node.

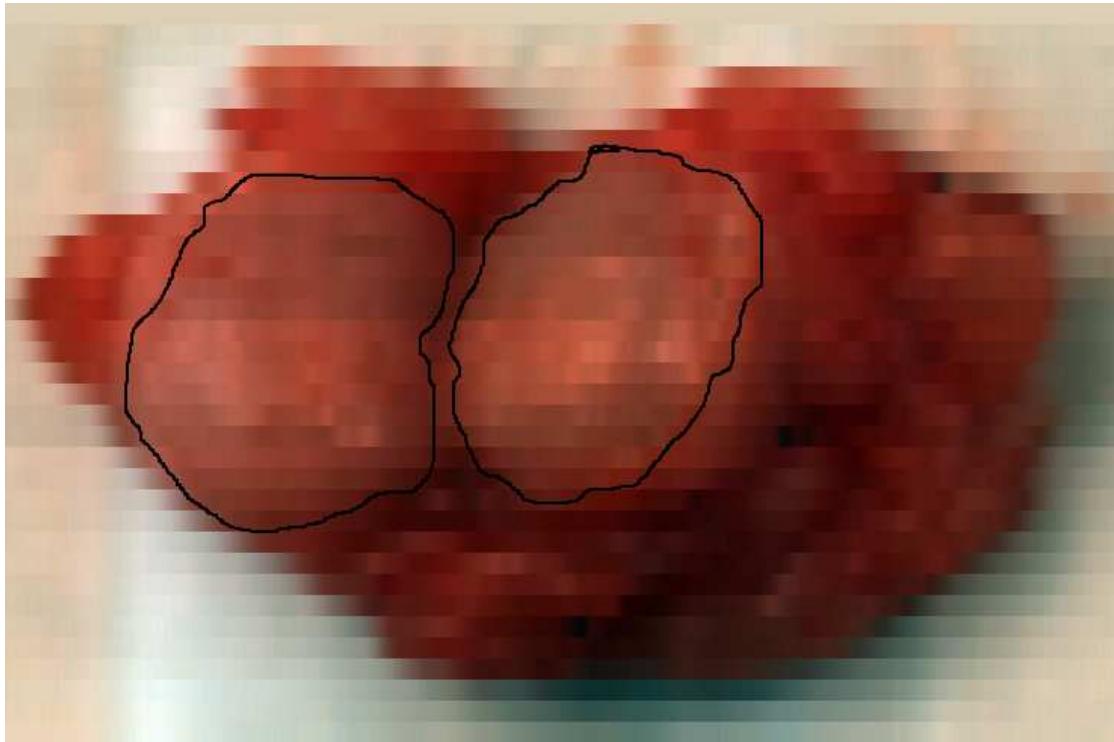


FIGURE 5.20: A pseudo-colour image generated from the hyperspectral image of a resected SLN. The approximate extent of the node is shown by black lines, marking the two cut surfaces of the bi-valved node, the remainder of tissue is the surrounding fatty tissue in which the node was originally embedded.

As with the test images in the previous section, an SAM was calculated and a histogram of the spectral angles generated. The histogram is shown in figure 5.21. This shows a significant peak at the left hand side (lower spectral angles) and a number of peaks on the right hand side (higher spectral angles). The peak at low spectral angle values represents the background pixels, the background being a piece of uniform, or near uniform, grey card, the peaks at high spectral angle values represents the biological tissue. The rest of the analysis focusses on the biological tissue, which in this case involves pixels with a spectral angle greater than 0.1 radians.

Peaks were automatically fitted to this section of the histogram using Peakfit. The software provides an automatic estimate of a line of best fit through the points and an initial fit with a number of peaks. The parameters (peak width, height and position) are adjusted and the fit iterated until no further improvement in fit quality is achieved. The width is kept the same for all peaks on the assumption that they are all the result of similar perturbations. In this case the software found 9 separate peaks, as shown in figure 5.22.

The spectral image was then classified using these peaks as class centres, resulting in the image shown in figure 5.23.

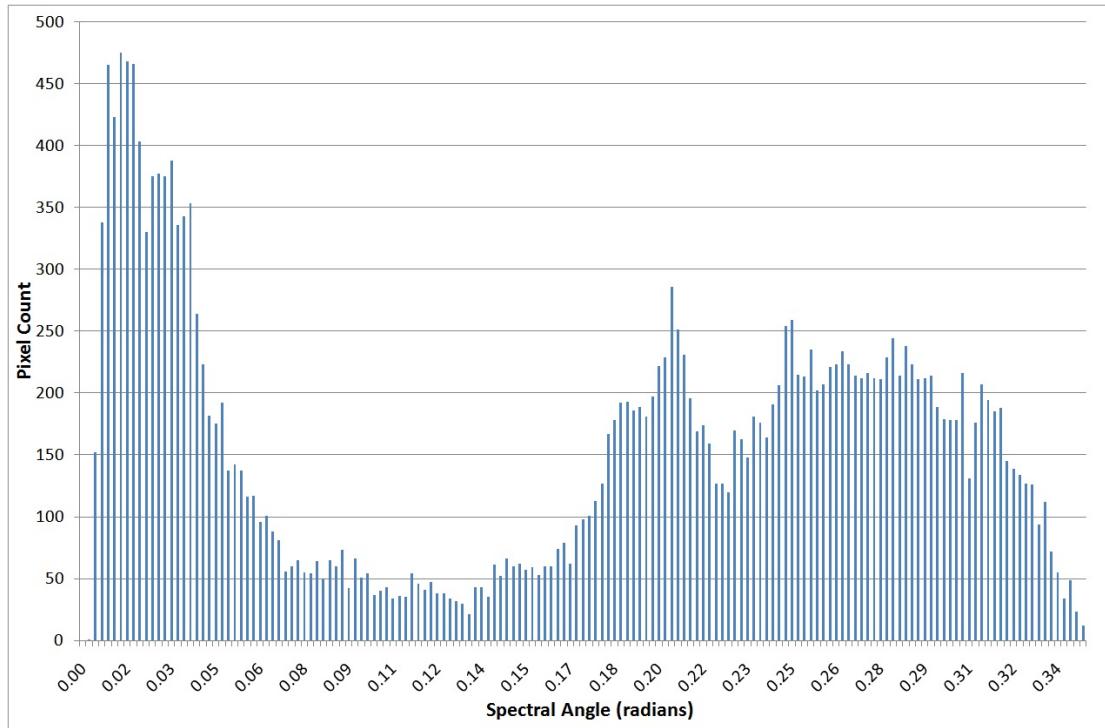


FIGURE 5.21: The histogram of spectral angles present in the hyperspectral image of the SLN from figure 5.20.

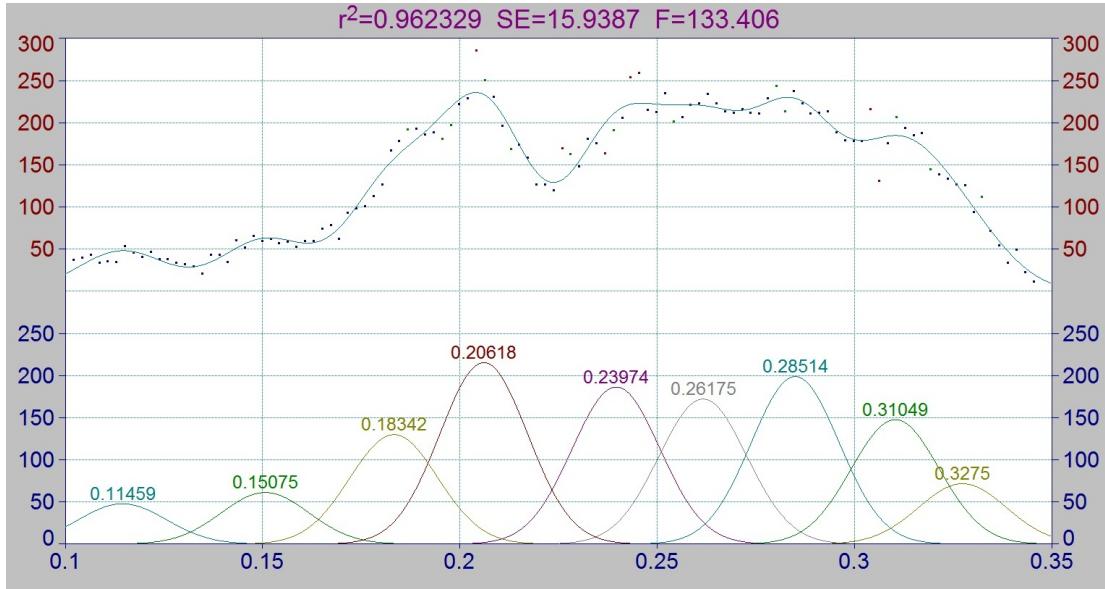


FIGURE 5.22: Gaussian peaks that have been automatically fitted to the histogram from figure 5.21.

From figure 5.23, and using the group colour key from figure 4.17, it can clearly be seen that the biological tissue has been separated from the background. It can further be observed that the tissue comprising the node itself (groups 3, 4 and 5) has been separated from the fatty tissue surrounding it (groups 7, 8, 9 and 10). The ellipsoidal shape of the node and the symmetry expected from a bi-valve cut are clearly evident.

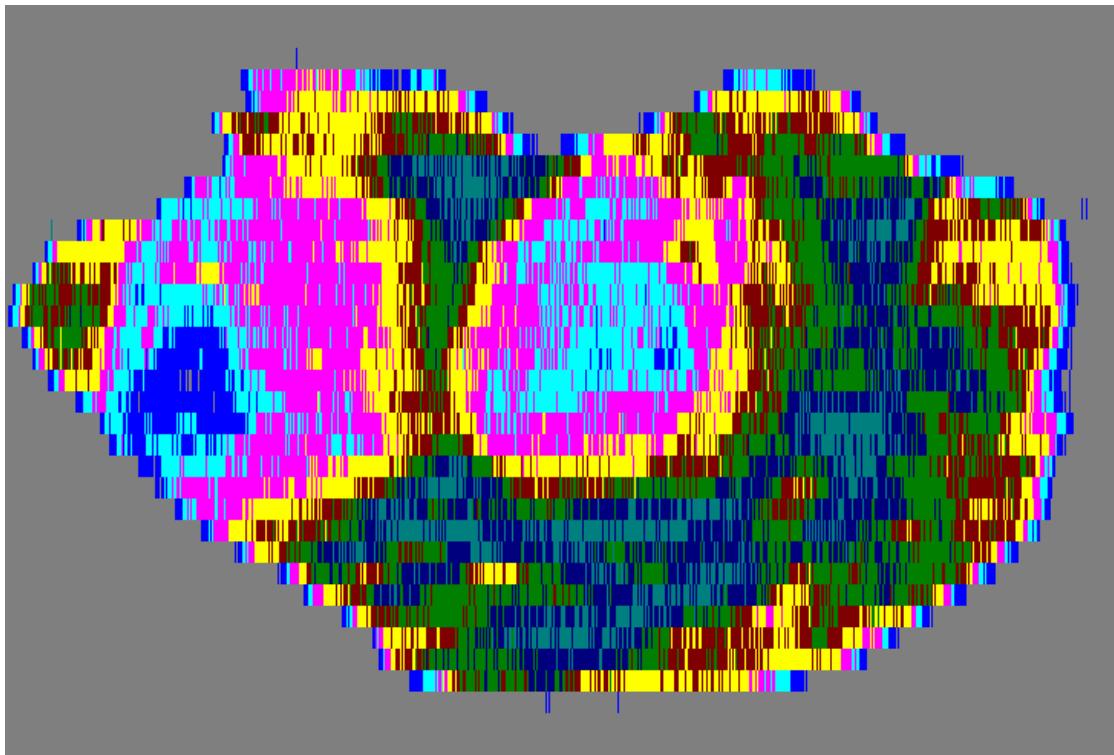


FIGURE 5.23: The hyperspectral image from figure 5.20 with pixels classified by spectral angle, using the Gaussian peaks from figure 5.22 as class centres. The grey background pixels were not considered during the classification process.

As well as this general success, there is evidence of the shortcomings of this classification, namely that group 6 seems to comprise both fatty tissue and also some of the edge of the nodal tissue, and also that the outline of the tissue as a whole is comprised of pixels classified to the same groups as the nodal tissue.

The SPM was then calculated and a 2D histogram of spectral angle against dominant plane generated. Some areas of interest have been highlighted within coloured boxes.

The spatial distribution of the pixels within these areas is shown in figure 5.25.

It is evident that pixels in group 3 in the left hand side of the node in figure 5.23 correspond well to the red pixels from figure 5.25, it is also notable that the pixels in group 3 around the edge of the tissue in figure 5.23 are not replicated in figure 5.25.

The green and blue pixels in figure 5.25 demonstrate a second area in which the dominant plane analysis has allowed a separation of pixels that was not possible using spectral angle alone. These areas were in groups 6 and 7 in figure 5.23, groups which also had a significant crossover into the surrounding fatty tissue. Whilst there are still a small number of blue and green pixels in the surrounding tissue, the predominant concentration is very much around the edge of the nodes. Again, the symmetry shown by these distributions confirms what would be expected from the known history of the tissue.

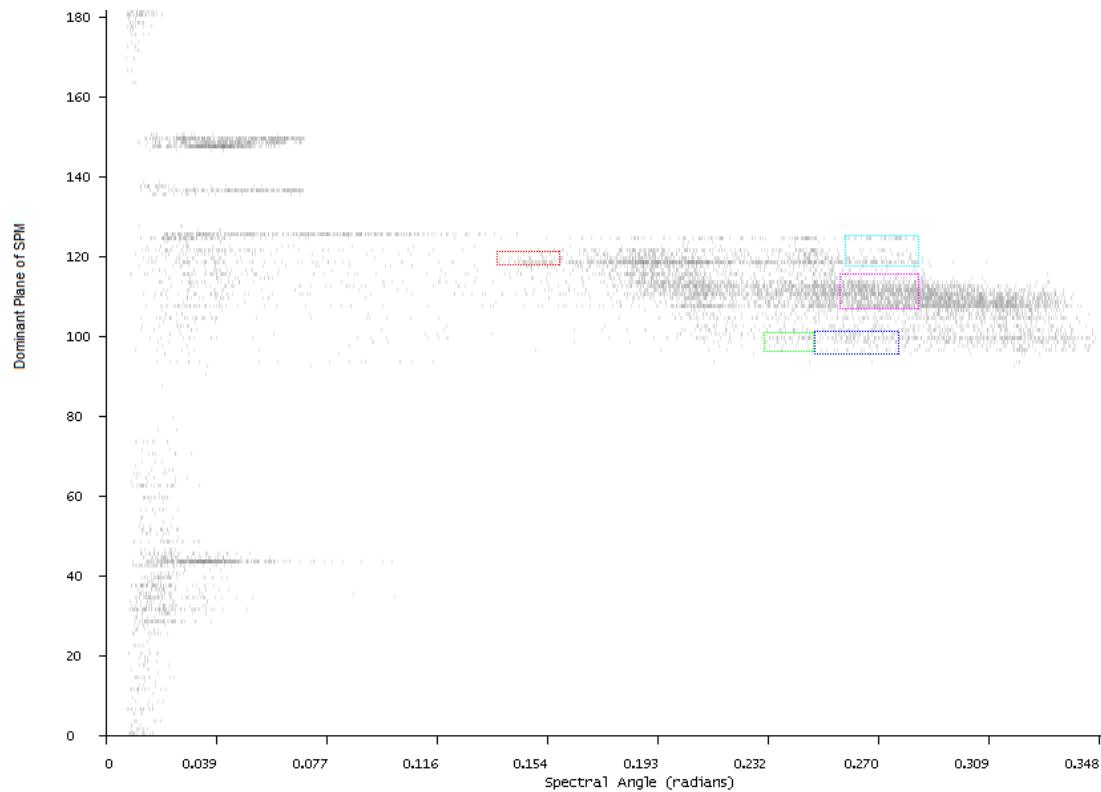


FIGURE 5.24: 2D histogram of spectral angle against the dominant plane of the SPM of the hyperspectral image from figure 5.20. Some areas of interest have been highlighted by coloured rectangles.

The cyan pixels, as can be seen from figure 5.24, cover a broadly similar range of spectral angles to the blue pixels. They are clearly separated in the dominant plane however, and on figure 5.25 it can be seen that they represent pixels that are clearly part of the surrounding tissue, rather than the nodal tissue represented by the blue. The magenta pixels also fall in the same spectral angle range and in terms of dominant plane, fall between the cyan and blue groups. The majority of these pixels are clearly in the surrounding tissue, however a significant number also form part of the edge of the node, particularly on the right hand side.

Average spectra for each of the classes from figure 5.23 are shown in figure 5.26, the colours of the lines corresponding to the colours of the classes as given by the scale bar in figure 4.17.

The spectrum corresponding to the background is largely uniform around 8000a.u. The dip below 650nm is most probably explained by the presence in the class of a small number of pixels that actually correspond to tissue. A peak is clearly visible at around 550nm and is replicated in each of the spectra, this is close to a strong mercury peak in the spectrum of the fluorescent lights used in the hospital (546nm), and it is possible that this peak is an artefact of this illumination. The shape of the spectra for

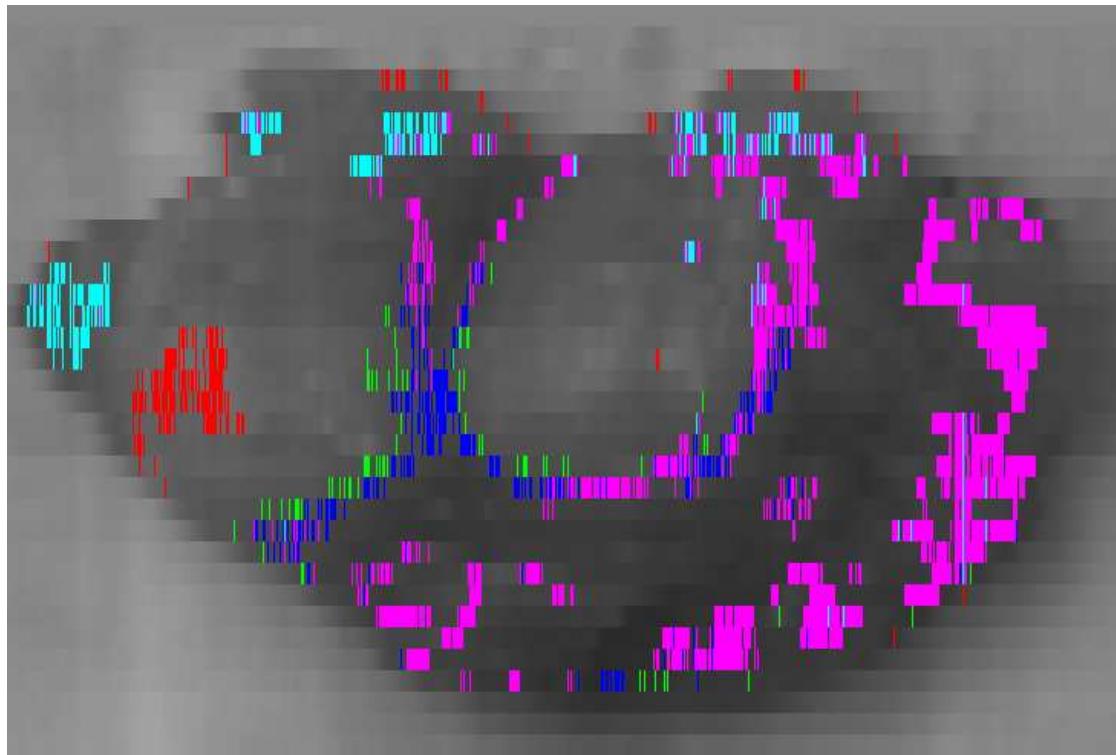


FIGURE 5.25: The hyperspectral image from figure 5.20 with pixels inside the coloured boxes from figure 5.24 shown in the same colours.

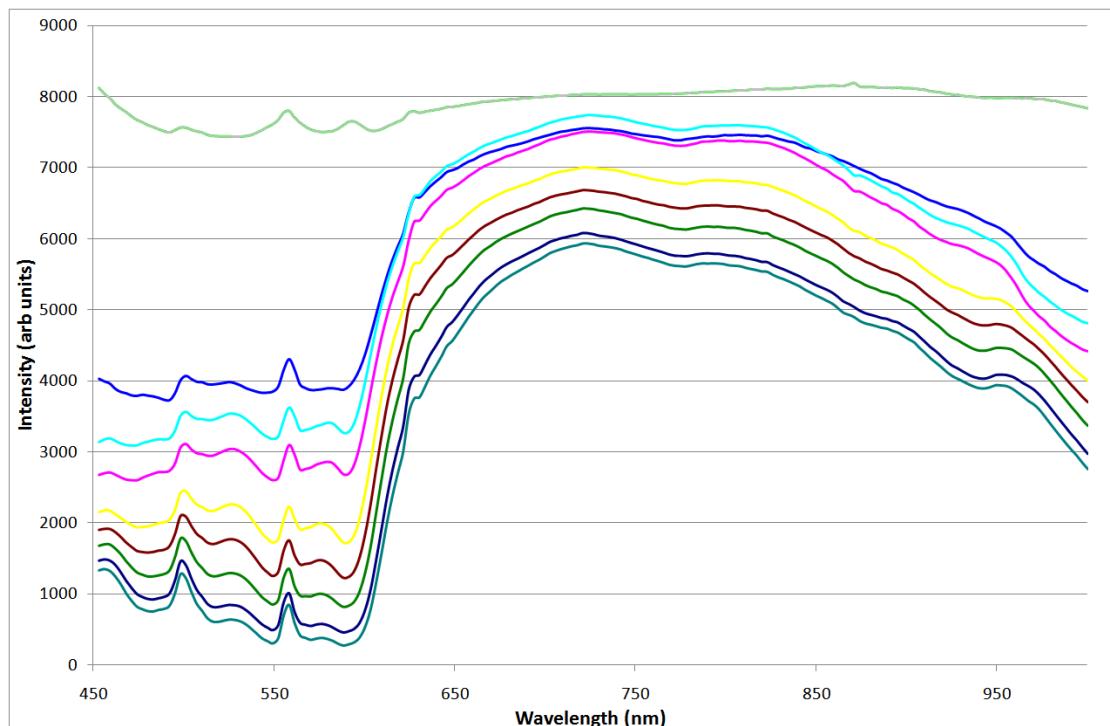


FIGURE 5.26: The average spectra of each class from the spectral angle classified image from figure 5.23.

the rest of the groups is typical for biological tissue, high intensity in the long wavelength visible and NIR, owing to the lack of chromophores in this region.

The spectra from groups 3, 4 and 5 are similar in shape, as might be expected as they all represent node tissue; there are detailed differences, in particular the spectra from group 4 twice crosses that from group 3 (at around 630nm and again at just below 850nm). The spectra from groups 7, 8, 9 and 10, also follow similar shapes with an apparent peak at just above 950nm. When considering that these spectra represent fatty tissue, this “peak” is actually likely to be the rise from a trough at just below 950nm, corresponding to the absorption peak seen in pork fat and detailed in figure 1.7. In the spectrum for group 6, this apparent peak is more akin to a shoulder, this would be explained by it being a mixture of node and fatty tissue.

The spectra from the groups show in figure 5.25 can be seen in figure 5.27, again with the colours of the lines corresponding to the colours of the groups.

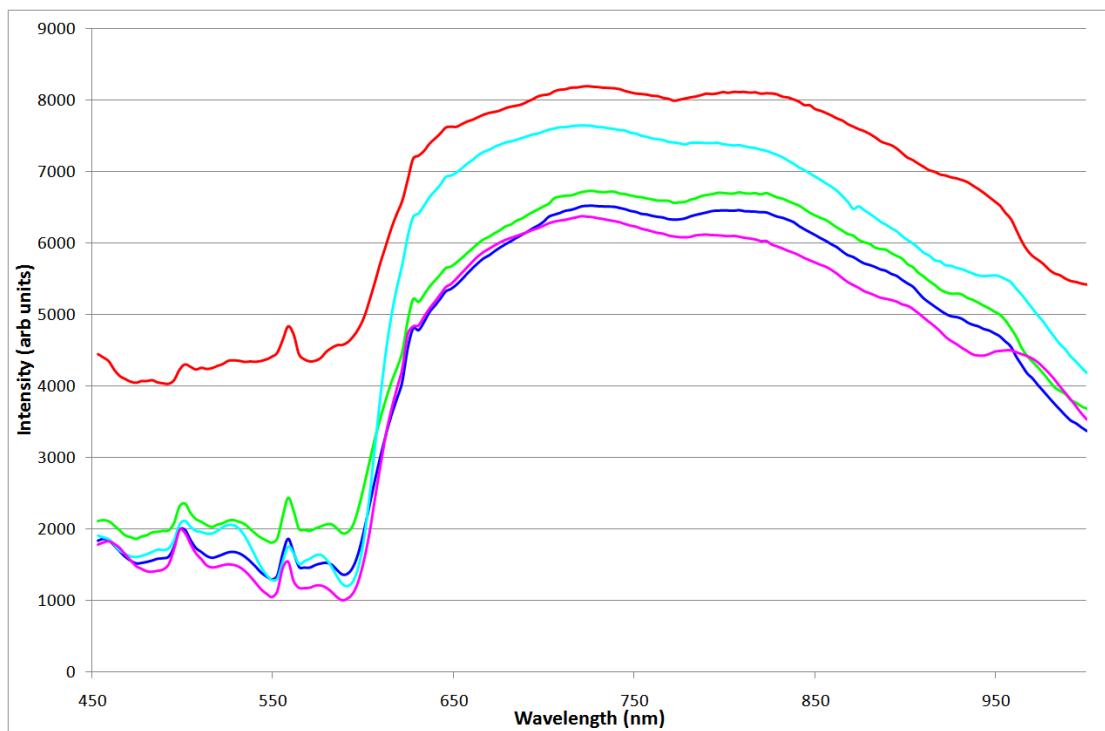


FIGURE 5.27: The average spectra of each of the groups of pixels shown in figure 5.25.

Here, the red spectrum clearly has significantly higher reflection up to 630nm and a pronounced dip beyond 950nm (this dip is evident on groups 3, 4 and 5 in figure 5.26, particular for group 4). This could be explained by a trough in the spectrum corresponding to the water absorption peak at 970nm.

In the magenta spectrum (consisting of fatty tissue) the trough just below 950nm is clearly recognisable as such, rather than there being a peak just above 950nm. The cyan spectrum, which otherwise has many differences from the magenta, also shows

strong evidence of this trough. The green and blue spectra much more closely resemble the shape of the red spectrum in this long wavelength region, supporting the separation of node and fatty tissue at the same spectral angle.

Average spectra from four small areas of the image were also calculated, the areas being shown as coloured squares on figure 5.28. The spectra themselves are shown in figure 5.29.

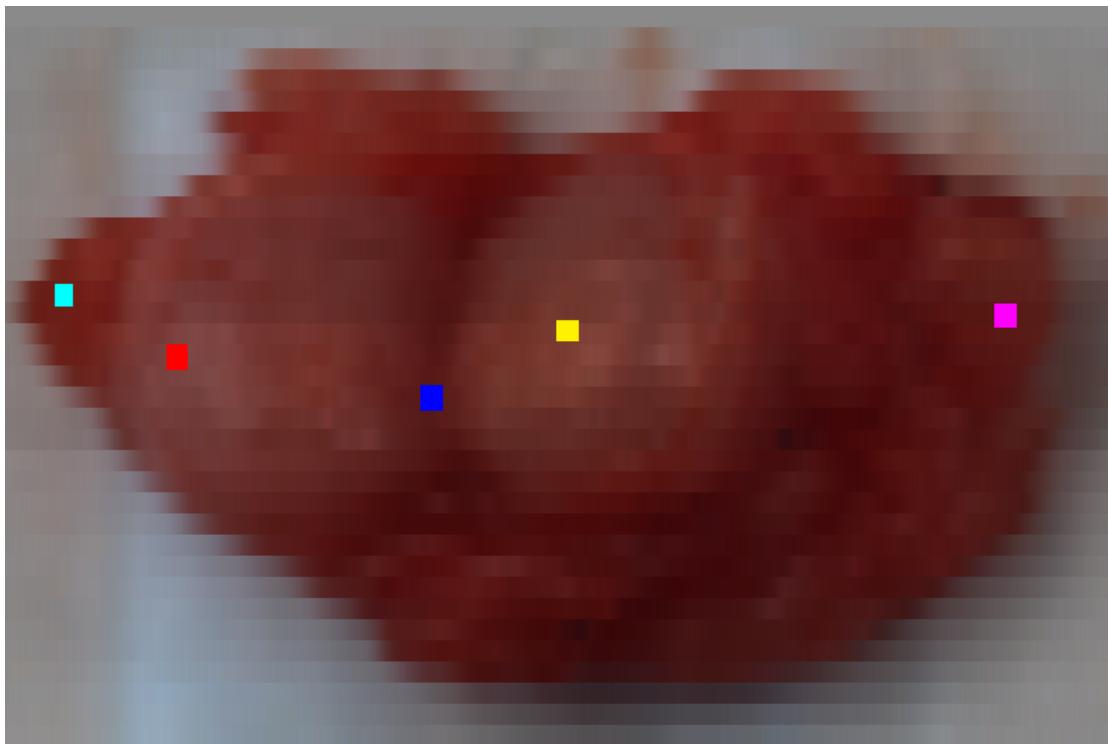


FIGURE 5.28: A pseudo-colour image of the hyperspectral image shown in figure 5.20 with four areas of interest highlighted, corresponding approximately to areas shown in the four groups in figure 5.25.

These coloured areas correspond closely to the equivalent coloured areas from figure 5.25, with the addition of a yellow area corresponding to some node tissue.

The trough just below 950nm is again strongly evident in the magenta spectrum, but interestingly not in the cyan spectrum. The strong dip above 950nm is again seen in the red and blue spectra, and there is some evidence of it also in the cyan. The yellow spectrum also shows some sign of this dip, although much less strongly than the red and blue.

The most important thing to note from figure 5.29 is that these spectra much more strongly resemble the results from figure 5.27 than they do the results from figure 5.26.

The node in this analysis is known to have tested positive for the presence of cancerous cells in cytology. Unfortunately, the preliminary nature of this study has meant that

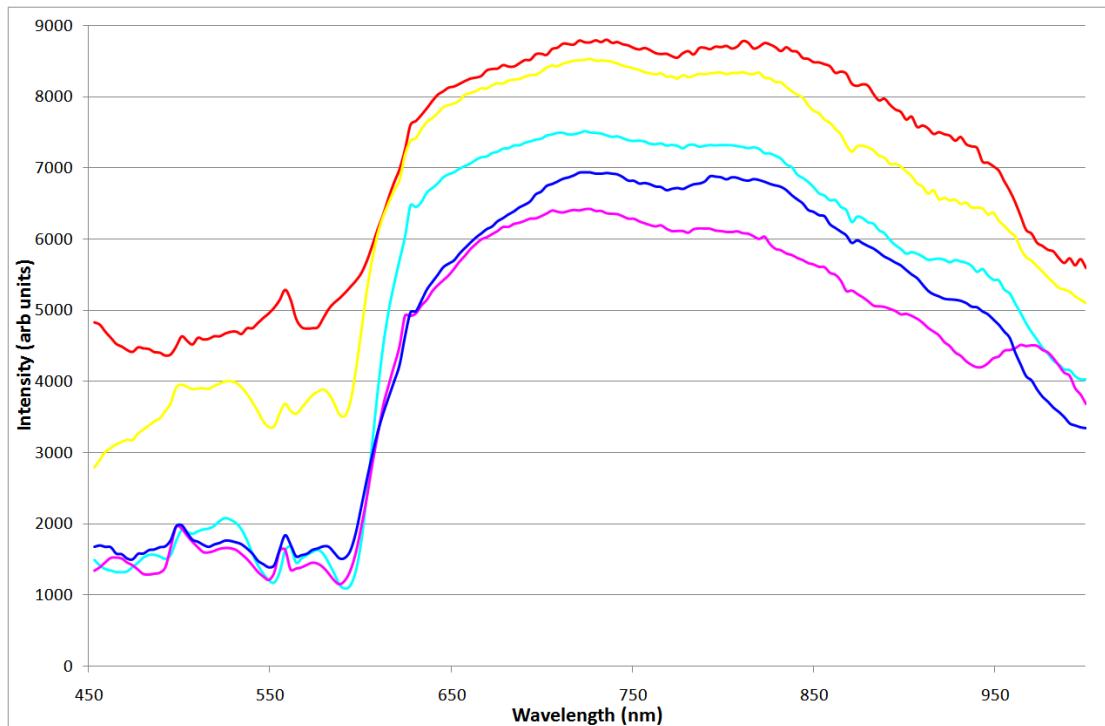


FIGURE 5.29: The average spectra of the four areas shown in figure 5.28.

once they had been imaged, the nodes were not followed with sufficient detail to allow comparisons between the cytology images and the spectral images, this point is discussed in more detail in chapter 6. It is known however that metastasis in the lymph node starts around the edges of the nodes, gradually moving in towards the centre. The blue and green groups in figure 5.25 are a tantalising hint that this may be what has been detected in this image, however it is not possible to verify this.

A second sentinel lymph node from another patient is shown below in figure 5.30. Again, this covers the spectral range of 450-1000nm with 183 channels and is a pseudo-colour image created in the same way as figure 5.20. The image is cropped to 981 pixels wide and 53 pixels high. Again, the approximate extent of the node is marked in black.

Again, a histogram was created and peaks fitted to the region representing the biological tissue, in this case 7 peaks were fitted. The classification according to these peaks gives the image shown in figure 5.31.

This image is harder to interpret than the previous example; in this case the node tissue lies along a diagonal, with the cut running from upper left to lower right. With this in mind, a symmetrical ellipsoidal shape starts to suggest itself, with group 4 being the clearest area of node, the discrimination is far from clear however. The average spectra for these classes are shown in figure 5.32.

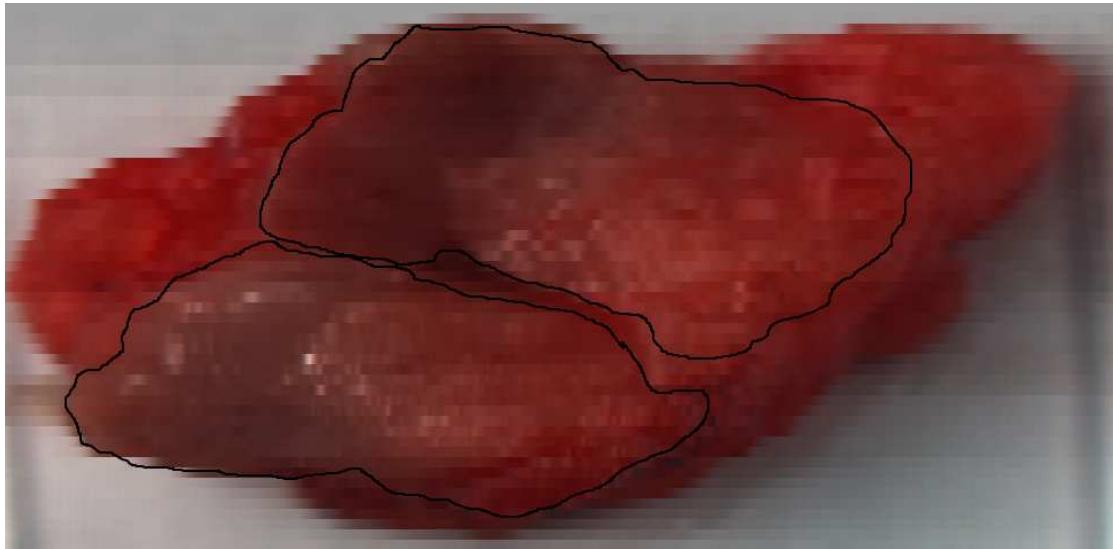


FIGURE 5.30: A pseudo-colour image generated from the hyperspectral image of a resected SLN. The approximate extent of the node is shown by black lines, marking the two cut surfaces of the bi-valved node, the remainder of tissue is the surrounding fatty tissue in which the node was originally embedded.

