TISSUE CHARACTERISATION IN MRI: AN ASSESSMENT OF THE EFFICACY OF TEXTURAL FEATURES

by
David W. M. Boyce

A thesis submitted to
The University of London
in partial fulfilment of the requirements
for the degree of PhD.

Department of Electrical and Electronic Engineering
Imperial College
London

ABSTRACT

This project considers the design and implementation of computational methods to assist clinicians in their interpretation of images from Magnetic Resonance Imaging (MRI).

The ultimate aim of any diagnostic procedure is tissue characterisation. Conventional radiological image interpretation suffers from significant diagnostic inaccuracies. The research undertaken as part of this project specifically explores the use of sophisticated texture analytic methods to derive robust pathological indicators from MRI images with the ultimate aim of minimising diagnostic errors.

A novel image visualisation and analysis environment is first developed utilising current graphics standards and user ergonomic features. This will form the basis of a clinically-oriented image processing platform for the evaluation and testing of the methods developed in this project.

Three sets of texture analytic tools have been developed. The first set of algorithms is based on the image statistics and includes first order moments, run length, cluster size, grey tone difference and co-occurrence matrices. The second comprises of image transform methods and includes Fourier parameters and an original approach to the Walsh and Slant transforms. The third set is based on the measurement of the fractal dimension and its interpretation as a textural measure.

The performance of the measures derived from each set of texture tools are evaluated when applied to sample sets of different texture classes from a library of clinical MRI images. The performance of the features is found from correlation and clustering measures as the image parameters are varied, for example; region of interest size, number of grey levels and data normalisation. Simple two feature tissue segmentation graphs are constructed to demonstrate the best performing texture measures. The results of this study enable the relative performances of the three sets of texture analytic tools to be compared.

TABLE OF CONTENTS

TITLE PAGE 1
ABSTRACT 2
TABLE OF CONTENTS 3
FIGURE LIST 7
GRAPH LIST 8
TABLE LIST 11
STATEMENT OF ORIGINALITY 13
ACKNOWLEDGEMENTS 14

CHAPTER 1 - INTRODUCTION

1.1 Project motivation 15
1.2 Diagnostic information 16
1.3 Texture 18
1.4 Non clinical applications of texture 25
1.5 Image data 28
1.6 Structure of thesis 29

CHAPTER 2 - TISSUE CHARACTERISATION

2.1 Introduction 33
2.2 Non textural tissue characterisation 35
2.3 Textural tissue characterisation 43
2.4 Summary 50

CHAPTER 3 - THE *MAIVIS* IMAGE VISUALISATION SYSTEM

3.1 Introduction 54
3.2 The role of *MAIVIS* in medical imaging 56
3.3 An overview of the system development 59
3.4 X-Windows concepts
3.4.1 Introduction 63
3.4.2 Window hierarchy 64
3.4.3 Mapping windows 64
3.4.4 Program control and events 65
3.5 The structure of *MAIVIS*
3.5.1 System organisation 65
3.5.2 The screen layout 69
3.5.3 Menu structure 70
3.5.4 User definable menu 72
3.5.5 Personal menus 73
3.5.6 Restrictive menu options 74
3.5.7 Result data and result images 75
3.5.8 Regions of interest 76

3.6 Basic image processing tools
3.6.1 General tools 76
3.6.2 Image statistics 77
3.6.3 Visual manipulation 77
3.6.4 Image orientation 78
3.6.5 Image arithmetic 78
3.6.6 Grey level manipulation 78
3.6.7 Edge detection masks 79
3.6.8 Spectral tools 79
3.7 Summary 79

CHAPTER 4 - STATISTICAL TEXTURAL FEATURES

4.1 Introduction 80
4.2 First order statistical approaches using moments 83
4.3 Run length coding 84
4.4 Cluster size 87
4.5 Grey tone difference method 94
4.6 Grey tone spatial dependence matrices 96
4.7 Sum and difference co-occurrence matrices 101
4.8 Summary 104

CHAPTER 5 - TRANSFORM TEXTURAL FEATURES

5.1 Introduction 107
5.2 Fourier power spectrum 111
5.3 Hadamard and Walsh transforms 115
5.4 Slant transform 124
5.5 Texture analysis using non-directional forms of Walsh and Slant transforms 130
5.6 Summary 137

CHAPTER 6 - FRACTAL TEXTURAL FEATURES

6.1 Introduction 139
6.2 The fractal dimension 143
6.3 The fractal dimension and the box method 147
6.4 The fractal dimension and the blanket method 149
6.5 The fractal dimension and the power spectral density
6.5.1 Relationship between PSD and fractal dimension 152
6.5.2 Manipulating the frequency information 155
6.5.3 Calculating the fractal dimension 157
6.6 The fractal dimension and automated region selection 160
6.7 Summary 161

CHAPTER 7 - EXPERIMENTAL DETAILS

7.1 Introduction 164
7.2 Data collection 164
7.3 Data normalisation 166
7.4 Data analysis and performance evaluation 168
7.5 Image data sets 171

CHAPTER 8 - RESULTS: STATISTICAL TEXTURE FEATURES

8.1 Introduction 176
8.2 First order statistical texture features
8.2.1 Discriminative performance overview 178
8.2.2 Correlation study 181
8.2.3 Optimised feature set 182
8.3 Run length analysis texture features
8.3.1 Discriminative performance overview 182
8.3.2 Correlation study 185
8.3.3 Optimised feature set 186
8.4 Grey tone difference texture features
8.4.1 Discriminative performance overview 187
8.4.2 Correlation study 189
8.4.3 Optimised feature set 189
8.5 Co-occurrence texture features
8.5.1 Discriminative performance overview 190
8.5.2 Correlation study 194
8.5.3 Optimised feature set 195
8.6 Comparison of optimal statistical texture features 196
8.7 Examples of tissue characterisation 198

CHAPTER 9 - RESULTS: TRANSFORM TEXTURE FEATURES

9.1 Introduction 206
9.2 Power spectrum texture features
9.2.1 Introduction to power spectrum texture features 207
9.2.2 Performance overview 208
9.2.3 Correlation study 210
9.2.4 Optimised feature set 210
9.2.5 Examples of tissue characterisation 211
9.3 Walsh and Slant texture features
9.3.1 Introduction to Walsh and slant texture features 215
9.3.2 Performance overview 217
9.3.3 Correlation study 222
9.3.4 Optimised feature set 225
9.3.5 Examples of tissue characterisation 225
9.4 Non-directional Walsh and Slant texture features
9.4.1 Introduction to non-directional Walsh & Slant 233
9.4.2 Performance overview 235
9.4.3 Correlation study 239
9.4.4 Optimised feature set 241
9.4.5 Examples of tissue characterisation 241

CHAPTER 10 - RESULTS: FRACTAL TEXTURE FEATURES

10.1 Introduction 245
10.2 Fractal features performance overview
10.2.1 Power spectral density method performance overview 246
10.2.2 Blanket method performance overview 248
10.2.3 Box method performance overview 250
10.2.4 Summary of conditions for optimal performance 251

10.3 Correlation between all fractal features 253
10.4 Examples of fractal tissue characterisation 254

CHAPTER 11 - DISCUSSION OF RESULTS

11.1 Introduction 263
11.2 Summary of texture feature evaluation
11.2.1 Statistical methods 264
11.2.2 Transform methods 266
11.2.3 Fractal methods 268
11.2.4 Optimised texture features 269
11.3 Comparison between optimised texture features
11.3.1 Performance comparison 271
11.3.2 Correlation 276
11.4 Discussion 278
11.5 Project conclusions 283
11.6 Future work 284

BIBLIOGRAPHY 286

APPENDICES

Appendix 1 An example of menu function 296
Appendix 2 Run length texture features 302
Appendix 3 Cluster identification flowchart 303
Appendix 4 Co-occurrence texture features 304

FIGURE LIST

Figure 1.1	Texture example No.1	23
Figure 1.2	Texture example No.2	23
Figure 1.3	Texture example No.3	23
Figure 1.4	Texture example No.4	24
Figure 1.5	Texture example No.5	24
Figure 1.6	Texture example No.6	24
Figure 1.7	Texture example No.7	24
Figure 3.1	Network overview at Biomedical systems group, Imperial College	62
Figure 3.2	X Windows initialisation (C-language code)	64
Figure 3.3	*MAIVIS* applications	66
Figure 3.4	*MAIVIS* main screen windows	67
Figure 3.5	*MAIVIS* image windows	67
Figure 3.6	*MAIVIS* actual screen view #1	68
Figure 3.7	*MAIVIS* actual screen view #2	68
Figure 3.8	Menu data structure	71
Figure 4.1	Cluster size method - original image	88
Figure 4.2	Cluster size method - padded image	88
Figure 4.3	Cluster size method - image mask	89
Figure 4.4	Cluster size method - example 1D histogram	92
Figure 4.5	Cluster size method - example 2D histogram	93
Figure 5.1	1D Hadamard basis functions	118
Figure 5.2	1D Walsh basis functions	120
Figure 5.3	2D Walsh basis functions	121
Figure 5.4	Axes of dyadic translations	122
Figure 5.5	1D Slant basis functions	127
Figure 5.6	2D Slant basis functions	128
Figure 5.7	Clipping image	132
Figure 5.8	Image padding	133
Figure 6.1	Fourier transformed image quadrants	155
Figure 6.2	Fourier values at radius $R=\text{sqrt}(U^2+V^2)$	157
Figure 7.1	Extraction from an exported texture feature file	165
Figure 7.2	Histogram from ROI	167
Figure 7.3	Equalised histogram from ROI	167
Figure 7.4	Texture class distributions	170
Figure 7.5	General clinical data sets	173
Figure 7.6	Brain clinical data sets	174
Figure 7.7	Texture sample images sets	175
Figure A1.1	A simple menu tree	296
Figure A1.2	MENU node data structure	297
Figure A1.3	MENU option selected and re-sixed	298
Figure A1.4	Option2 selected and re-sized	298
Figure A1.5	Option2 node data structure	299
Figure A1.6	Option10 node data structure	299

Figure A1.7 Option10 selected and re-sized 300
Figure A1.8 Option21 selected and function called 300
Figure A1.9 Option21 node data structure 301
Figure A1.10 C-language for function selection 301