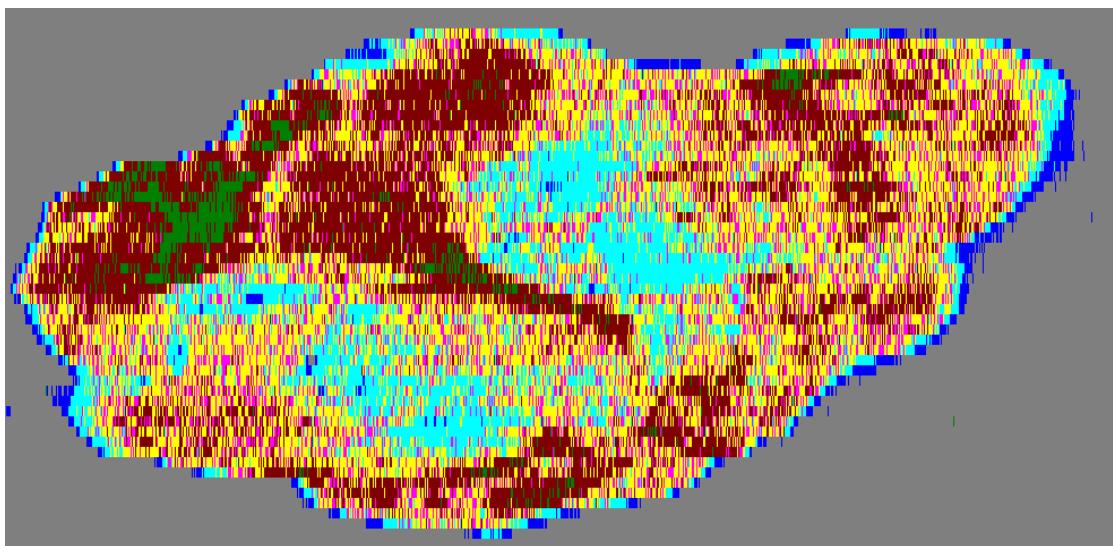


FIGURE 5.31: The hyperspectral image from figure 5.30 with pixels classified by spectral angle, using the Gaussian peaks calculated from the histogram of spectral angles as class centres. The grey background pixels were not considered during the classification process.

It is immediately apparent that the spectra for groups 5 and 6 are almost identical, which suggests that an extra peak has been fitted where it is not actually needed¹. This makes sense also in figure 5.31 where the pixels for these groups are not spatially resolved from each at all. The lipid trough just below 950nm is evident to some degree in all of the spectra suggesting that none of these groups properly resolves the node

¹In iterating this case one of the peaks gradually shifted until it was very nearly coincident with another.

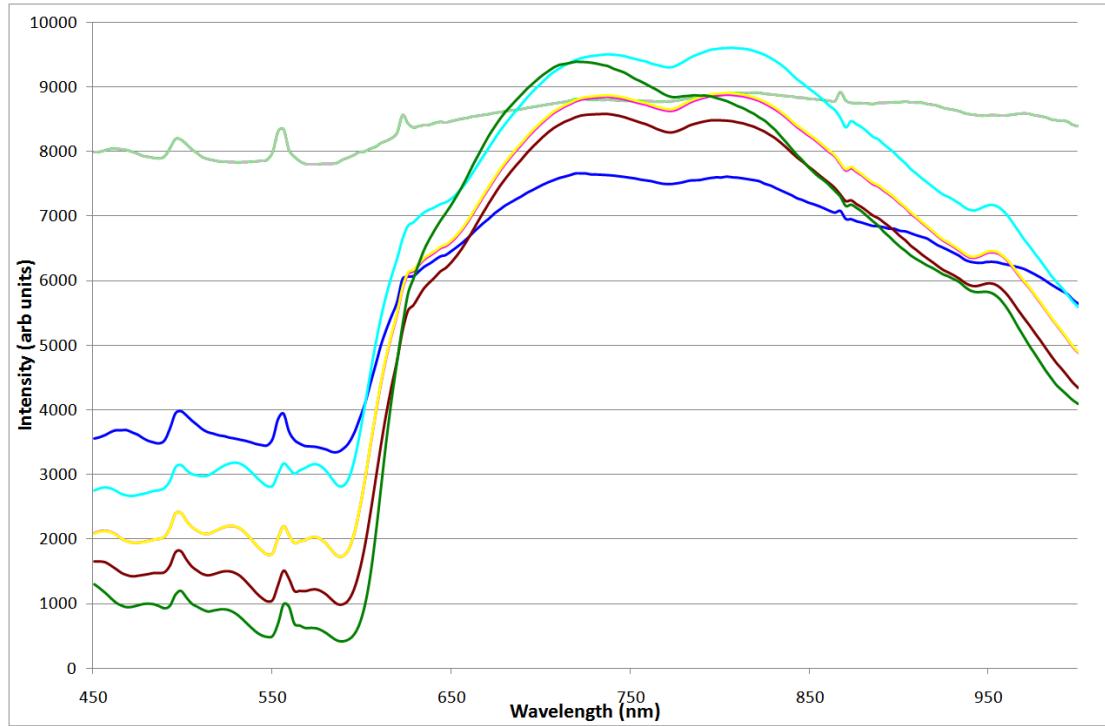


FIGURE 5.32: The average spectra of each class from the spectral angle classified image from figure 5.31.

from the surrounding fat. Group 4 exhibits a significant shoulder at around 630nm which, whilst present is much less pronounced in the other groups. This feature is not as evident in figure 5.29 where all of the spectra rise more sharply in this region.

Again, a 2D histogram of spectral angle and dominant plane was plotted, and this is shown in figure 5.33, with some coloured boxes drawn in to highlight certain areas of the plot. At the right hand side of this plot, two bands appear in the dominant plane, and as suggested by the shape of the coloured boxes, these are significant.

The spatial arrangement of these pixels is shown in figure 5.34.

From figure 5.34 it can be seen that the green pixels now cover most of the surrounding fatty tissue, the blue pixels form an approximate boundary around the node tissue and the red and cyan pixels come from the node tissue itself. The average spectra for these regions are shown in figure 5.35.

In these the lipid trough just below 950nm is strongly evident in the blue and green spectra and much less so for the red and cyan, which supports the fact that the green pixels cover the surrounding tissue and suggests that the blue, whilst demarcating the node, actually have significant crossover into the surrounding tissue. The red and cyan spectra also show stronger evidence of a sharp dip towards the water trough at 970nm. The sharp shoulder at around 630nm is very evident here in the blue, red and cyan spectra, whilst the green peak has significantly higher intensity from around 630nm up

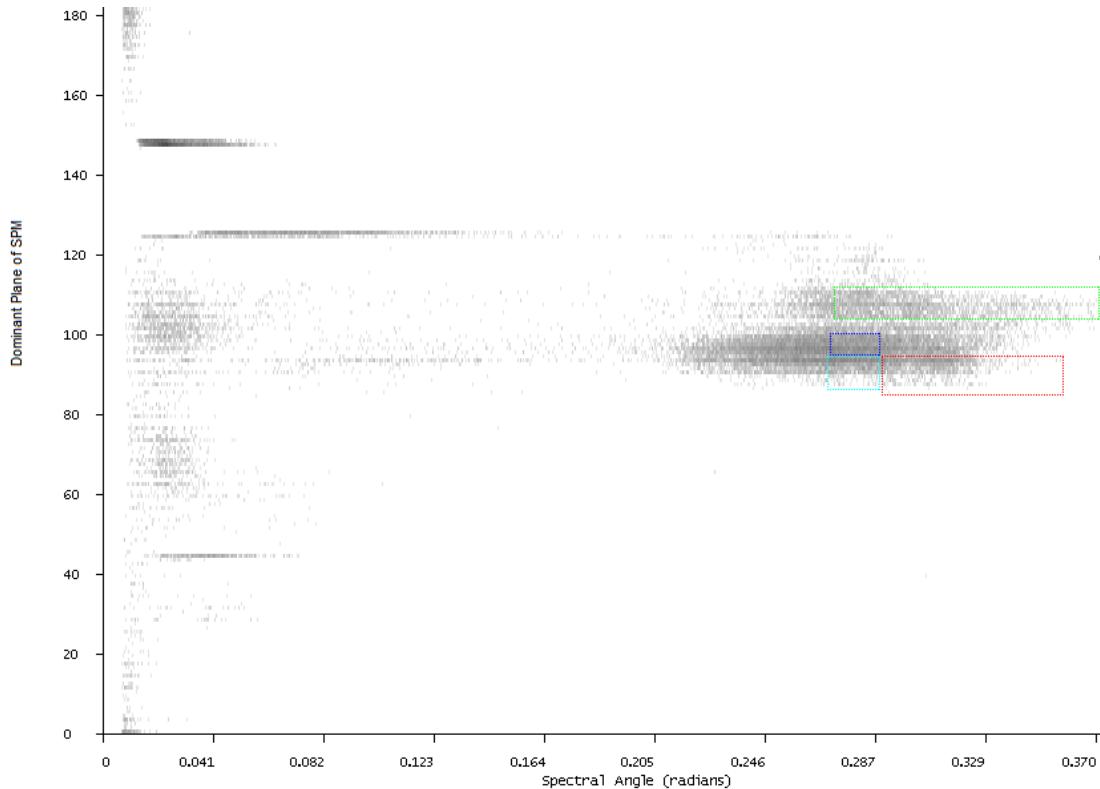


FIGURE 5.33: 2D histogram of spectral angle against the dominant plane of the SPM of the hyperspectral image from figure 5.30. Some areas of interest have been highlighted by coloured rectangles.

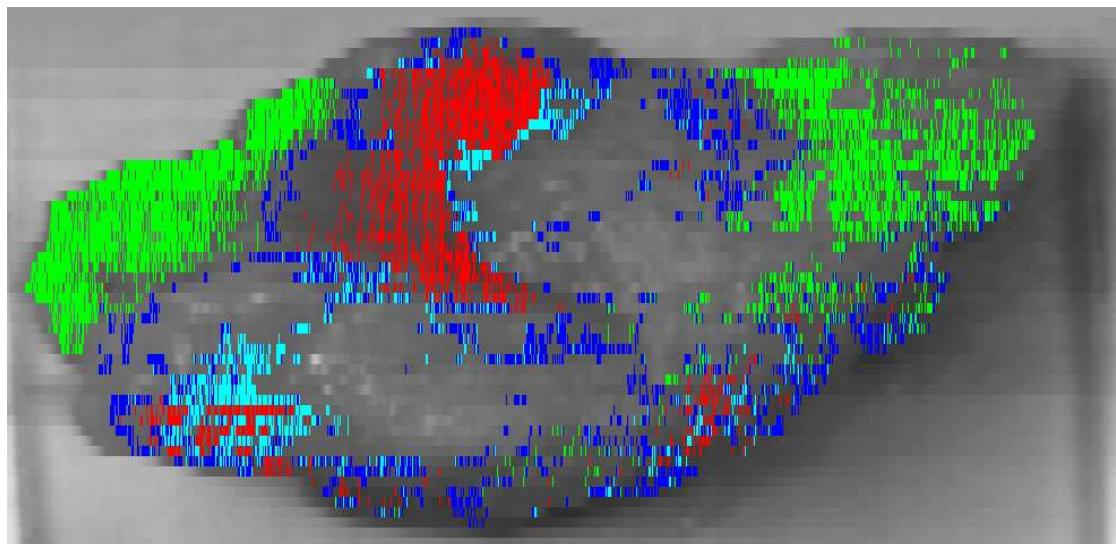


FIGURE 5.34: The hyperspectral image from figure 5.30 with pixels inside the coloured boxes from figure 5.33 shown in the same colours.

to around 800nm. From figure 5.33 it is clear that the green pixels cover the same spectral angle range as all of the others here, thus demonstrating that major differences in the shape of the spectra can be missed by use of the spectral angle alone.

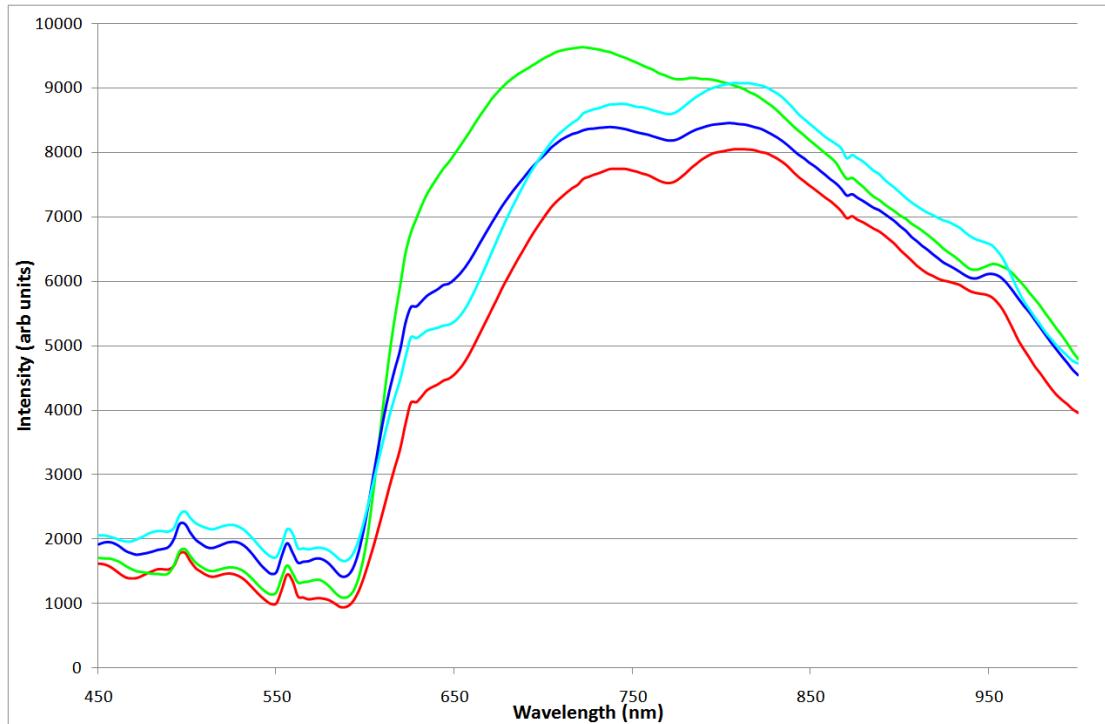


FIGURE 5.35: The average spectra of each of the groups of pixels shown in figure 5.34.

These two examples both demonstrate an enhanced capability to distinguish between different spectra when the spectral angle is used alongside information from the SPV compared to when it used on its own. Whilst only partial 2D classifications have been presented here, there is evidence that this combination of measure will allow for a complete classification that is more accurate than can be achieved when only working with spectral angle. The fact that the groups in figures 5.25 and 5.34 are clustered in spatially coherent groups, rather than being randomly distributed about the image, also suggests that the extra discrimination that can be achieved with the spectral position is real.

These two examples also demonstrate the variability between patients. When working with these nodes, this is most obvious spatially, in the array of sizes and shapes in which they exist. From the spectra presented above, it appears that there is spectral variation also, and whilst some features, like the easily explained lipid and water troughs, appear to be consistent, the as yet unexplained shoulder at 630nm appears to be less so. It must be stressed that a sample of 2 is far too small to draw any conclusions, but this variability is an important consideration when carrying this work forward. Analogous images and spectra for several more nodes can be found in appendix B.

5.3.3 Multispectral Results

The multispectral imager described in section 5.2.1.3 above was used with narrowband (10nm FWHM) interference filters at 470nm, 527nm, 610nm, 900nm and 970nm in order to record 5D multispectral images of the sentinel lymph nodes. The pair of NIR filters at 900nm and 970nm were selected for their sensitivity to the water absorption peak, the three visible filters were chosen partly to cover the full range of the visible spectrum and partly on the basis of the spectra reported by Bigio et al. (2000).

Figure 5.36 shows a monochrome image of the multispectral image of the same sentinel node as shown in figure 5.20. The image was taken through a neutral density filter. The image has been cropped from its original 512×512 pixels to 400×330 pixels, in order to reduce the number of background pixels. The approximate extent of the node is outlined in black.



FIGURE 5.36: A monochrome image generated from the multispectral image of a resected SLN. The approximate extent of the node is shown by black lines, marking the two cut surfaces of the bi-valved node, the remainder of tissue is the surrounding fatty tissue in which the node was originally embedded.

The same approximate geometry of the node and the surrounding tissue can be seen in both the hyperspectral and multispectral images; the apparent distortion in the height can be explained by the fact that the horizontal and vertical spatial scales are the

same in the multispectral image, whereas they differ in the hyperspectral image as a result of the differences in spatial resolutions in the horizontal and vertical dimensions.

As with the hyperspectral images, the SAM was calculated and a histogram of spectral angles formed. The histogram is shown in figure 5.37. There are a cluster of peaks here at low spectral angles, representing the background pixels, with a further set of peaks at higher spectral angles representing tissue. Again, the biological tissue starts at around 0.1 radians.

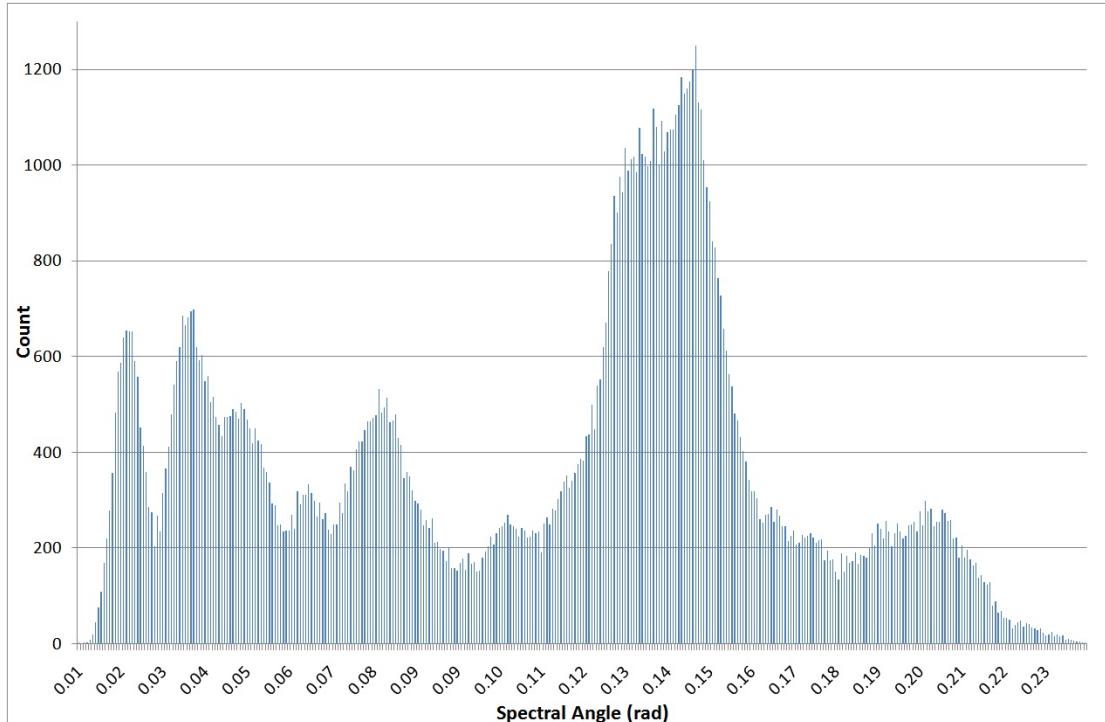


FIGURE 5.37: The histogram of spectral angles present in the multispectral image of the SLN from figure 5.36.

Focussing on the region representing tissue, 7 peaks were automatically fitted to this histogram as shown in figure 5.38.

The spectral image was then classified using these peaks, as shown in figure 5.39. Again, the colour scale from figure 4.17 is used.

Again, the first thing to note is that the background can be well separated from the tissue, although there does appear to be some tissue on the lower right hand side that is classified as background. The general shape of the nodes can be seen, primarily in group 4 on the right hand side and groups 4 and 5 on the left hand side. The node is well delineated from the surrounding tissue at the top, with group 6 representing the edges of the node, and some fatty tissue, but at the bottom, although the general shape is perceptible, the same two groups, 4 and 5, represent the node and all the surrounding tissue. This creates the impression of horizontal banding across the sample.

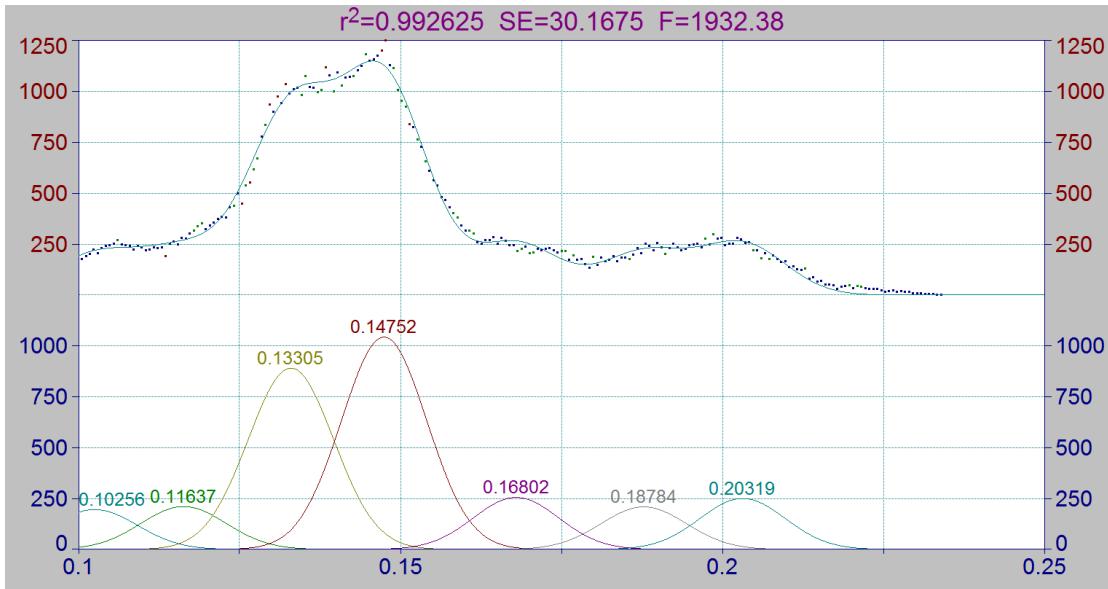


FIGURE 5.38: Gaussian peaks that have been automatically fitted to the histogram from figure 5.38.

2D histograms of the spectral angle with each of the 4 planes of the SPM are shown in figure 5.40. The histograms with planes 1 and 2 show the greatest variability in spectral position, and of these the histogram with plane 1 shows more areas of high concentrations of pixels, as such this plane was selected for further analysis. The histogram, with some boxes drawn on to highlight some key areas is shown in figure 5.41 and the spatial distribution of the pixels in each box is shown in figure 5.42.

The red box on figure 5.41 clearly covers a large range of spectral angles, but is limited in spectral position. The green and blue boxes both cover smaller spectral angle ranges, but both are within the range of the red box and as such would be indistinguishable by spectral angle alone, however they both highlight peaks in the 2D histogram. From figure 5.42, it is clear that the blue and green cover pixels from the node tissue itself, and the red box is limited to the surrounding tissue. Clearly all three boxes cover the range that mapped to groups 4 and 5 in figure 5.39. At higher spectral angles, the yellow, cyan and magenta boxes all also cover the same spectral angle range. Their distribution on figure 5.42 suggests that the yellow boxes cover more of the fatty tissue, whereas the cyan and magenta are both on the periphery of the node tissue. Since the yellow pixels are closer to the spectral position range of the red pixels than the green and blue, this could suggest that the spectral position is as important, if not more important in this case, for separating the node from the surroundings.

Figure 5.43 shows the monochrome, neutral density image of the same sentinel lymph node as shown in figure 5.30. This image has been cropped from its original 512×512 pixels to 436×497 pixels. Again, the spatial distortion compared to the hyperspectral image (in fact, the hyperspectral image is distorted rather than this multispectral

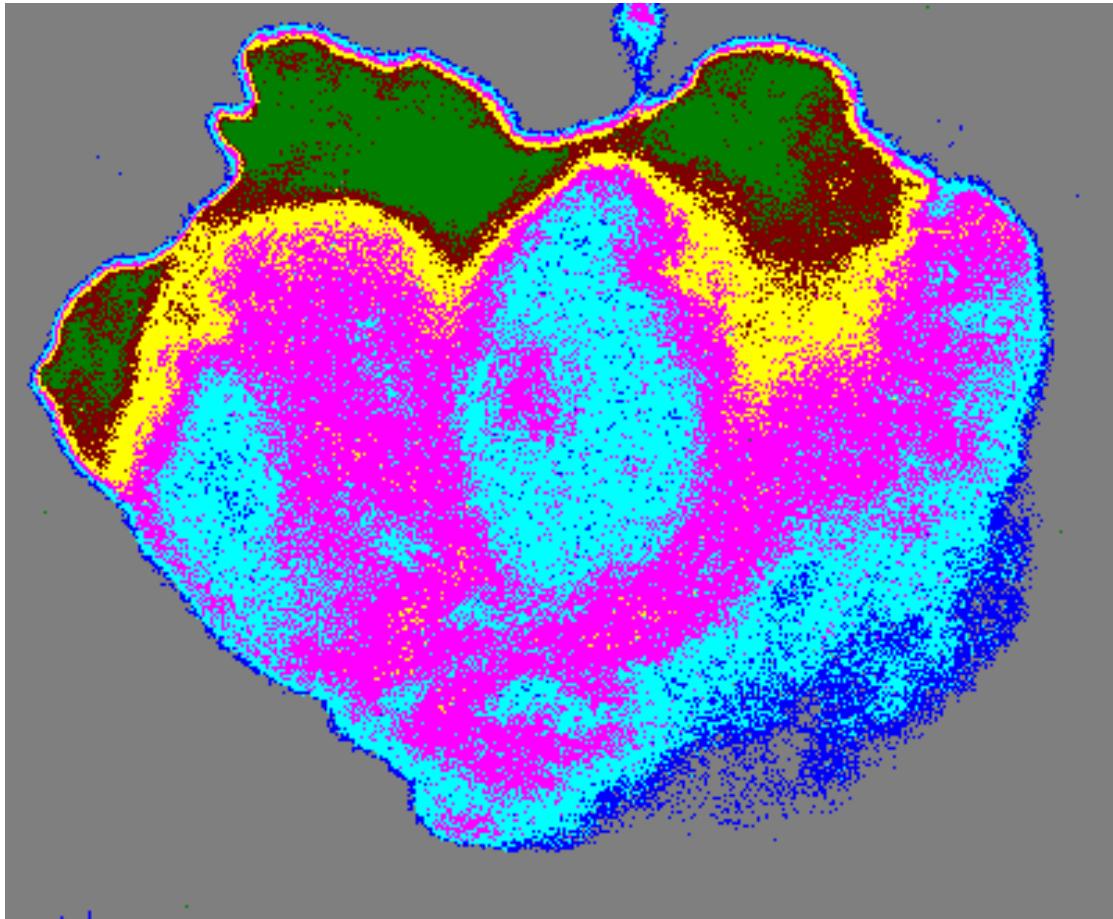


FIGURE 5.39: The multispectral image from figure 5.36 with pixels classified by spectral angle, using the Gaussian peaks from figure 5.38 as class centres. The grey background pixels were not considered during the classification process.

image) can be seen, but again the geometry is recognisably the same. The approximate extent of the node itself is marked in black.

Again, a histogram of spectral angles was created, and this is shown in figure 5.44.

It is apparent from figure 5.44 that there is a strong peak at low spectral angle (centred around 0.035 radians) and this was taken to be the background, the peak fit was thus performed on the region of the histogram greater than 0.05 radians. Seven peaks were fitted to the histogram and the pixels were classified accordingly. The result of the classification can be seen in figure 5.45.

It is very apparent from figure 5.45 that the background pixels in the bottom right hand corner have been classified as biological (node) tissue (group 3). There is a further misclassification regarding the feature at the top centre of the image, which has been classified as groups 7 and 8, consistent with fatty tissue. It was seen in the hyperspectral example, and indeed it can be seen from figure 5.43 that this should actually be node tissue, along with the area to its lower left, which has been classified as groups 5 and 6 (fatty tissue). In other respects this classification has been successful

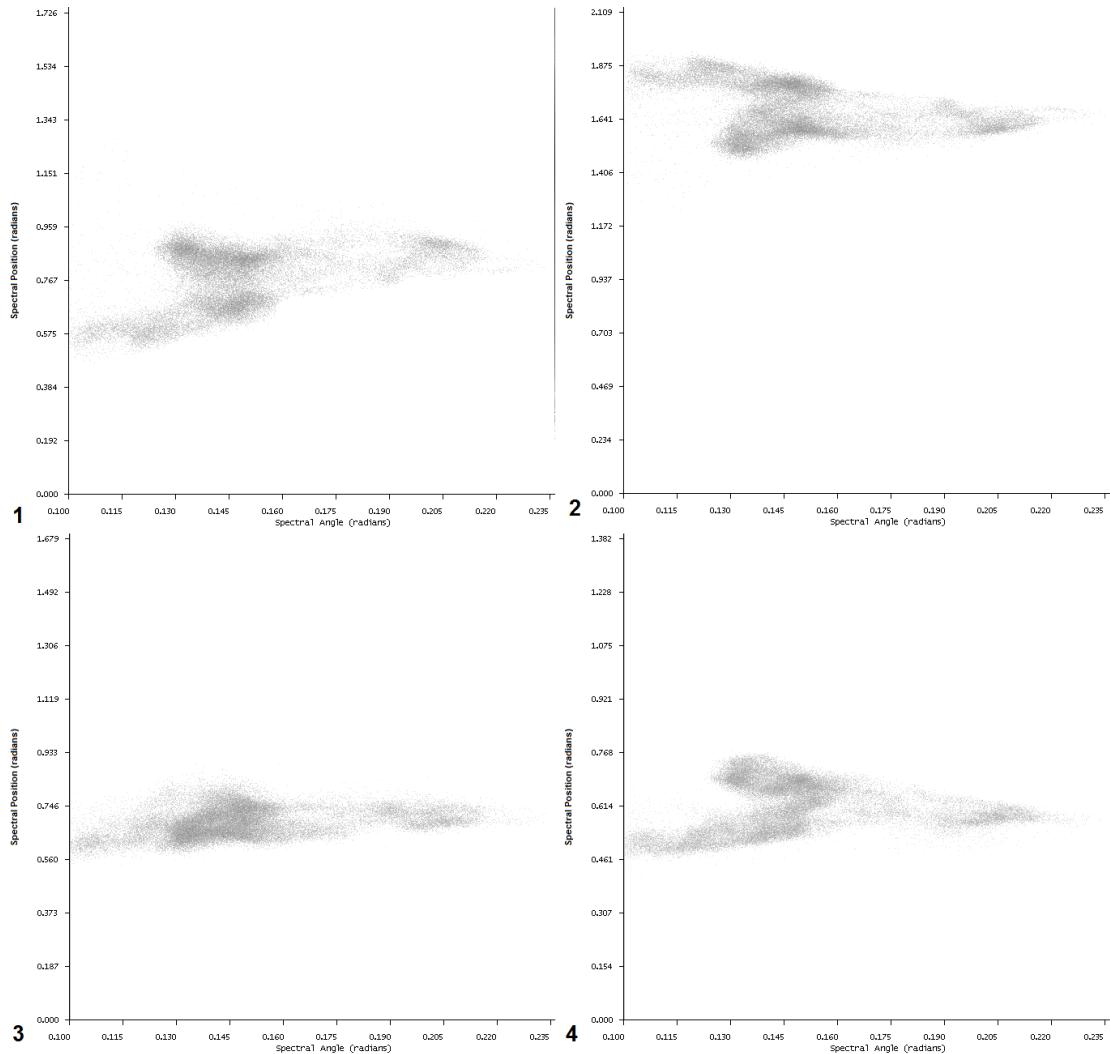


FIGURE 5.40: 2D histograms of spectral angle against each of the the planes of the SPM for the multispectral image in figure 5.36. The plane is noted in the corner of each image.

in that the node tissue (group 4) is well separated from the surrounding fatty tissue (groups 6-8) in a way that is consistent with the separation seen in the hyperspectral image example.

The 2D histograms of spectral angle against the planes of the SPM are shown in figure 5.46. From this it can bee seen that plots 1 and 2 have a high spread in the spectral position dimension at low spectral angles, whereas plots 3 and 4 have higher spreads at higher spectral angles. The misclassifications in this case came from two separate areas, the background misclassification is a low spectral angle problem, and the top central misclassification is a high spectral angle problem. It therefore would appear to be beneficial to approach these two areas using two separate plots; plot 1 was selected to try to resolve the background issue, and plot 4 to resolve the top central issue.

The 2D histogram of spectral angle against the first plane of spectral position is

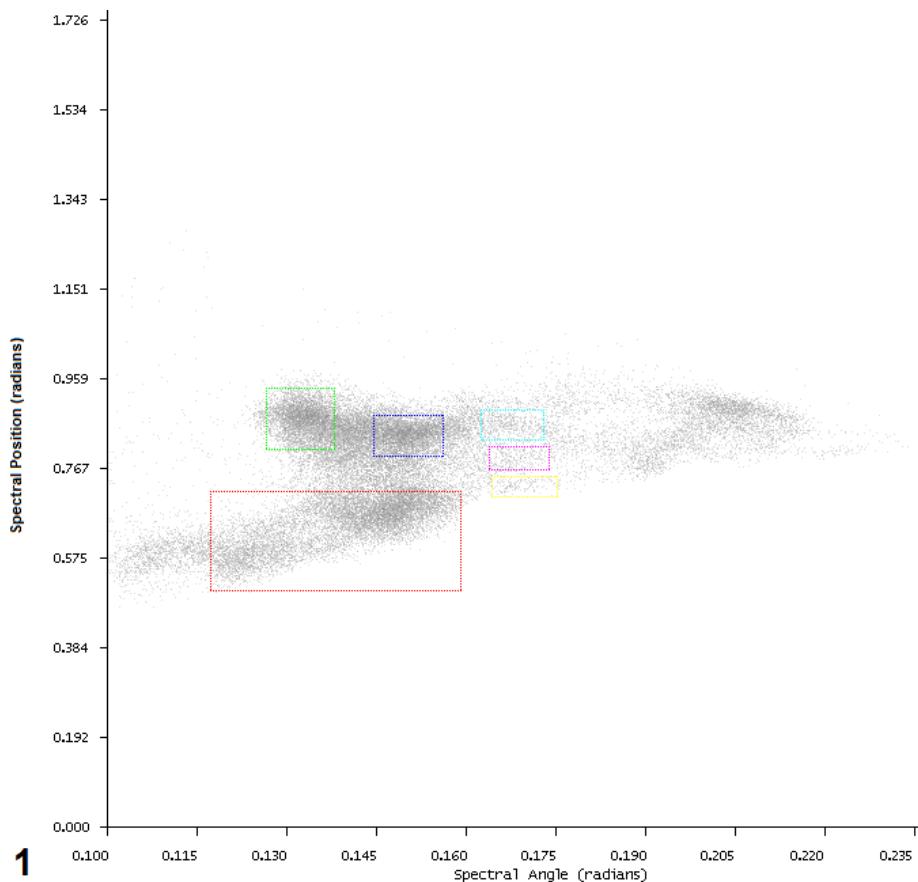


FIGURE 5.41: The 2D histogram of spectral angle against the first plane of the SPM of the multispectral image from figure 5.36 with some regions of interest highlighted by coloured rectangles.

shown, with some boxes highlighting key areas, in figure 5.47 and the spatial distribution of these pixels is shown in figure 5.48.

It is clear from figure 5.48 that almost all of the background erroneously classified as biological tissue is represented by the red pixels, with the green pixels marking the lower right border of the tissue. There is still a small patch of misclassified background in the lower left of the image however, demonstrating that the separation is not perfect. The blue and cyan boxes from figure 5.47 clearly cover the spectral angle range in which the other major misclassification occurs, at different spectral positions. The misclassification is not able to be resolved using this plot.

Figure 5.49 shows the 2D histogram of spectral angle against the fourth plane of the spectral position with some coloured boxes drawn in to highlight the key areas. The spatial distributions of these pixels are shown in figure 5.50.

The red, green and blue boxes in figure 5.49 trace what appears to be a single curving feature in the histogram, covering a large range of spectral angles but a comparatively small range of spectral position. The cyan and magenta boxes cover some of the same

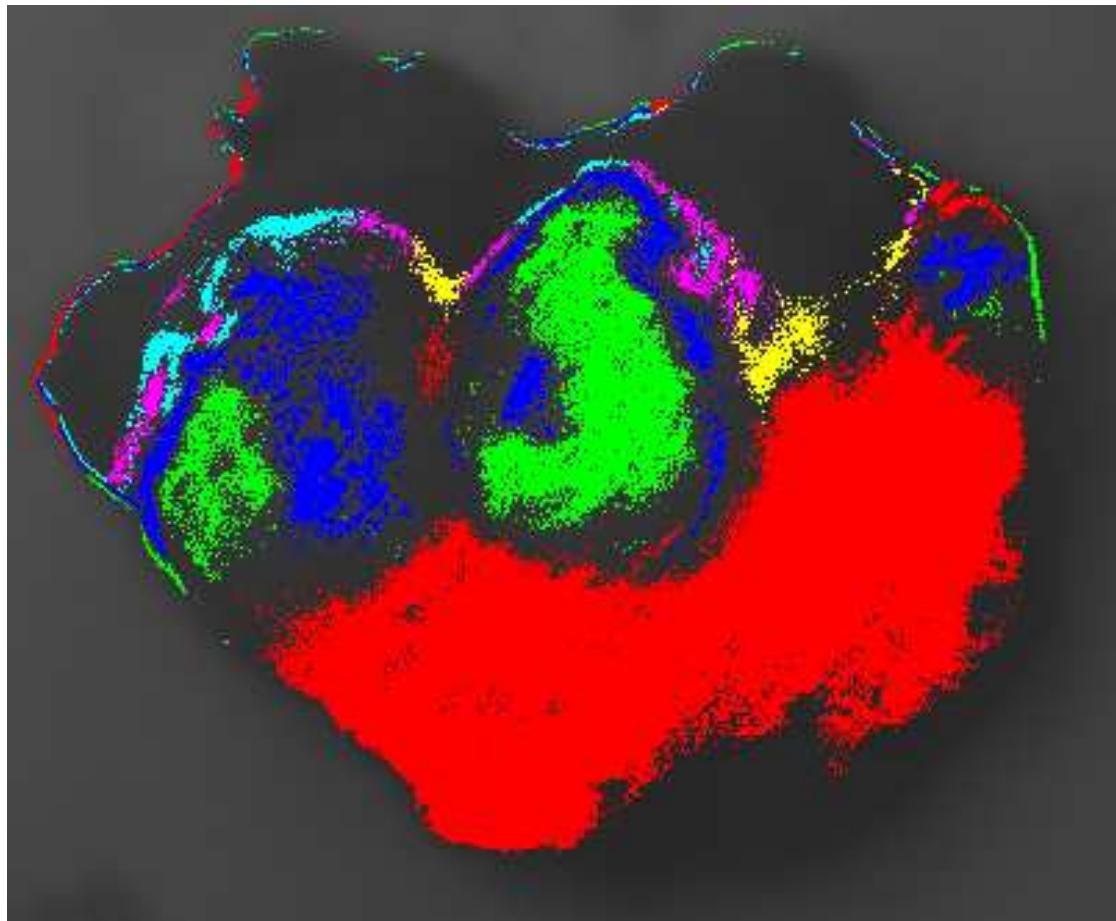


FIGURE 5.42: The multispectral image from figure 5.36 with pixels inside the coloured boxes from figure 5.41 shown in the same colours.

spectral angle range and trace what again appears to be a single feature at a different spectral position. From figure 5.50 it is clear that the cyan and magenta pixels correspond to the fatty tissue whereas the red, green and blue pixels correspond to the node tissue (c.f. the red pixels from figure 5.34).

This example strongly suggests that there is relevant but different information in the various spectral position dimensions. Again, the spectral position has been used in conjunction with the spectral angle to increase the classification ability, and in this example, two different spectral position dimensions have allowed two different misclassification problems to be resolved.

5.4 Conclusions

The test results presented in this chapter again demonstrate how use of the SPM can increase the discriminatory ability of unsupervised classification of hyperspectral images. These data also demonstrate that this is true for the kinds of spectra typically

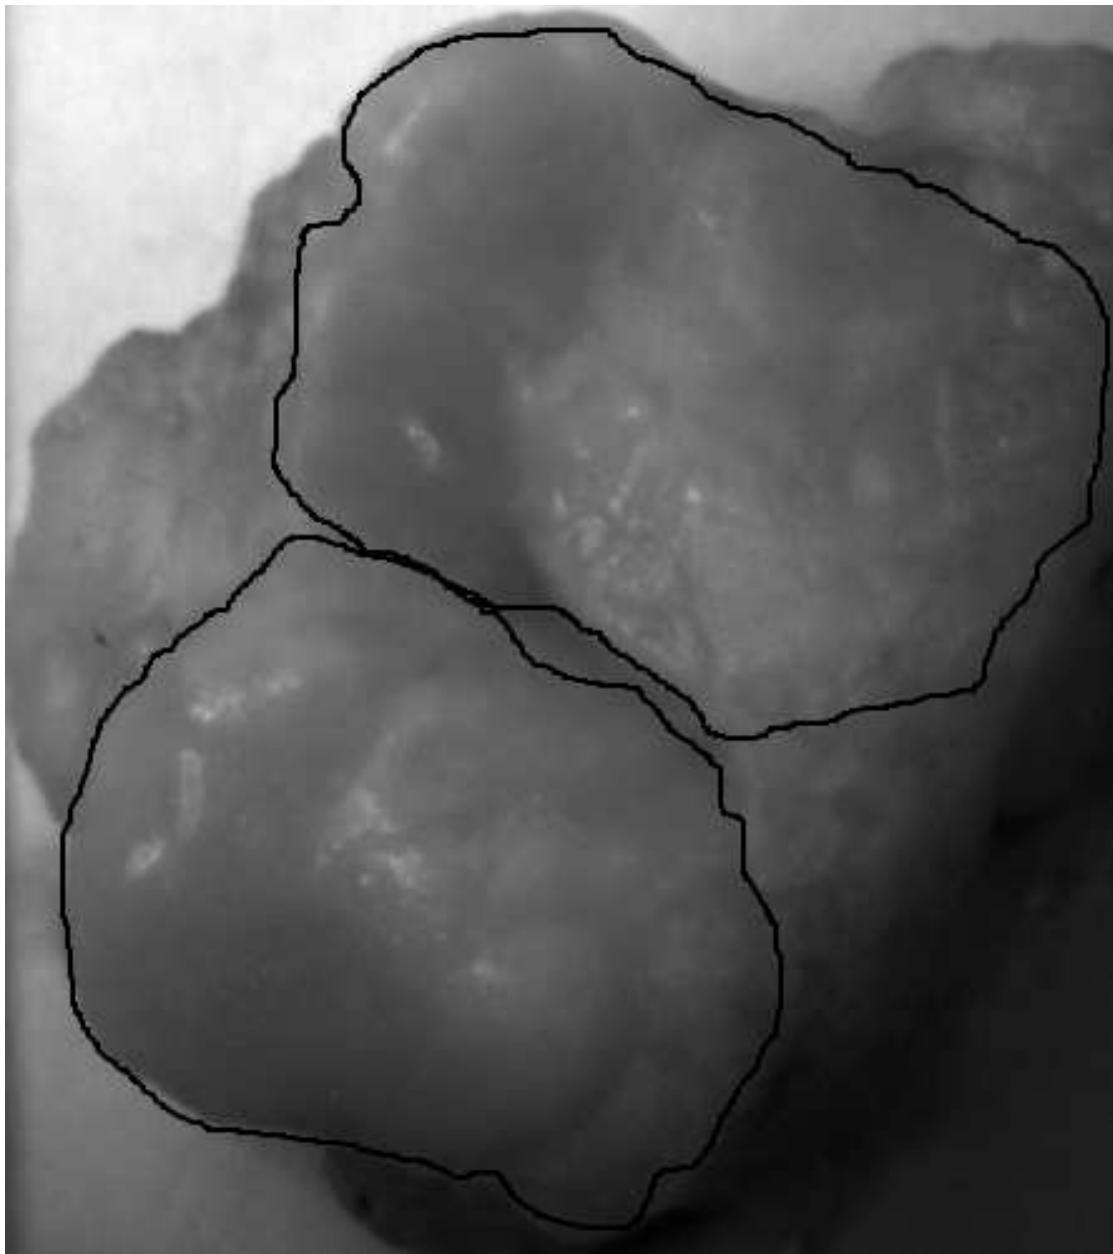


FIGURE 5.43: A monochrome image generated from the multispectral image of a resected SLN. The approximate extent of the node is shown by black lines, marking the two cut surfaces of the bi-valved node, the remainder of tissue is the surrounding fatty tissue in which the node was originally embedded.

exhibited by biological tissue, and more importantly, the specific types of tissue on which the later analysis is performed.

It is further shown that the SPM introduces another strength over spectral angle alone, namely that it appears to be more robust in the presence of noise. In real world applications this noise could exhibit from the imaging system itself, or more importantly, from distortions to the spectra caused by “pixel-mixing”, that is where one pixel in the image covers more than one material type.

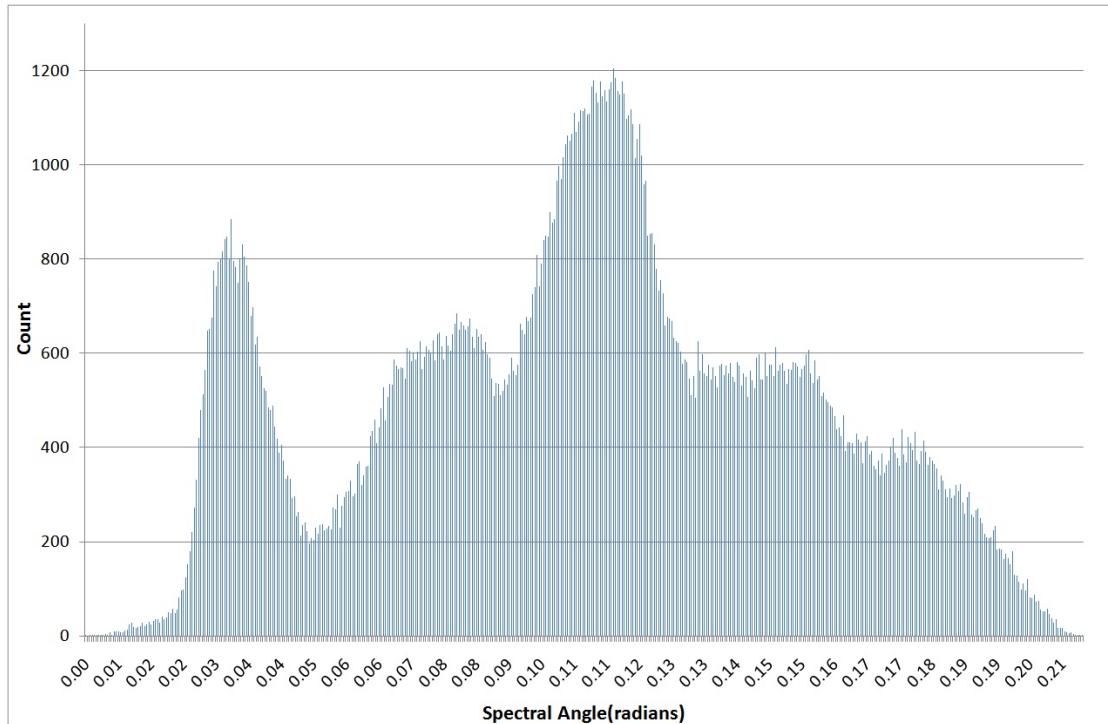


FIGURE 5.44: The histogram of spectral angles present in the multispectral image of the SLN from figure 5.43.

The real results presented in this chapter represent the preliminary work undertaken on actual tissue samples and do not represent the ideal experimental setup or analysis, points which are more fully discussed in chapter 6. They do however demonstrate that the improvement in classification that can be achieved with the use of the SPM, seen with so many test images, is also present for the biological tissue acquired for this work. The SPM is clearly able to distinguish between different types of biological tissue, and subsequent analysis of the average class spectra demonstrate that the spectral differences being detected are real and conform to expectations of the different materials which should be present in the different tissue types.

The SPM analysis method, although clearly in its infancy, is a promising method for the unsupervised analysis of spectral images.

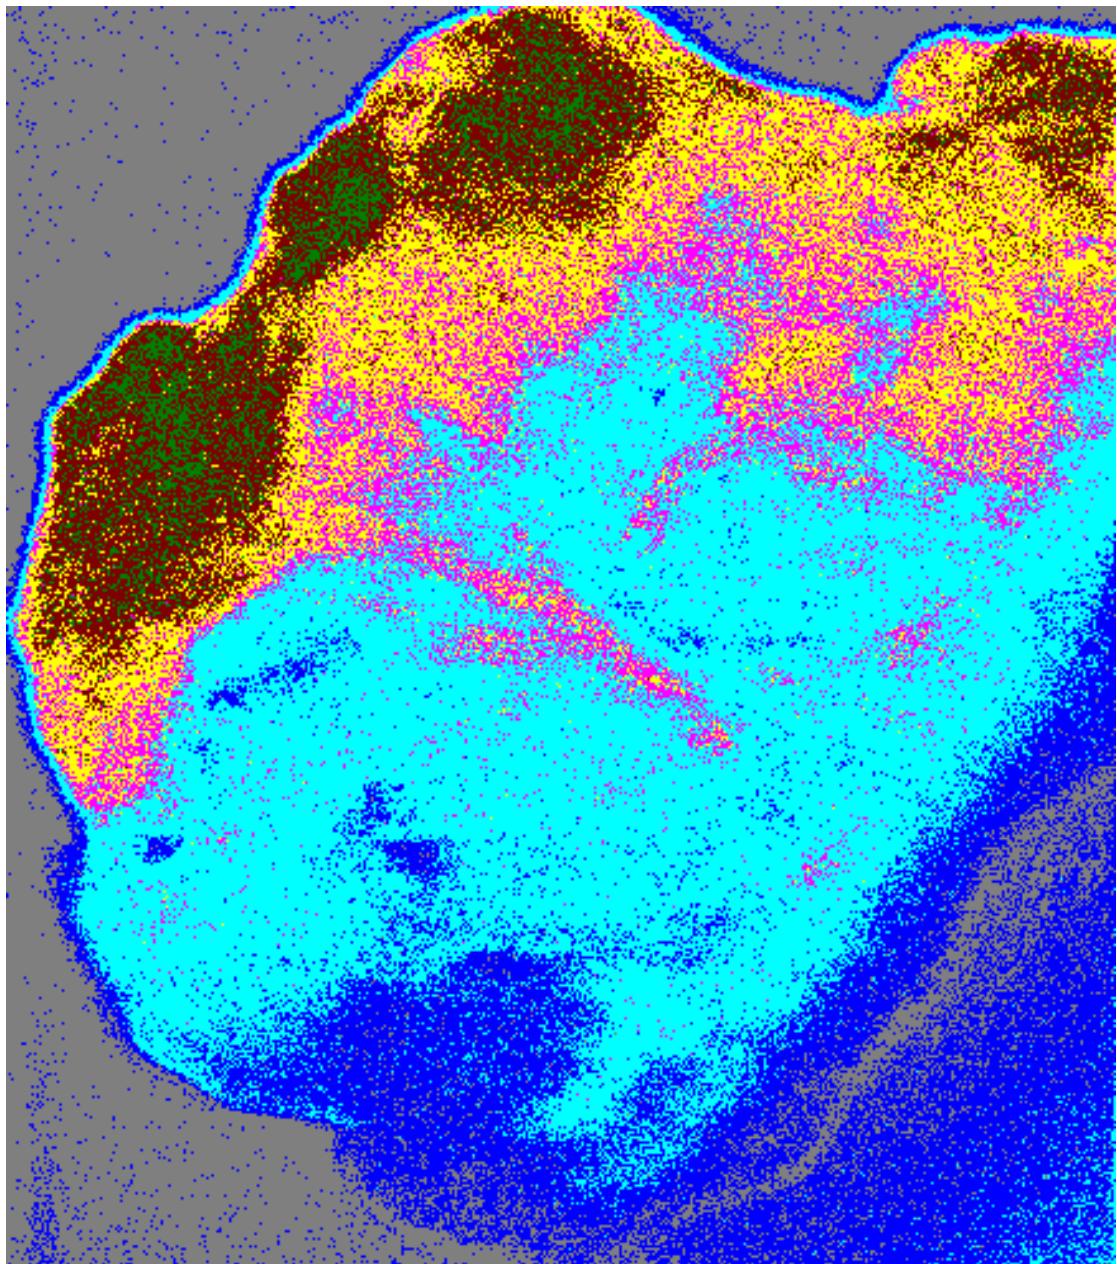


FIGURE 5.45: The multispectral image from figure 5.43 with pixels classified by spectral angle, using the Gaussian peaks fitted to the histogram of spectral angles as class centres. The grey background pixels were not considered during the classification process.

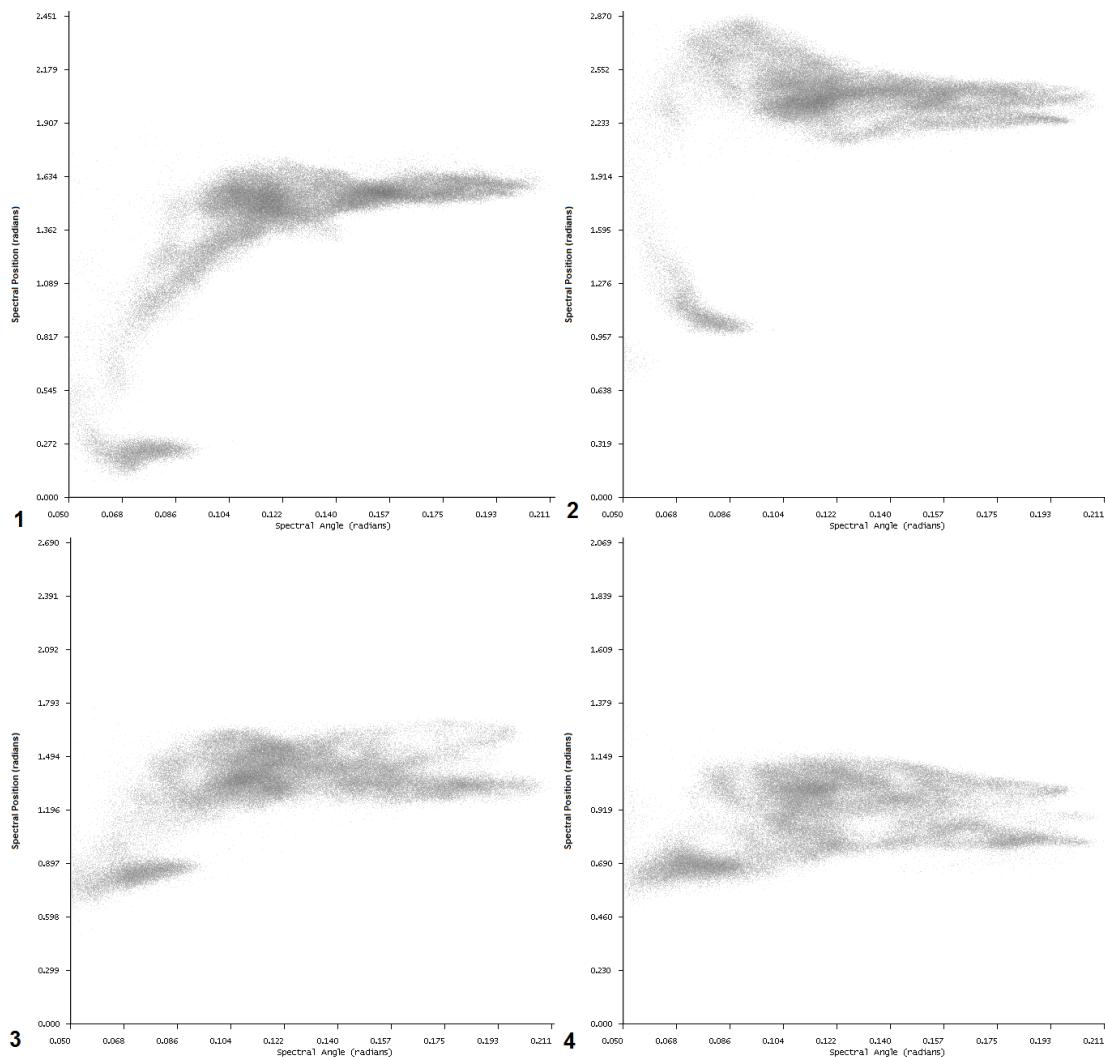


FIGURE 5.46: 2D histograms of spectral angle against each of the the planes of the SPM for the multispectral image in figure 5.43. The plane is noted in the corner of each image.

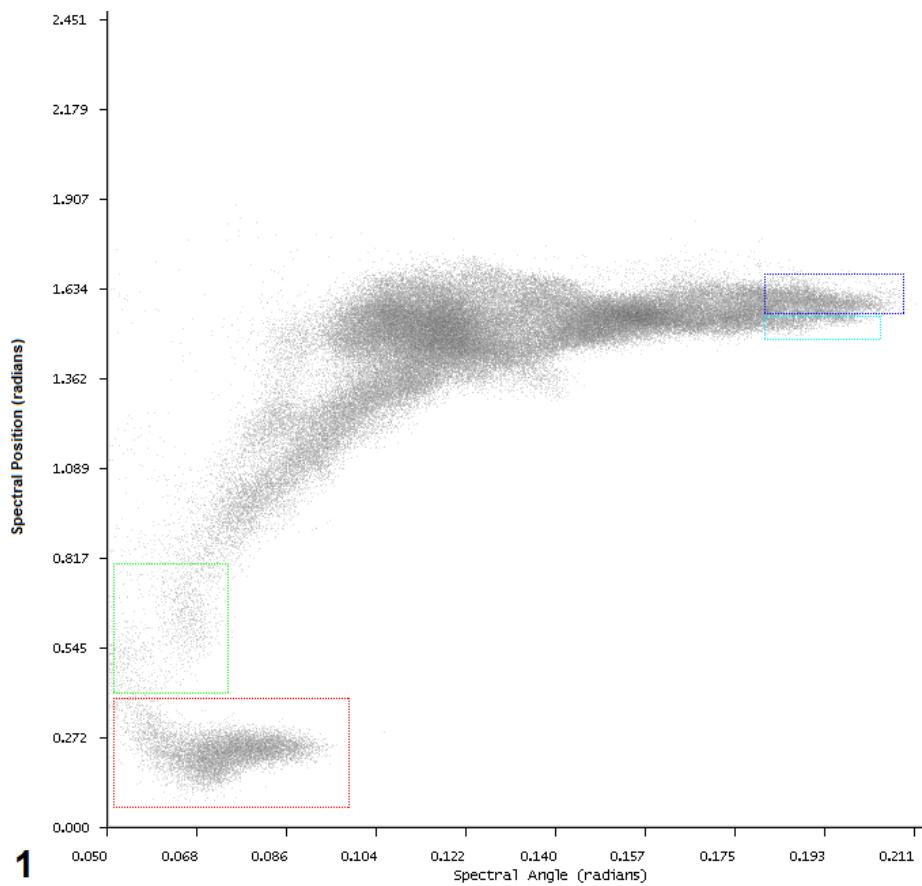


FIGURE 5.47: The 2D histogram of spectral angle against the first plane of the SPM of the multispectral image from figure 5.43 with some regions of interest highlighted by coloured rectangles.

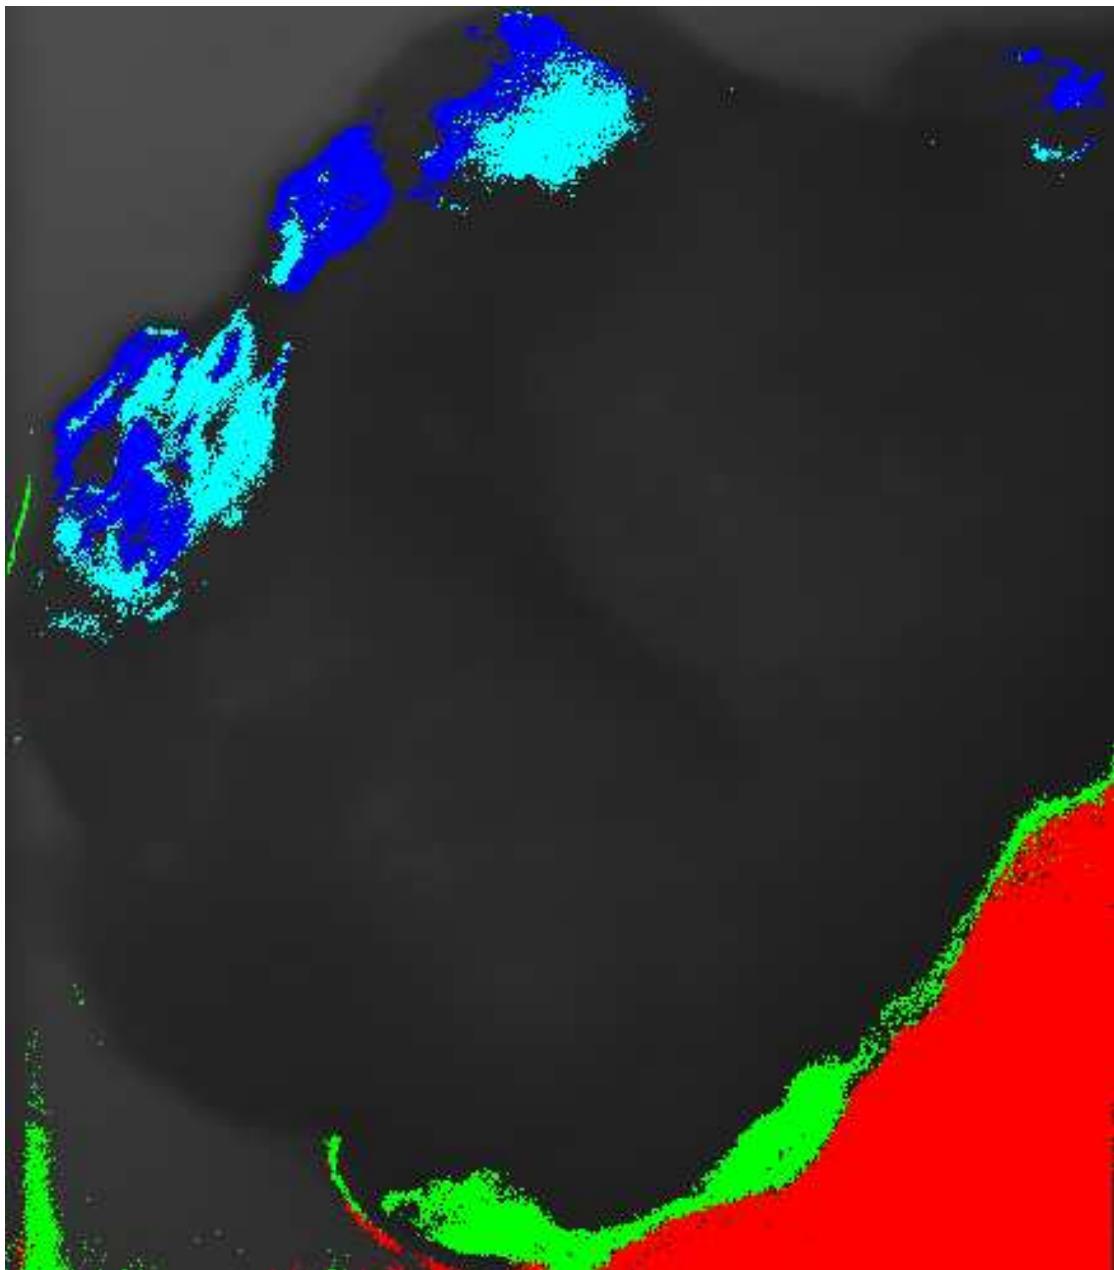


FIGURE 5.48: The multispectral image from figure 5.43 with pixels inside the coloured boxes from figure 5.47 shown in the same colours.

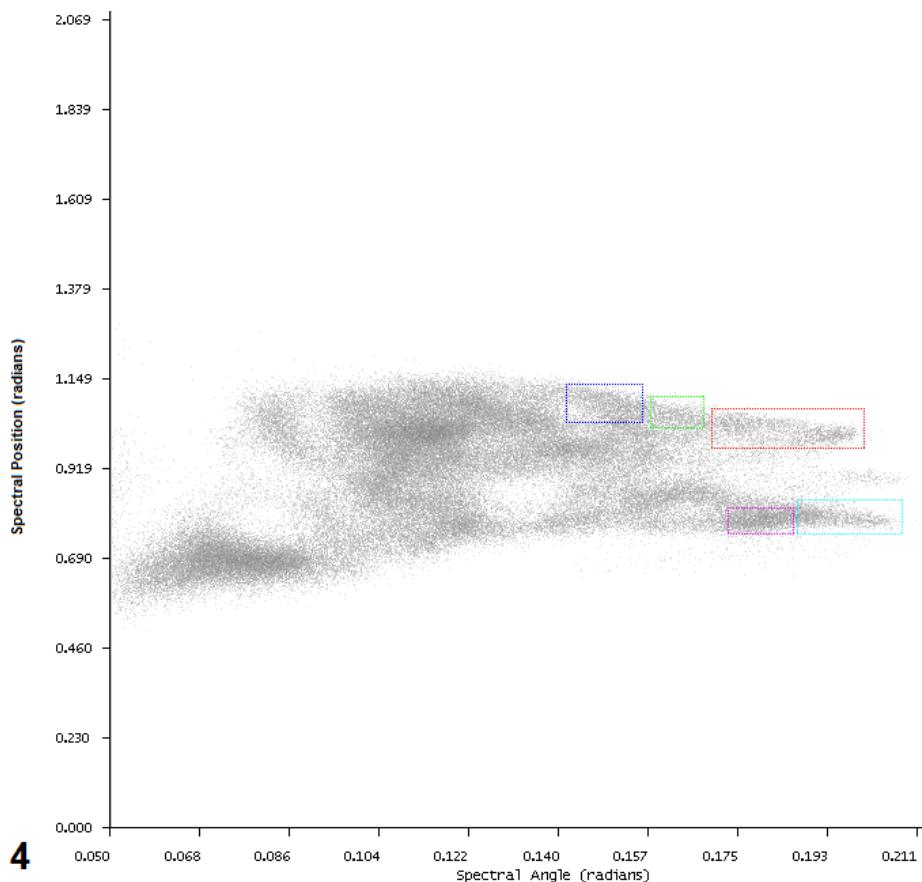


FIGURE 5.49: The 2D histogram of spectral angle against the fourth plane of the SPM of the multispectral image from figure 5.43 with some regions of interest highlighted by coloured rectangles.

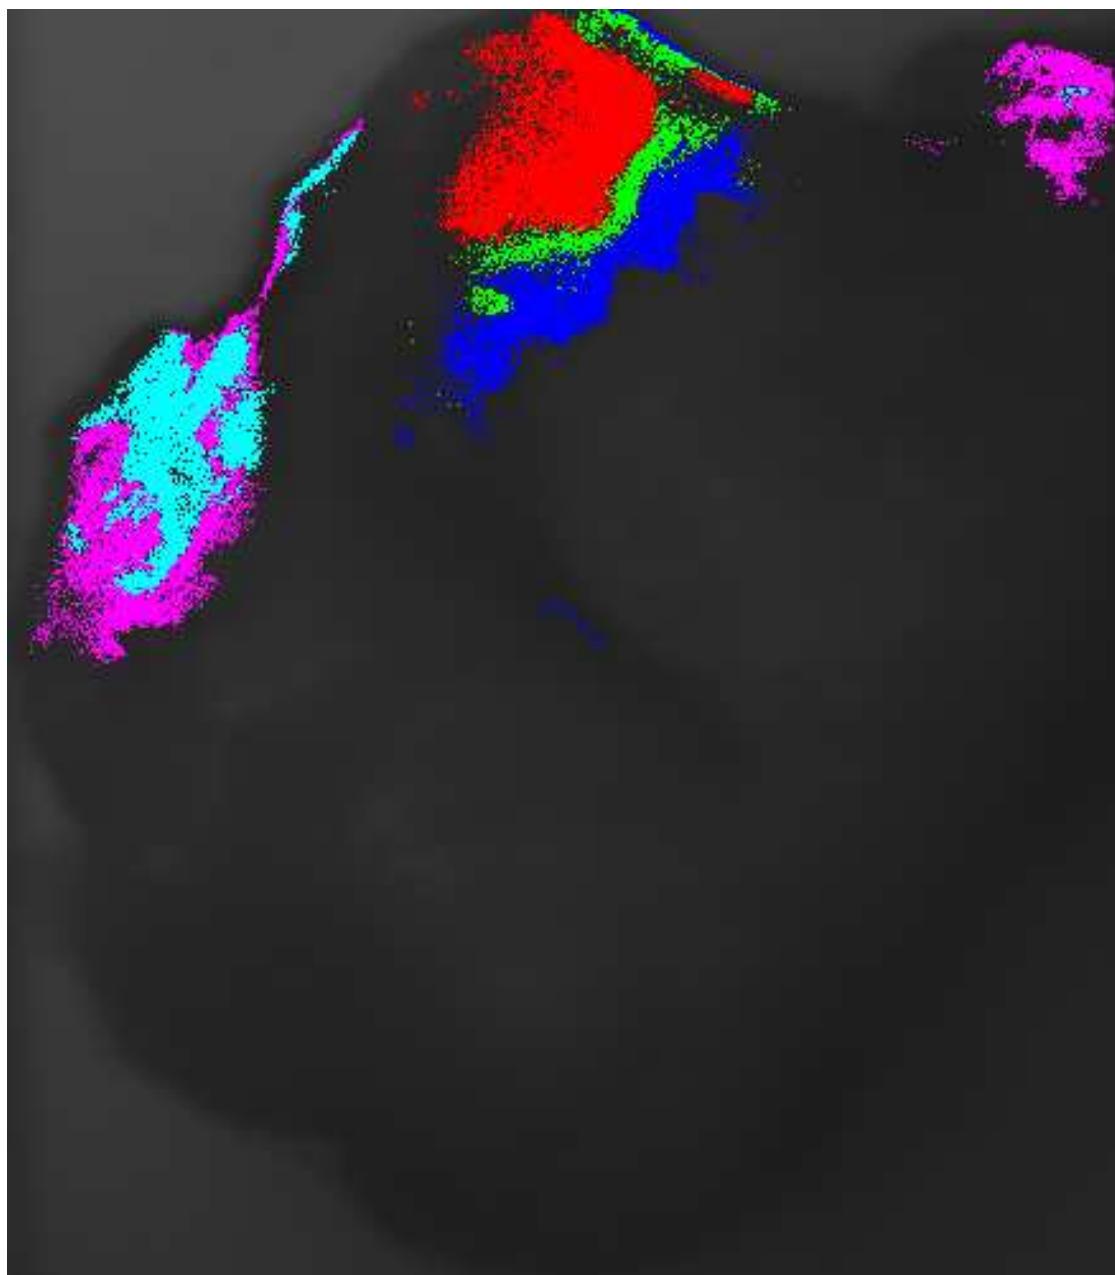


FIGURE 5.50: The multispectral image from figure 5.43 with pixels inside the coloured boxes from figure 5.49 shown in the same colours.

Chapter 6

Discussion

As stated in chapter 1, at the core of this project has always been the hypothesis that if different types of visually similar biological tissue can be distinguished by their VNIR spectrum, then a spectral imaging system could provide benefits to patients and treatments centres alike that are far in excess of the cost of the components required to build it. The purpose of this project has thus been to test the assertion that different types of tissue are indeed distinguishable by their VNIR spectra.

This project was not intended, nor conducted as a detailed research project into the VNIR spectral properties of different types of biological tissue, rather it was intended as a demonstration of principle. The key question to be addressed was not whether or not differences in the spectra of different tissue types exist, or what the physical basis for such differences may be, but rather whether any such differences could be detected by a spectral imager in such a way as to allow a classification of different areas on the image; in short, is spectral imaging a suitable candidate for development of an optical biopsy?