Figure A3.1 Cluster identification 303

GRAPH LIST

Graph 8.1 Behaviour of first order features with grey level bits (N0) #1 179
Graph 8.2 Behaviour of first order features with grey level bits (N0) #2 179
Graph 8.3 Behaviour of first order features with normalisation (B8) #1 180
Graph 8.4 Behaviour of first order features with normalisation (B8) #2 180
Graph 8.5 Comparison of optimised first order statistical features 182
Graph 8.6 Behaviour of vertical run length features with grey level bits (N0) 183
Graph 8.7 Behaviour of horizontal run length features with grey level bits (N0) 183
Graph 8.8 Behaviour of cluster size features with grey level bits (N0) 184
Graph 8.9 Behaviour of vertical run length features with normalisation (B5) 184
Graph 8.10 Behaviour of horizontal run length features with normalisation (B5) 185
Graph 8.11 Behaviour of cluster size features with normalisation (B5) 185
Graph 8.12 Comparison of optimised run length features 186
Graph 8.13 Behaviour of grey tone difference with grey level bits (N0) 187
Graph 8.14 Behaviour of grey tone difference with normalisation (B7) 188
Graph 8.15 Behaviour of grey tone difference with direction dependence 188
Graph 8.16 Comparison of optimised grey tone difference features 190
Graph 8.17 Behaviour of co-occurrence features with grey level bits (N0) #1 191
Graph 8.18 Behaviour of co-occurrence features with grey level bits (N0) #2 191
Graph 8.19 Behaviour of co-occurrence features with normalisation (B8) #1 191
Graph 8.20 Behaviour of co-occurrence features with normalisation (B8) #2 192
Graph 8.21 Comparison of all co-occurrence features under optimum conditions 192
Graph 8.22 Behaviour of co-occurrence features with direction dependence 193
Graph 8.23 Behaviour of co-occurrence features with vector length dependence 194
Graph 8.24 Comparison of optimised co-occurrence features 195
Graph 8.25 Comparison of optimised statistical texture features 196
Graph 8.26 Segmentation of 'Thigh' data - features M2vM3 (ROI 16x16) 201
Graph 8.27 Segmentation of 'Thigh' data - features HGLNvGT1 (ROI 16x16) 201
Graph 8.28 Segmentation of 'Thigh' data - features F1vF8 (ROI 16x16) 201
Graph 8.29 Segmentation of 'Thigh' data - features F9vF11 (ROI 16x16) 201
Graph 8.30 Segmentation of 'Abdomen' data - features M2vM3 (ROI 12x12) 202
Graph 8.31 Segmentation of 'Abdomen' data - features M3vHGLN (ROI 12x12) 202
Graph 8.32 Segmentation of 'Abdomen' data - features F1vF8 (ROI 12x12) 202
Graph 8.33 Segmentation of 'Abdomen' data - features F9vF11 (ROI 12x12) 202
Graph 8.34 Segmentation of 'Leg104' data - features M2vM3 (ROI 16x16) 203
Graph 8.35 Segmentation of 'Leg104' data - features M3vHGLN (ROI 16x16) 203
Graph 8.36 Segmentation of 'Leg104' data - features F1vF8 (ROI 16x16) 203
Graph 8.37 Segmentation of 'Leg104' data - features F9vF11 (ROI 16x16) 203
Graph 8.38 Segmentation of 'Leg192' data - features M2vM3 (ROI 16x16) 204
Graph 8.39 Segmentation of 'Leg192' data - features M3vHGLN (ROI 16x16) 204
Graph 8.40 Segmentation of 'Leg192' data - features F1vF8 (ROI 16x16) 204
Graph 8.41 Segmentation of 'Leg192' data - features F9vF11 (ROI 16x16) 204

Graph 8.42	Segmentation of 'Brain' data - features M2vM3 (ROI 16x16)	205
Graph 8.43	Segmentation of 'Brain' data - features HGLNvGT1 (ROI 16x16)	205
Graph 8.44	Segmentation of 'Brain' data - features F1vF8 (ROI 16x16)	205
Graph 8.45	Segmentation of 'Brain' data - features F9vF11 (ROI 16x16)	205
Graph 9.1	Behaviour of power spectrum features with grey level bits (N0) #1	208
Graph 9.2	Behaviour of power spectrum features with grey level bits (N0) #2	209
Graph 9.3	Behaviour of power spectrum features with normalisation (B8) #1	209
Graph 9.4	Behaviour of power spectrum features with normalisation (B8) #2	209
Graph 9.5	Comparison of optimised power spectrum features	211
Graph 9.6	Segmentation of 'Thigh' data - features FT1vFT5 (ROI 8x8)	213
Graph 9.7	Segmentation of 'Thigh' data - features FT1vFT5 (ROI 16x16)	213
Graph 9.8	Segmentation of 'Abdomen' data - features FT1vFT5 (ROI 16x16)	213
Graph 9.9	Segmentation of 'Leg104' data - features FT1vFT5 (ROI 16x16)	214
Graph 9.10	Segmentation of 'Leg192' data - features FT1vFT5 (ROI 16x16)	214
Graph 9.11	Segmentation of 'Brain' data - features FT1vFT5 (ROI 16x16)	214
Graph 9.12	Behaviour of Walsh vertical run length features (8x8)	218
Graph 9.13	Behaviour of Walsh horizontal run length features (8x8)	218
Graph 9.14	Behaviour of Walsh cluster size features (8x8)	218
Graph 9.15	Behaviour of Walsh Fourier and axial symmetry features (8x8)	218
Graph 9.16	Behaviour of Walsh co-occurrence features WF1-WF6' (8x8)	219
Graph 9.17	Behaviour of Walsh co-occurrence features WF7-WF14 (8x8)	219
Graph 9.18	Behaviour of Slant vertical run length features (8x8)	219
Graph 9.19	Behaviour of Slant horizontal run length features (8x8)	219
Graph 9.20	Behaviour of Slant cluster size features (8x8)	220
Graph 9.21	Behaviour of Slant Fourier features (8x8)	220
Graph 9.22	Behaviour of Slant co-occurrence features SF1-SF6' (8x8)	220
Graph 9.23	Behaviour of Slant co-occurrence features SF7-SF14 (8x8)	220
Graph 9.24	Behaviour of best performing Walsh features	221
Graph 9.25	Behaviour of best performing Slant features	221
Graph 9.26	Comparison of optimised Walsh and Slant features	225
Graph 9.27	Segmentation of 'Thigh' data - features WVRL3 v WF11 (No FFT)	228
Graph 9.28	Segmentation of 'Thigh' data - features WF6' v WFT1 (FFT)	228
Graph 9.29	Segmentation of 'Thigh' data - features SF6 v SF8 (No FFT)	228
Graph 9.30	Segmentation of 'Thigh' data - features SF9 v SF11 (FFT)	228
Graph 9.31	Segmentation of 'Abdomen' data - features WVRL3 v WF11 (No FFT)	229
Graph 9.32	Segmentation of 'Abdomen' data - features WF6' v WFT1 (FFT)	229
Graph 9.33	Segmentation of 'Abdomen' data - features SF6 v SF8 (No FFT)	229
Graph 9.34	Segmentation of 'Abdomen' data - features SF9 v SF11 (FFT)	229
Graph 9.35	Segmentation of 'Leg104' data - features WVRL3 v WF11 (No FFT)	230
Graph 9.36	Segmentation of 'Leg104' data - features WF6' v WFT1 (FFT)	230
Graph 9.37	Segmentation of 'Leg104' data - features SF6 v SF8 (No FFT)	230
Graph 9.38	Segmentation of 'Leg104' data - features SF9 v SF11 (FFT)	230
Graph 9.39	Segmentation of 'Leg192' data - features WVRL3 v WF11 (No FFT)	231
Graph 9.40	Segmentation of 'Leg192' data - features WF6' v WFT1 (FFT)	231
Graph 9.41	Segmentation of 'Leg192' data - features SF6 v SF8 (No FFT)	231
Graph 9.42	Segmentation of 'Leg192' data - features SF9 v SF11 (FFT)	231
Graph 9.43	Segmentation of 'Brain' data - features WVRL3 v WF11 (No FFT)	232
Graph 9.44	Segmentation of 'Brain' data - features WF6' v WFT1 (FFT)	232
Graph 9.45	Segmentation of 'Brain' data - features SF6 v SF8 (No FFT)	232
Graph 9.46	Segmentation of 'Brain' data - features SF9 v SF11 (FFT)	232
Graph 9.47	Behaviour of non-dir^{n} Walsh vertical run length features (8x8)	236