This project has successfully demonstrated the separation of areas of different types of biological tissue on both multispectral and hyperspectral images acquired in the VNIR range. In doing so it has provided strong evidence to support the assertion that spectral imaging can be used as the basis of an optical alternative to the traditional biopsy. This is very much a preliminary conclusion and much work is yet to be done even before such an optical biopsy could be taken to a clinical trial; the discussion of the limitations of the work in this project and the future research that would need to be conducted, is presented here in two parts, technique and application. Section 6.2, discussing technique relates mostly to the work presented in chapters 3 and 4, section 6.3 addresses the application, namely an optical sentinel lymph node biopsy, corresponding to the work presented in chapter 5.

With regards to the detection and treatment of skin abnormalities, it became apparent that the main problem in this work was going to be in the treatment system rather

than the detection system and so this work was brought to a close at an early stage, as discussed in chapter 2. Given the limited nature of this work, it is not intended to present a second discussion on it here.

6.1 A Review of Procedure

The image acquisition procedure for this project was changed very little from the initial procedure laid out by Dr. Hoy for his own project. The focus of the author has been primarily on the data analysis, and in consequence the data acquisition process was somewhat overlooked. This allowed a number of areas for further development to persist; these shall be discussed in this section.

6.1.1 Illumination Issues

6.1.1.1 Use of an AC Lamp

Illumination for all images in this project was provided by standard desk lamps running directly from the alternating current (AC) mains power supply. AC power supplies have current flowing in alternating directions, in the UK this occurs at a rate of 50Hz, in consequence, the current output is a sinusoid at this frequency. The result of this is that the power (proportional to the square of the current) supplied to the lamps in this project varied sinusoidally at a rate of 100Hz; this correction to the power supply is often referred to as the 100Hz ripple. This ripple means that the energy emitted by the lamps in two equal time intervals occurring non-concurrently is not guaranteed to be equal. Only if the time periods are a whole integer multiple of the period of the ripple will the lamps be guaranteed to emit the same total energy.

In order to determine how much of an effect this may have had on the results presented in this thesis, it is necessary to consider two factors, firstly the size of the ripple compared to the total power output of the lamps, and secondly, the period over which any two exposures may be subject to differences.

The first part was achieved by measuring the output of the lamps using a silicon based photodiode sensitive in the visible range of the EM spectrum. The lamp and the photodiode were DC coupled, and the equipment was arranged such that the lamp was the only source of illumination visible to the detector and did not saturate the detector's response. The detector's response is shown in figure 6.1 below. This shows an average illumination level of 9.1V with a clear sinusoidal variation, with maxima and minima at 9.9V and 8.3V respectively. The period of this ripple is 10ms, which is in accordance with a frequency of 100Hz. The ripple is thus shown to cause a variability of up to 8.8% of the average illumination intensity.

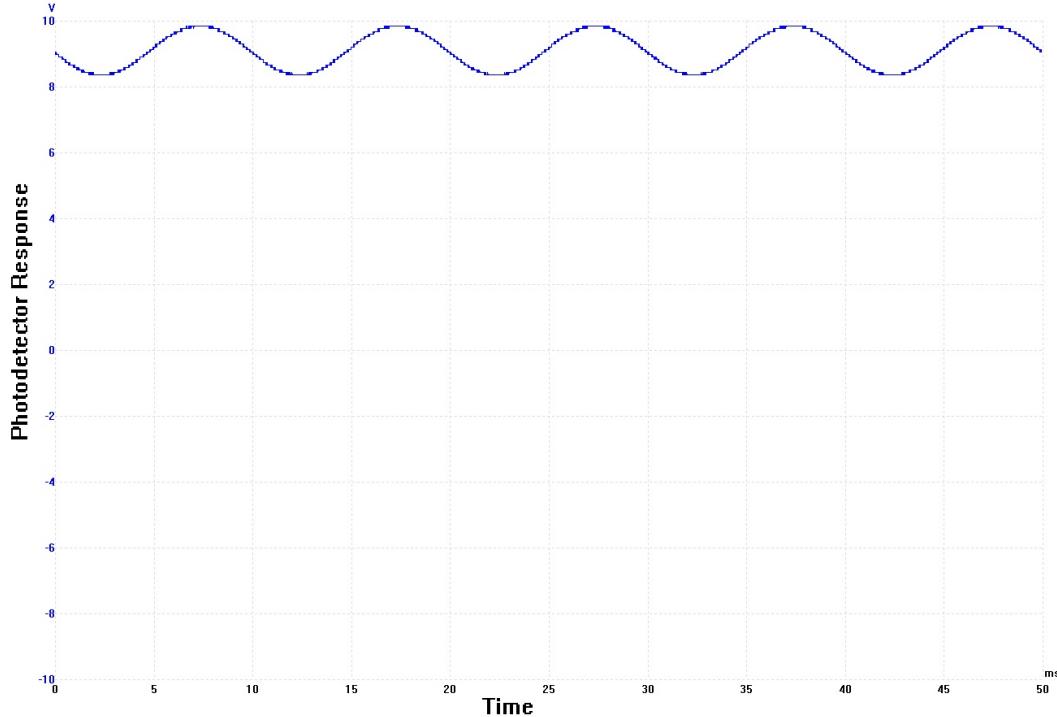


FIGURE 6.1

The second consideration, that of the period over which this ripple will have an influence, is calculated here using representative values as the exact exposure lengths were variable. The majority of exposures were in the range 200-300ms. Using 205ms as a representative value will produce a conservative estimate as to the size of any effect. This exposure length represents 20 complete periods of ripple and a 5ms period over which the illumination can vary. Using a worst case scenario, this 5ms period would occur entirely in the above average half of the period for one exposure and entirely in the below average half of the period for the next.

Using figure 6.1, the curve to integrate to calculate the energy, E (in arbitrary units) can be described as:

$$E = 9.1 + 0.8 \sin(\omega t) \quad (6.1)$$

where t is the time in milliseconds and $\omega = \frac{2\pi}{10}$, and 9.1 and 0.8 are the average and peak illumination voltages.

Calculating this integral between 0ms and 5ms, and then between 5ms and 10ms gives $\sim 48\text{msV}$ and $\sim 43\text{msV}$ respectively. That is to say, that compared to the 45.5msV average illumination, there is a 5msV ($\sim 11\%$) variation. Of course, it must be remembered that this is an 11% variation in a period representing $\sim 5\%$ of the overall exposure length. Over the two 205ms periods, the total energies are

$(20 \times 91) + 48 = 1868\text{msV}$ and $(20 \times 91) + 43 = 1863\text{msV}$, the total difference is thus $\sim 0.26\%$ of the total energy. Again, it must be remembered that this is the worst case scenario, for a conservative representative exposure length, and this is before considering that for the hyperspectral imager, each exposure was the result of co-adding 16 separate exposures.

Given these considerations, it can be safely assumed that the 100Hz had a negligible effect on the results presented in this thesis; nevertheless, for increased scientific rigour in future experiments, it would be desirable to use DC powered lamps.

6.1.1.2 Fluorescent Lighting Artefacts

As mentioned in section 5.3.2, the images presented in this thesis suffer from artefacts of the hospital's fluorescent lighting system. Efforts were made to remove the influence of the light source from the images by normalising them against a standard background, however this removes only multiplicative effects. The artefacts remaining are additive spectra resulting from a small amount of the light being specularly reflected by the sample. This can, and in future work, should be eliminated by constructing a dark box in which to place the sample, carefully controlling the light that enters to illuminate the sample.

Another way to eliminate, or at least reduce such specular reflections is the use of polarising filters. Figure 6.2 demonstrates how specular reflections can be greatly reduced by the use of a polarising filter. Moreover, light that is specularly reflected at a surface maintains its original linear polarisation state (Anderson, 1991), so by using a linearly polarised illumination source, and by imaging through a polarising lens perpendicular to that polarisation, surface reflections can be eliminated entirely (or at least, to the effectiveness of the cross-polarising filters). This has been demonstrated and reported by Morgan and Stockford (2003).

In this work, polarising filters were available for the multispectral imaging system and these were used chiefly in early images where regions appeared to be saturating on the image, or where there was suspected glisten. In these cases, differencing the images taken through two, orthogonally polarised filters, showed up specular reflection as the only, or principal, difference between the images. This problem appeared to have been resolved by using more carefully diffused light to illuminate the samples, and so unfortunately the polarised images were largely neglected thereafter. The quality of the data available could be improved by taking all images through cross-polarisers as described above and it is recommended that any further work implement this improvement.



FIGURE 6.2: An image of a car windscreen with and without a polarising filter. It can clearly be seen that without a filter there is reflection of the illuminating source (in this case, the sky in general) from the windscreens, when a filter is used the reflections are no longer evident and the interior of the car can be seen. Further, the window in the background, which is approximately perpendicular to the windscreens, behaves in the opposite manner, with reflections only visible through the polarising filter (original image available at <http://www.flickr.com/photos/aidanwojtas/2252215670/?q=polarising%20filter%20examples>, images reproduced with permission of the owner through a creative commons licence).

6.1.2 Image Calibration

A further area in which the image acquisition for this project was out of step with what would be considered as best scientific practice is the gain and offset calibration of the acquired images. As detailed in section 5.2.2.1, all hyperspectral images were pre-processed by removing the influence of known “hot” pixels and then further by normalising with respect to a neutral background. The multispectral images were registered spatially and cropped, but otherwise underwent no pre-processing.

The pre-processing would benefit from the measurement of a dark frame at the beginning and end of each set of measurements. These dark frames can be compared to assess any change in the background noise of the detector over the measurement time. Interpolating between the two measurements, an estimate of the noise level at the acquisition time for each exposure can then be calculated and subtracted from the image.

The main source of error that this extra processing accounts for is changes in the background noise of the CCD with temperature. The multispectral imaging system used a camera with a cooled CCD, which kept the detector at a constant temperature throughout operation, thus the absence of such dark frame measurements is unlikely to have had a significant effect on the multispectral images. Although these dark frame images were not taken for the hyperspectral images, some effort was made to minimise the effect of CCD temperature in that the camera was only powered on for the minimum period required to take a set of exposures. There will, however, inevitably have been a benefit to measuring and accounting for this noise directly.

As detailed in section 5.2.2.1, measurements were made of the readings of the pixels in dark frames, and an average count of around 0.3% of the maximum signal value was found to be typical. The images in this thesis were exposed so as to try to produce an average pixel value of around 25% of the maximum signal, so that the brightest pixels would not be saturating the detector. This means that the noise level due to the CCD was in fact higher than the 0.3% and would typically have been more likely to have been in the 1-1.5% range. Again, this would have been mitigated to some degree by recording the average of a number of exposures, however, the uncertainty of exactly what noise level may have been present is an inherent weakness in the acquisition process that should be corrected in any further work.

The images from the multispectral imager would also have benefitted from a full frame normalisation against a neutral background under the same illumination conditions, as was performed for the hyperspectral images. This would remove any fixed pattern noise from the detector and illumination. Again, this should be implemented in any future work.

6.1.3 Appraisal of the System with the Analysis Algorithm

For best practice in the scientific method, the use of the analysis algorithm in conjunction with the imaging system used should have appraised and characterised with a well understood target. Unfortunately, the images to do this were not acquired during the project and it has not been possible to acquire them in the intervening period. However, it is not felt by the author that this represents a significant detriment to the work presented in this thesis.

The imaging systems used for this project were themselves well understood and well characterised. The cameras used were commercial, “off the shelf” silicon CCD sensors, with well known and published spectral profiles. As detailed in section 5.2.2.1, any fixed pattern noise from the Prosilica CCD was removed from the final images by a full frame normalisation, and whilst this was not the case for the Alta CCD, observationally this did not appear to be a problem.

The Specim hyperspectral imager was also a commercially available, standard piece of equipment, whose spectral response was characterised by the author, as described in section 5.2.2.1. The spectra recorded using this system were entirely in keeping with what was being imaged, as demonstrated in figure 5.8 (the illumination from a standard fluorescent light source and an argon lamp being well understood).

In short, the images produced for this thesis originated from sources whose performances are well understood.

Chapter 4 is spent trying to understand and characterise the performance of the analysis algorithm developed for this project. The images used in that chapter are of well understood targets and originate from a number of different hardware and software sources. The software developed for this project is entirely agnostic as to the source of images, all images are merely regarded as vectors of data to be manipulated, and there is no indication that the source of the images has any bearing on the performance of the algorithm. Given that the algorithm and the image acquisition systems are both well understood, whilst it would be preferable to have some separate appraisal of the two working together, it is highly unlikely that any new or interesting information would be presented by this.

6.2 The Extended Spectral Angle Metric and Classification System

The separation of different tissue types was achieved by developing a classification system based on a well established distance metric, the spectral angle. This classification system was unsupervised to the extent no assumptions were made about either the number of classes or the sizes of the classes (both absolute and relative). These were blindly calculated by the software on the assumptions that different spectra would subtend different spectral angles, and that noise in the system would exhibit as Gaussian curves in a histogram where lines would be expected in an ideal case. This second assumption was tested using dummy spectral images and found to hold true (see section 5.3.1).

Spectral image processing covers a diverse range of techniques and there are a number of ways in which to proceed. The choice of technique is largely determined by the application and for this project some form of classification, i.e. labelling every pixel in the image, was an obvious approach. This project was further constrained to requiring an unsupervised approach because, unlike with some applications of spectral imaging, a reliable library of spectra that would be anticipated to be required, does not appear to exist in the literature. To obtain such a library would involve a project that was very disruptive to the normal course of surgery and, as discussed below in section 6.3, this work sought to avoid such disruption wherever possible. It is possible that a future project could address this.

In order to perform a classification there has to be a measure that is used to compare pixels, either to each other or to references. In basing the classification on a distance metric, the multidimensional dataset created by the spectral imager was reduced to a dataset in a single dimension. Such a reduction is necessarily lossy (i.e. the original data cannot be reconstructed from the reduced data) and is a many-to-one transform, inherently creating degeneracy in the metric space. Reducing a 183D hyperspectral

vector to a single number represents a significant decrease in the data content and inevitably a significant loss of information. The data reduction is easily quantified, the SAM in this case is 0.55% of the size of the full spectral image, however it is not so easy to quantify the information loss because the data has not simply been discarded to create the SAM, each datum has been considered in a calculation resulting in a single number. This loss of information presents in the system as misclassifications.

The clear simplicity of using a single metric is ultimately outweighed by the limited performance of the classifier, as can be seen from the examples in this thesis. The extension to the spectral angle metric proposed in this project (and referred to as the spectral position) is the multidimensional analogue of the azimuthal angle in three dimensions and reduces the dimensionality of the dataset by one (when considering it separately from the spectral angle). Whilst it is simple to conceptualise and calculate, it is much less clear exactly how it should be used, particularly in hyperspectral cases where the dimensional reduction is much less significant than in a low dimensional multispectral image.

The obvious extension to the classifier already developed is to extract a single measure from the SPV as use it a complement to the spectral angle in calculating a 2D histogram from which to discover classes. This has been done in this project in two different ways dependent on whether the spectral image was originally multi- or hyper-. The multispectral method is based on the examination of principal component score plots (for an explanation of which see Grahn and Geladi (2007)) and currently requires supervision in so far as the final histogram to be used for classification is chosen by the user. In an attempt to remove this supervised step however, it is possible to base this decision on an objective and quantifiable measure of overall contrast (see section 4.3). The hyperspectral method uses the high dimensionality of the SPV itself to create this secondary measure, which is the plane of the SPV that has the largest absolute value.

6.2.1 Future Directions

The 2D histograms created in each case appear to be informative. Structures such as peaks and “watersheds”, as described by Grahn and Geladi in their discussion on score-plots (Grahn and Geladi, 2007), appear in the histograms and it is clear that the data points are not randomly scattered in the non-spectral angle dimension. Further, these structures correspond to pixels that are spatially proximate, or from areas of similar origin when mapped back into the original spectral image.

These 2D histograms provide enhanced separation of different areas of the spectral image compared to the classifier based on spectral angle, significantly improving the classifier’s performance. To date this enhanced classifier is supervised to the extent that only areas selected by the user are mapped from the histogram to the spectral

image, and these areas are limited to rectangles. The logical next step in the development would be to automate this process to provide complete and unsupervised classification. There are two ways in which to do this, which would need to be investigated. The first would be to attempt to fit 2D Gaussian peaks to the histogram in a straightforward extension of the current method, the second would be to attempt to use the apparent structures as the basis for classes.

The second option here would seem likely to be the most difficult, however there are suggestions in the results presented in this thesis that it would also be the most profitable. In figure 5.49 the groups selected trace structures in the histogram rather than attempting to pick out neat peaks and it can be seen in figure 5.50 that similar spectra cover a wider of spectral angles than might naively be assumed but that they are similar in respect in spectral position.

An alternative development of this method would be to use it only where misclassification is suspected when using spectral angle alone. In this system the spectral angle range of the histogram that needed to be considered would be vastly reduced and far fewer structures would need to be considered or identified. Indeed, it is even possible that a group suspected to contain multiple types of spectra could be classified in a similar way to the whole image, with a 1D histogram of spectral position being generated for the small spectral angle range in question, and fitted with Gaussian peaks as before. This would seem to be a quicker and simpler fix for misclassifications, however if the structures in the 2D histogram are indeed discovered to be relevant and have an underlying cause, then this method would suffer from not being capable of using the entire structure in many cases. Further, it would introduce a supervised stage into the classifier, running contrary to the long term aims of the project as a whole.

There is also an open question as to when images should start being considered and treated as hyper- rather than multi- spectral. As mentioned in the opening chapter, the distinction between the two is not precise or clear cut, however it seems clear that in a small number of dimensions the “dominant plane” method outlined for hyperspectral images would simply not provide the requisite range for interesting structure to emerge. In other words, this thesis suggests that there is a real difference between multi- and hyper- spectral images, particularly pertaining to their analysis, that goes beyond merely nomenclature. If indeed this is the case, then it seems likely that the spectral position would be important to understanding the difference and maybe to providing an indication of where the difference occurs. It is apparent that the spectral position should be examined mathematically in a way that is much more rigorous than the time constraints of this project have so far allowed.

The further work in this section has been presented in a pragmatic order, the spectral position appears to enhance the performance of the classifier presented in this thesis and so developing it to provide a more automated, less subjective classification system

appears to be a sensible course of action. However, it might be that a more thorough understanding of it would alter the way in which it is applied and in planning any future development of this work it would seem sensible to give this a high priority.

The unsupervised classification using the extended spectral angle measure has been developed in response to a specific problem posed in a medical context, however it is an entirely generic technique and may be of use to researchers in spectral imaging in any field and in any spectral region. As such it is not important that any future research discussed above be carried out in conjunction with or even with reference to the specific application that will be discussed in the next section.

6.3 The Optical Sentinel Lymph Node Biopsy

The central problem posed by the SLNB is that the time taken to perform a standard lumpectomy is determined not by the difficulty of the surgery itself, but by the external factor of the determination of lymphatic involvement. This forms part of a wider discussion in medicine today on the efficacy and reliability of histopathological assessment in general. The job of the pathologist is to look for certain, well defined characteristics of the cells on a slide that indicate their health, or otherwise; these characteristics typically being the shape of cells, their distribution and/or level of order, and the colour which they stain. It is acknowledged that the decisions made about a given slide can show not only inter-observer, but also intra-observer differences. As is often the case when human interpretation is required, two different, equally trained and skilled observers may come to a different conclusion about the same piece of evidence, but worse, a single observer may come to different conclusions when viewing the same evidence at different times, depending on a range of variables including their tiredness and the context of previously viewed cases. It is easy to imagine that a borderline case may be seen as normal if it comes after a succession of diseased cases, but diseased if it comes after a succession of normal cases.

It is now possible to capture microscope slides as very high resolution digital images, known as virtual slides, which can be up to 200,000 dots per inch (for comparison an ink-jet printer will produce prints at between 300 and 600 dots per inch). In one exciting research project, the University of Leeds has built a 50 megapixel “powerwall” display to allow these slides to be examined (http://www.virtualpathology.leeds.ac.uk/research/HCI/Powerwall/virtual_reality_powerwall.php). The obvious development, particularly given the digital nature of the slides, is to automate the analysis so that a suitably developed computer algorithm searches for the well-defined characteristics and makes a diagnosis by comparison against some metric. The cost in terms of skilled operative time could be drastically reduced if such a system could be

introduced, however no such system is yet able to match the performance of a trained pathologist. The aim and purpose of this work is substantially similar.

The results presented in this thesis represent work from the very early stages of development of an optical biopsy and as such they are rather limited in detail. Both the data acquisition and the analysis have been at experimental phases for this project and it is clear now that there is a significant limitation in the conclusions that can be drawn from the results so far.

Each lymph node examined for this project was identified only by reference to an anonymised patient number, starting at “01” and increasing sequentially, and the date and time at which the images were captured. This system has allowed the data to be stored anonymously, thus protecting patient confidentiality, whilst also providing enough information to the clinical staff at the hospital to be able to identify the appropriate patient records. It was not felt that there was a need to maintain a strict blindness to the actual metastatic state of the nodes in this project due to the exploratory nature of this work, and so results from cytology are available for each node examined, and have been throughout the project. However, this result comes in the form of a binary indicator, the node is either regarded as having tested positive for the presence of cancerous cells, or negative. Details such as the extent of the metastasis or the spatial distribution throughout the node are not available.

This lack of detailed information presents a problem when assessing the performance of the classifier developed, the “ground-truth” of the images is not sufficiently detailed to fully appraise the classification.

At the outset of this work there was no clear roadmap as to how the images would be analysed for the presence of cancerous cells. The work with the dummy images (see section 5.3.1) demonstrated that in the binary case of having just two spectral types, the differences were clear and easily identifiable. It might have been hoped that carrying this over into the real images some clear indication of cancerous cells might have manifested as a readily identifiable peak or structure in the 2D histograms that only occurred in those images relating to nodes that tested positive. This appears not to be the case and the histograms for all images look broadly similar. It is still possible that actually some identifiable feature may still occur, it is anticipated that in any given positive node the number of pixels covering cancerous cells would be small compared with the size of the image and as such any feature in the histograms would be faint and easily missed.

6.3.1 Future Directions

The next key question to be answered is can this method definitively distinguish between normal node tissue and cancerous tissue? It is clear that in order to answer

that, this spectral imaging process will need to be far more integrated with the current pathological tests than it has been thus far. This means that this work will become more disruptive to the overall current procedure in that it will require close involvement with the pathology staff likely to require significant investments of their time. This is something that has been avoided up until now on the basis that with such an experimental and untested system, there should be an absolute minimum of disruption to the normal levels of care and operations of the hospital. Further, there is an extra burden of ethics approval for work of this nature, and the more disruptive the process, the more difficult it becomes to gain approval. However, for any further progress to be made with this work it is vital that the spectral images can be directly correlated with the diagnoses made from the cytology examinations. A direct mapping between spectral images and these examinations is the only way to establish fully labelled images to assess the classification.

Exactly how this would develop would need to be discussed in more detail with the clinical and pathology staff in order to find a system that could be easily accommodated, however some possibilities can be outlined already.

The spectral images taken for this project are essentially the reverse of the slide imprints as they are already assessed and so it is possible that this mapping could be easily achieved by simply providing the images and some means of labelling them to the pathologist along with the slide. However, human tissue is not rigid and so there will be significant spatial deformations as it is imprinted against the slide meaning the mapping is not exact.

It is also possible that after running an unsupervised classifier based on the spectral angle/spectral position methods outlined in this thesis the classified images would be examinable and interpretable by the pathologist in much the same way as the slides are currently. This would essentially bring this project to the same point as the University of Leeds' virtual pathology project in that the pathologist is performing their traditional job but on a digital format image.

It is possible that the best approach might be to not image the node directly as is done now but to image the slide that is examined by the pathologist. Again this is very much related to the virtual pathology project.

A final alternative is that the spectral image be taken of the tissue through the slide as it is being imprinted. This would mean that the spectral image and the assessed slide would have a direct and exact mapping.

Of course, it is possible that an entirely new way may present itself in the course of discussions with clinical experts. Whatever course of action is ultimately pursued, the key development is to know with confidence exactly where different types of tissue are expected on the image.

This thesis has discussed working with both multi- and hyper- spectral images and a natural question to consider is which of these systems should be pursued further. Both systems represent a compromise from the ideal situation; in the case of multispectral imaging, the compromise is an increase in image acquisition speed and two-dimensional spatial resolution, for the cost of a decrease in spectral resolution, whereas for hyperspectral imaging the spectral resolution is attained at the cost of spatial resolution in one dimension and acquisition time. In a production system, it is the author's opinion that the spatial resolution and image acquisition time offered by a multispectral system makes this the preferred method. It is likely that continued research using hyperspectral images will be the best way in which to determine which spectral lines or bands should be used in a multispectral system. Indeed, if a small enough number of spectral filters were required (this is likely to need to be four or fewer), these could form a mask directly onto a camera face in much the same way as filters are applied in a Bayer pattern to form colour cameras today. This would create genuinely snap-shot spectral images.

As discussed in section 6.2 above, the techniques presented in this thesis are presented as generic spectral imaging techniques, which begs the wider question of whether, within a clinical/medical setting, and in the context of being able to distinguish different types of biological tissue, this technique has wider applicability. Lymph nodes are prognostic indicators in cancers other than breast cancers and it is likely that by opening up collaborations with specialists in other oncological surgeries the number of cases that could be included in future research could be greatly increased.

Beyond lymph nodes, even beyond cancer, it seems entirely plausible that there are many areas of medicine where spectral imaging could open up new avenues of research.

Appendix A

Publications

A.1 CLEO 2011 - Baltimore, Maryland, May 2011

An Extended Spectral Angle Map for Hyperspectral and Multispectral Imaging

O'Sullivan, J.D.¹, Hoy, P.R.², Rutt, H.N.¹

¹Faculty of Physical and Applied Sciences, University of Southampton, Highfield, Southampton, SO17 1BJ, UK

²School of Medicine, University of Southampton, Highfield, Southampton, SO17 1BJ, UK

Email:jdo@orc.soton.ac.uk, Tel: +44 2830 593163, Fax: +44 2380 593142

Abstract: We present an extension to the widely used spectral angle metric, calculating an azimuthal angle around a reference vector. We demonstrate that it provides additional information, thus improving the classification ability of the spectral angle.

©2010 Optical Society of America

OCIS codes: (280.0280) Remote Sensing and Sensors; (280.4788) Optical sensing and sensors; (110.4234) Multispectral and Hyperspectral Imaging

1. Introduction

The spectral angle[1] is a common metric in spectral imaging. It treats spectra as vectors in a high-dimensional space allowing the scalar product of, and hence the angle between, a test vector and reference vector to be calculated (see equation 1 where the test spectrum, t , and the reference spectrum, r , are each considered as N dimensional vectors). Smaller angles indicate more similar vectors.

$$\theta = \arccos \left(\frac{\sum_{i=1}^N t_i r_i}{\sqrt{\sum_{i=1}^N t_i^2} \sqrt{\sum_{i=1}^N r_i^2}} \right) \quad (1)$$

This metric is simple and computationally inexpensive to calculate. It is easily understood as an analogue of a generalised distance between two points, indeed for small angles it is approximately equal to the Euclidean distance between two points[2]. Further, it is intensity independent as all test vectors kt subtend the same angle from r ; as such illumination changes across an image do not affect this measurement.

For a given reference vector however, the angle θ does not specify a unique test vector, or even a unique test unit-vector. This is clear in two dimensions where there are exactly two test vectors for each θ , one between the reference vector and each axis. In three dimensions the vectors lying on the surface of a cone around the reference vector all subtend the same angle, in principle an infinite number of vectors. In a discrete space, as the number of dimensions increases, the number of degenerate test vectors also increases, although in more than three dimensions this is very difficult to visualise.

We suggest then that as well as calculating the spectral angle, a set of azimuthal angles should also be calculated so as to reduce this degeneracy. In section 2 we set out the method for this and in section 3 we show how this can be applied to a real dataset.

2. Method

In a three dimensional space, we have a set of vectors \mathbf{T} , at an angle θ from a reference vector, \mathbf{R} , forming a cone. We can rotate the co-ordinate space $(x,y,z) \rightarrow (x', y', z')$, such that $\mathbf{R} = R\mathbf{z}'$, the projection of \mathbf{T} onto the $x'y'$ plane is now a circle. Each member of \mathbf{T} can now be specified by considering the angles that this projection subtends from the x' and y' axes.

In an arbitrary $N+1$ dimensions, $\mathbf{T}_i = \sum_{j=1}^{N+1} T_{ij} \mathbf{x}_j$, and if we rotate the co-ordinate space such that $\mathbf{R} = R\mathbf{x}'_1$, then the projection of \mathbf{T} onto the N dimensional subspace orthogonal to \mathbf{R} is $\mathbf{T}'_i = \sum_{j=2}^N T'_{ij} \mathbf{x}'_j$. For each \mathbf{T}'_i , we can now calculate the set of angles φ_k between the vector and relevant axis x'_k (equation 2).

$$\varphi_{ik} = \arccos \left(\frac{T'_{ik}}{\sqrt{\sum_{j=2}^{N+1} T'^2_{ij}}} \right) \quad (2)$$

3. Application and Discussion

We demonstrate this technique with an example using a publicly available multispectral image from The University of Eastern Finland[3]. The image in question is available as “Landscape.zip”. The image is a multispectral image with seven channels of data, and shows a landscape scene of a tree against a backdrop of a field. A monochrome version of the image is shown in figure 1.



Figure 1 – Monochrome landscape image

We calculated the spectral angle of each pixel in the image with respect to a common reference (a “grey” pixel, i.e. with the same value in each channel). We then calculated a histogram of the spectral angle values and fitted Gaussian peaks to the histogram. We then classified each pixel by assigning it to the peak to which it most strongly belonged, the result of this is shown in figure 2. We can see from this image that the road/track, which shows up most as red, blue and green in the lower portion of the image, and the longer grass in the central third of the image, which shows up as predominantly pink, are clearly distinguished. However, this classification assigns the tree and the areas of shorter grass the same category (the yellow areas of the image).

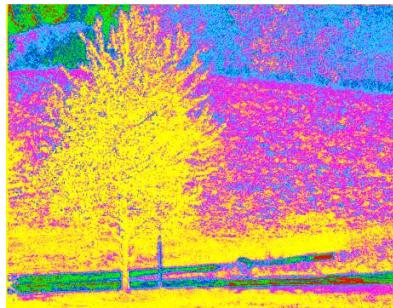


Figure 2 – Classification of pixels using the spectral angle metric

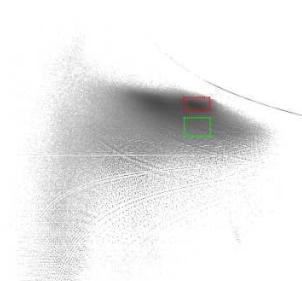


Figure 3 – A plot showing spectral angle (horizontal axis) against azimuthal angle φ_2 (vertical axis)



Figure 4 – Spatial configuration of pixels from the red and green boxes shown on figure 3.

We then plotted the distribution of spectral angle against the azimuthal angle φ_2 (the angle from the x' axis), as shown in figure 3. From this we can see that at any given spectral angle, there is a spread of azimuthal angles. The darker areas on this plot show areas of higher pixel densities and so these are the interesting cases to examine further. The red box on the image highlights one such area, the green box highlights pixels with the same spectral angle as those in the red box, but a different azimuthal position. Figure 4 shows the pixels in these highlighted areas on the spatial image. On this image we can clearly see the pixels in the red box (red areas of the image) correspond well to the areas of shorter grass, and those in the green box correspond well to the leaves of the tree.

4. Conclusion

We have presented an extension to the spectral angle metric in order to address the limits placed on this measure by the fact that a single angle does not uniquely specify a spectrum. By applying this technique to a real multispectral image we have demonstrated that pixels corresponding to different materials which were inseparable by spectral angle alone can be separated by this azimuthal position metric.

5. References

- [1] Kruse, F.A., et al., *THE SPECTRAL IMAGE-PROCESSING SYSTEM (SIPS) - INTERACTIVE VISUALIZATION AND ANALYSIS OF IMAGING SPECTROMETER DATA*. *Remote Sensing of Environment*, 1993. **44**(2-3): p. 145-163.
- [2] Chang, C.-I., *An information-theoretic approach to spectral variability, similarity, and discrimination for hyperspectral image analysis*. *Information Theory, IEEE Transactions on*, 2000. **46**(5): p. 1927-1932.
- [3] *Spectral Image Database*. University of Joensuu Color Group [cited 2010 26/11/2010]; Database of spectral images in different formats]. Available from: <http://spectral.joensuu.fi/multispectral/spectralimages.php>.

A.2 OSA/SPIE European Conferences on Biomedical Optics - Munich, Germany, May 2011

A.2.1 Conference Proceedings Paper for ECBO, Munich, May 2011

Spectral Imaging as a Potential Tool for Optical Sentinel Lymph Node Biopsies

Jack D. O'Sullivan^{*a}, Paul R. Hoy^b, Harvey N. Rutt^a

^aFaculty of Physical & Applied Sciences, University of Southampton, Southampton, SO17 1BJ, UK;

^bSchool of Medicine, University of Southampton, Southampton, SO17 1BJ, UK

ABSTRACT

Sentinel Lymph Node Biopsy (SLNB) is an increasingly standard procedure to help oncologists accurately stage cancers. It is performed as an alternative to full axillary lymph node dissection in breast cancer patients, reducing the risk of long-term health problems associated with lymph node removal. Intraoperative analysis is currently performed using touch-print cytology, which can introduce significant delay into the procedure. Spectral imaging is forming a multi-plane image where reflected intensities from a number of spectral bands are recorded at each pixel in the spatial plane. We investigate the possibility of using spectral imaging to assess sentinel lymph nodes of breast cancer patients with a view to eventually developing an optical technique that could significantly reduce the time required to perform this procedure. We investigate previously reported spectra of normal and metastatic tissue in the visible and near infrared region, using them as the basis of dummy spectral images. We analyse these images using the spectral angle map (SAM), a tool routinely used in other fields where spectral imaging is prevalent. We simulate random noise in these images in order to determine whether the SAM can discriminate between normal and metastatic pixels as the quality of the images deteriorates. We show that even in cases where noise levels are up to 20% of the maximum signal, the spectral angle map can distinguish healthy pixels from metastatic. We believe that this makes spectral imaging a good candidate for further study in the development of an optical SLNB.

Keywords: Sentinel Lymph Node Biopsy, Multispectral Imaging, Hyperspectral Imaging, Spectral Angle Map, Optical Biopsy

1. INTRODUCTION

1.1 Sentinel lymph node biopsy procedure

In order to effectively plan treatment and determine prognosis for cancer patients, it is important that the stage of the cancer is correctly assessed. Cancer spreads through the body using the lymphatic system; as such this system plays a crucial role in this staging process. In the case of breast cancer patients, fluids from the tumour region enter the lymphatic system through the axillary lymph nodes. During a standard lumpectomy, it is these nodes which must be assessed in order to determine the level of metastasis.

The Sentinel Lymph Node Biopsy (SLNB) is a standard alternative to the Axillary Lymph Node Dissection (ALND) procedure for determining the extent of this involvement. The ALND is a full removal of all lymph nodes in the axillary region, these are then tested for the presence of cancerous cells. The Sentinel Lymph Node (SLN) is the first node to which fluids from the tumour region will drain and as such is the primary gateway for metastases and the purpose of the SLNB is to determine whether cancerous cells are present in the SLN.

If it is determined that the SLN contains cancerous cells then a full ALND will be performed to assess the extent of lymphatic involvement. If no cancer is found in the SLN then the remaining axillary lymph nodes can be left intact as there is low probability that metastasis has started to occur¹. This reduces the risk of long term health problems associated with the removal of lymph nodes, such as lymphedema².

Currently the assessment of the SLN is performed by way of touch print or touch prep (TP) cytology³. This cytological test is presently the rate-limiting step in the tumour resection surgery. Where the facilities for performing this test are available onsite, the results can be obtained during the course of the surgery, although a delay of up to an hour is not uncommon and this is time that the patient is anaesthetised and “open” in theatre. If no facilities are available onsite,

* jdo@orc.soton.ac.uk; phone +44 2380 593163

2. SPECTRAL ANGLE MAP

The spectral angle requires a geometric interpretation of a spectrum. Each wavelength or waveband is treated as a coordinate axis in a high-dimensional space; each possible spectrum thus specifies a point in this space and can be thought of as a high-dimensional vector.

As is familiar for three dimensional space, any two vectors, \mathbf{a} and \mathbf{b} , in this space will be co-planar, and will diverge from each other at a fixed angle, θ , which can be calculated from the scalar product:

$$\theta = \arccos\left(\frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}\right). \quad (1)$$

This angle, θ , is the spectral angle between two spectra. For exactly coincident spectra, this angle will be 0 and it will increase as the spectra become more dissimilar. Since all values in a spectral image are non-negative, we are only considering arccos in the domain $[0,1]$, this places an upper limit on the spectral angle of $\pi/2$ radians. It is worth noting that all vectors $k\mathbf{a}$ will subtend the same angle from \mathbf{b} . In terms of spectral imaging, this means that the spectral angle is illumination independent, i.e. changes in illumination across the image will not affect the results.

This angle can be calculated for each image pixel (against a standard reference spectrum), thus creating a spectral angle map (SAM).

3. APPLICATION

We are working on spectral imaging in the visible and near infrared range (400-1000nm), and so to test the viability of the SAM, we have constructed a dummy image using visible spectra of healthy and metastatic sentinel nodes of breast cancer sufferers reported by Bigio et al⁹ (see Figure 2).

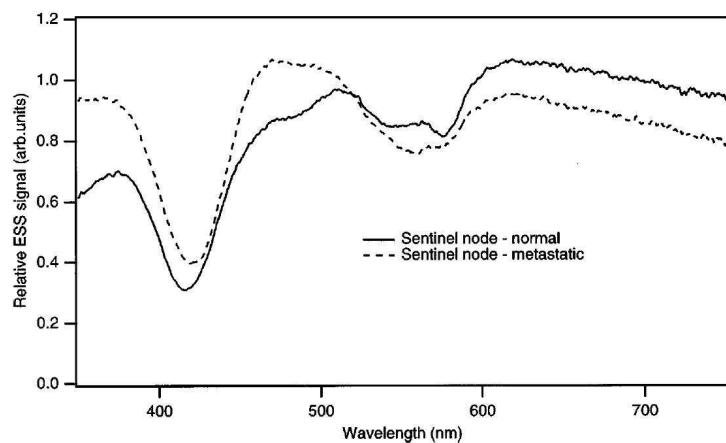


Figure 2 - Diffuse reflectance spectra in the range ~350-750nm for normal and metastatic sentinel lymph nodes.

We created a hyperspectral image in which each pixel was either assigned the normal spectrum or the metastatic spectrum according to a binary mask image. Further hyperspectral images were created such that random deviations from the spectrum were added to simulate noise in the measurements. Each measurement of each pixel was randomly adjusted by adding or subtracting a number randomly selected between zero and a given percentage of the maximum signal value. This 'noise level' was increased from 1% to 20% in integral increments, each with a new hyperspectral image.

For each image we then created an SAM by calculating the spectral angle of each pixel from a grey reference spectrum, i.e. a spectrum in which every value is the same. Figure 3 below shows the SAMs for two of these hyperspectral images; the image on the left shows the original (all pixels are either filled with the normal, or metastatic spectrum as reported), the image on the right shows the hyperspectral image with the addition of a 20% noise level.



Figure 3 - SAMs of the dummy hyperspectral images, on the left, the original, on the right the image with a 20% noise level.

It is clear that increasing the noise level decreases the separation of the normal and metastatic pixels.

It is evident from Figure 3 that the background (normal) pixels are significantly less well separated from the foreground (metastatic) pixels in the noisy image than they are in the original binary image. In particular the finer spatial details in the foreground are less well defined.

Figures 4-6 show histograms of spectral angle from three of these images, the original image, a 10% noise level image and a 20% noise level. On the histogram of the original image we have just two lines at the values of the spectral angles of the normal and metastatic spectra. As we increase the noise, these lines are smeared out into peaks, which gradually merge into each other. This is confirmation of what we would expect to be the case. In the presence noise on the signal, the normal pixels exhibit a range of closely correlated spectra, and the metastatic pixels also exhibit this range. It also confirms what we observed visually from the SAMs, i.e. that the normal and metastatic pixels become less separable in the spectral angle space as the noise was increased.

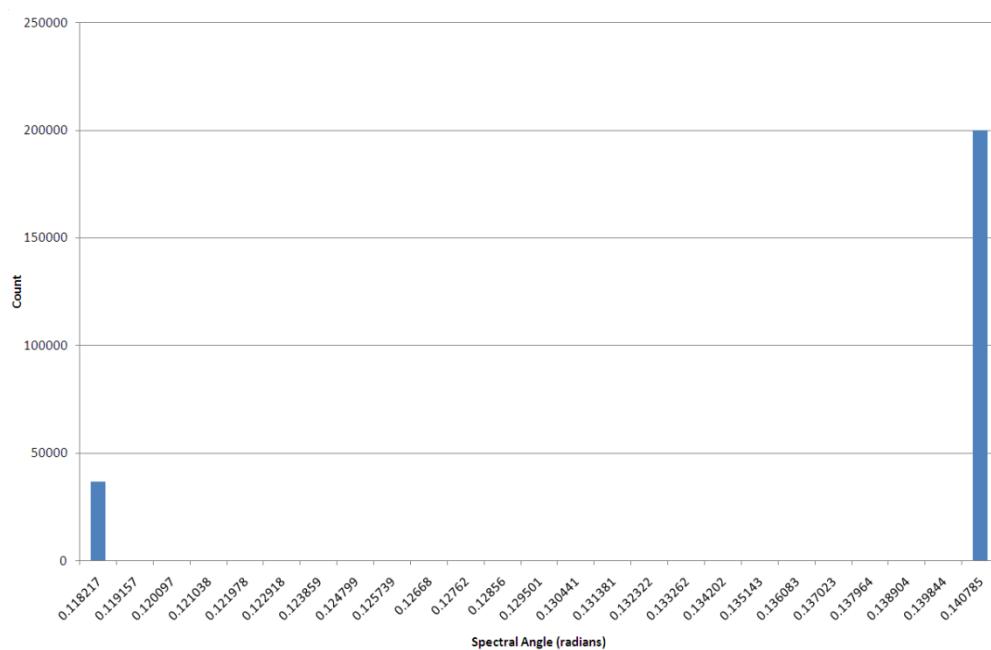


Figure 4 - Histogram of the spectral angles for the original (no noise) dummy hyperspectral image.

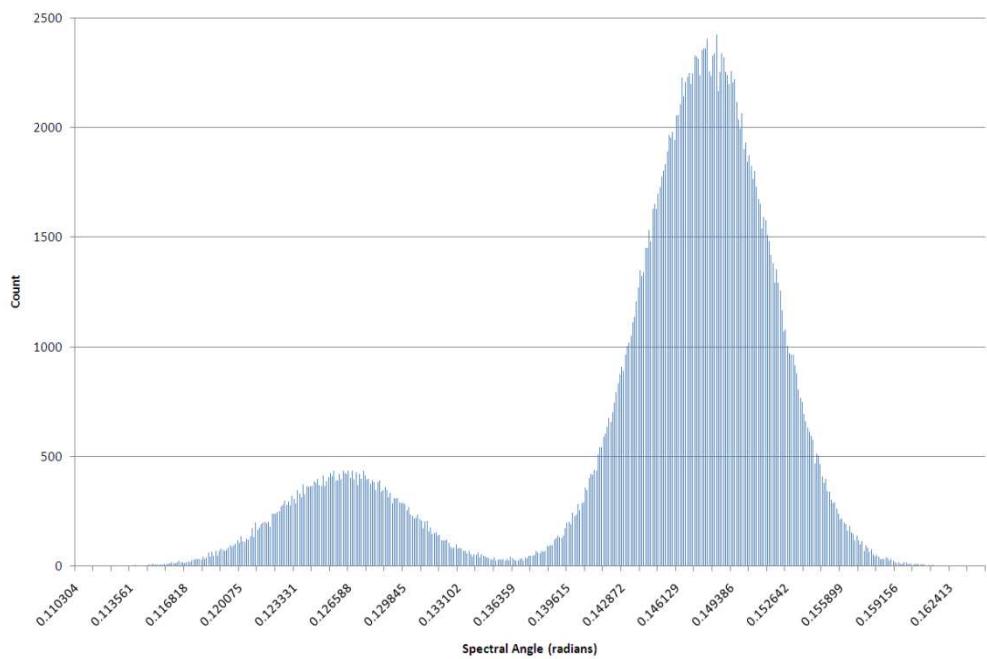


Figure 5 - Histogram of the spectral angles for the 10% noise level hyperspectral image.

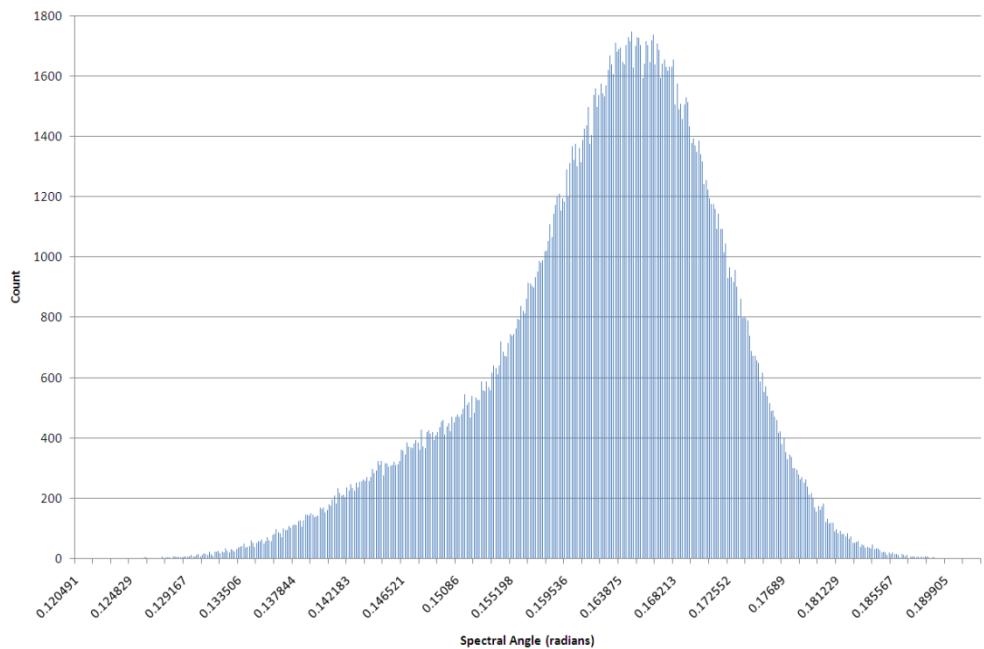


Figure 6 - Histogram of the spectral angles for the 20% noise level hyperspectral image.

We fitted Gaussian peaks to each of the histograms, the values of the peaks and their widths (shown as error bars) for increasing noise levels, are shown in Figure 7.

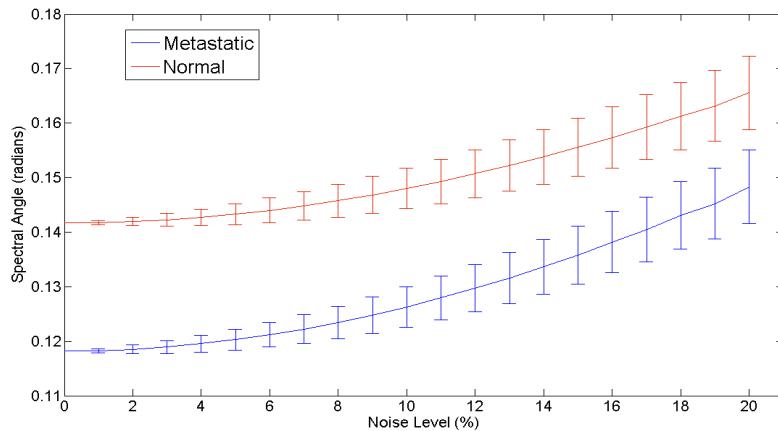


Figure 7 - Graph showing positions and widths of the Gaussian peaks fitted to the histograms of spectral angles for all dummy hyperspectral images. The width of the peak is shown as the error bar at that point.

Using these Gaussian profiles, we classified each pixel as normal or metastatic purely on the basis of its spectral angle. For each pixel, we calculated the probability of it being of part of each peak, it was then randomly assigned to one of the peaks, based on these probabilities. For example, a pixel with a 75% probability of being a part of the normal peak would have a 75% chance of being assigned to the normal peak. Figure 8 below shows the results of classification for a low noise level (5%) and the 20% noise level image.



Figure 8 - Classification of pixels as being normal (black) or metastatic (white) based on spectral angle. The left image shows a 5% noise level spectral angle, the right image shows a 20% noise level.

The 5% noise level image shows virtually no misclassification whereas the 20% noise level image shows a great deal. The foreground is still distinguishable from the background, and the classification seems to be correct for the majority of pixels. We counted the number of misclassified pixels and this is shown below for increasing noise levels in Figure 9. This shows that the number of misclassifications was steady and very low up until noise levels of around 9-10% before rising steadily. It is around the 9-10% level at which the peaks on the histograms begin to merge significantly.

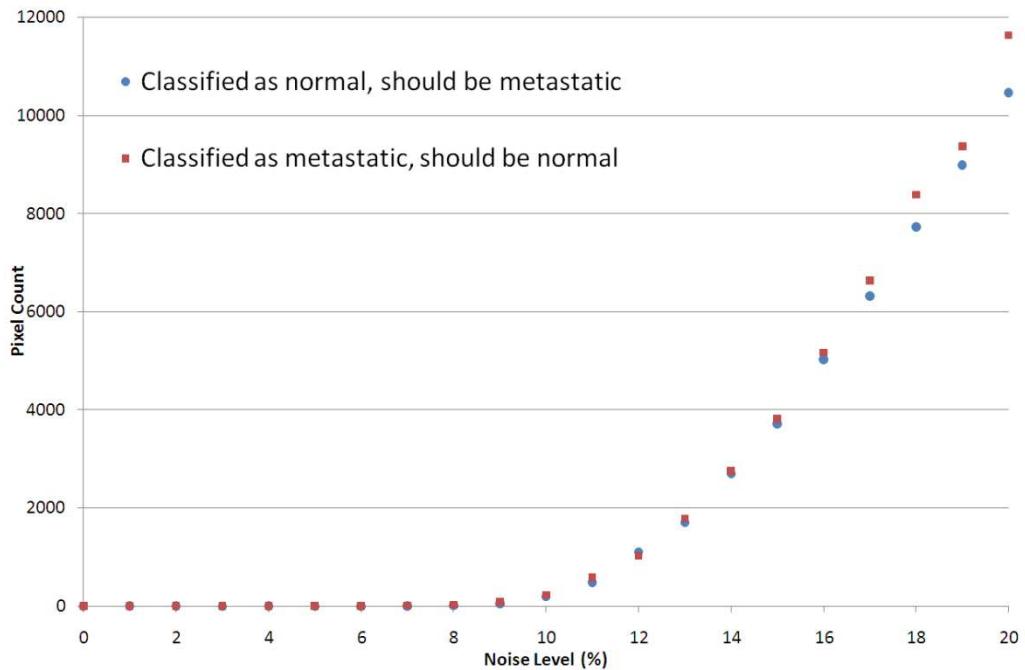


Figure 9 - A graph showing the number of misclassified pixels as a function of noise level, both for pixels that were misclassified as normal and those that were misclassified as metastatic. There are 199853 normal and 36768 metastatic pixels in the original.

Further, we calculated the average spectra for those pixels classified as normal and those classified as metastatic for the 5% and 20% noise level image; these are shown together in Figure 10 below.

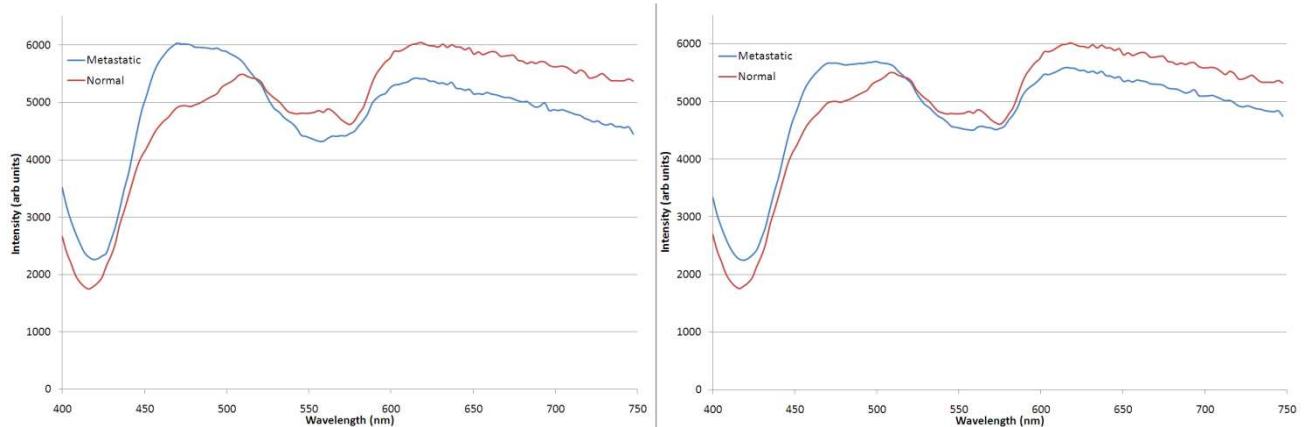


Figure 10 - Average spectra for pixels classified as normal and metastatic. The graph on the left shows the 5% noise level image, the right hand graph shows those from the 20% noise level image.

Figure 10 shows that, although still clearly distinguishable, the differences between the spectra are beginning to diminish at high noise levels. In particular, the peak in the metastatic spectrum at around 460nm is significantly reduced, and in the area between that and the start of the decline at around 510nm the spectrum becomes much flatter. The separation between the peaks at higher wavelengths is also markedly reduced. It is noted that the normal spectrum is much less changed than the metastatic.

4. CONCLUSION

The spectral angle map clearly allows us to distinguish between the normal and metastatic spectra. Increasing the noise level on the image makes the distinction between normal and metastatic pixels less clear, as the histograms in figures 4-6 demonstrate. However, Figure 7 clearly shows that the spectral angles remain quantifiably separable, even at high noise levels.

The limits on this separability have been explored somewhat further by assessing how well the pixels are classified as being normal or metastatic purely on the basis of their distance in metric space from the peaks of the Gaussian curves present on the histogram. We find that up to around 10%, the misclassification rate is very low but that it climbs steadily as the image gets noisier thereafter. The number of misclassifications is broadly similar for both normal and metastatic (the difference can be accounted for by pixels no longer being classified as either as they have spectral angles far removed from either peak), however this has more impact on the average calculated metastatic spectrum.

There are approximately 5 normal pixels for every metastatic pixel in the original image; we would anticipate that in a real world situation there would be less cancerous tissue in a node than normal tissue. This means however, that every misclassified pixel affects the metastatic spectrum a lot more than the normal spectrum, so whilst the differences between them are changing, as demonstrated in Figure 10, most of this can be attributed to the metastatic spectrum becoming more like the normal.

These results, based as they are on “dummy” data, are very preliminary, and much work remains to be done; however, we believe that they demonstrate that the SAM shows much potential for being the basis of a rigorous and automated classification system, which itself could lead to the development of an optical system for performing SLNBs.

ACKNOWLEDGEMENT

The authors would like to thank Mr Richard Sainsbury at the School of Medicine, University of Southampton, for sharing his expertise on the subject of breast cancer treatment and specifically the sentinel lymph node biopsy.

REFERENCES

- [1] Howard, S., et al., "Sentinel node biopsy in the staging of breast cancer". *American journal of surgery*. 176(4): p. 305-310.(1998)
- [2] McMasters, K.M., et al., "Sentinel Lymph Node Biopsy for Breast Cancer: A Suitable Alternative to Routine Axillary Dissection in Multi-Institutional Practice When Optimal Technique Is Used". *J Clin Oncol*. 18(13): p. 2560-2566.(2000)
- [3] Rubio, I., et al., "Use of touch preps for intraoperative diagnosis of sentinel lymph node metastases in breast cancer". *Annals of Surgical Oncology*. 5(8): p. 689-694.(1998)
- [4] Park, B., et al., "Contaminant Classification of Poultry Hyperspectral Imagery using a Spectral Angle Mapper Algorithm". *Biosystems Engineering*. 96(3): p. 323-333.(2007)
- [5] Li, Q.L., et al., "Detection of physical defects in solar cells by hyperspectral imaging technology". *Optics and Laser Technology*. 42(6): p. 1010-1013.(2010)
- [6] Ng, P.H.R., et al., "Detection of illicit substances in fingerprints by infrared spectral imaging". *Analytical and Bioanalytical Chemistry*. 394(8): p. 2039-2048.(2009)
- [7] Jusoff, H.K., "Pixel-based airborne hyperspectral sensing technique for search-and-rescue of the missing RMAF NURI helicopter in Genting-Sempah, Malaysia". *Disaster Prevention and Management*. 19(1): p. 88-102.(2010)
- [8] Kruse, F.A., et al., "The Spectral Image-Processing System (SIPS) - Interactive Visualization and Analysis of Imaging Spectrometer Data". *Remote Sensing of Environment*. 44(2-3): p. 145-163.(1993)
- [9] Bigio, I.J., et al., "Diagnosis of breast cancer using elastic-scattering spectroscopy: preliminary clinical results". *Journal of Biomedical Optics*. 5(2): p. 221-228.(2000)

A.2.2 Poster Presented at ECBO, Munich, May 2011

Spectral imaging as a potential tool for optical sentinel node biopsies

Jack D. O'Sullivan¹, Paul R. Hoy², Harvey N. Rutt¹ **UNIVERSITY OF Southampton**
¹Faculty of Physical & Applied Sciences, University of Southampton,
Southampton, SO17 1BJ, UK
²School of Medicine, University of Southampton, Southampton, SO17 1BJ, UK **Optoelectronics Research Centre**

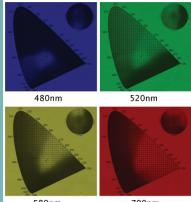
We investigate the potential use of spectral imaging for performing an optical Sentinel Lymph Node Biopsy (SLNB). Dummy images generated from previously published data are analysed using a spectral angle map. We conclude that there is scope for spectral imaging to form the basis of an optical biopsy technique.

Sentinel Lymph Node Biopsy Background

. SLNB is a routine surgical procedure to determine the extent of metastasis
. The sentinel node is the first to which fluids from the tumour site will drain
. During tumour resection, the node is removed and tested for cancerous cells by cytology ^[1]
. If the node is clear there is a low probability that metastasis has started to occur ^[2]

. Testing facilities are not available at all treatment centres; this can necessitate further surgery
. If on-site facilities exist, the test occurs whilst the patient is under anaesthetic and is often the longest part of the surgery
. An automatic optical based assessment could reduce the delay currently experienced

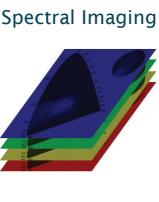
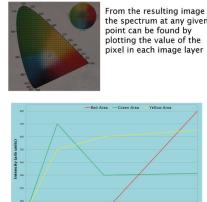
Spectral Imaging and Spectral Angle Background



A scene is imaged a number of times through filters at various wavelengths or bands of wavelengths

Spectral Imaging

These images are stacked together to form an "image cube"

From the resulting image the spectrum at any given point can be found by plotting the value of the pixel in each image layer

To compare pixels, treat each spectrum as a high dimensional vector.

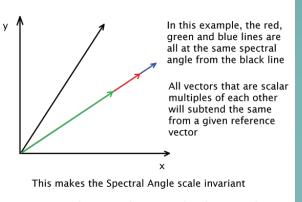
Red Pixel = [10, 12, 15, 800]
Green Pixel = [12, 700, 300, 315]

Then calculate the angle between them using the inner product:

$$\cos \theta = \frac{\mathbf{a} \cdot \mathbf{b}}{|\mathbf{a}| |\mathbf{b}|}$$

This angle is called the "Spectral Angle"

It is zero for identical spectra and increases as the spectra become less similar



In this example, the red, green and blue lines are all at the same spectral angle from the black line

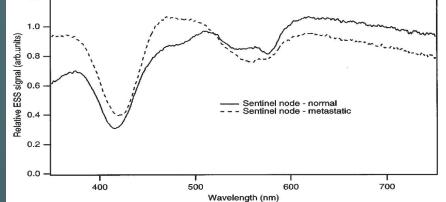
All vectors that are scalar multiples of another will subtend the same angle from a given reference vector

This makes the Spectral Angle scale invariant

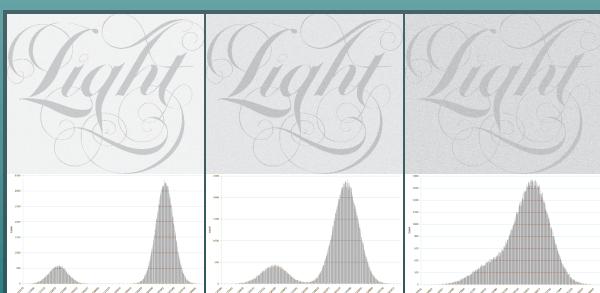
In spectral imaging this means that the Spectral Angle is invariant to changes in overall illumination

Application

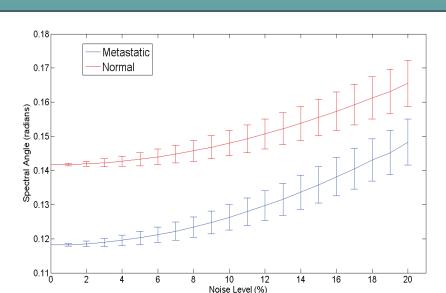
. To test spectral imaging in the visible spectrum (400-700nm) we made dummy images using the spectra reported by Bigio et al ^[3]
. Pixels were assigned either the normal or the metastatic spectrum according to a binary image mask
. Random deviations from the spectra (i.e. noise) were added, with increasing maximum amplitudes, from 0-20% of the maximum signal
. A "Spectral Angle Map" (SAM) was then generated for each noise level, the spectral angle calculated with reference to a uniform "grey" spectrum
. Histograms of the spectral angle were then created for each noise level
. Gaussian peaks were fitted to each histogram and the means and peak widths plotted against noise level



The spectra in the wavelength range of visible light for normal and metastatic sentinel lymph nodes of breast cancer patients, as reported by Bigio et al ^[3]



SAM (top) and histogram of spectral angles for noise levels of 5%, 10% and 20% respectively. The clear separation becomes increasingly blurred as noise is increased, and the peaks in the histograms move closer and begin to merge.



A plot of the positions and widths of the peaks fitted to the histograms. The decrease in separation is apparent, as is the spreading of the peaks. However the widths do not overlap even at a noise level of 20%

Conclusion

Increasing the noise level of the spectra decreases the separability of normal and metastatic spectra that exists with the SAM. However, even at a high noise level, the Gaussian peaks are still quantifiably distinct. This suggests that a spectral imaging system could potentially be developed for rapid SLNBs.

References:

[1] Rubio et al, Annals of Surgical Oncology, 5(8), 1998
[2] Howard et al, American Journal of Surgery, 176(4), 1998
[3] Bigio et al, Journal of Biomedical Optics, 5(2), 2000

Contact: jdo@orc.soton.ac.uk

Appendix B

Further Hyperspectral Images



FIGURE B.1: Pseudo-colour image of a resected SLN. Staining from the blue dye can clearly be seen at the edges of the node and in some of the surrounding tissue on the right hand side. The approximate extent of the node is shown by white lines, marking the two cut surfaces of the bi-valved node, the remainder of tissue is the surrounding fatty tissue in which the node was originally embedded.

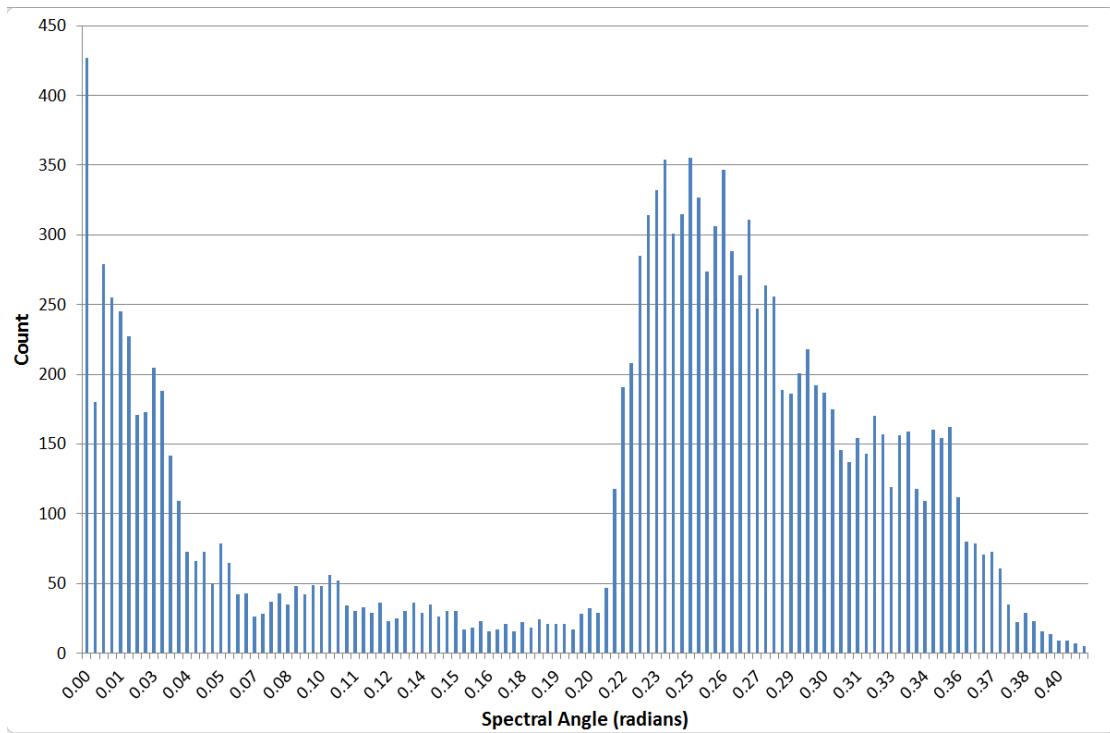


FIGURE B.2: The histogram of spectral angles present in the SAM for the node in figure B.1.

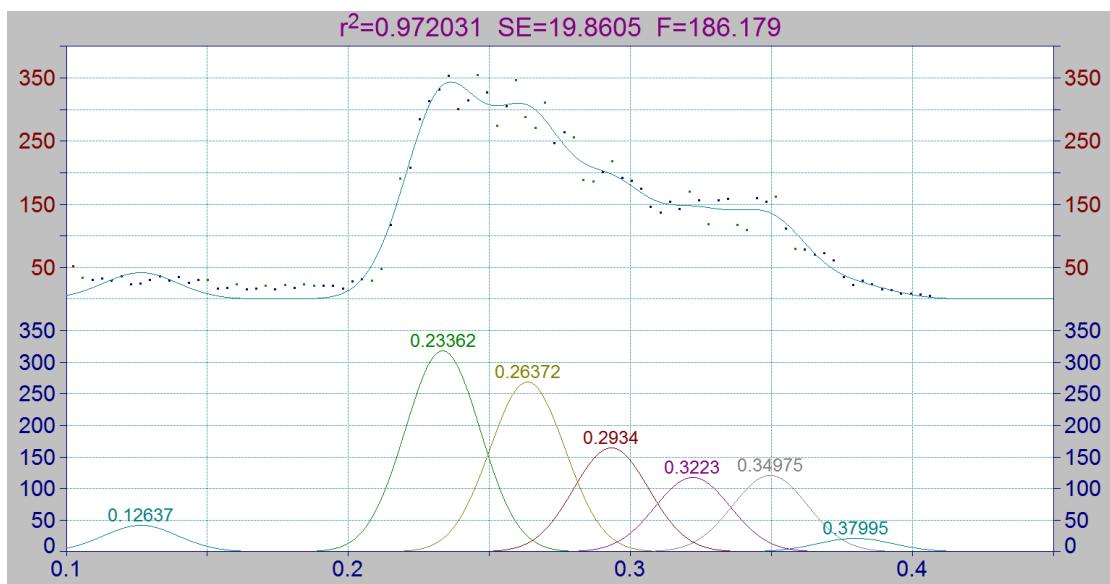


FIGURE B.3: Gaussian peaks automatically fitted to the histogram of spectral angles.

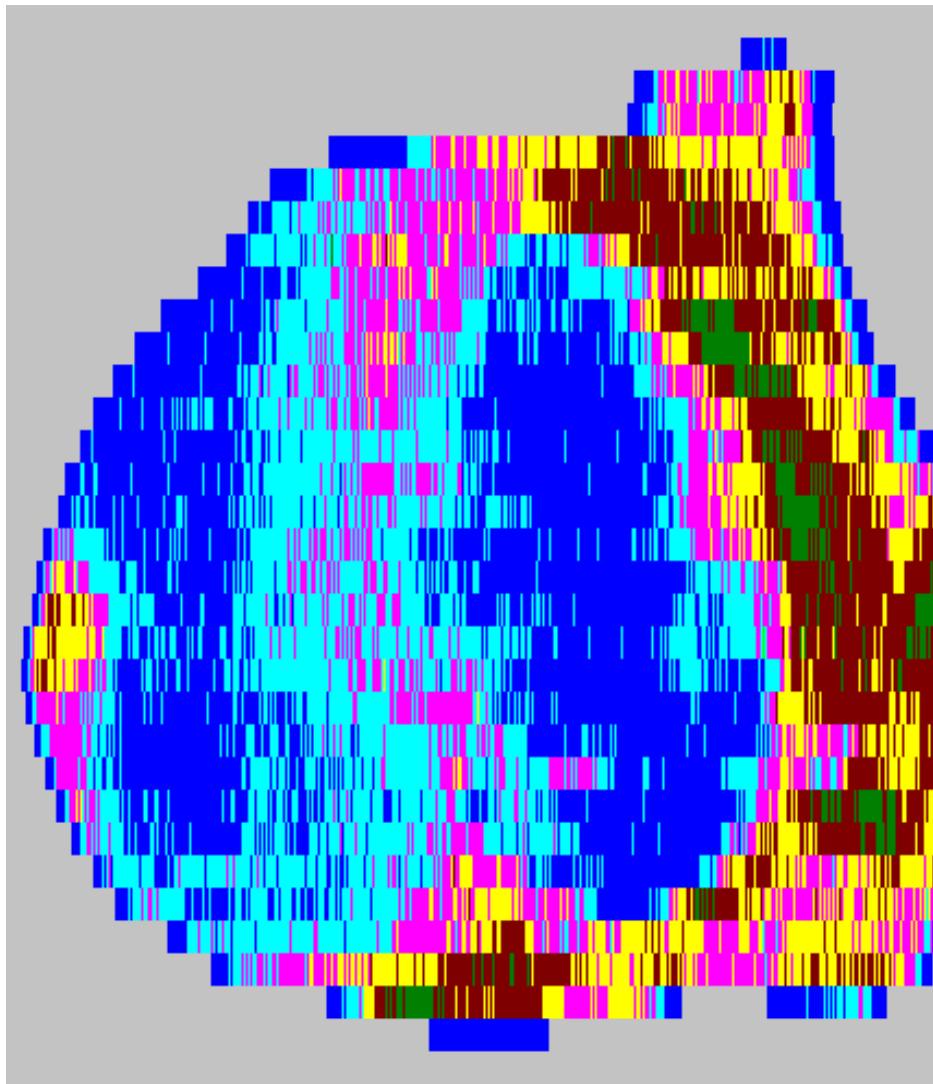


FIGURE B.4: Spectral image classified by spectral angle. The node is clearly identified as groups 3 and 4. The area of node on the left hand that was stained blue is classed along with the surrounding tissue as groups 5, 6 and 7.