Graph 9.48	Behaviour of non-dir^{n} Walsh horizontal run length features (8x8)	236
Graph 9.49	Behaviour of non-dir^{n} Walsh cluster size features (8x8)	236
Graph 9.50	Behaviour of non-dir^{n} Walsh Fourier and axial symmetry features (8x8)	236
Graph 9.51	Behaviour of non-dir^{n} Walsh co-occurrence features NSF1-NSF6' (8x8)	237
Graph 9.52	Behaviour of non-dir^{n} Walsh co-occurrence features NSF7-NSF14 (8x8)	237
Graph 9.53	Behaviour of non-dir^{n} Slant vertical run length features (8x8)	237
Graph 9.54	Behaviour of non-dir^{n} Slant horizontal run length features (8x8)	237
Graph 9.55	Behaviour of non-dir^{n} Slant cluster size features (8x8)	238
Graph 9.56	Behaviour of non-dir^{n} Slant Fourier features (8x8)	238
Graph 9.57	Behaviour of non-dir^{n} Slant co-occurrence features NSF1-NSF6' (8x8)	238
Graph 9.58	Behaviour of non-dir^{n} Slant co-occurrence features NSF7-NSF14 (8x8)	238
Graph 9.59	Behaviour of best performing non-directional Walsh features	239
Graph 9.60	Comparison of optimised non-directional Walsh features	241
Graph 9.61	Segmentation of 'Thigh' data - features NWF4 v NWFT1	243
Graph 9.62	Segmentation of 'Abdomen' data - features NWF4 v NWFT1	243
Graph 9.63	Segmentation of 'Leg104' data - features NWF4 v NWFT1	243
Graph 9.64	Segmentation of 'Leg192' data - features NWF4 v NWFT1	244
Graph 9.65	Segmentation of 'Brain' data - features NWF4 v NWFT1	244
Graph 10.1	Behaviour of power spectral density fractal feature	247
Graph 10.2	Behaviour of blanket fractal feature with grey level bits (12x12 & N0)	248
Graph 10.3	Behaviour of blanket fractal feature with grey level bits (24x24 & N0)	249
Graph 10.4	Behaviour of blanket fractal feature with normalisation (12x12 & B4)	249
Graph 10.5	Behaviour of blanket fractal feature with normalisation (24x24 & B5)	249
Graph 10.6	Behaviour of box fractal feature	250
Graph 10.7	Behaviour of average box fractal feature	251
Graph 10.8	Behaviour comparison of features FR1&FR2 (ROI=16x16)	252
Graph 10.9	Behaviour comparison of features FR2,FR3&FR4 (ROI=24x24)	252
Graph 10.10	Segmentation of 'Thigh' data - features FR1vFR2 (ROI 8x8)	257
Graph 10.11	Segmentation of 'Thigh' data - features FR1vFR2 (ROI 16x16)	257
Graph 10.12	Segmentation of 'Thigh' data - features FR2vFR3 (ROI 12x12)	257
Graph 10.13	Segmentation of 'Thigh' data - features FR2vFR3 (ROI 24x24)	257
Graph 10.14	Segmentation of 'Abdomen' data - features FR1vFR2 (ROI 8x8)	258
Graph 10.15	Segmentation of 'Abdomen' data - features FR1vFR2 (ROI 16x16)	258
Graph 10.16	Segmentation of 'Abdomen' data - features FR2vFR3 (ROI 12x12)	258
Graph 10.17	Segmentation of 'Leg104' data - features FR1vFR2 (ROI 8x8)	259
Graph 10.18	Segmentation of 'Leg104' data - features FR1vFR2 (ROI 16x16)	259
Graph 10.19	Segmentation of 'Leg104' data - features FR2vFR3 (ROI 12x12)	259
Graph 10.20	Segmentation of 'Leg104' data - features FR2vFR3 (ROI 24x24)	259
Graph 10.21	Segmentation of 'Leg192' data - features FR1vFR2 (ROI 8x8)	260
Graph 10.22	Segmentation of 'Leg192' data - features FR1vFR2 (ROI 16x16)	260
Graph 10.23	Segmentation of 'Leg192' data - features FR2vFR3 (ROI 12x12)	260
Graph 10.24	Segmentation of 'Leg192' data - features FR2vFR3 (ROI 24x24)	260
Graph 10.25	Segmentation of 'Brain' data - features FR1vFR2 (ROI 8x8)	261
Graph 10.26	Segmentation of 'Brain' data - features FR1vFR2 (ROI 16x16)	261
Graph 10.27	Segmentation of 'Brain' data - features FR2vFR3 (ROI 12x12)	261
Graph 10.28	Segmentation of 'Brain' data - features FR2vFR3 (ROI 24x24)	261
Graph 10.29	Segmentation of 'Brodatz' data - features FR1vFR2 (ROI 32x32)	262
Graph 10.30	Segmentation of 'Brodatz' data - features FR1vFR2 (ROI 64x64)	262
Graph 11.1	DB performance between all three tissue classes	273
Graph 11.2	DB performance between tissue 1 (muscle) and tissue 2 (fat)	273

Graph 11.3	DB performance between tissue 2 (fat) and tissue 3 (tumour)	273
Graph 11.4	DB performance between tissue 1 (muscle) and tissue 3 (tumour)	274
Graph 11.5	Fisher linear classifier between tissue 1 (muscle) and tissue 2 (fat)	274
Graph 11.6	Fisher linear classifier between tissue 2 (fat) and tissue 3 (tumour)	274
Graph 11.7	Fisher linear classifier between tissue 1 (muscle) and tissue 3 (tumour)	274
Graph 11.8	Student's t-test between tissue 1 (muscle) and tissue 2 (fat)	275
Graph 11.9	Student's t-test between tissue 2 (fat) and tissue 3 (tumour)	275
Graph 11.10	Student's t-test between tissue 1 (muscle) and tissue 3 (tumour)	275

TABLE LIST

Table 2.1	Texture methods - (S) statistical, (T) transform and (F) fractal	52
Table 4.1	Values of 2D histogram *Cluster[i,j]*	93
Table 7.1	Clinical data sets (256x256 images, 12 bit)	172
Table 7.2	Clinical tissue classes	172
Table 7.3	Non-clinical data sets (256x256 images, 8-bit digitised)	172
Table 8.1	Statistical texture features	177
Table 8.2	First order feature correlation	181
Table 8.3	Run length feature correlation	186
Table 8.4	Grey tone difference feature correlation	189
Table 8.5	Co-occurrence feature correlation	194
Table 8.6	Grouping of most correlated co-occurrence features	195
Table 8.7	Correlation between best performing statistical texture features	197
Table 8.8	Statistical texture features - graph key	198
Table 8.9	Description of image data	198
Table 8.10	Examples of tissue segmentation with statistical texture features	199
Table 9.1	Power spectrum texture features	207
Table 9.2	Power spectrum feature correlation (B8N0 ROI 8x8)	210
Table 9.3	Power spectrum feature correlation (B8N0 ROI 16x16)	210
Table 9.4	Power spectrum texture features - graph key	211
Table 9.5	Description of image data	212
Table 9.6	Examples of tissue segmentation with power spectrum texture features	212
Table 9.7	Texture features derived from Walsh transformed data	216
Table 9.8	Texture features derived from Slant transformed data	216
Table 9.9	Best performing Walsh and Slant features	221
Table 9.10	Walsh correlation (8x8)	222
Table 9.11	Walsh correlation (16x16)	222
Table 9.12	Walsh correlation (8x8 FFT)	222
Table 9.13	Walsh correlation (16x16 FFT)	222
Table 9.14	Slant correlation (8x8)	223
Table 9.15	Slant correlation (16x16)	223
Table 9.16	Slant correlation (8x8 FFT)	224
Table 9.17	Slant correlation (16x16 FFT)	224
Table 9.18	Description of image data	226
Table 9.19	Examples of tissue segmentation with Walsh and Slant texture features	226
Table 9.20	Texture features derived from non-dir^n Walsh transformed data	234

Table 9.21	Texture features derived from non-dir^n Slant transformed data	234
Table 9.22	Best performing non-directional Walsh features	239
Table 9.23	Feature correlation (ROI 8x8)	240
Table 9.24	Feature correlation (ROI 16x16)	240
Table 9.25	Description of image data	242
Table 9.26	Examples of tissue segmentation with non-dir^n Walsh texture features	242
Table 10.1	Fractal texture features	245
Table 10.2	Optimal feature performance parameters	251
Table 10.3	Feature correlation (ROI 8x8)	253
Table 10.4	Feature correlation (ROI 12x12)	253
Table 10.5	Feature correlation (ROI 16x16)	253
Table 10.6	Feature correlation (ROI 24x24)	253
Table 10.7	Fractal texture features - graph key	255
Table 10.8	Description of image data	255
Table 10.9	Examples of tissue segmentation with fractal texture features	255
Table 10.10	Comparison between PSD and blanket fractal features (ROI 32x32)	262
Table 10.11	Comparison between PSD and blanket fractal features (ROI 64x64)	262
Table 11.1	Optimised texture features	270
Table 11.2	Correlation between performance measures	272
Table 11.3	Correlation between optimised texture features	277

STATEMENT OF ORIGINALITY

The original aspects of this thesis are considered to be:

IMAGE PROCESSING ENVIRONMENT

- The design and implementation of *MAIVIS* (image visualisation and analysis environment)
- The implementation of a large number of general image processing tools
- The development and implementation of the textural analysis algorithms
- The implementation of the histogram equalisation normalisation algorithm
- The design and implementation of the fixed statistical moments normalisation algorithm

TEXTURE ANALYSIS THEORY

- The proposal of an original 2D run length texture method called cluster size
- The proposal of an orientation in-sensitive version of the grey tone difference method
- The proposal of an additional co-occurrence matrix texture feature
- The proposal of additional Fourier spectrum texture features
- The proposal of the use of Walsh and Slant transforms for texture analysis
- The proposal of original Walsh and Slant transform texture features
- The proposal of orientation in-sensitive Walsh and Slant transforms
- The development of a methodology for deriving the Fourier spectrum fractal dimension

TEXTURE ANALYSIS EXPERIMENTS

- The proposal of a simple discriminative performance measure
- The identification of optimum image parameters for each texture feature
- The identification of the best performing texture features for each method
- A relative performance comparison of statistical, transform and fractal methods
- The measurement of correlation between all the texture features
- The identification of redundant texture methods

OVERALL PROJECT

- The compehensive study of the performance of an extensive range of texture features collected under different experimental conditions
- The discussion and conclusions of this study

ACKNOWLEDGEMENTS

I would firstly like to express my gratitude to my project supervisor Dr Keith Straughan for his enthusiasm in this field of work. Dr Straughan has been a constant source of inspiration, and gave me a great deal of encouragement to explore and develop many of my own ideas within the bounds of the project subject area. I would also like to thank him for his continued support since he has left Imperial College.

I would like to thank Professor Richard Kitney for continuing the supervision of my project since the departure of Dr Straughan. Professor Kitney is a committed advocate of the use of computation methods in diagnostic imaging, and I have benefited from his enthusiasm for this subject.

I would like to acknowledge the staff and students of the Biomedical Systems group at Imperial College for the general support for my work. I would also like to thank my colleagues on the MagNET project for the flexibility forded to me to enable the completion of this thesis. I am especially grateful to Dr Janet De Wilde for her friendship and for the moral support she has offered me over the years I have known her.

I would like to thank the staff at St. Mary's Hospital MR scanning unit for providing me with access to their MR image data archive, from which much test data for this project was obtained. I am grateful to Dr Samantha Free for her assistance in this task.

A very special thanks must go to my friends and family for their eternal patience and relentless encouragement. In particular, I would like to individually thank my Mother and Father for the special support they both gave me. I am very grateful for everything they have done for me.

The greatest of all thanks, however, must go to my wife Catherine, without whom this work could not have been completed. She is my friend, companion and partner, and I owe her a large debt of gratitude for the sacrifices she has made over the years this work has been continuing. I therefore dedicate this thesis to her and to our son Harry.