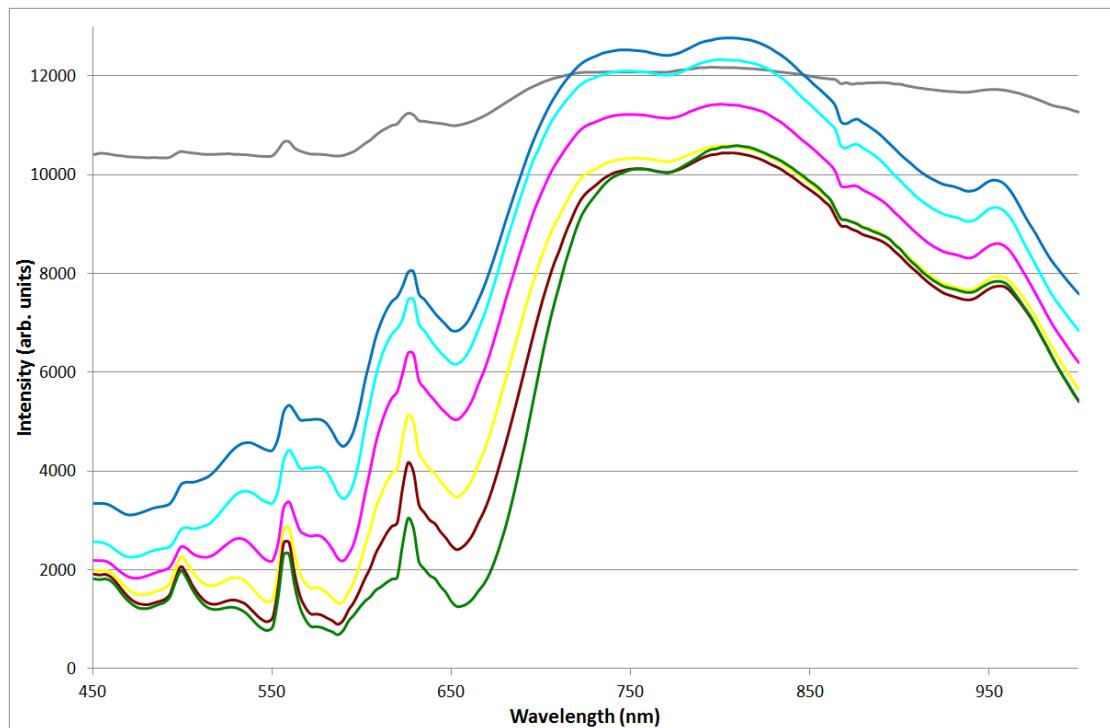


FIGURE B.5: Average spectra of the groups in figure B.4.

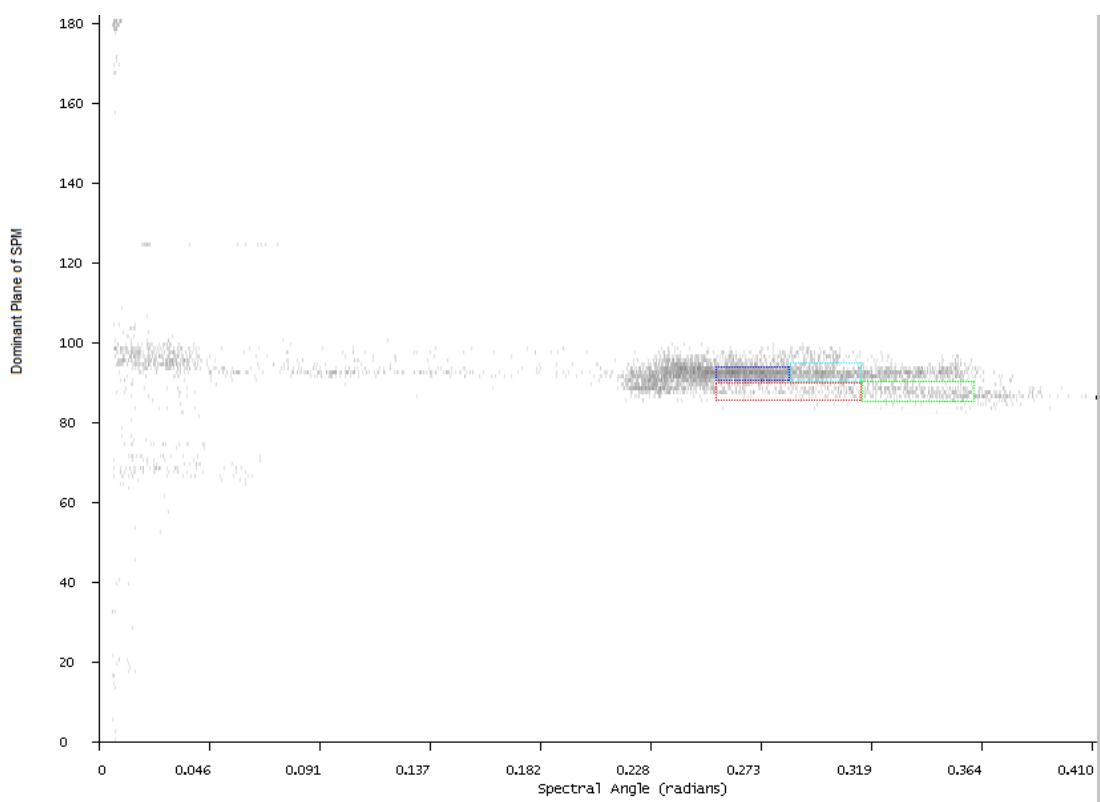


FIGURE B.6: 2D histogram of spectral angle against dominant plane of SPM, with some areas of interest marked in coloured rectangles.

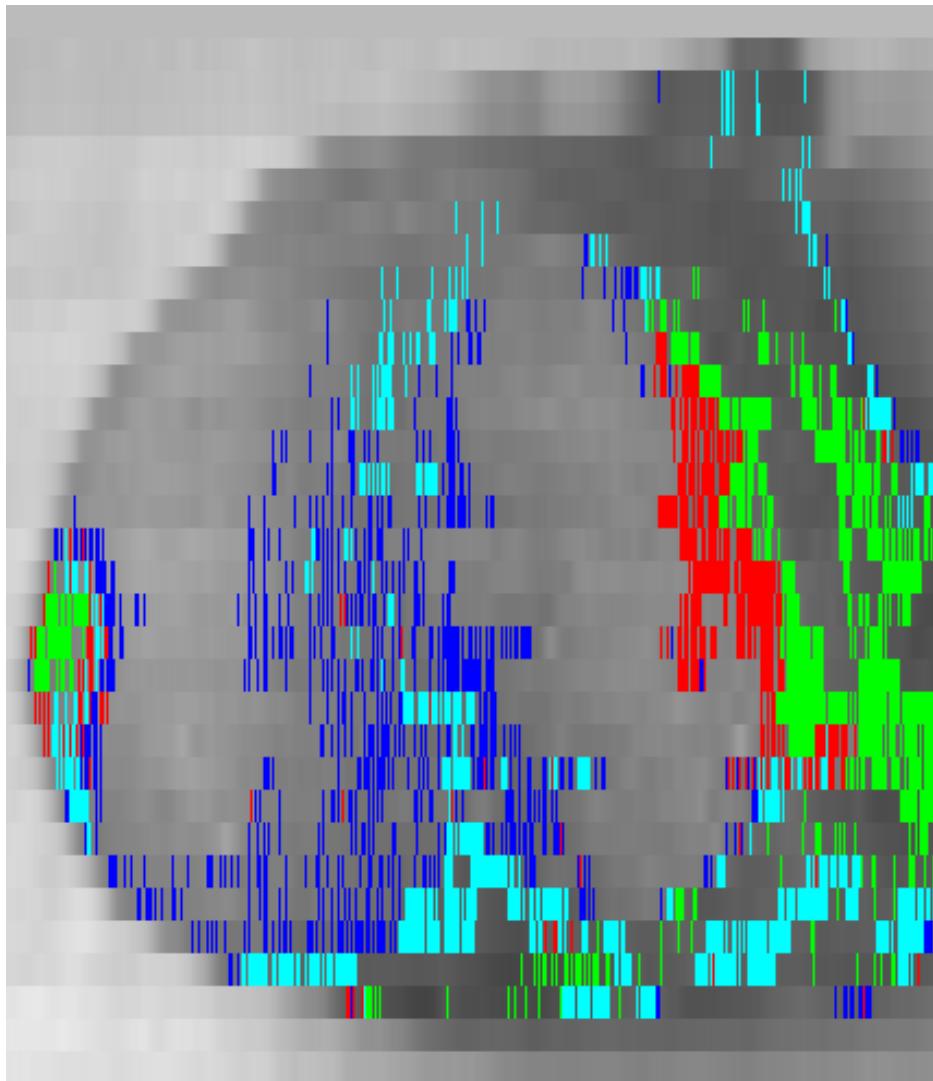


FIGURE B.7: Partially classified hyperspectral image showing the pixels inside the coloured boxes from figure B.6. The blue, red and cyan pixels cover the same spectral angle range but identify node tissue, stained tissue and surrounding tissue respectively. The green pixels cover the blue stained surrounding tissue on the right hand side, and also the stained tissue on the left hand side, suggesting this may actually not be node as assumed.

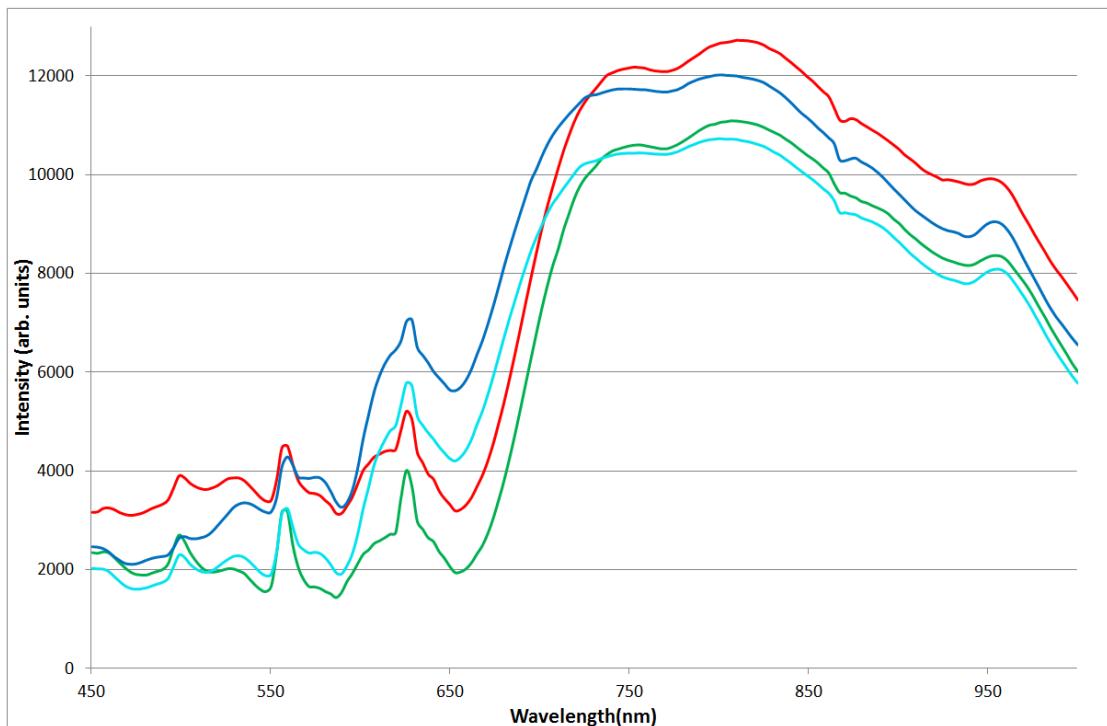


FIGURE B.8: Average spectra of the classes from figure B.7. The red and green spectra (both representing stained tissue) show lower intensities from around 500-700nm, and much steeper rises in the area above 650nm. This is the influence of the blue dye, which has high absorption in the long wavelength range of the visible spectrum. All show evidence of the lipid trough below 950nm.

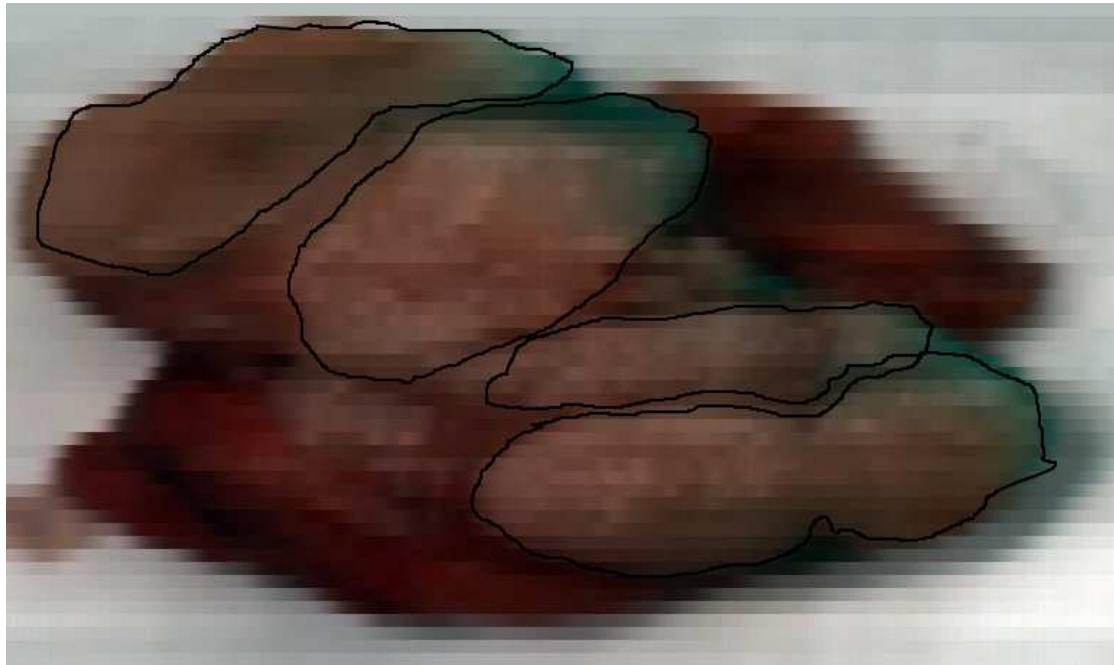


FIGURE B.9: Pseudo-colour image of two resected and butterflied lymph nodes. Some blue staining is visible at the top edge of both nodes. The approximate extent of the node is shown by black lines, marking the two cut surfaces of the bi-valved node, the remainder of tissue is the surrounding fatty tissue in which the node was originally embedded.

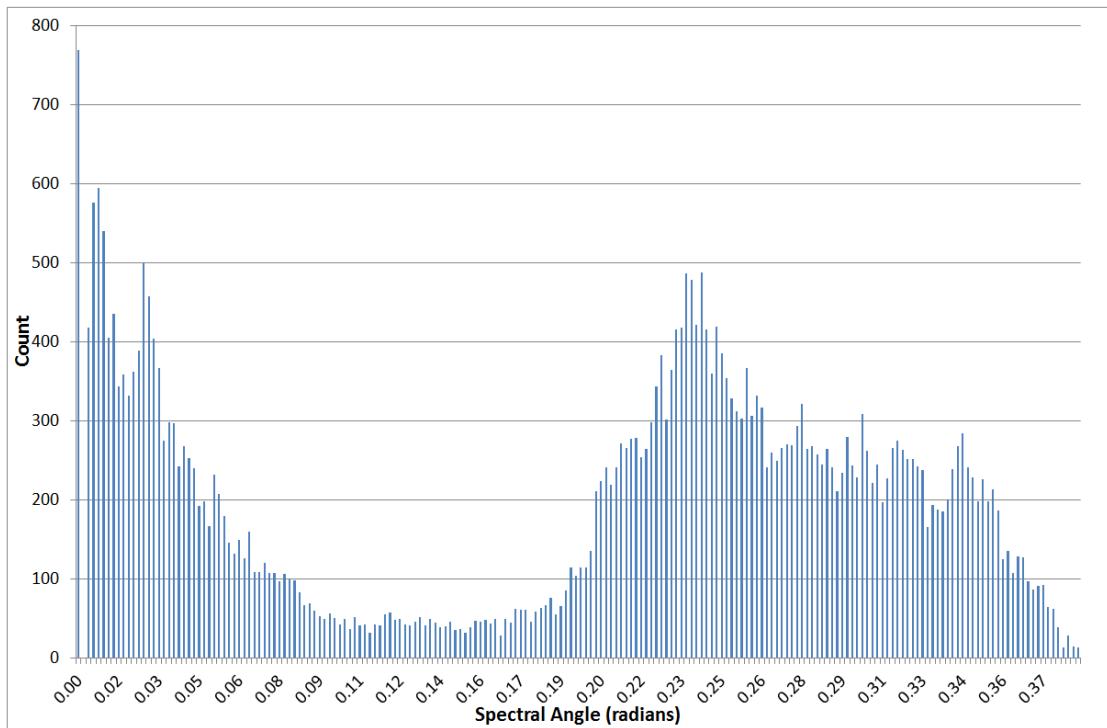


FIGURE B.10: The histogram of spectral angles present in the SAM for figure B.9.

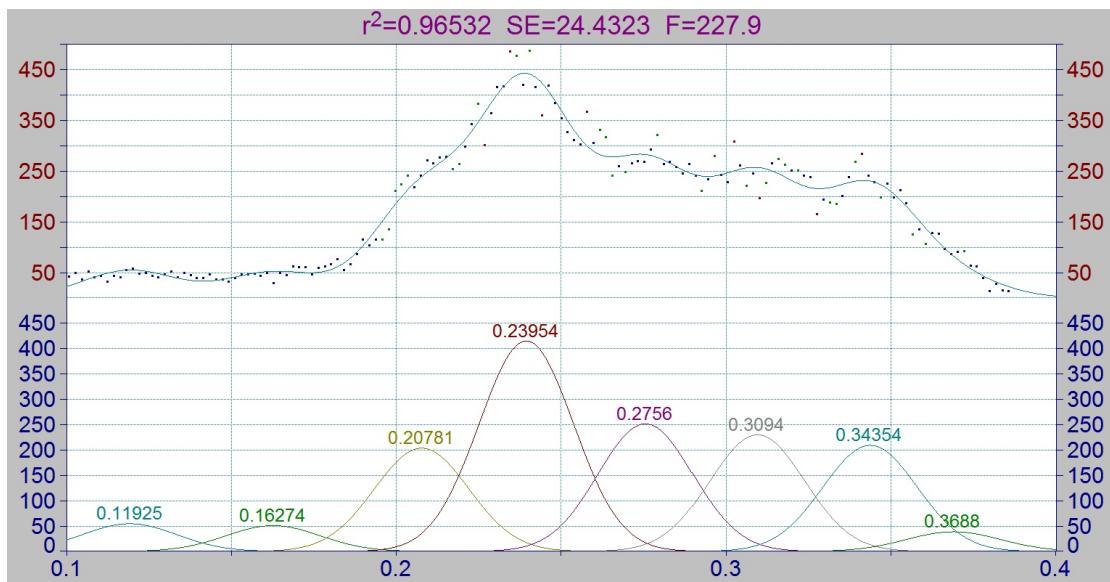


FIGURE B.11: Gaussian peaks automatically fitted to the histogram of spectral angles.

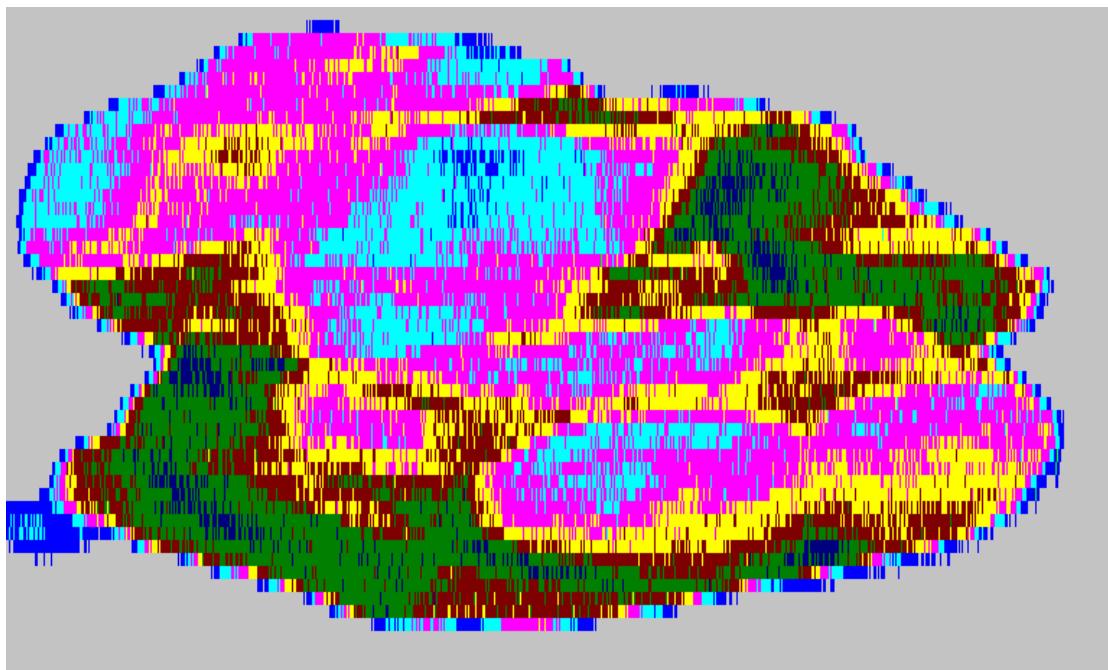


FIGURE B.12: Spectral image classified by spectral angle. The nodes are clearly distinguished as groups 4 and 5, with group 6 seemingly forming the margins and some of the surrounding tissue.

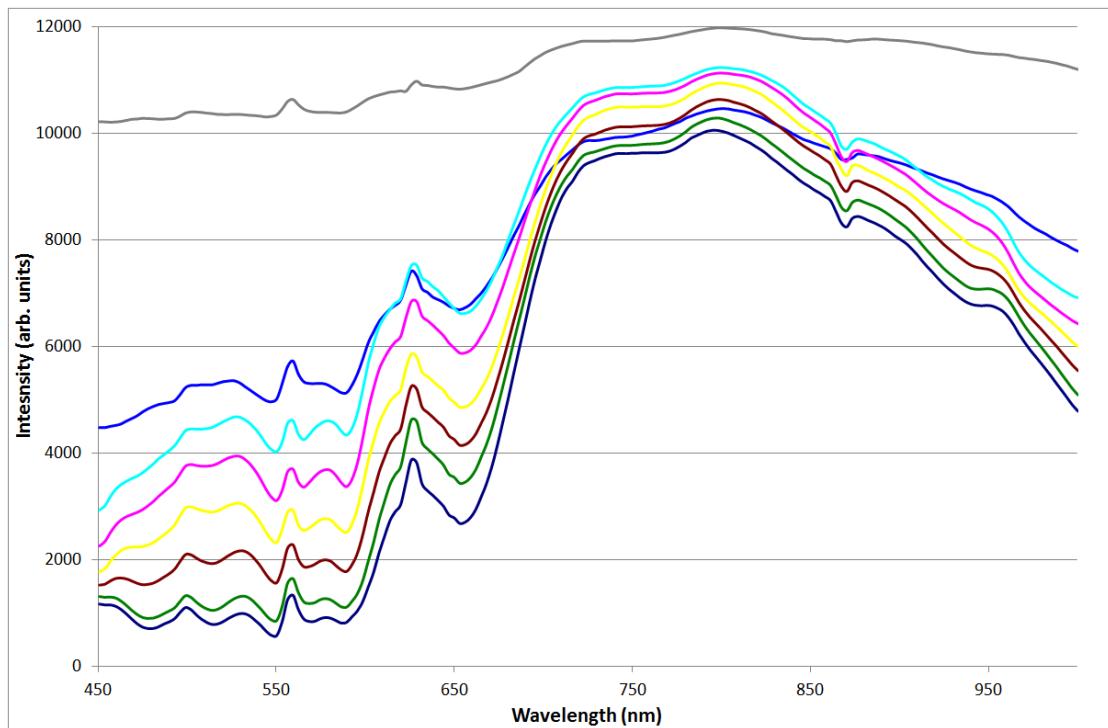


FIGURE B.13: Average spectra of the classes from figure B.12.

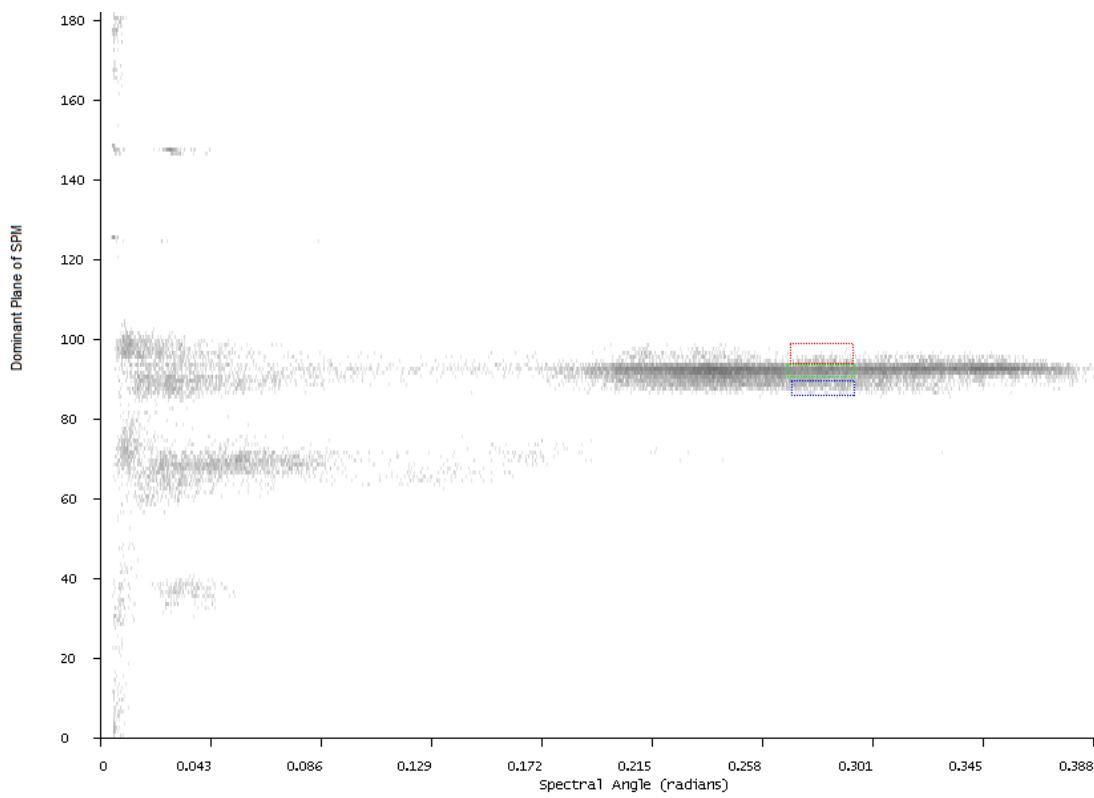


FIGURE B.14: 2D histogram of spectral angle against dominant plane of SPM, with some areas of interest marked in coloured rectangles.

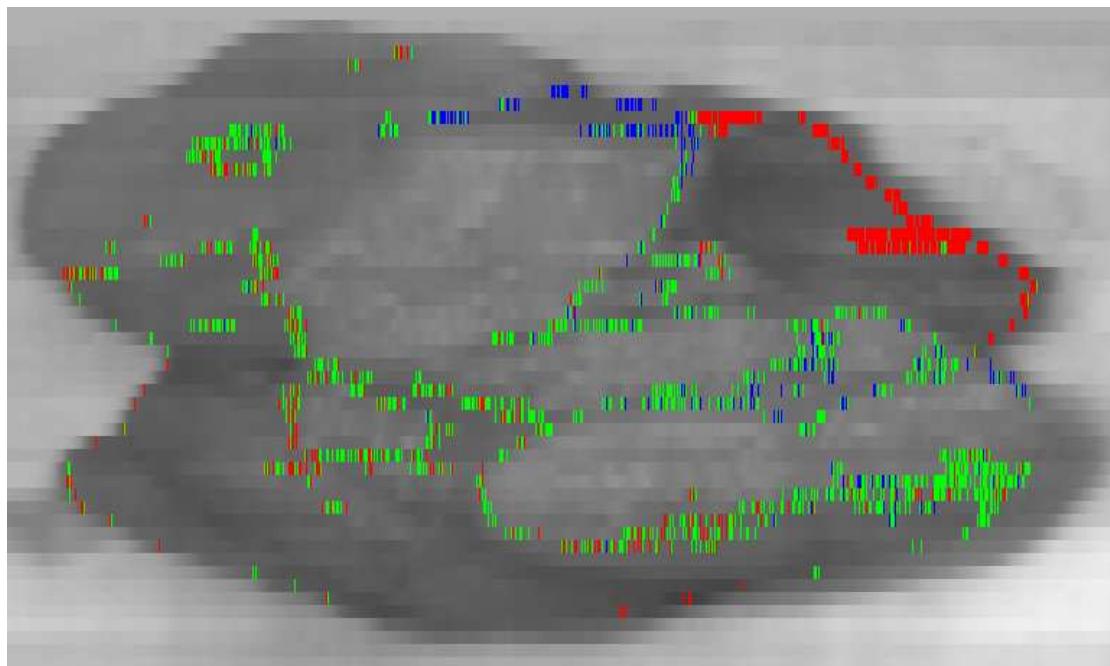


FIGURE B.15: Partially classified spectral image showing the pixels from figure B.14. The three boxes cover the same spectral angle range, roughly corresponding to group 6 from figure B.12. Here the red seems to correspond much more strongly to the surrounding tissue, with the green forming the node margins. The blue pixels cover some of the stained tissue.

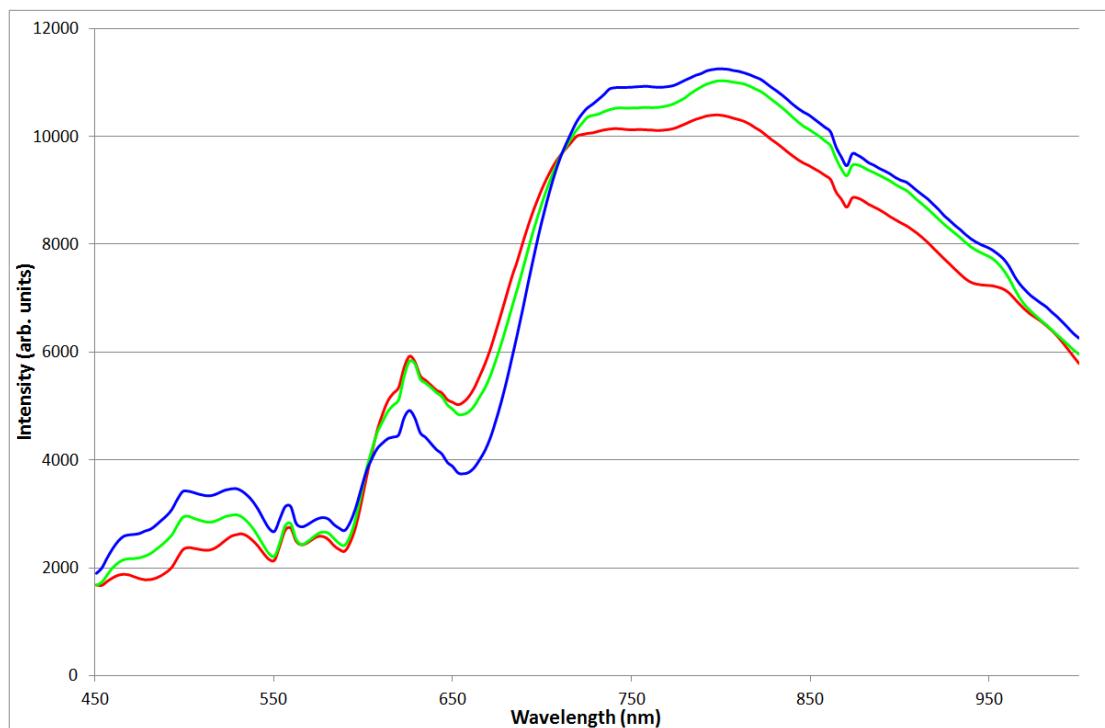


FIGURE B.16: Average spectra of the classes from figure B.15. The blue line shows a steeper rise above 650nm consistent with the stained tissue spectra from figure B.8. The lipid trough is much more pronounced on the red spectra here, with the water trough being more apparent on the green and blue.

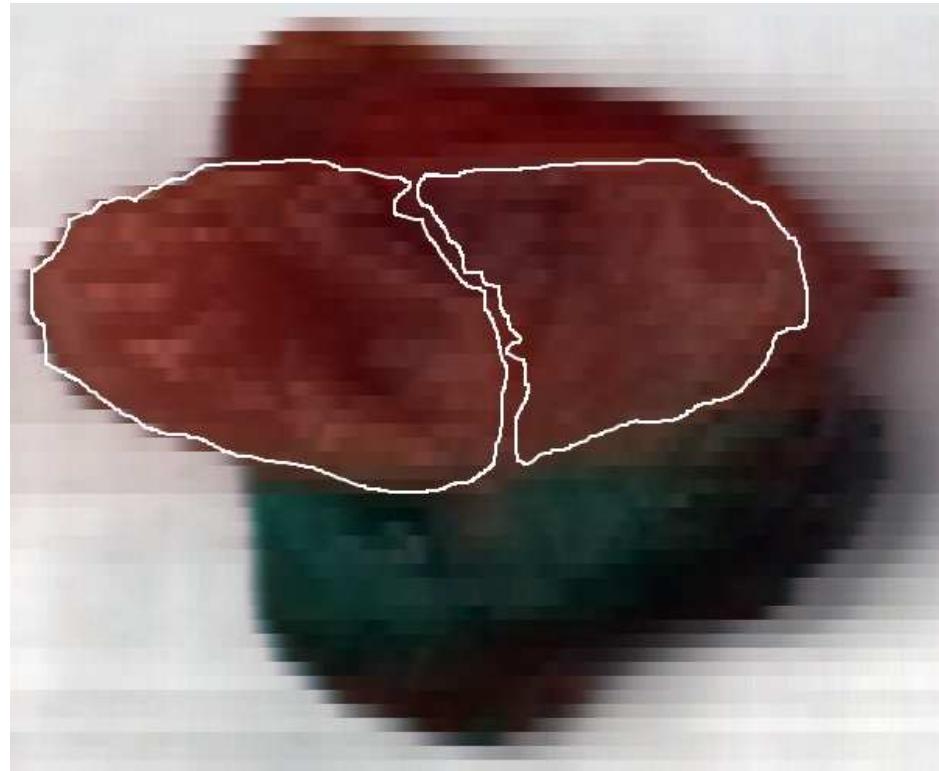


FIGURE B.17: Pseudo-colour image of a resected SLN. Here the surrounding tissue (at the bottom) is clearly heavily stained by the blue dye. The approximate extent of the node is shown by white lines, marking the two cut surfaces of the bi-valved node, the remainder of tissue is the surrounding fatty tissue in which the node was originally embedded.

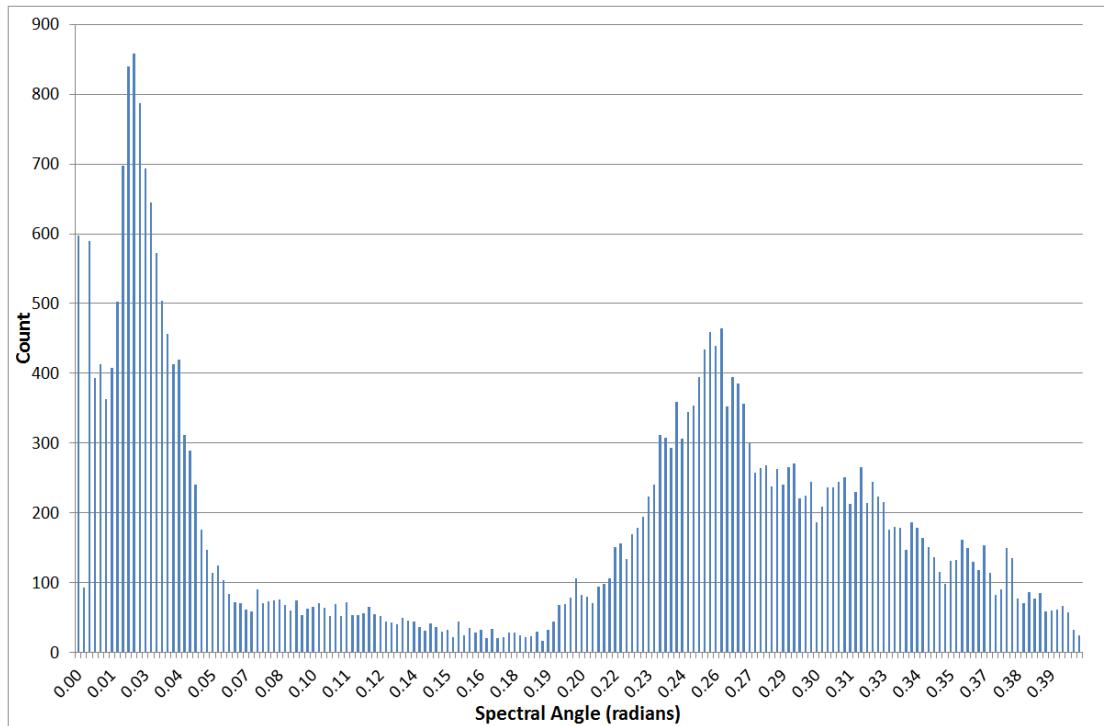


FIGURE B.18: The histogram of spectral angles present in the SAM for the node in figure B.17.

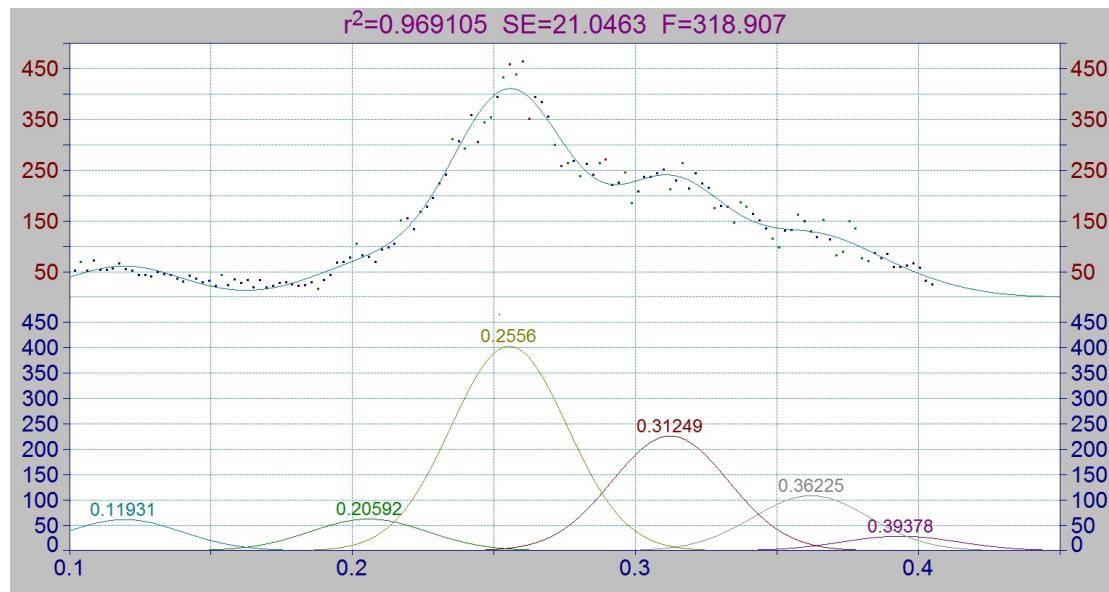


FIGURE B.19: Gaussian peaks automatically fitted to the histogram of spectral angles.

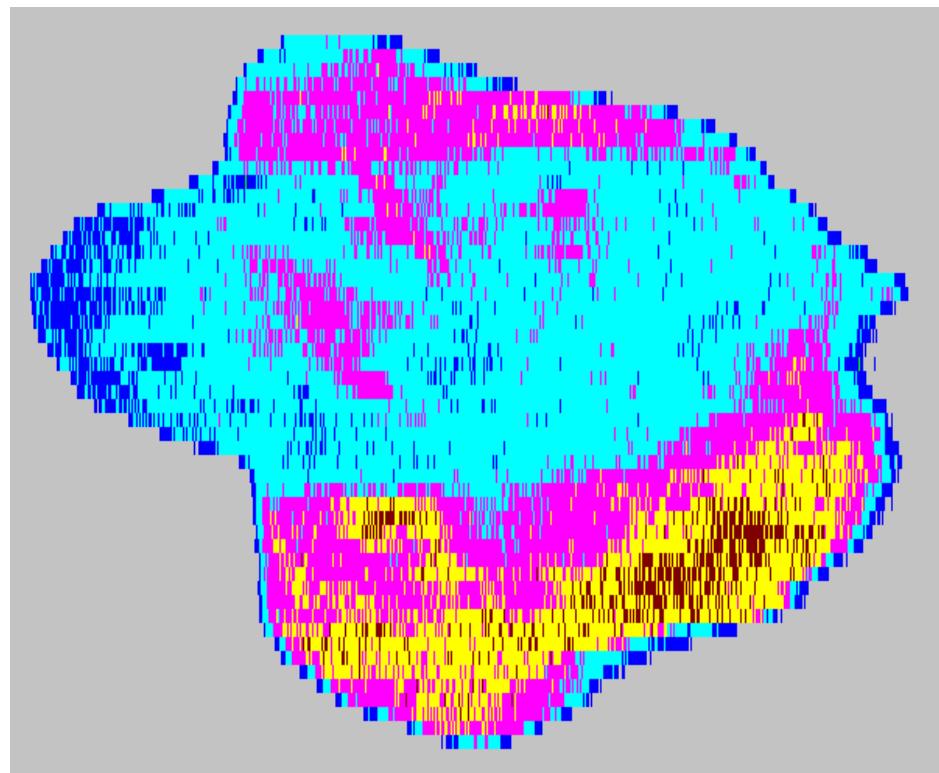


FIGURE B.20: Spectral image classified by spectral angle. Here the node is clearly marked by groups 3 and 4, and although group 5 largely covers the blue stained surrounding tissue, it also covers some of the node tissue.

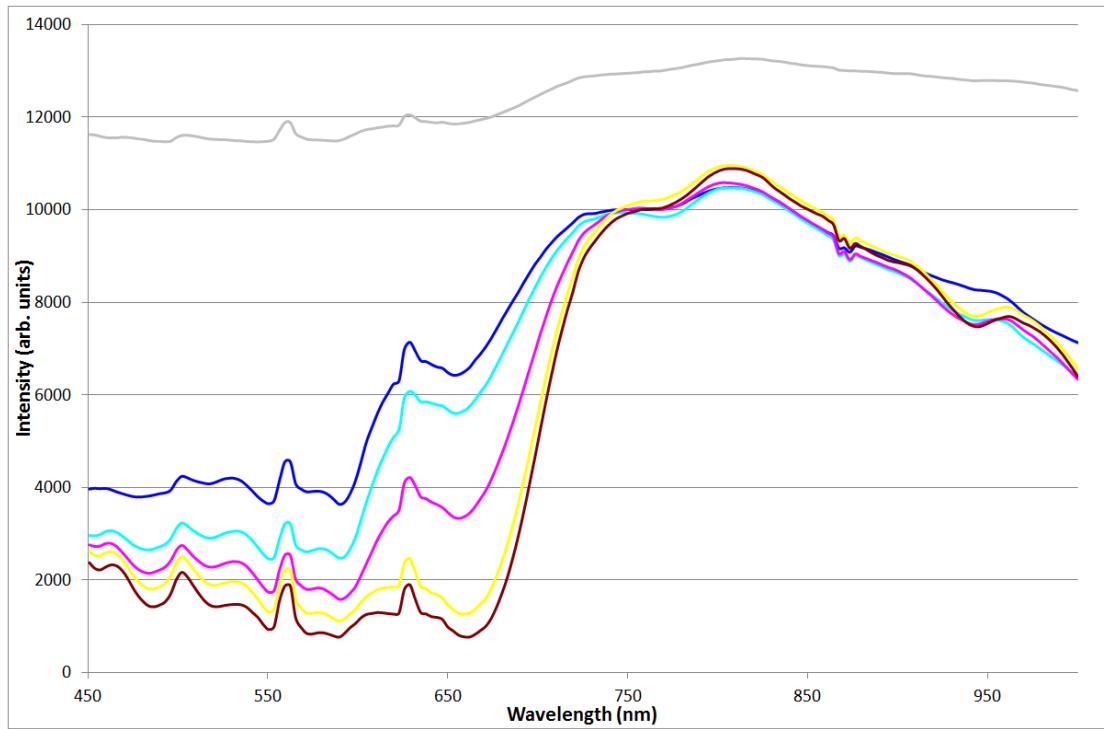


FIGURE B.21: Average spectra of the classes from figure B.20. The spectra for groups 6 and 7 show the steep rise above 650nm that is consistent with blue staining. As might be expected, for group 5 this rise is somewhere between that for 6 and 7 and that for 4.

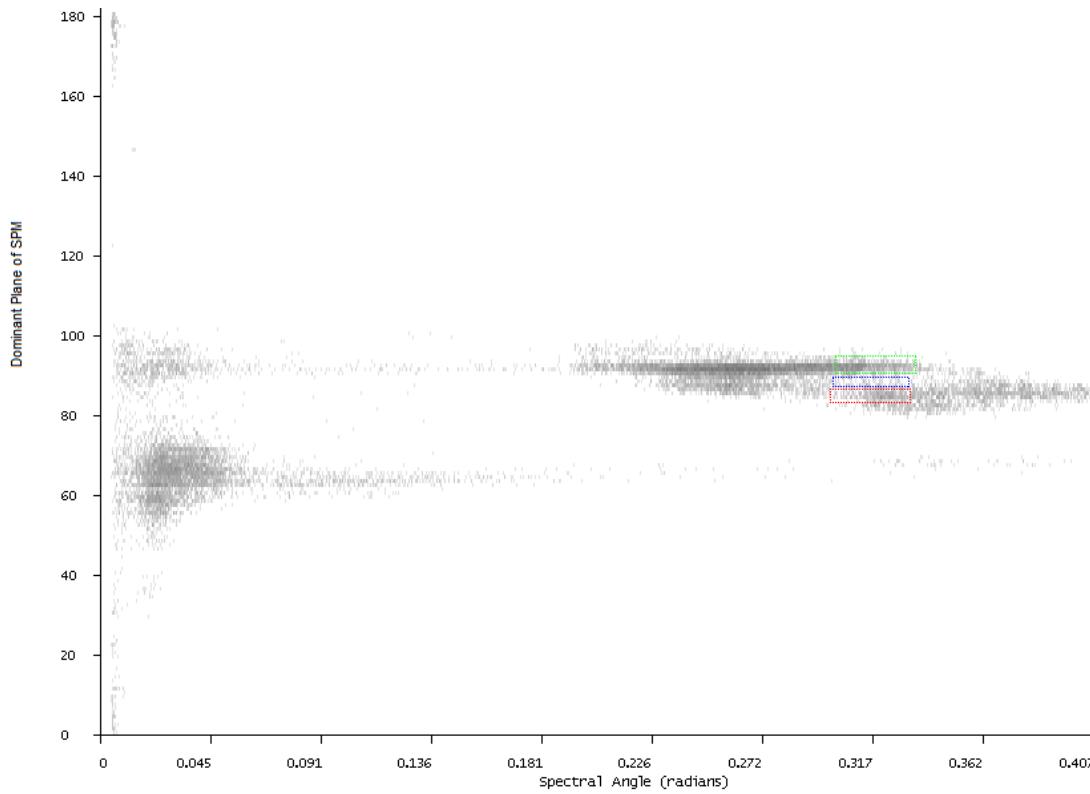


FIGURE B.22: 2D histogram of spectral angle against dominant plane of SPM, with some areas of interest marked in coloured rectangles.

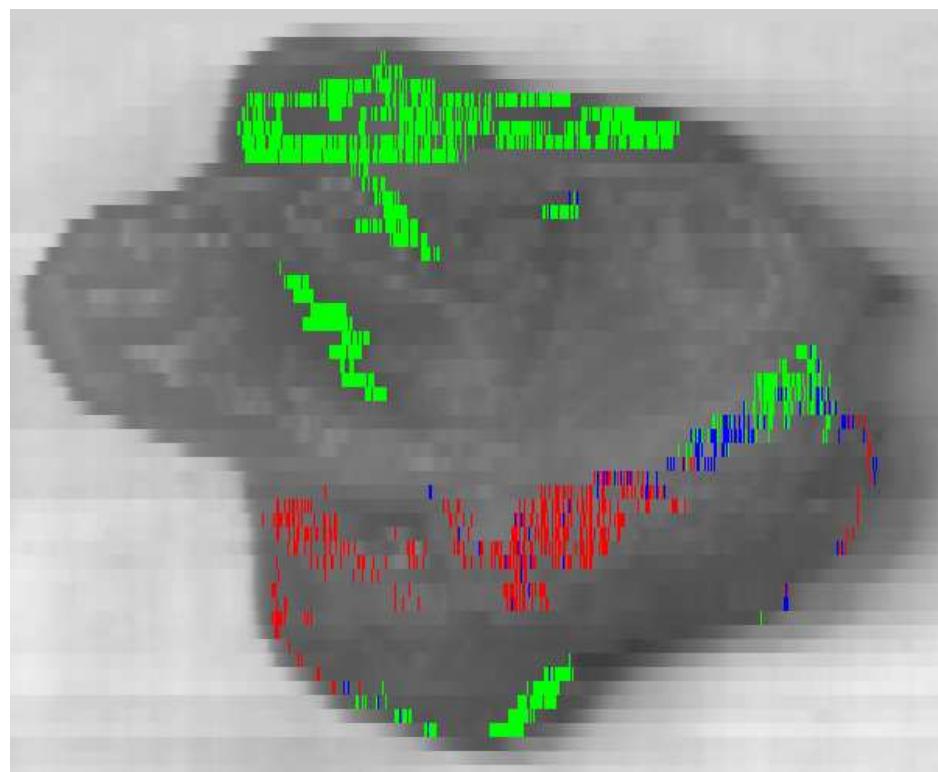


FIGURE B.23: Partially classified spectral image showing the pixels from the boxes in figure B.22. All three groups cover the same spectral angle range, roughly corresponding to group 5 from figure B.20. Here the red and blue largely cover the blue stained tissue whereas the green pixels come from the unstained node and that of the surrounding tissue which is unstained.

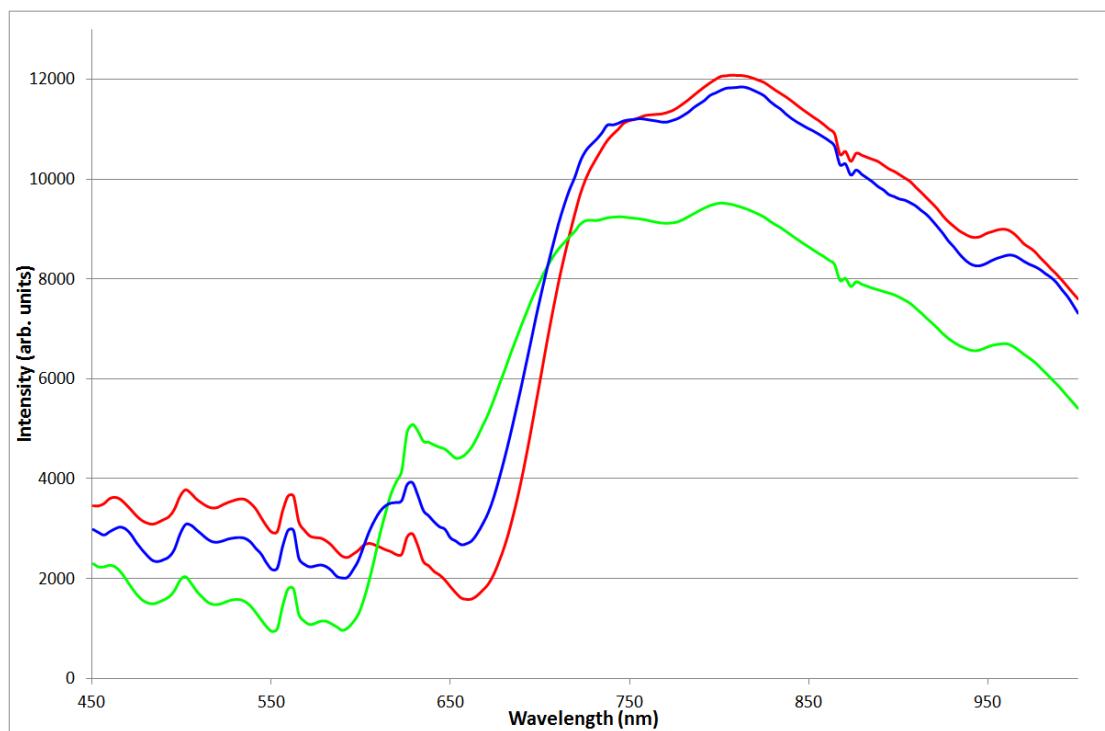


FIGURE B.24: Average spectra of the classes from figure B.23. The steep rise above 650nm for the blue and red spectra again are consistent with the blue staining. All three spectra show strong evidence of the lipid trough below 950nm.

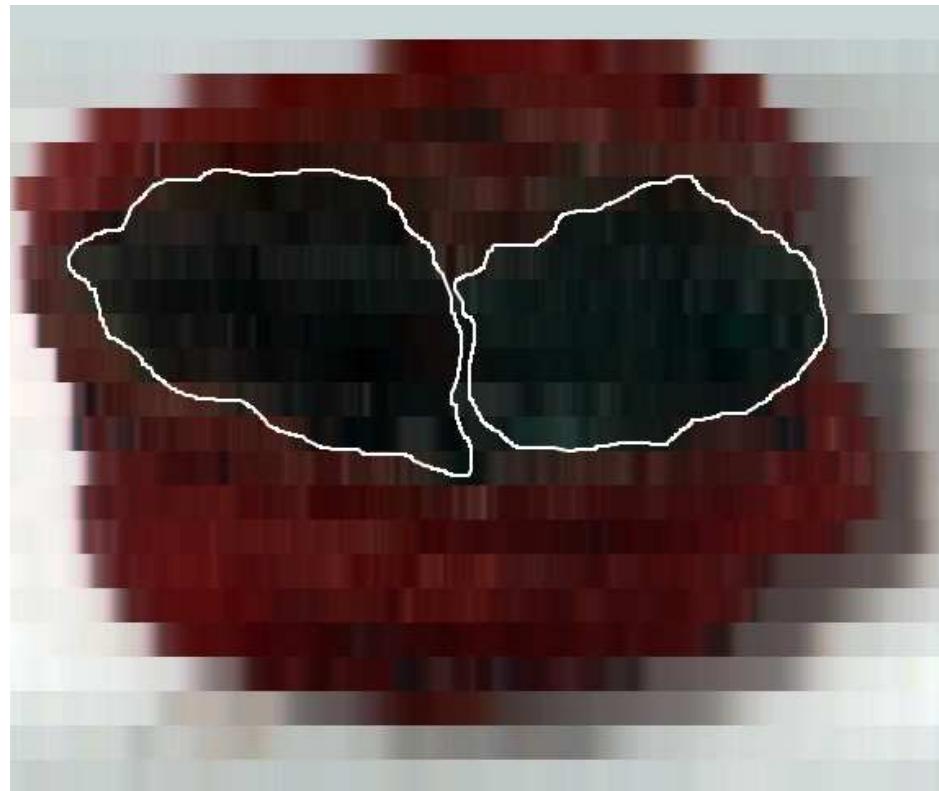


FIGURE B.25: Pseudo-colour image of a resected SLN. Here the lymph node itself is heavily stained blue with the surrounding tissue largely unstained. The approximate extent of the node is shown by white lines, marking the two cut surfaces of the bi-valved node, the remainder of tissue is the surrounding fatty tissue in which the node was originally embedded.

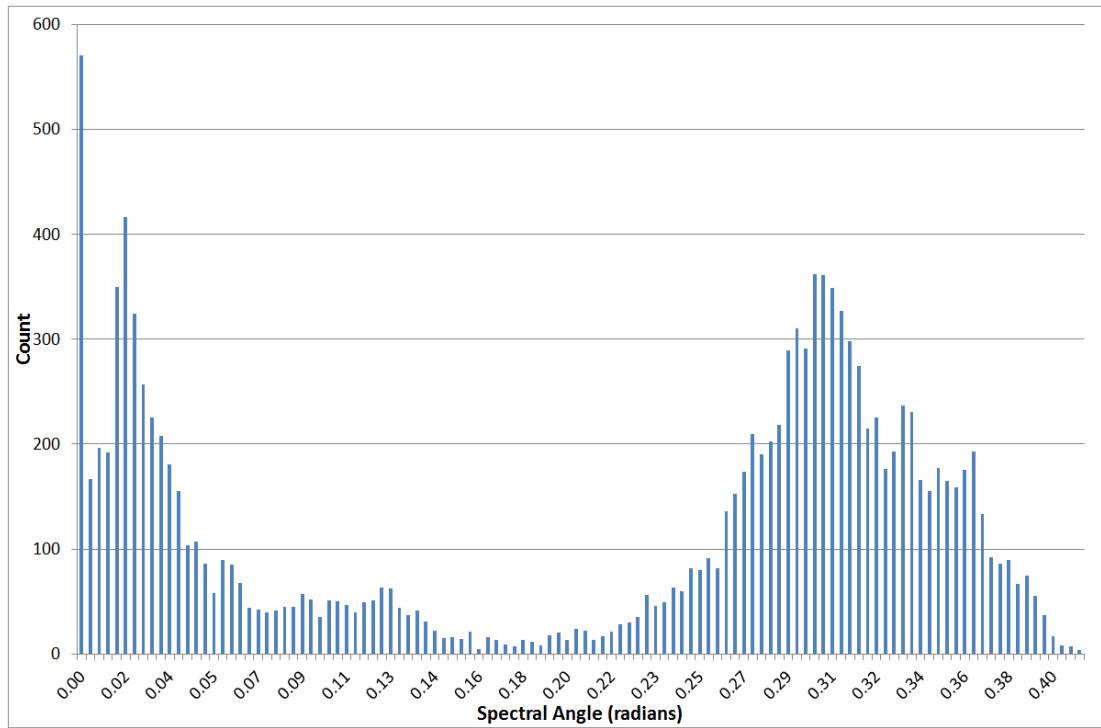


FIGURE B.26: The histogram of spectral angles present in the SAM for the node in figure B.25.

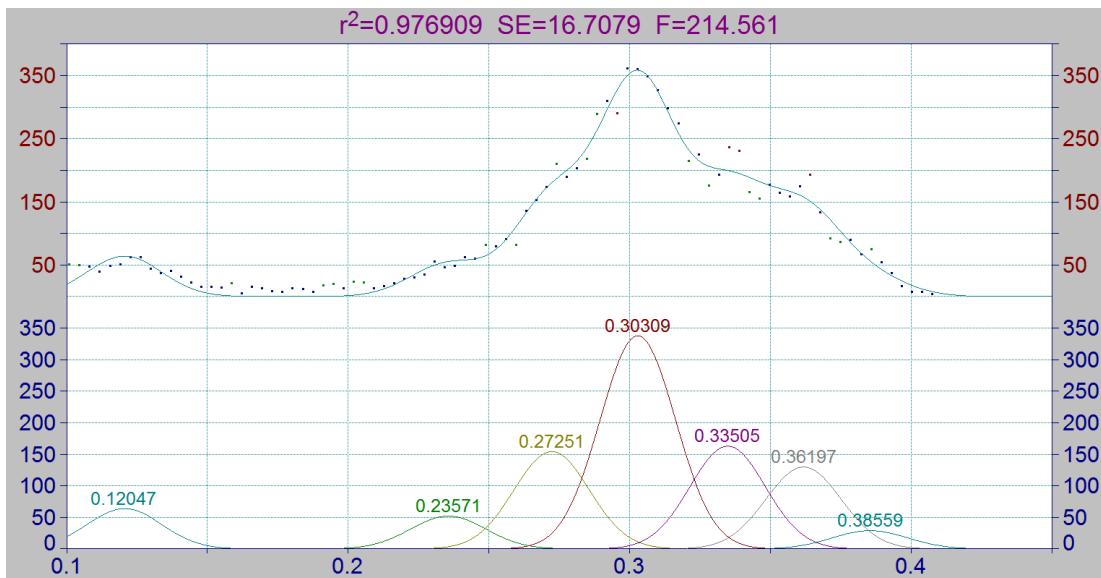


FIGURE B.27: Gaussian peaks automatically fitted to the histogram of spectral angles.

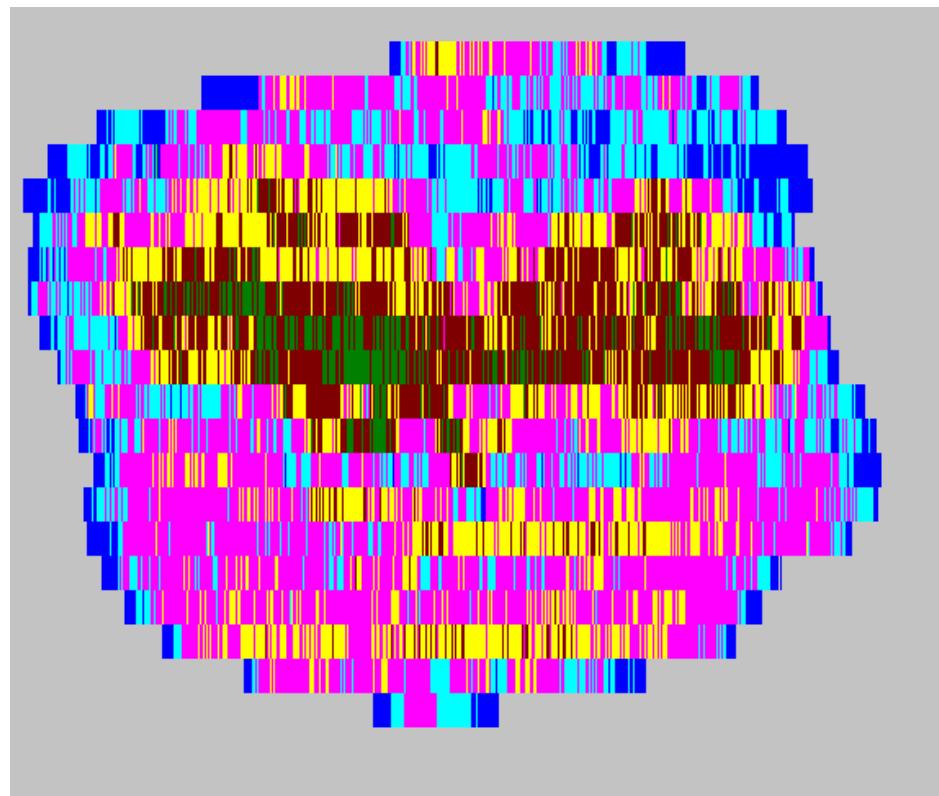


FIGURE B.28: Spectral image classified by spectral angle. The node tissue here is largely covered by groups 6, 7 and 8, although group 6 is also quite heavily present in the surrounding tissue.

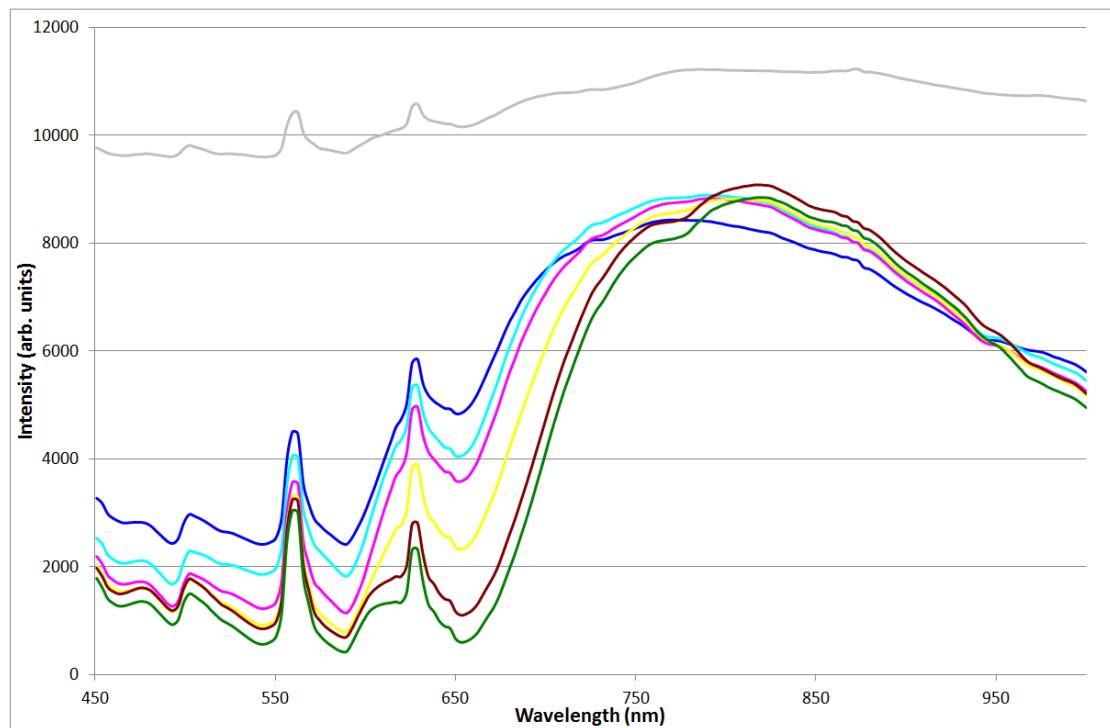


FIGURE B.29: Average spectra of the classes found in figure B.28. The spectra for groups 7 and 8 again show a steep rise above 650nm consistent with blue staining.

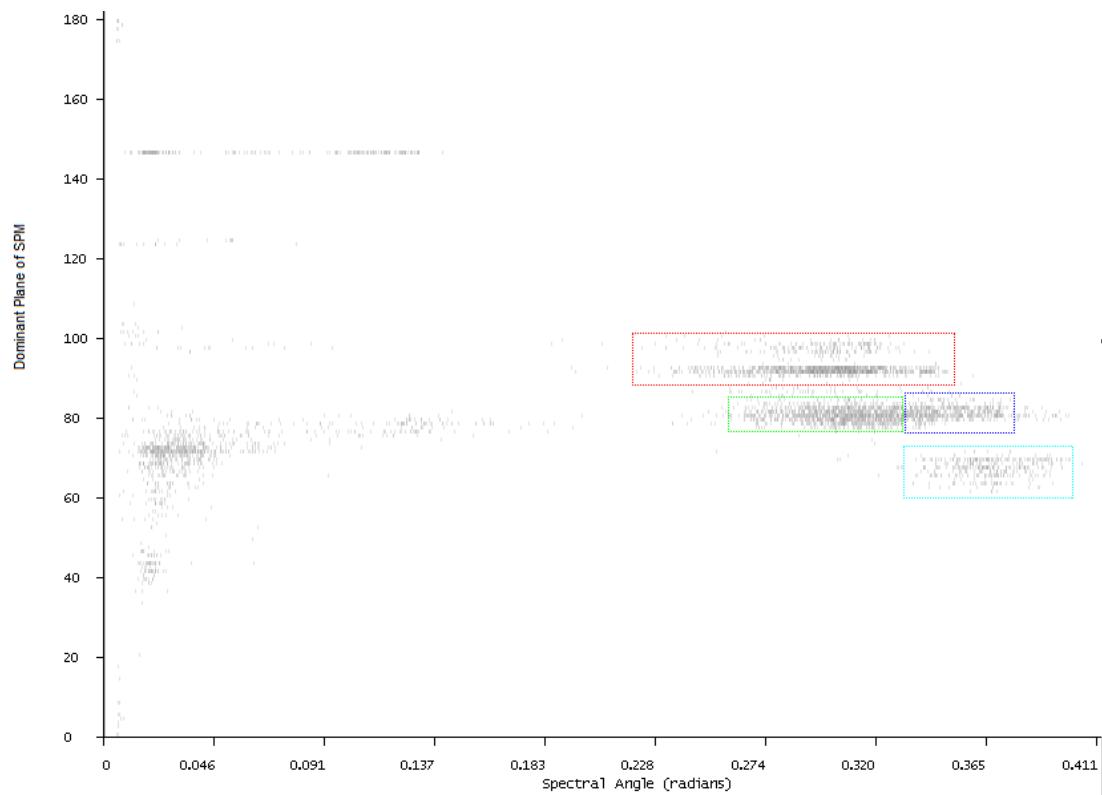


FIGURE B.30: 2D histogram of spectral angle against dominant plane of SPM, with some areas of interest marked in coloured rectangles. The areas here cover most of the biological tissue in rather large crude groups that are well separated in dominant plane of SPM but not so well in spectral angle.

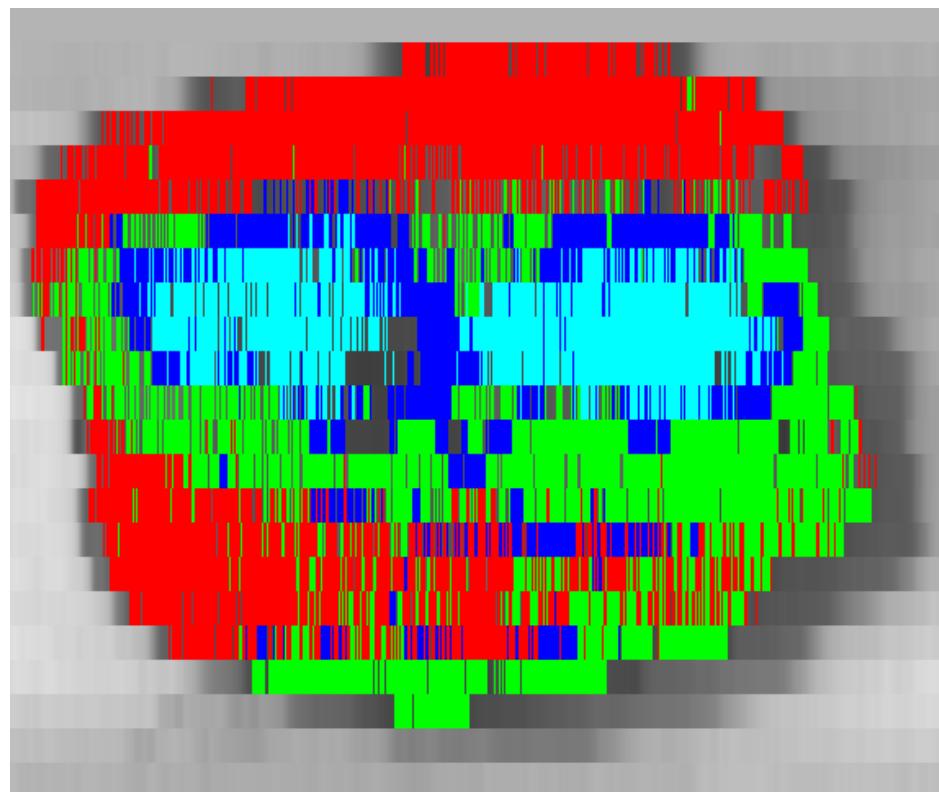


FIGURE B.31: Partially classified spectral image showing the pixels in the boxes from figure B.30. The red pixels here are purely drawn from the surrounding tissue and cover areas that were classified as node in figure B.28. The blue pixels cover some of those that were classified in group 6 from figure B.28 and whilst there are still blue pixels in the surrounding tissue here, the crossover is reduced compared to group 6.