CHAPTER 1 - INTRODUCTION

1.1 PROJECT MOTIVATION

The motivation for this research project is derived from the potential that computational methods can offer to assist hospital clinicians and radiologists in their interpretation of patient images from Magnetic Resonance Imaging (MRI) imaging. Accurate computer analysis of medical images can help the clinicians to make more efficient use of their skills and resources, and consequently can provide a better service and a faster throughput of patients.

When MRI imaging was first introduced into the clinical environment, the technology was very new and the scanners were expensive prototypes. In general the functionality was very limited, typically only single slice acquisitions were possible with just a handful of sequences. The radiologists had at their disposal a limited set of variable scanning parameters and very basic sets of post-processing tools such as for grey level windowing. The clinical significance of the images was judged on prior knowledge of the appearance of normal and abnormal tissue. There were no system tools to aid the radiologists in the evaluation of images.

The imaging equipment manufacturers initially designed their systems from a purely engineering stance. The thrust of early MRI development was in the technical performance and efficiency of the data retrieval process. With such rapid advancements in the scanning technology, there was little apparent interest in the diversity of user requirements. It is only now, when technology, clinical expertise and the clinical needs are becoming integrated properly, that the enormous potential of MRI is being realised.

In addition to the latest imaging techniques, the modern MRI system also provides the user with a sophisticated image review capability. Depending on the package, this may also include sophisticated post-processing tools such as maximum intensity projections (MIPS) for angiographic studies, multi-planar reformatting (MPR), or T1/T2 parametric image mapping for example. Apart from T1/T2 maps, which are not offered on all commercial systems, there are currently no tools available on these systems for tissue characterisation. The modern MRI installation can potentially be networked to a large number of in-expensive workstations, providing the clinicians with enormous post-processing resources. In this technological climate there is clearly great scope for new tissue characterisation methods, however sophisticated.

It is no coincidence, therefore, that this research project is directed towards the development of sophisticated tissue characterisation post-processing techniques, and in particular towards the use of texture analysis. The motivation for using texture analysis for tissue characterisation comes from reports made by a number of authors (**chapter 2**) that texture analysis has the potential to provide improved specificity over indices derived directly from image or parametric image data.

Because the project is completely software orientated, a software platform was required on which images could be loaded and the analysis methods designed and tested. The development of such an image visualisation environment formed a major part of this project and required a considerable amount of time to be dedicated to its inception. This system, more commonly known as *MAIVIS* (*M*edical *A*pplication *I*mage *V*isualisation and *I*nterpretation *S*ystem), is a general purpose image visualisation package. The basic system contains a set of general image processing tools with the specific function, or application, defined by a choice of user definable algorithms. In addition to serving the needs of this project, *MAIVIS* has successfully been employed as the image processing platform for a number of other research projects (**chapter 3, section 3.1**).

Another important motivation for the development of image visualisation systems like *MAIVIS* is the advent of sophisticated databases such as Picture Archiving and Communication Systems (PACS) into the clinical environment **(chapter 3, section 3.2)**. PACS is a multi-media database system that allows text, pictures, speech and other media to be stored in single documents or records. In the hospital PACS would provide digital patient records combining images from a whole range of scanning modalities. Image visualisation packages like *MAIVIS* running on low cost workstations could easily be networked into PACS providing the clinicians with powerful image review platforms. With such a system, additional diagnostic tools such as for tissue characterisation could easily be introduced as and when required. In the context of this project, texture analysis has been investigated as a valid method for tissue characterisation.

1.2 DIAGNOSTIC INFORMATION

In Magnetic Resonance imaging the image acquired provides a true representation of the object of the scan. The image features correspond directly to shape, position, and size of

anatomical structure. Because the grey level (intensity) distribution in these images depends on the contrast mechanisms adopted by the scan, subjectivity can arise in the characterisation of the true nature of tissue ie. detecting the changes in grey level associated with disease. Unless the clinician has *a priori* knowledge of the characteristic appearance of pathology, then this subjectivity will increase the probability of mistakes in diagnostic interpretation.

Motivated by the subjectivity suffered by conventional radiological image interpretation, much work has been aimed at quantitative diagnostic techniques to derive precise tissue characteristic indices. The relative successes and failures of a large number of researchers with quantitative tissue characterisation methods based on the physical properties of a range of imaging modalities including MRI are discussed in **chapter 2, section 2**. The body of research undertaken as part of this project falls within this general strategy, and specifically explores the use of sophisticated analytic methods to derive robust pathological indicators from clinical MRI images. These methods can be used to determine if the tissue in a region of interest possesses normal or pathological characteristics. The tissue can be classified by assimilation with a knowledge base of the typical appearances of various tissue characteristics.

In order to develop appropriate analytic techniques, the image features which the radiologist explicitly or implicitly uses in the appraisal of tissue appearance must be considered. The most commonly cited image features used by clinicians are intensity, morphology and texture. These features correspond directly with the idea of 'textons' described by (JULESZ, 1975; JULESZ, 1981) ie. local conspicuous features that trigger the human perceptive system to discriminate between different image subjects. Texture has been chosen as the method of tissue characterisation for this project because of its potential sensitivity to pathological detail. This approach is supported in the conclusions drawn by many of the authors in **chapter 2**.

Visual texture is a subjective feature that the human observer finds it difficult to define precisely. It is only using mathematical techniques that we can quantify and objectively define texture. It is to the development and implementation of these techniques that this project makes a large contribution. The subsequent experiments carried out with the tissue indices derived indicate the efficacy of the texture methods for tissue characterisation.

The texture analysis techniques investigated as part of this project have been developed to run on the *MAIVIS* computer platform. A large number of specialised algorithms have been

developed specifically for this purpose. The additional functionality on *MAIVIS* enables the user to interrogate image regions of interest and extract tissue characterisation indices using a large diversity of texture analysis methods. The main sub-groups of texture analysis employed are statistical and syntactic (or structural) methods. The statistical texture methods are discussed in **chapter 4**. The syntactic methods can be further sub-divided into transform methods and fractal methods, and are described in **chapters 5** and **6**, respectively.

This overall aim of this research project is to take a fundamental look at texture analytic approaches to tissue characterisation in MRI. The long term aim is to produce a reliable set of computational tools that are easily used by clinicians that can accurately classify tissue types from medical images. Such precise tools are likely to enhance the skills of the radiologist in addition to assisting them in their work. The short term result, which is the specific aim of this project, is to perform a feasibility study on a vast number of existing and new texture techniques. The study is structured in such a way that initially the three types of texture methods will be examined separately. The best performing features from all groups will subsequently be compared. The texture features are measured on sample sets of different textures from a large library of MRI images. Performance is measured in terms of the discriminative power of the features under varying experimental conditions. Features will be excluded on the basis of poor performance and poor stability, and feature sets will be reduced where correlation occurs.

In the final analysis, the best performing texture features will be used as the basis for simple tissue characterisation schemes based on selected regions of interest. More advanced schemes involving classification and segmentation are outside the scope of this project.

1.3 TEXTURE

The image post-processing tools available on an MRI scanner have evolved through three distinct stages. These naturally group themselves under the headings of visualisation, enhancement and interpretation. On modern scanners these functions can be performed on both the scanner console and additional networked workstations.

The MRI process can either be used to collect sets of successive 2D image slices in any plane (or oblique) through the object being scanned, or to collect a 3D image data set from which

2D images are derived. The acquisition and display of image data, generally known as image visualisation, is the primary post-processing function of the scanner. Image acquisition is normally only possible on the scanner console (host computer).

Enhancement post-processing tools enable critical image details to be examined more carefully to provide the clinician with more diagnostic information. The most widely used enhancement tools are radiometric functions to vary the image window level and width (image contrast and brightness). Other commonly used tools include zooming, edge detection and image filtering. Most of these methods are described in the standard image processing texts of (PRATT, 1978; CLARKE, 1985; GONZALEZ and WINTZ, 1987; SCHALKOFF, 1989).

The third group of image post-processing techniques perform data interpretation and image analysis. Because of the computing power required for these methods, they are normally only performed on additional networked workstations. Large amounts of useful information can be extracted from images by examining the image statistics or by applying a large diversity of image transforms and operators. These tools can be very effective in assisting the clinician in the characterisation of pathology or abnormality. Texture analysis comes under this category of image post-processing techniques.

Texture is a term that can be used to characterise the surface of a given object, providing a valuable description of the variation of structure over the object's surface. Texture has proved to be an important feature in image processing and image pattern recognition applications (**section 1.4**, this chapter). In this context, the texture describes the spatial distribution of tonal (grey level) variations. It is suggested by (JULESZ, 1975; HARALICK, 1979) that texture can be completely described by its tonal primitives, or 'textons'. Tonal primitives are defined by the relationship of the grey levels of neighbouring pixels. Image texture is essentially a neighbourhood property.

In simple terms, image texture can be described as being either fine, coarse or smooth. The description could be further elaborated to describe the texture as being either rippled, mottled, irregular or lineated. Although the human observer can easily recognise and describe textures using these terms, texture is extremely difficult to define precisely. A formal approach or precise definition of texture does not exist.

The most difficult step in characterising pictorial information, such as the tissue groups in MRI images, is that of defining a set of meaningful features to describe the regions of interest. **Chapter 2, section 2.3** describes work in a variety of imaging modalities of those who have pioneered the use of texture analysis for quantitative tissue characterisation. These computational methods attempt to describe texture quantitatively from the underlying pixel-to-pixel relationships across the image. Most of these studies have been entirely circumspect, and only a few conclusive results have been reached. In order that some of these conclusions can be substantiated, a more comprehensive study examining the efficacy of quantitative texture tools in their application to MRI images is needed. This study should also establish if these methods can be used to characterise tissue more reliably than the human observer. The results of this project attempt to provide the foundations for these studies by evaluating the strengths and weaknesses of this quantitative approach with a large range of texture analysis methods under varying experimental conditions.

Computational texture analysis methods can be grouped into either statistical methods or syntactic (structural) methods. Statistical texture features are derived from probability density functions of image pixel intensities or of the relationship of intensities between neighbouring pixels. Statistical methods account for the majority of all texture methods. **Chapter 4** describes a large diversity of approaches to statistical texture analysis and identifies those methods that have been implemented as part of this project. By definition, syntactic methods characterise texture by its basic structural components. For example, the Fourier transform can be used to examine the frequency composition of texture. The Fourier transform and several additional transforms not previously used for texture analysis have been explored as part of this project. Another syntactic approach that has been investigated is the characterisation of texture by its fractal properties. **Chapters 5** and **6**, respectively, discuss transform and fractal methods.

There are two main benefits of using computational methods to assist the human observer in the analysis of texture, and in the subsequent characterisation of tissue. The first gain is that computational methods provide quantitative measurements, where the human observer's perception is subjective and often variable. The second benefit of these methods is their ability to measure changes in image statistics. Subtle difference between textures are often cause by small variations in image statistics which the human observer is not always able to perceive. The extent of the human observer's perception of image statistics can be understood from its

fundamental definition.

The first order statistics of an image are found from the frequency histogram of pixel brightness (or grey levels). The performance of the human observer in discriminating first order texture differences is extremely good (JULESZ, 1975; JULESZ, 1981). This is easily demonstrated in **figure 1.1,** which shows two textures with different first order statistics (indicated by different element size). The textures are clearly different.

The definition of second order image statistics given by (WAGNER *et al*, 1985) is the joint distribution of brightness found at the pixels at the two ends of a line or 'dipole', as a function of all lengths and orientations of the 'dipole'. In simple terms the second order statistics can be though of as the spatial distribution of brightness across the image. The human observer's pre-attentive system performs poorly at discriminating second order texture differences (JULESZ, 1975). The human pre-attentive system is the eye-brain perceptive system which allows unconscious recognition.

WAGNER *et al* describe the third order image statistics as the joint distribution of the brightness found at the three vertices of a triangle, as a function of all the possible triangle side lengths and angles. Using this model, the second order statistics correspond to the brightness at the two ends of a degenerate triangle. The information conveyed by the third order statistics relates to the structure of image patterns. This can be explained by examining two textures with identical second order, but with different third order statistics. The textures in **figures 1.2** and **1.3** are both good examples of this. In each case, the two patterns have similar dipolar structure, but the overall structure of the patterns are different. This is demonstrated by the mirrored primitives in **figure 1.2**, and the differently arranged primitives in **figure 1.3**. The two textures in each of the figures are, at first glance, indistinguishable and require some scrutiny to distinguish them. JULESZ concluded that the human pre-attentive system cannot globally compute third or higher order statistics, and that attentive scrutiny was required to differentiate textures of these levels.

JULESZ describes how the human observer discriminates between texture as a result of local conspicuous features, called *textons*. *Textons* cue the human pre-attentive system for discrimination. These prompts are caused by differences in the texture primitives eg. size, shape, orientation and distribution. **Figures 1.4** to **1.7** illustrate how *textons* help the human

observer to discriminate between different textures. **Figures 1.4** and **1.5** contain textures with the same first order, but different second order statistics. The textures in **figure 1.4** are easily distinguishable because of the variation in element distribution. The textures in **figure 1.5** are easily distinguishable because their elements have different orientations. **Figures 1.6** and **1.7** contain textures with identical second order statistics. **Figure 1.6** is below the pre-attentive threshold and the textures are, at first glance, indistinguishable. Due to the shape of the primitives, the two textures in **figure 1.7** are above the pre-attentive threshold and are easily distinguishable.

The previous examples demonstrate that the human observer can easily differentiate between textures with differences in first order statistics, but relies on conspicuous features to discriminate between textures with differences in higher order statistics. Changes in image statistics to the n^{th} order, however, can be measured very proficiently using quantitative computational methods. These findings fully justify the study of the efficacy of statistical texture analysis methods for tissue characterisation carried out as part of this project. The syntactic texture methods tackled are equally justified on this basis because changes in third and higher order image statistics relate to differences in texture structure.

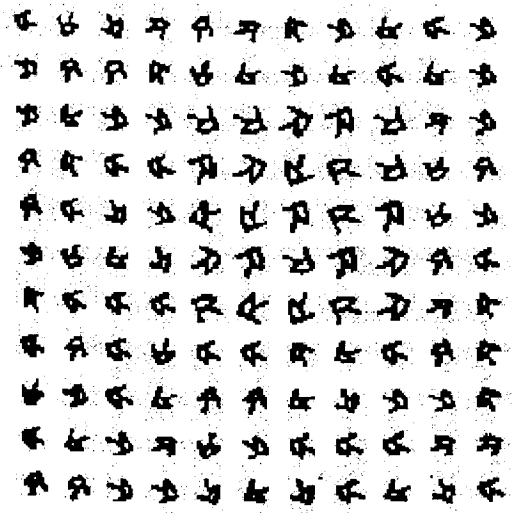
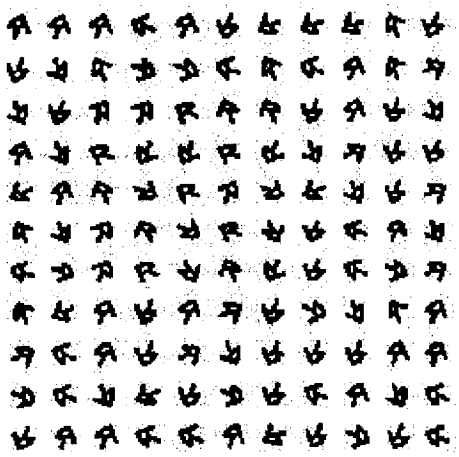


Figure 1.1 : Texture example No.1
(Different 1st order statistics)

From (JULESZ, 1975)



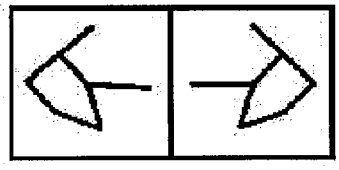
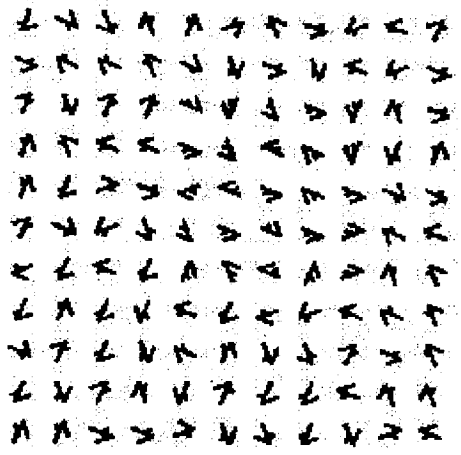


Figure 1.2 : Texture example No.2
(Same 2nd order but different higher order)

From (JULESZ, 1975)



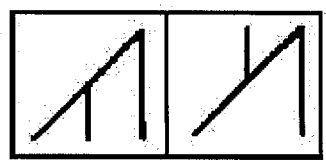


Figure 1.3 : Texture example No.3
(Same 2nd order but different higher order)

From (JULESZ, 1975)

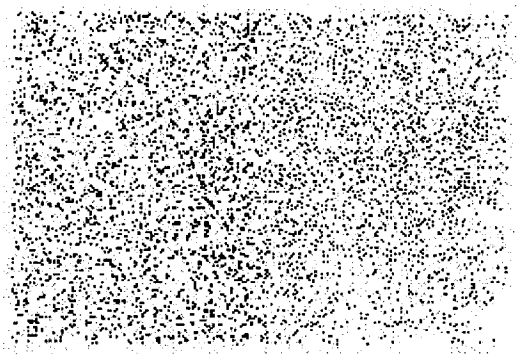


Figure 1.4 : Texture example No.4
(Same 1st order but different 2nd order)
From (JULESZ, 1975)

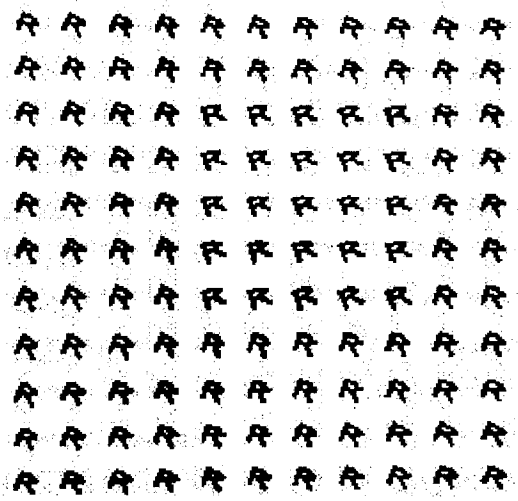
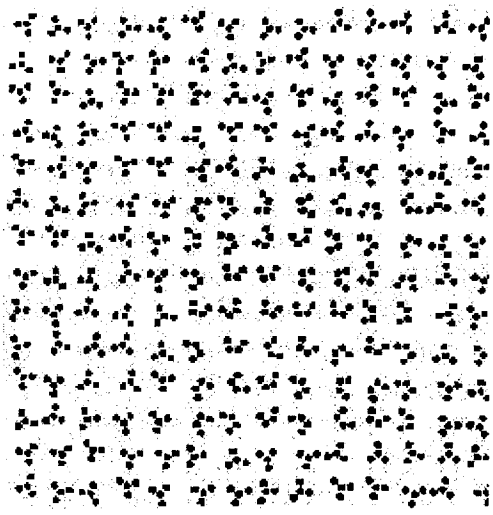


Figure 1.5 : Texture example No.5
(Same 1st order but different 2nd order)
From (JULESZ, 1975)



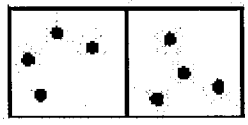
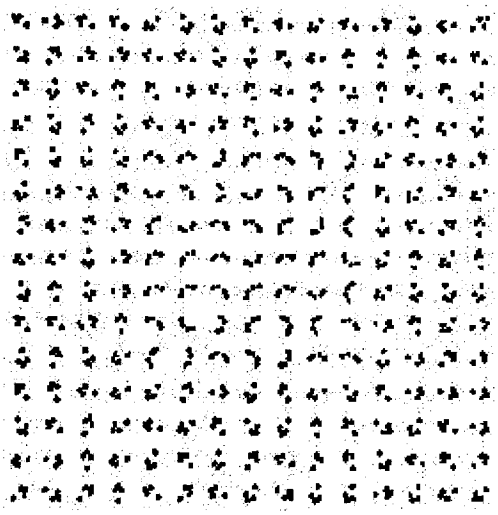


Figure 1.6 : Texture example No.6
(Same 2nd order but textures are inconspicuous)
From (JULESZ, 1975)



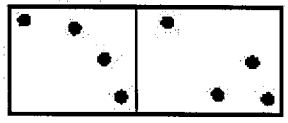


Figure 1.7 : Texture example No.7
(Same 2nd order but textures are conspicuous)
From (JULESZ, 1975)

1.4 NON CLINICAL APPLICATIONS OF TEXTURE

The first application that employed the use of texture analysis was terrain analysis ie. the interpretation of aerial photographs or satellite pictures. The general motivation for this development was to achieve more objective and consistent texture classification with computational methods than was possible with the human observer.