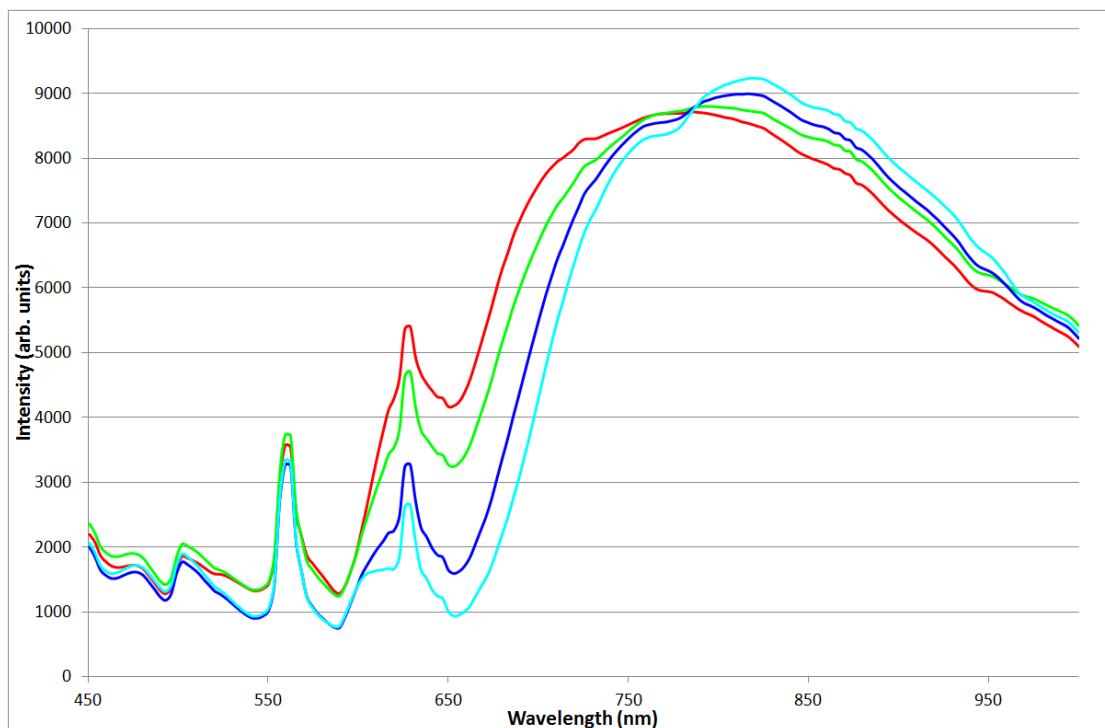


FIGURE B.32: Average spectra of the groups shown in figure B.31. The blue and cyan pixels clearly show the steep rise above 650nm that is consistent with blue staining.

Appendix C

High Dimensional Image Formats

The High Dimensional Image (HDI) format is a generic data container for both multi- and hyper-spectral images. There are two versions of the format, the HDI, and the HDIx (High Dimensional Image eXtension). The HDI is a 32bit format and as such can only deal with file sizes up to around 4GB, the HDIx is a 64bit format and for all practical purposes has no upper size limit (the theoretical upper limit is around 20 exabytes (20×10^{18} bytes)). Both format consist of a header, a table of contents and a number of blocks of data.

Where 0x prefixes a number it signifies that the number is written in hexadecimal format. This is a convenient way to write numbers when dealing with computer storage as one hexadecimal digit represents 4 bits, 1 byte is therefore a 2 digit hexadecimal. A 32bit integer can thus be written as 4 two digit numbers, for example the number 1,000,000, becomes 0x000F4240 .

The `ints` in these files, are stored in little-endian format, i.e. the least significant bit is stored first. When viewing the file in hex mode, the number 1,000,000 would appear as 0x40 42 0F 00.

C.1 HDI

C.1.1 File Header

The header length is given here as a checksum value by which the reading program can ensure that the header has been correctly written.

Bytes	Format	Contents	Comments
0-3	char	“HDI ”	ASCII characters identifying the file as HDI format
4-7	int32	0x18	Length of header in bytes
8-11	int32		Image height in pixels
12-15	int32		Image width in pixels
16-19	int32		Number of spectral channels
20-23	int32		Bit depth of the image

TABLE C.1: Layout of the HDI file header

C.1.2 Table of Contents

The table of contents begins at byte 24. For a valid image, at least the image data must be stored, there are then four further possible data blocks that can be written, one for metric information, one for classification information, one for a mask layer, and one for a list of channels. All further values shall be given as offsets from the beginning of the data block in which they are present, assuming a file that has been fully populated, and lengths (in bytes). Any number which is not fixed, but instead given here a default value is considered volatile and will be suffixed v.

Offset	Length	Format	Contents	Comments
0	4	int32	0x2Cv	Table length in bytes
4	4	int32	0x01	Identifier for image data block
8	4	int32		Address of the image data block (in bytes from the beginning of the file)
12	4	int32	0x02	Identifier for the metric data block, i.e. SAM
16v	4	int32		Address of the metric data block (in bytes from the beginning of the file)
20v	4	int32	0x03	Identifier for the class data block
24v	4	int32		Address of the class data block (in bytes from the beginning of the file)
28v	4	int32	0x04	Identifier for the mask data block
32v	4	int32		Address of the mask data block (in bytes from the beginning of the file)
36v	4	int32	0x05	Identifier for the channel data block
40v	4	int32		Address of the channel data block (in bytes from the beginning of the file)

TABLE C.2: Layout of the HDI table of contents

Additional identifiers can be defined to extend the number of data blocks that the file can store, thus making the file an extensible format.

Offset	Length	Format	Contents	Comments
0	4	int32	0x01	Identifier for the image data block
4	4	int32	0x0C	Length of the image data block header
8	4	int32	$h \times w \times c \times 4$	Length of the image data block
12	$h \times w \times c \times 4$	Uint32 array		Image data

TABLE C.3: Layout of the HDI image data block

Offset	Length	Format	Contents	Comments
0	4	int32	0x02	Identifier for the metric data block
4	4	int32	0x10	Length of the metric data block header
8	4	int32		Length of the metric data block
12	4	int32		Identifier for type of metric stored
16		float array		Metric data

TABLE C.4: Layout of the HDI metric data block

C.1.3 Image Data Block

h = image height, w = image width, c = number of channels.

Image data is written pixel wise, in rows. For an N channel image, the first N 4 byte integers represent the values of the first pixel in each channel, the second N 4 byte integers represent the values of the second pixel in each channel, and so on. The first w sets of N integers are the pixels from the first row of the image, and so on.

C.1.4 Metric Data Block

There are three fundamental types of metric data arrays. The first deals is a single plane of metric values, one per image pixel, written in rows. The second has three planes, each of which measures a metric with respect to a different reference and the third has c planes and stores the full SAM and SPM for each image pixel. For the second and third type, metric data is written plane wise, in rows.

C.1.5 Class Data Block

The class number for each pixel is stored by row.

Offset	Length	Format	Contents	Comments
0	4	int32	0x03	Identifier for the class data block
4	4	int32	0x10	Length of the class data block header
8	4	int32	$h \times w \times 4$	Length of the class data block
12	4	int32		Identifier for the type of classification performed
16	$h \times w \times 4$	int32 array		Class data

TABLE C.5: Layout of the HDI class data block

Offset	Length	Format	Contents	Comments
0	4	int32	0x04	Identifier for the mask data block
4	4	int32	0x0C	Length of the mask data block header
8	4	int32	$h \times w \times 4$	Length of the mask data block
12	$h \times w \times 4$	Uint32 array		Image mask data

TABLE C.6: Layout of the HDI image mask data block

C.1.6 Mask Data Block

A single plane, binary image mask is stored, by row.

C.1.7 Channel Data Block

Offset	Length	Format	Contents	Comments
0	4	int32	0x05	Identifier for the channel data block
4	4	int32	0x0C	Length of the channel data block header
8	4	int32	c	Length of the channel data block
12	$c \times 4$	float array		Channel data

TABLE C.7: Layout of the HDI channel data block

A label for each channel is stored in the same order as the channels of data.

C.2 HDIx

C.2.1 File Header

The header length is given here as a checksum value by which the reading program can ensure that the header has been correctly written.

Bytes	Format	Contents	Comments
0-3	char	“HDI ”	ASCII characters identifying the file as HDI format
4-7	Uint32	0x18	Length of header in bytes
8-11	int32		Image height in pixels
12-15	int32		Image width in pixels
16-19	int32		Number of spectral channels
20-23	int32		Bit depth of the image

TABLE C.8: Layout of the HDIx file header

C.2.2 Table of Contents

The table of contents begins at byte 24. For a valid image, at least the image data must be stored, there are then four further possible data blocks that can be written, one for metric information, one for classification information, one for a mask layer, and one for a list of channels. All further values shall be given as offsets from the beginning of the data block in which they are present, assuming a file that has been fully populated, and lengths (in bytes). Any number which is not fixed, but instead given here a default value is considered volatile and will be suffixed v.

Offset	Length	Format	Contents	Comments
0	4	int32	0x40v	Table length in bytes
4	4	Uint32	0x01	Identifier for image data block
8	8	int64		Address of the image data block (in bytes from the beginning of the file)
16	4	Uint32	0x02	Identifier for the metric data block, i.e. SAM
20v	8	int64		Address of the metric data block (in bytes from the beginning of the file)
28v	4	Uint32	0x03	Identifier for the class data block
32v	8	int64		Address of the class data block (in bytes from the beginning of the file)
40v	4	Uint32	0x04	Identifier for the mask data block
44v	8	int64		Address of the mask data block (in bytes from the beginning of the file)
52v	4	Uint32	0x05	Identifier for the channel data block
56v	8	int64		Address of the channel data block (in bytes from the beginning of the file)

TABLE C.9: Layout of the HDIx table of contents

Additional identifiers can be defined to extend the number of data blocks that the file can store, thus making the file an extensible format.

Offset	Length	Format	Contents	Comments
0	4	Uint32	0x01	Identifier for the image data block
4	4	Uint32	0x10	Length of the image data block header
8	8	int64	$h \times w \times c \times 4$	Length of the image data block
16	$h \times w \times c \times 4$	Uint32 array		Image data

TABLE C.10: Layout of the HDIx image data block

Offset	Length	Format	Contents	Comments
0	4	Uint32	0x02	Identifier for the metric data block
4	4	Uint32	0x14	Length of the metric data block header
8	8	int64		Length of the metric data block
16	4	int32		Identifier for type of metric stored
20		float array		Metric data

TABLE C.11: Layout of the HDIx metric data block

C.2.3 Image Data Block

h = image height, w = image width, c = number of channels.

Image data is written pixel wise, in rows. For an N channel image, the first N 4 byte integers represent the values of the first pixel in each channel, the second N 4 byte integers represent the values of the second pixel in each channel, and so on. The first w sets of N integers are the pixels from the first row of the image, and so on.

C.2.4 Metric Data Block

There are three fundamental types of metric data arrays. The first deals is a single plane of metric values, one per image pixel, written in rows. The second has three planes, each of which measures a metric with respect to a different reference and the third has c planes and stores the full SAM and SPM for each image pixel. For the second and third type, metric data is written plane wise, in rows.

C.2.5 Class Data Block

The class number for each pixel is stored by row.

Offset	Length	Format	Contents	Comments
0	4	Uint32	0x03	Identifier for the class data block
4	4	Uint32	0x14	Length of the class data block header
8	8	int64	$h \times w \times 4$	Length of the class data block
16	4	int32		Identifier for the type of classification performed
20	$h \times w \times 4$	int32 array		Class data

TABLE C.12: Layout of the HDIx class data block

Offset	Length	Format	Contents	Comments
0	4	Uint32	0x04	Identifier for the mask data block
4	4	Uint32	0x10	Length of the mask data block header
8	8	int64	$h \times w \times 4$	Length of the mask data block
16	$h \times w \times 4$	Uint32 array		Image mask data

TABLE C.13: Layout of the HDIx image mask data block

C.2.6 Mask Data Block

A single plane, binary image mask is stored, by row.

C.2.7 Channel Data Block

Offset	Length	Format	Contents	Comments
0	4	Uint32	0x05	Identifier for the channel data block
4	4	Uint32	0x10	Length of the channel data block header
8	8	int64	c	Length of the channel data block
16	$c \times 4$	float array		Channel data

TABLE C.14: Layout of the HDIx channel data block

A label for each channel is stored in the same order as the channels of data.

Bibliography

R. Rox Anderson. Polarized light examination and photography of the skin. *Arch Dermatol*, 127(7):1000–1005, 1991.

Paul G.L. Andrus and Robert D. Strickland. Cancer grading by fourier transform infrared spectroscopy. *Biospectroscopy*, 4(1):37–46, 1998.

P. Asawanonda, R. Rox Anderson, Y. Chang, and C. R. Taylor. 308-nm excimer laser for the treatment of psoriasis. *Archives of Dermatology*, 136:619–624, 2000.

M. R. Austwick, B. Clark, C. A. Mosse, K. Johnson, D. W. Chicken, S. K. Somasundaram, K. W. Calabro, Y. Zhu, M. Falzon, G. Kocjan, T. Fearn, S. G. Bown, I. J. Bigio, and M. R. S. Keshtgar. Scanning elastic scattering spectroscopy detects metastatic breast cancer in sentinel lymph nodes. *Journal of Biomedical Optics*, 15(4), 2010.

M.J. Bernhardt. *Laser Resurfacing*. 1998. <http://www.dcmsonline.org/jax-medicine/1998journals/february98/laserresurfacing.htm>. *date accessed, 19/08/2011*.

A. Bhattacharyya. On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society*, 35:99–109, 1943.

I. J. Bigio, S. G. Bown, G. Briggs, C. Kelley, S. Lakhani, D. Pickard, P. M. Ripley, I. G. Rose, and C. Saunders. Diagnosis of breast cancer using elastic-scattering spectroscopy: preliminary clinical results. *Journal of Biomedical Optics*, 5(2):221–228, 2000.

B. Bird, M. Romeo, N. Laver, and M. Diem. Spectral detection of micro-metastases in lymph node histo-pathology. *Journal of Biophotonics*, 2(1-2):37–46, 2009.

Benjamin Bird, Milos Miljkovic, Melissa J. Romeo, Jennifer Smith, Nicholas Stone, Michael W. George, and Max Diem. Infrared micro-spectral imaging: distinction of tissue types in axillary lymph node histology. *BMC Clin Pathol*, 8:8, 2008.

R. Bissonnette, H. S. Zeng, D. I. McLean, W. E. Schreiber, D. L. Roscoe, and H. Lui. Psoriatic plaques exhibit red autofluorescence that is due to protoporphyrin IX. *Journal of Investigative Dermatology*, 111(4):586–591, 1998.

S. Blit, E. G. Weaver, T. A. Rabson, and F. K. Tittel. Continuous wave UV radiation tunable from 285 nm to 400 nm by harmonic and sum frequency generation. *Appl. Opt.*, 17(5):721–723, 1978.

Cancer Research UK. *CancerStats*. 2011.
<http://info.cancerresearchuk.org/cancerstats/>. *date accessed, 11/08/2011*.

Malavika Chandra, Robert H. Wilson, James Scheiman, Diane Simeone, Barbara McKenna, Julianne Purdy, and Mary-Ann Mycek. Optical spectroscopy for clinical detection of pancreatic cancer. In Georgakoudi Irene, Popp Jurgen, and Svanberg Katarina, editors, *SPIE European Conferences on Biomedical Optics*, volume 7368, page 73681G. SPIE, 2009.

Chein I. Chang. Spectral information divergence for hyperspectral image analysis. In *Geoscience and Remote Sensing Symposium, 1999. IGARSS '99 Proceedings. IEEE 1999 International*, volume 1, pages 509–511 vol.1, 1999.

G. Cheng, S. Kurita, D. A. Torigian, and A. Alavi. Current status of sentinel lymph-node biopsy in patients with breast cancer. *European Journal of Nuclear Medicine and Molecular Imaging*, 38(3):562–575, 2011.

H.S. Chung, A.E. Cerussi, S.I. Merritt, J. Ruth, and B. J. Tromberg. Non-invasive tissue temperature measurements based on quantitative diffuse optical spectroscopy (dos) of water. *Physics in Medicine and Biology*, 55(13):3753, 2010.

Clinuvel Pharmaceuticals. *Melanin: Function Skin Pigment*. 2009.
<http://www.clinuvel.com/dermatology/melanin/function-skin-pigment>. *date accessed, 02/01/2009*.

J. M. Conway, K. H. Norris, and C. E. Bodwell. A new approach for the estimation of body composition: infrared interactance. *Am J Clin Nutr*, 40(6):1123–1130, 1984.

Alvaro P. Crósta, Charles Sabine, and James V. Taranik. Hydrothermal alteration mapping at Bodie, California, using AVIRIS hyperspectral data. *Remote Sensing of Environment*, 65(3):309–319, 1998.

S. C. Davison, A. Ballsdon, M. H. Allen, and J. N. W. N. Barker. Early migration of cutaneous lymphocyte-associated antigen (CLA) positive T cells into evolving psoriatic plaques. *Experimental Dermatology*, 10(4):280–285, 2001.

O.A. De Carvalho and P.R. Meneses. Spectral correlation mapper: An improvement on the spectral angle mapper. Technical report, JPL Publication 00-18, 2000.

Laurent de Schoulenikoff, Valentin Mitev, Valentin Simeonov, Bertrand Calpini, and Hubert van den Bergh. Experimental investigation of high-power single-pass Raman shifters in the ultraviolet with Nd:YAG and KrF lasers. *Appl. Opt.*, 36(21): 5026–5043, 1997.

David Delgado Gomez, Constantine Butakoff, Bjarne Ersbll, and Jens Michael Carstensen. Automatic change detection and quantification of dermatological diseases with an application to psoriasis images. *Pattern Recognition Letters*, 28(9): 1012–1018, 2007.

B. Diffey. Conformal phototherapy (private correspondence), 2008.

C. Elwell and J. Hebden. *NIR Research - Biomedical Optics Research Laboratory*. 2008. <http://www.brainhelpers.com/NIR-Research.html>. *date accessed, 02/01/2009*.

V. V. Ermolenkov, V. A. Lisinetskii, Y. I. Mishkel, A. S. Grabchikov, A. P. Chaikovskii, and V. A. Orlovich. A radiation source based on a solid-state raman laser for diagnosing tropospheric ozone. *Journal of Optical Technology*, 72(1):32–36, 2005.

H. Fabian, P. Lasch, M. Boese, and W. Haensch. Infrared microspectroscopic imaging of benign breast tumor tissue sections. *Journal of Molecular Structure*, 661-662: 411–417, 2003.

Heinz Fabian, Ngoc Anh Ngo Thi, Michael Eiden, Peter Lasch, Jurgen Schmitt, and Dieter Naumann. Diagnosing benign and malignant lesions in breast tissue sections by using IR-microspectroscopy. *Biochimica et Biophysica Acta (BBA) - Biomembranes*, 1758(7):874–882, 2006.

Naoko Fujioka, Yuji Morimoto, Tsunenori Arai, and Makoto Kikuchi. Discrimination between normal and malignant human gastric tissues by fourier transform infrared spectroscopy. *Cancer Detection and Prevention*, 28(1):32–36, 2004.

Alejandro Garcia-Uribe, Nasser Kehtarnavaz, Guillermo Marquez, Victor Prieto, Madeleine Duvic, and Lihong V. Wang. Skin cancer detection by spectroscopic oblique-incidence reflectometry: Classification and physiological origins. *Appl. Opt.*, 43(13):2643–2650, 2004.

P.U. Giacomoni. *Biophysical and Physiological Effects of Solar Radiation on Human Skin*, volume 1. Royal Society of Chemistry, London, 2007.

M. Gniadecka, O. F. Nielsen, and H. C. Wulf. Water content and structure in malignant and benign skin tumours. *Journal of Molecular Structure*, 661-662: 405–410, 2003.

Monika Gniadecka, Peter Alshede Philipsen, Sigurdur Sigurdsson, Sonja Wessel, Ole Faurskov Nielsen, Daniel Hojgaard Christensen, Jana Hercogova, Kristian Rossen, Henrik Klem Thomsen, Robert Gniadecki, Lars Kai Hansen, and

Hans Christian Wulf. Melanoma diagnosis by raman spectroscopy and neural networks: Structure alterations in proteins and lipids in intact cancer tissue. *J Investig Dermatol*, 122(2):443–449, 2004.

A. Goyal and R. E. Mansel. Recent advances in sentinel lymph node biopsy for breast cancer. *Current Opinion in Oncology*, 20(6):621–626, 2008.

H.F. Grahn and P. Geladi. *Techniques and Applications of Hyperspectral Image Analysis*. John Wiley and Sons Ltd, Chichester, 1st edition, 2007.

Dirk Grosenick, Heidrun Wabnitz, K. Thomas Moesta, Jorg Mucke, Peter M. Schlag, and Herbert Rinneberg. Time-domain scanning optical mammography: II. Optical properties and tissue parameters of 87 carcinomas. *Physics in Medicine and Biology*, 50(11):2451–2468, 2005.

R. Gush and T. King. Discrimination of capillary and arterio-venular blood flow in skin by laser doppler flowmetry. *Medical and Biological Engineering and Computing*, 29(4):387–392, 1991.

D. A. Haner and I. S. McDermid. Stimulated Raman shifting of the Nd:YAG fourth harmonic (266 nm) in H2, HD, and D2. *Quantum Electronics, IEEE Journal of*, 26 (7):1292–1298, 1990.

P. J. Hendra. What is Raman spectroscopy? *Internet Journal of Vibrational Spectroscopy*, 1(5):6–16, 1998.

Snider Howard, Dowlatshahi Kambiz, Fan Ming, M. Bridger William, G. Rayudu, and Oleske Denise. Sentinel node biopsy in the staging of breast cancer. *American Journal of Surgery*, 176(4):305–310, 1998.

P.R. Hoy. *Imaging of Mice and Men; Adventures in Multispectral Imaging*. PhD thesis, University of Southampton, 2009.

Susan Macdonald Hull, Mark Goodfield, Edward John Wood, D. Phil, and William John Cunliffe. Active and inactive edges of psoriatic plaques: Identification by tracing and investigation by laser-doppler flowmetry and immunocytochemical techniques. *J Investig Dermatol*, 92(6):782–785, 1989.

Michael Jackson, Michael G. Sowa, and Henry H. Mantsch. Infrared spectroscopy: a new frontier in medicine. *Biophysical Chemistry*, 68(1-3):109–125, 1997.

K. S. Johnson, D. W. Chicken, D. C. O. Pickard, A. C. Lee, G. Briggs, M. Falzon, I. J. Bigio, M. R. Keshtgar, and S. G. Bown. Elastic scattering spectroscopy for intraoperative determination of sentinel lymph node status in the breast. *Journal of Biomedical Optics*, 9(6):1122–1128, 2004.

H. K. Jusoff. Pixel-based airborne hyperspectral sensing technique for search-and-rescue of the missing RMAF NURI helicopter in Genting-Sempah, Malaysia. *Disaster Prevention and Management*, 19(1):88–102, 2010.

K.S. Kane, J. Bissonette, H.P. Baden, R.A. Johnson, and A.J. Stratigos. *Color Atlas & Synopsis of Pediatric Dermatology*. McGraw-Hill Medical, New York, international edition, 2002.

M. R. S. Keshtgar, D. W. Chicken, M. R. Austwick, S. K. Somasundaram, C. A. Mosse, Y. Zhu, I. J. Bigio, and S. G. Bown. Optical scanning for rapid intraoperative diagnosis of sentinel node metastases in breast cancer. *British Journal of Surgery*, 97(8):1232–1239, 2010.

K. Köllner, M.B. Wimmershof, C. Hintz, M. Landthaler, and U. Hohenleutner. Comparison of the 308-nm excimer laser and a 308-nm excimer lamp with 311-nm narrowband ultraviolet B in the treatment of psoriasis. *British Journal of Dermatology*, 152:750–754, 2005.

C. Krafft, S. B. Sobottka, K. D. Geiger, G. Schackert, and R. Salzer. Classification of malignant gliomas by infrared spectroscopic imaging and linear discriminant analysis. *Analytical and Bioanalytical Chemistry*, 387(5):1669–1677, 2007.

C. Krafft, K. Thummel, S. B. Sobottka, G. Schackert, and R. Salzer. Classification of malignant gliomas by infrared spectroscopy and linear discriminant analysis. *Biopolymers*, 82(4):301–305, 2006.

O. Kreylos. *Oliver Kreylos' homepage*. 2010. <http://www.idav.ucdavis.edu/okreylos/>. date accessed, 11/08/2011.

F. A. Kruse, A. B. Lefkoff, J. W. Boardman, K. B. Heidebrecht, A. T. Shapiro, P. J. Barloon, and A. F. H. Goetz. The spectral image-processing system (SIPS) - interactive visualization and analysis of imaging spectrometer data. *Remote Sensing of Environment*, 44(2-3):145–163, 1993.

S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

D. A Landgrebe. *Signal Theory Methods in Multispectral Remote Sensing*. Wiley Series in Remote Sensing and Image Processing. John Wiley and Sons Ltd, Hoboken, NJ, 1st edition, 2003.

Q. L. Li, W. S. Wang, C. Ma, and Z. Q. Zhu. Detection of physical defects in solar cells by hyperspectral imaging technology. *Optics and Laser Technology*, 42(6):1010–1013, 2010.

T. Loree, R. Sze, D. Barker, and P. Scott. New lines in the UV: SRS of excimer laser wavelengths. *Quantum Electronics, IEEE Journal of*, 15(5):337–342, 1979.

Anna Matas, Michael G. Sowa, Geoff Taylor, and Henry H. Mantsch. Melanin as a confounding factor in near infrared spectroscopy of skin. *Vibrational Spectroscopy*, 28(1):45–52, 2002.

Laura M. McIntosh, Mike Jackson, Henry H. Mantsch, Miroslaw F. Stranc, Dragana Pilavdzic, and A. Neil Crowson. Infrared spectra of basal cell carcinomas are distinct from non-tumor-bearing skin components. *Journal of Investigative Dermatology*, 112(6):951–956, 1999.

Laura M. McIntosh, Randy Summers, Michael Jackson, Henry H. Mantsch, James R. Mansfield, Marilyn Howlett, A. Neil Crowson, and John W. P. Toole. Towards non-invasive screening of skin lesions by near-infrared spectroscopy. *Journal of Investigative Dermatology*, 116(1):175–181, 2001.

M. J. T. Milton, G. Ancellet, A. Apituley, J. Bsenberg, W. Carnuth, F. Castagnoli, T. Trickl, H. Edner, L. Stefanutti, T. Schaberl, A. Sunesson, and C. Weitkamp. Raman-shifted laser sources suitable for differentialabsorption lidar measurements of ozone in the troposphere. *Applied Physics B: Lasers and Optics*, 66(1):105–113, 1998.

Emmanuel Mimoun, Luigi De Sarlo, Jean-Jacques Zondy, Jean Dalibard, and Fabrice Gerbier. Sum-frequency generation of 589 nm light with near-unit efficiency. *Opt. Express*, 16(23):18684–18691, 2008.

Stephen P. Morgan and IanM. Stockford. Surface-reflection elimination in polarization imaging of superficial tissue. *Opt. Lett.*, 28(2):114–116, Jan 2003.

A. K. Murray, A. L. Herrick, T. L. Moore, T. A. King, and C. E. M. Griffiths. Dual wavelength (532 and 633 nm) laser doppler imaging of plaque psoriasis. *British Journal of Dermatology*, 152(6):1182–1186, 2005.

P. H. R. Ng, S. Walker, M. Tahtouh, and B. Reedy. Detection of illicit substances in fingerprints by infrared spectral imaging. *Analytical and Bioanalytical Chemistry*, 394(8):2039–2048, 2009.

M.S. Nixon and A.S. Aguado. *Feature Extraction and Image Processing*. Elsevier Science and Technology, 1st edition, 2002.

D. Pantalone, F. Andreoli, F. Fusi, V. Basile, G. Romano, G. Giustozzi, L. Rigacci, R. Alterini, and M. Monici. Multispectral imaging autofluorescence microscopy in colonic and gastric cancer metastatic lymph nodes. *Clinical Gastroenterology and Hepatology*, 5(2):230–236, 2007.

B. Park, W. R. Windham, K. C. Lawrence, and D. P. Smith. Contaminant classification of poultry hyperspectral imagery using a spectral angle mapper algorithm. *Biosystems Engineering*, 96(3):323–333, 2007a.

Seok Chan Park, Sung Jae Lee, Hankyu Namkung, Hoeil Chung, Sung-Hwan Han, Moon-Young Yoon, Jong-Jae Park, Ju-Han Lee, Chil-Hwan Oh, and Young-Ah Woo. Feasibility study for diagnosis of stomach adenoma and cancer using IR spectroscopy. *Vibrational Spectroscopy*, 44(2):279–285, 2007b.

John A. Parrish and Kurt F. Jaenicke. Action spectrum for phototherapy of psoriasis. *J Investig Dermatol*, 76(5):359–362, 1981.

M. J. Pepels, Jhmj Vestjens, M. de Boer, M. Smidt, P. J. van Diest, G. F. Borm, and V. C. G. Tjan-Heijnen. Safety of avoiding routine use of axillary dissection in early stage breast cancer: a systematic review. *Breast Cancer Research and Treatment*, 125(2):301–313, 2011.

Shazza Rehman, Zanyar Movasaghi, Art T. Tucker, Simon P. Joel, Jawaad A. Darr, Alexander V. Ruban, and Ihtesham U. Rehman. Raman spectroscopic analysis of breast cancer tissues: identifying differences between normal, invasive ductal carcinoma and ductal carcinoma in situ of the breast tissue. *Journal of Raman Spectroscopy*, 38(10):1345–1351, 2007.

Luigi Rigacci, Renato Alterini, Pietro Antonio Bernabei, Pierluigi Rossi Ferrini, Giovanni Agati, Franco Fusi, and Monica Monici. Multispectral imaging autofluorescence microscopy for the analysis of lymph-node tissues. *Photochemistry and Photobiology*, 71(6):737–742, 2000.

Melissa J. Romeo and Max Diem. Infrared spectral imaging of lymph nodes: Strategies for analysis and artifact reduction. *Vibrational Spectroscopy*, 38(1-2):115–119, 2005.

Valentin Simeonov, Valentin Mitev, Hubert van den Bergh, and Bertrand Calpini. Raman frequency shifting in a CH₄:H₂:Ar mixture pumped by the fourth harmonic of a Nd:YAG laser. *Appl. Opt.*, 37(30):7112–7115, 1998.

Youngsinn Sohn, Emilio Moran, and Francesco Gurri. Deforestation in North-Central Yucatan (1985-1995): Mapping secondary succession of forest and agricultural land use in Sotuta using the cosine of the angle concept. *Photogrammetric Engineering and Remote Sensing*, 65(8):947 – 958, 1999.

Youngsinn Sohn and N. Sanjay Rebello. Supervised and unsupervised spectral angle classifiers. *Photogrammetric Engineering and Remote Sensing*, 68(12):1271–1280, 2002.

Sony Electronics. ICX285AL CCD specification sheet. Technical report, Sony Corporation, 2003.

G. G. M. Stoffels, P. Schmidt, N. Dam, and J. J. terMeulen. Generation of 224-nm radiation by stimulated raman scattering of ArF excimer laser radiation in a mixture of H-2 and D-2. *Applied Optics*, 36(27):6797–6801, 1997.

H.A. Sturges. The choice of a class interval. *J. American Statistical Association*, 21: 65–66, 1926.

BJ Tromberg, A Cerussi, N Shah, M Compton, A Durkin, D Hsiang, J Butler, and R. Mehta. Imaging in breast cancer: Diffuse optics in breast cancer: detecting tumors in pre-menopausal women and monitoring neoadjuvant chemotherapy. *Breast Cancer Research*, 7(6):279–285, 2005.

Vo-Dinh Tuan. A hyperspectral imaging system for in vivo optical diagnostics. *Engineering in Medicine and Biology Magazine, IEEE*, 23(5):40–49, 2004.

UCL Biomedical Optics Research Laboratory. *Shedding Light on the Human Body*. 2006. <http://www.medphys.ucl.ac.uk/research/borl/sheddinglight/background.htm>. *date accessed, 02/01/2009*.

Y. Uesugi, Y. Mizutani, S.G. Kruglik, A. G. Shvedko, V.A. Orlovich, and T. Kitagawa. Characterization of stimulated raman scattering of hydrogen and methane gases as a light source for picosecond time-resolved raman spectroscopy. *Journal of Raman Spectroscopy*, 31(4):339–348, 2000.

University of Joensuu Color Group. *Spectral Database*. 2006. <http://spectral.joensuu.fi/multispectral/spectralimages.php>. *date accessed, 26/11/2010*.

US Food and Drug Administration. FDA Executive Summary for the GeneSearch BLN assay. Technical report, United States Food and Drug Administration, 2006.

I. M. C. van der Ploeg, O. E. Nieweg, M. C. van Rijk, R. A. Valds Olmos, and B. B. R. Kroon. Axillary recurrence after a tumour-negative sentinel node biopsy in breast cancer patients: A systematic review and meta-analysis of the literature. *European Journal of Surgical Oncology (EJSO)*, 34(12):1277–1284, 2008.

Umberto Veronesi, Giovanni Paganelli, Giuseppe Viale, Alberto Luini, Stefano Zurruda, Viviana Galimberti, Mattia Intra, Paolo Veronesi, Chris Robertson, Patrick Maisonneuve, Giuseppe Renne, Concetta De Cicco, Francesca De Lucia, and Roberto Gennari. A randomized comparison of sentinel-node biopsy with routine axillary dissection in breast cancer. *N Engl J Med*, 349(6):546–553, 2003.

V. P. Wallace, A. J. Fitzgerald, S. Shankar, N. Flanagan, R. Pye, J. Cluff, and D. D. Arnone. Terahertz pulsed imaging of basal cell carcinoma ex vivo and in vivo. *British Journal of Dermatology*, 151(2):424–432, 2004.

Robert H. Wilson, Malavika Chandra, Leng-Chun Chen, William R. Lloyd, James Scheiman, Diane Simeone, Julianne Purdy, Barbara McKenna, and Mary-Ann Mycek. Photon-tissue interaction model enables quantitative optical analysis of human pancreatic tissues. *Optics Express*, 18(21):21612–21621, 2010.

Robert H. Wilson, Malavika Chandra, William Lloyd, Leng-Chun Chen, James Scheiman, Diane Simeone, Barbara McKenna, and Mary-Ann Mycek. Optical spectroscopy for quantitative sensing in human pancreatic tissues. In Nirmala Ramanujam and Jürgen Popp, editors, *Clinical and Biomedical Spectroscopy and Imaging II*, volume 8087. SPIE, 2011.

Robert H. Wilson, Malavika Chandra, James Scheiman, Diane Simeone, Barbara McKenna, Julianne Purdy, and Mary-Ann Mycek. Mathematical modeling of reflectance and intrinsic fluorescence for cancer detection in human pancreatic tissue. In Wax Adam and Backman Vadim, editors, *Photonics West: Biomedical Optics*, volume 7187, page 71870H. SPIE, 2009.

J. C. F. Wong and Alfio Parisi. Assessment of ultraviolet radiation exposures in photobiological experiments, 1999.

Patrick T. T. Wong, Sanford M. Goldstein, Roy C. Grekin, Thomas A. Godwin, Chris Pivik, and Basil Rigas. Distinct infrared spectroscopic patterns of human basal cell carcinoma of the skin. *Cancer Res*, 53(4):762–765, 1993.

R. M. Woodward, B. E. Cole, V. P. Wallace, R. J. Pye, D. D. Arnone, E. H. Linfield, and M. Pepper. Terahertz pulse imaging in reflection geometry of human skin cancer and skin tissue. *Physics in Medicine and Biology*, 47(21):3853, 2002.

R. M. Woodward, V. P. Wallace, D. D. Arnone, E. H. Linfield, and M. Pepper. Terahertz pulsed imaging of skin cancer in the time and frequency domain. *Journal of Biological Physics*, 29(2):257–259, 2003.

B. Xu, G. M. Yue, Y. C. Zhang, H. L. Hu, J. Zhou, and S. X. Hu. Generation of UV laser light by stimulated raman scattering in D-2, D-2/Ar and D-2/He using a pulsed Nd : YAG laser at 355nm. *Chinese Physics*, 12(9):1021–1025, 2003.

W. G. Zijlstra, A. Buursma, and W. P. Meeuwesen-van der Roest. Absorption spectra of human fetal and adult oxyhemoglobin, de- oxyhemoglobin, carboxyhemoglobin, and methemoglobin. *Clin Chem*, 37(9):1633–1638, 1991.

K.J. Zuzak, M.D. Schaeberle, E.N. Lewis, and I. W. Levin. Visible reflectance hyperspectral imaging: Characterization of a noninvasive, *in vivo* system for determining tissue perfusion. *Analytical Chemistry*, 74(9):2021–2028, 2002.