One of the earliest experiments with texture analysis was carried out by (KAIZER, 1955) who assessed the performance of the auto-correlation function as a source of texture information. KAIZER asked a group of subjects (photo-interpreters) to rank seven aerial photographs of different Arctic terrain on a scale from fine detail to coarse detail. KAIZER found a high correlation (coefficient of 0.99) between the rankings and the results of the auto-correlation function. The results of this experiment clearly demonstrated the potential effectiveness of computational methods for texture analysis.

The most important contribution to the field of texture analysis was probably made by (HARALICK *et al*, 1973). HARALICK *et al* presented 14 easily computable texture features based on the grey-tone spatial dependencies (co-occurrence) of neighbouring image points. This method is now almost universally accepted as the standard approach to texture analysis. In the experiments reported in his publication, HARALICK *et al* tested the efficacy of these texture features with three different sources of image texture; photo-micrographs of sandstone, aerial photographs of a variety of terrains, and satellite pictures over the California coastline (composites from radiance at four narrow bands of wavelengths). The results from the category identification tasks using the 14 texture features showed high accuracy with all three data sets (at least 82 percent).

Another major contribution to the science of texture analysis was made by (GALLOWAY, 1975), who presented a set of texture features she had derived from image grey-level run lengths. GALLOWAY obtained promising results in identifying different classes of terrain samples from a set of aerial photographs with this method.

One of the first general reviews of texture analysis was carried out by (WESZKA *et al*, 1976) in which he examined the comparative performance of a number of different texture methods for terrain classification. For this study, texture features were derived from the (HARALICK

et al, 1973) co-occurrence method, the (GALLOWAY, 1975) run length method, the Fourier power spectrum and the first order statistics of local grey level differences. The performance of these features was examined in the classification of aerial photographic terrain samples belonging to nine land use classes and LANDSAT imagery samples belonging to three geological terrain types. WESZKA *et al* found that apart from the Fourier power spectrum features, which performed more poorly, the performance of all feature sets was comparative.

WESZKA *et al* results demonstrated that local grey level differences are a valuable source of texture information. A similar approach was proposed by (MITCHELL *et al*, 1977), who described a method that uses the relative frequency of local extremes in grey level. This method is independent of image illumination and invariant under sampling rate changes. Mitchell also demonstrated comparative performance between his method and the (HARALICK *et al*, 1973) co-occurrence method for a selection of image texture samples.

Other review articles published include the survey by (HARALICK, 1979) of statistical and structural approaches to texture and the theoretical comparison of texture algorithms carried out by (CONNERS and HARLOW, 1980). HARALICK *et al* described successful implementations of statistical texture approaches using run length, co-occurrence, auto-correlation, digital transforms such as the power spectrum, and several methods based on the relationships between grey levels of neighbouring pixels. He reported little use of pure structural approaches based on more complex primitives than grey level. Using statistically generated textures, such as Markov textures, CONNERS and HARLOW examined the comparative performance of co-occurrence, run length, grey level difference and power spectrum texture methods. The co-occurrence method was found to give the best results.

The most successful texture methods are those that exploit the second order statistics of the texture. Of these approaches, the (HARALICK *et al*, 1973) co-occurrence method has proved to be the most effective. This method has subsequently become the focus of a number of investigators' work with texture analysis.

The 14 co-occurrence texture features introduced by HARALICK *et al* were used by (GOTLIEB and KREYSZIG, 1990) to classify sets of homogenous texture samples digitised from the (BRODATZ, 1966) book of photographs of natural textures. GOTLIEB and KREYSZIG suggested that there would be correlation between some of the 14 texture

features, and consequently redundant information. Combinations of features were explored, and four groups were established according to the nature of the features. GOTLIEB and KREYSZIG found no information was lost by representing the texture by just one feature from each group.

The sum and difference histograms were introduced by (UNSER, 1986) as an alternative to the co-occurrence method for texture analysis to over-come the computational overheads required for its calculation. UNSER's method used sum and difference matrices derived from the two de-correlated probability density functions defined by the principal axis of the co-occurrence joint probability density function. This method requires much less computation than the co-occurrence method for a small loss in performance. UNSER tested his method with the (BRODATZ, 1966) natural textures and found the classification percentages to be very high, and increasing with region size.

Fractal geometry has received a surge of interest in the last decade due to the work of (MANDELBROT, 1977). MANDELBROT introduced the concept of the fractal dimension as a measure of the fractal properties of a line, surface or volume, and (PENTLAND, 1984) subsequently proposed that this measure could be used to classify and segment image texture. The feasibility of using the fractal dimension for texture analysis has been pioneered by several investigator including (PENTLAND, 1984; PELEG *et al*, 1984; MEDIONI and YASUMOTO, 1984; ALLINSON and LAWSON, 1990; LINNETT and CLARKE, 1990), with selection of image subjects such as aerial photographs, sonar images and (BRODATZ, 1966) natural textures. The overall results indicate that the fractal dimension can be used as an effective measure of texture, despite instances when textures with some similar features have proved to be indistinguishable.

Texture analysis was pioneered in the 1970's with terrain analysis applications. The texture features developed were based predominantly on the statistics of local grey level variations, or on the digital transforms such as the Fourier power spectrum. It took another decade before fractal texture methods were also being used. Texture analysis has subsequently been implemented in newer and more diverse applications, such as for tissue characterisation in medical imaging, or more specifically for tissue characterisation in MRI.

The commentary presented in this section, coupled with the review in **chapter 2** of the

application of texture analysis for tissue characterisation in medical imaging, provides an excellent focus for the selection of texture methods for this project's experiments with MRI tissue characterisation.

1.5 IMAGE DATA

The texture analysis algorithms designed and coded as part of this project were tested on image data collected from several sources. Testing was carried out on both non-clinical texture samples and on clinical data, which was initially very difficult to obtain.

Non-clinical texture images were digitized from the book of texture samples by (BRODATZ, 1966). A set of largely homogenous texture images resulted from this capture process. The image loading format was trivial ie. only one byte per pixel, because the data were limited to 8-bits by the digitisation process.

The first sets of clinical data investigated came on reels of 1/2 inch magnetic tape. This was the only method by which data could be archived at that time. Although this media could be read on standard tape drives, the data encoding was proprietary. The whole data set was read off in one block. From subsequent examination of both ASCII and Hexadecimal dumps of this data, the 12-bit images could be individually extracted. Additional examination of the data was often required to establish the image loading format. Examples of the problems encountered are finding the correct byte order for images with more that one byte per pixel and identifying patterns of image data zero padding. This whole process was very time consuming, and depending on the format, could take as long as a several weeks with a new source of data.

During the latter stages of algorithm testing process, data became available from the GE Signa 1.5 T MRI scanner at St Mary's Hospital, Paddington. The scanner was networked to a Sun workstation, which in turn was networked with the Sun workstations at Imperial College. Suitable clinical examples were selected from the scanning records and re-loaded onto the scanner from the optical disk archive. These images were then transferred over the network to Imperial College. The image format was trivial and therefore no problems were experienced in the loading images into *MAIVIS*.

Once the texture analysis algorithms and associated functions in *MAIVIS* had been fully tested,

the studies presented in the results chapters of this thesis were initiated. All the clinical data used in these studies was from taken from the GE scanner at St Mary's Hospital, details of which are provided in **chapter 7**. The texture samples digitised from (BRODATZ, 1966) were also used in these experiments.

1.6 STRUCTURE OF THESIS

Chapter 1 provides an introduction to the thesis. The chapter begins with a section describing the motivation for this research project. This section concludes that computational tissue characterisation methods such as texture analysis may prove to be valuable diagnostic tools to assist the MRI user in clinical practice. The importance of image visualisation packages such as *MAIVIS* (developed as part of this project) are also discussed in this context.

The diagnostic information provided by MRI images is discussed, and texture is identified as a key feature that the clinician uses in the appraisal of tissue appearance. The concept of image texture, including examples, and its potential for the quantitative description of pictorial information is also discussed. The three methods of texture analysis that have been investigated as part of this project are introduced at this point.

The earliest texture analysis applications were non clinical, such as for terrain analysis from satellite pictures. The successes and failures of many of these early applications are described in this section. These examples demonstrate the potential of texture analysis for tissue characterisation, and have provided much of the motivation for subsequent applications, many of which are discussed in **chapter 2**.

In addition to the development of the *MAIVIS* image visualisation package, considerable project time was dedicated to obtaining image data. Some of the issues relating to image decoding and loading are discussed along with details of the image data used in the studies presented in this thesis.

The most basic form of tissue characterisation is when the clinician recognises tissue structures by their shape, boundary definition, relative location, and relative contrast. This approach, however, is subjective and relies on *a priori* knowledge of the characteristic appearance of tissue. **Chapter 2** describes how more accurate tissue characterisation can be

achieved from quantitative analysis of the tissue parameters related the physical properties associated with the imaging modality. A large review of work published on these quantitative tissue parameters in ultrasound, X-ray, CT and MRI is presented in this chapter. The common conclusion reached by many of the authors is that parameters derived purely from physical properties lack the specificity required to perform accurate quantitative tissue characterisation. It is suggested that additional information can be provided from histology, morphological characteristics, and texture analysis.

A major part of **chapter 2** is dedicated to a review of clinical examples of the use of texture analysis in tissue characterisation. Most of this work has been carried out in ultrasound, with some examples in X-ray, CT and MRI. It is clear from the few examples of texture analysis in MRI that this approach offers great potential in providing the clinician with additional diagnostic information. Many of the texture analysis methods summarised at the end of this chapter have been examined as part of this project. Methods previously applied to MRI are examined to find the conditions under which they perform optimally. Methods previously only implemented in either ultrasound, X-ray or CT are examined to see if they offer the same potential for MRI.

Chapter 3 provides an overview of the *MAIVIS* image visualisation system. A significant amount of project time was dedicated to the design and implementation of this software platform. *MAIVIS* was originally developed not only to serve the needs of this project, but also to provide image visualisation resources to all subsequent projects and research programmes associated with the Biomedical Systems group at Imperial College. The role of *MAIVIS* is also discussed in the context of digital radiology.

A proportion of this chapter is dedicated to the technical aspects of the inception and subsequent development of *MAIVIS*. A description of the software structure, function and the basic image processing tools available is also included.

Chapters 4, **5** and **6** are dedicated to the three approaches to texture analysis examined as part of this project. These three chapters derive texture features from statistical, transform, and fractal analysis methods.

Chapter 4 describes five methods of extracting textural information from image statistics. The

five sets of texture features are derived from first order statistics, run length analysis, cluster size analysis, grey-tone difference probability density functions, and grey-tone spatial dependence matrices (co-occurrence). Cluster size analysis, a two dimensional extension of the run length analysis concept, was developed as part of this project. Modifications were also made to the grey tone difference method (WESZKA *et al*, 1976; CONNERS and HARLOW, 1980) as part of this project. These alterations make this method insensitive to texture orientation, a significant source of errors with clinical MRI image data.

Chapter 5 describes transform methods from which texture features can be derived. These methods include the Fourier, Walsh and Slant transforms. It has be shown (KAIZER, 1955) that useful texture information can be extracted from the Fourier power spectrum of image regions, although this has not been demonstrated with clinical MRI images. The Walsh and Slant transforms have been proposed as similar sources of texture information as part of this project. Several methods of extracting texture features from the results of these transforms have been considered. Non-directional implementations of the Walsh and Slant transforms have also been proposed with the same texture features.

Chapter 6 introduces the concept of the fractal dimension and describes how this parameter can be used to extract useful image texture information. Three methods of obtaining the fractal dimension are discussed in this chapter. Although these methods are based on existing fractal theory, many of the practical issues of implementation have been addressed as part of this project.

Chapter 7 addresses some of the important experimental details that are associated with the studies of texture features reported in later chapters. Information about the MRI image data sets and (BRODATZ, 1966) texture samples used in these studies is provided as a reference. Samples of both sources of images are given to indicate the texture classes investigated.

Also included in this chapter is a description of the process of data collection in the *MAIVIS* environment. This account makes specific reference to the variable parameters associated with each texture feature, ROI selection, data pre-processing and data normalisation. The mechanics of the normalisation methods implemented as part of this project are discussed further in a separate section. Having collected texture feature values for different tissue classes, this data is exported for performance evaluation using PC based spreadsheets. A

description of this evaluation process, including details of the discriminative performance measure adopted, is also provided in **chapter 7**.

Chapters 8, 9 and **10** provide the results of studies of statistical, transform and fractal texture features, respectively. Each chapter comprises a general discriminative performance overview, a correlation study between features, and examples of tissue characterisation with the best performing texture features.

Chapter 11 is the discussion of results and conclusions chapter. The chapter begins with a review of the findings from all three results chapters, including interpretation and a discussion of possible sources of error. A comparison of performance between the features from each of these sets of results, including a correlation study, enables the relationship between these approaches to texture analysis to be established. The most useful method and the best performing texture features are also identified during this study.

The project specifically explores the use of texture analysis to derive robust pathological indicators for MRI images. The conclusion contemplates whether the results obtained as part of this project indicate that texture analysis methods have the potential to provide such diagnostic information. This discussion also considers the influence of image data limitations from the modality. The conclusion questions if this project reaches all the goals set out in the introduction and to what degree this project has been successful. Possible improvements to the both the texture methods examined and the experimental method are discussed, and future work leading from this project is also proposed.

The thesis concludes with a bibliography and additional material in the form of appendices.

CHAPTER 2 - TISSUE CHARACTERISATION

2.1 INTRODUCTION

The aim of any developments in tissue characterisation must be to assist the clinician by increasing the accuracy and sensitivity of identification of structure or function in a particular region of the human body. In a sense, the contrast seen between different tissue structures as a result of the imaging process may be considered as a primitive form of tissue characterisation. The relative contrast between tissue structures under qualitative examination conveys information about the inherent physiology. The clinician recognises tissue structures by their shape, boundary definition, and relative location. Qualitative classification of pathology, however, is dependent on *a priori* knowledge of its characteristic appearance. Human visual perception provides a basic level of tissue characterisation based upon the contrast generated by the interaction of the modality with the tissue structures. For X-ray and Computed Tomography (CT), image contrast is determined by tissue X-ray attenuation. In ultrasound the tissue structure can only be visualised with the B-scan image. This image is a spatial realisation of the information in a series of A-scan echograms. The acoustic attenuation behaviour of the tissue determines the image contrast. Magnetic Resonance Imaging (MRI), however, is a much more complex modality than either X-ray or ultrasound. MRI images reflect multiple tissue parameters, and the visual display is dependent on the image acquisition protocol. However subjective, it is possible for clinicians to implement basic tissue characterisation in all four modalities from the visual display with an understanding of the fundamental imaging process. The main source of error in this approach to tissue characterisation lies in the inherent operator dependency, the outcome of which imposes a strong subjectivity in the results, and in many cases is responsible for poor precision.

More accurate tissue characterisation can be provided from quantitative analysis of the parameters generated by the modality. With such techniques, operator variation can be reduced. It is also possible to include normalisation schemes with to remove systematic errors and machine dependency, thus improving repeatability. Given the disparate nature of each modality, a different approach is required for each.

The interaction of ultrasound with biological tissues depends on the tissue structure and the relative acoustic impedances of its constituent materials. The detection of tissue disease is

possible when the disease process alters the tissue's acoustic properties. In quantitative ultrasound potential tissue characterisation parameters are those which describe physical properties of the tissue, such as ultrasonic attenuation, velocity of sound and scattering properties. For simplicity these properties are normally quantified from the A-scan RF echo amplitude or frequency spectra, although the B-mode image represents the same information.

Traditional X-ray imaging depends on visualising the *shadows* produced by structures within the body with different X-ray attenuation coefficients. Computed tomography provides trans-axial sectional X-ray images, with high density resolution and high accuracy. Both X-ray and CT provide a excellent soft tissue contrast from the tissue attenuation properties, but CT often requires contrast enhancement. Quantitative information for tissue characterisation is provided by the pixel grey level distribution in the digitised images.

The focus of Magnetic Resonance Imaging (MRI) tissue characterisation is those tissue characteristics which affect and influence the magnetic resonance signal. These characteristics relate mainly to water content and the distribution of protein and lipid molecules in the tissue. In some cases, other less commonly found molecules can provide valuable tissue information. Perfusion, diffusion, and flow mechanisms can also provide clinical information to aid tissue characterisation. The sensitivity of MRI to relatively minor changes in the state of water provides good soft tissue contrast. Quantitative MRI tissue characterisation is carried out from parametric images based primarily on calculated T1 (spin-lattice) and T2 (spin-spin) relaxation times and proton density.

Section 2.2 in this chapter highlights the extensive use of quantitative tissue characterisation methods implemented by a number of authors in clinical ultrasound, X-ray, CT and MRI. In each example, the tissue parameters derived relate to the physical properties associated with the modality. A common conclusion that has been drawn from this body of work is that parameters derived purely from physical properties lack the specificity required to perform accurate quantitative tissue characterisation. This result holds despite the efforts by some investigators to use carefully calibrated equipment and optimised methodologies.

To improve the accuracy of quantitative tissue characterisation, it has been suggested that additional information such as morphological characteristics and texture analysis should be included. Such techniques provide additional information about the structural

composition of tissue not always apparent in the basic display. **Section 2.3** provides clinical examples of the use of texture analysis in tissue characterisation. Most of this work has been carried out in ultrasound, where image texture results from the interactions between the ultrasonic pulses and the structure of the tissue. Some use of texture analysis in X-ray and CT has been demonstrated. In MRI there are considerably less examples of the successful use of texture analysis. This is partly due to the comparatively recent introduction of MRI into clinical practice, and also the complex nature of the modality. Nevertheless some successes have been reported in this field.

The review material presented in this chapter enables the most common and most effective texture methods employed in clinical tissue characterisation to be identified. It is clear from the few examples in MRI that such methods offer great potential in providing additional useful clinical information. The main objective of this project is to examine the efficacy of texture analysis methods in MRI tissue characterisation. Clearly the first step is to find some corroboration with the findings of authors with texture methods in MRI. The second step is to investigate additional texture methods that have been successfully implemented in ultrasound, X-ray and CT. In the implementation of both of these sets of texture methods as part of this project, consideration has been given to the influences of data normalisation, the explicit constraints of region of interest size, and sensitivity to tissue orientation. The quantitative tissue measurements derived from these texture methods must also be repeatable between different systems, patients, and images within the same study.

2.2 NON TEXTURAL TISSUE CHARACTERISATION

The earliest attempts to perform quantitative studies in ultrasound used a one dimensional display of echo amplitude versus time. This is termed A-scan ultrasound. It has been possible to differentiate between *in vitro* normal and diseased heart muscle using the acoustic attenuation information derived from A-scan ultrasound (JOYNT *et al*, 1979). JOYNT suggested that such characterisation would not be achievable *in vivo* because of the limited bandwidth, and the lack of control over factors such as overlying tissue and cardiac motion. However, JOYNT showed that diseased heart muscle could be detected by observing the state of contraction of cardiac muscle from fluctuations of the mean and variance of the RF echo amplitude and frequency spectra in the cardiac cycle. The shape of the echo amplitude distribution has also been found to provide a reliable indicator of tissue characteristics,

(JOYNT *et al*, 1980). JOYNT *et al* were also able to demonstrate the differentiation between *in vivo* normal and damaged myocardium.

The relationship between the acoustic attenuation and frequency was found by (KUC, 1980) to be dependent on the tissue. The slope of the acoustic attenuation coefficient plotted against frequency was shown to correlate well with the disease state of *in vivo* liver. The acoustic attenuation was also found to be dependent on temperature (PRICE *et al*, 1980). In addition to acoustic attenuation, PRICE *et al* suggested using the velocity of sound and the acoustic impedance for ultrasonic tissue characterisation. The attenuation parameters are mainly dependent on collagen content, while the speed of sound and acoustic impedance are mainly influenced by water content. The work of (CLOOSTERMANS *et al*, 1986) showed both the velocity of sound and acoustic impedance for *in vivo* liver tissue to give useful trends for the histological differences associated with diseased and normal tissue. CLOOSTERMANS *et al* were unable, however, to demonstrate quantitative tissue characterisation with these methods.

The more commonly used ultrasound mode is the B-scan. This mode provides a two dimensional image format. The position and direction of the ultrasound transducer determines the lines in the image, where the intensity values in the lines originate from the echo amplitude of separate A-scans. Instead of examining the attenuation coefficients for separate A-scan echoes, (WALACH *et al*, 1986) developed a novel technique for estimating ultrasonic attenuation coefficients from B-scan images. With this method he was able to segment the human liver from the surrounding tissue types, and produce parametric images indicating the different tissue groups present.

In general, little success has been demonstrated in quantitative ultrasonic tissue characterisation with parameters that describe physical properties of the tissue, such as acoustic attenuation, velocity of sound and acoustic impedance. However, considerably more success has been achieved by a large number of authors in tissue characterisation from analysis of the inherent texture within ultrasonic images that results from the interactions between the ultrasound and tissue structure.

In X-ray the most widely reported application of tissue characterisation comes with mammography, where breast tissue is characterised from measured X-ray attenuation coefficients. For example, good image contrast can be achieved between fat tissue and normal

fibrous tissue because they have markedly different attenuating properties. The degree of tissue attenuation can also vary with X-ray energy. This tissue property was utilized using dual energy X-ray imaging to suppress the fat/fibrous contrast by (JOHNS and YAFFE, 1987), enhancing the contrast between pathology and normal tissue. It was suggested by (WATMOUGH *et al*, 1988) that X-ray mammography was not sufficiently reliable for the evaluation of suspicious breast lesions, and proposed that a multi-pronged approach combining mammography, ultrasound and transillumination would be more successful. Nevertheless, X-ray mammography has been used as a method for cancer detection by the UK breast screening programme. The success of this work has been reported by (WELLS, 1990).

Because conventional X-ray yields only attenuation information, its use in tissue characterisation in all but a few cases is very limited. This modality gives primarily structural information and lacks sensitivity to the subtle pathological changes that occur in the early stages of disease. In computed tomography (CT), like X-ray, the CT-numbers represent the X-ray attenuation of the corresponding tissue at a given average energy. Unlike X-ray, however, CT is able to present the three dimensional internal structure of the body in the form of a series of slices, and with the sensitivity to show soft tissue detail.

Both (STEPHENS *et al*, 1977) and (LEVITT *et al*, 1977) found CT to be highly accurate in detecting and defining space-occupying lesions and in detecting fatty infiltration in the liver. However, both authors reported difficulty in the detection of diffuse liver disease. This result was latter confirmed by (COLEMAN *et al*, 1982) who found no significant difference in the CT-numbers of normal and cirrhosis tissue in the liver. The first-order statistical tissue characterisation parameters derived from the amplitude distribution of CT-numbers were reviewed by (PULLAN *et al*, 1978) and (DUERINCKX and MACOVSKI, 1979). The mean, standard deviation, coefficient of variation, skewness, range, and mode were all found to be promising tissue descriptors. However, with the dose limitations of clinical CT, DUERINCKX and MACOVSKI demonstrated that these parameters are dominated by artifacts caused by quantum noise. Successful characterisation of benign pulmonary nodules from the average CT-number was achieved by (SIEGELMAN *et al*, 1980). Clinical success was also demonstrated with histograms of CT-numbers by (RAITHEL *et al*, 1988). Pulmonary changes characteristic of asbestosis could be differentiated in accordance with their degree of severity, and progress of fibrotic pulmonary processes could be sensitively monitored.

The use of absolute CT-numbers for *in vivo* tissue characterisation is compromised by a number of technical and geometrical factors. The main causes of variation in CT-numbers are scan noise, scan technique, patient geometry, and scanner performance variability. The influence of these factors was demonstrated by (McCULLOUGH and MORIN, 1983). McCULLOUGH and MORIN observed significant intra- and inter-scanner variability in CT-numbers for thoracic geometry, indicating the need for quantities which are independent of the CT scanner and depend only on tissue composition. A description of how dual energy CT can be used to eliminate the energy dependency due to scattering processes from attenuation information to provide quantitative tissue parameters was provided by (ALVAREZ *et al*, 1981). Tissue characterisation of white and grey matter, which have extremely close attenuation values, was possible with this method where conventional CT failed. ALVAREZ *et al* do not, however, define the accuracy of this method. The successful application of this approach to the measurement of bone mineral density for the assessment of mineral content in osteoporosis is described by (CRAWLEY, 1990). This application is possible because the linear attenuation coefficient of a tissue is a function of it elemental composition. CRAWLEY also described lung pathology characterisation from density measurements.

Magnetic Resonance Imaging (MRI) tissue characterisation is based primarily on the T1 (spin-lattice) and T2 (spin-spin) relaxation times and proton density. Contrary to other imaging methods, MRI images reflect multiple tissue parameters. The degree to which these physical properties influence the image contrast depends on the acquisition parameters. Although not routine, it is common practise for these properties to be mapped separately as parametric images. Because these parameters provide excellent contrast between soft tissue, MRI is a potentially powerful method for the identification and characterisation of pathology. Early experiments with T1 and T2 relaxation rates by investigators such as (KOUTCHER *et al*, 1978) established MRI as a modality capable of distinguishing normal from diseased tissue with a high degree of consistency. KOUTCHER *et al* used relaxation rates to derive a malignancy index to discriminate *in vitro* normal and malignant samples of colon, lung and breast tissue. This result was significant in the light of *in vivo* experiments that were being carried out with animals demonstrating the feasibility of non-invasive tissue characterisation.

The possibility of tissue separability using images weighted towards different physical processes was demonstrated by (ADAMS *et al*, 1987). Statistically significant tissue separability was achieved between normal breast tissue, cancerous lesions, and benign lesions

based on tissue intensity levels from T1-weighted, T2-weighted, and proton density pulse sequences. It was suggested by (TAYLOR and BUSHELL, 1986) that substantially more information could be gained from MRI images if tissue could be characterised uniquely by their MRI parameters. Characteristic T1 and T2 time constants were calculated for six different tissues of male rat with good accuracy and with high correlation from individuals of different sex and species.

T2-selective parametric images of the human body were generated by (GERSONDE *et al*, 1985). T2 images were created by assigning grey levels to image pixels according to a linear mapping of T2 value. The histogram of T2 values of an image was shown to exhibit a characteristic pattern due to the composition and distribution of molecules in the tissue. GERSONDE *et al* were able to roughly characterise three primary classes of protons by their T2 value; water, lipid and protein protons. Secondary T2-selective images are created by windowing the histogram for an appropriate tissue type and mapping with grey levels or pseudo-colour.

The work of (GLAZER *et al*, 1986) examined a series of quantitative parameters based on absolute signal intensity, intensity ratios, and calculated T1 and T2 relaxation rates for adrenal gland tissue characterisation. Intensity ratios (calculated for the adrenal lesions to liver, kidney, and skeletal muscle) proved more accurate than either calculated T1 and T2 relaxation times or absolute signal intensity. GLAZER *et al* suggested that this result was due to the system dependency of signal intensities and calculated relaxation times. This result was corroborated by (JENKINS *et al*, 1987), who found T1 and T2 relaxation rates to be heavily dependent on magnetic field strength, the imaging equipment, and their method of measurement. The spatial resolution of the images was also inferior to those for equivalent contrast enhanced CT images. The resolution was further reduced in examinations at the site of the pancreas by pancreatic excursion during respiration and cardiac motion. A comparison of pancreatic disease studies using low-field MRI with CT and ultrasound was made by (SMITH *et al*, 1989). The conclusion of this work was that although MRI was limited by its relatively poor spatial resolution, it was neither better nor worse than either CT or ultrasound in accurately demonstrating tumours. These results demonstrate the potential of MRI as an alternative modality to ultrasound and CT in existing tissue characterisation applications.

A study of a group of patients with a wide spectrum of spinal dysraphic lesions by both CT

myelography and MRI with T1-weighted and T2-weighted sequences was carried out by (JASPAN *et al*, 1988). JASPAN *et al* found that MRI tissue characterisation was more accurate than CT for the spinal canal, and proposed that MRI could challenge CT-assisted myelography as the "gold standard". In addition to demonstrating superior characterisation and anatomical definition, MRI also offers the attractive benefits of being non-invasive and without radiation load.

Measurements of T1, T2 and proton density tissue parameters on a large group of patients with brain tumours were carried out by (JUST *et al*, 1988a; JUST and THELEN, 1988b). JUST *et al* used the brain as the subject of their studies because its structure offers high signal-to-noise ratio and minimal motion induced artifacts, and also offers a variety of different tumours. The proton density was normalised to the value of white matter in each patient. The results indicated that even with simultaneous calculation of MRI tissue parameters, broad overlapping of values occurs between different tumour groups. JUST *et al* concluded that reliable tissue characterisation of most solid tumours by means of their tissue parameters alone cannot be made, and that additional information like morphological characteristics, clinical data, and histology continue to play a dominant role in correct diagnosis. The low specificity of MRI tissue parameters can partially be explained by their dependence on water content, regressive changes in the tissue, and the relatively nonspecific features that can be found in varying degrees in many tumours independent of their type. It was suggested by (KJAER *et al*, 1991) that the low specificity of MRI tissue parameters reported by previous authors were, in general, due to shortcomings in experimental strategy ie. too few subjects, too few data points and mono-exponential T1 and T2 calculations. KJAER *et al* investigated a large number of cases of CT-proven intracranial tumours, carefully observing optimal experimental conditions. T1 and T2 relaxation times were calculated with both mono- and bi-exponential analysis methods with a statistically significant number of data points. No diagnostic improvements were provided by the extended experiments in this study. Significant differences in the T1 and T2 relaxation times were demonstrated, but biological scatter and overlap between tumour types seriously impeded the specificity of the results. In conclusion, KJAER *et al* suggested that the low specificity of the tissue parameters was due primarily to tissue heterogeneity and not experimental errors and methodology. KJAER *et al* proposed that a multi-parameter model of voxel-wise T1 and T2 distribution, morphological characteristics, and texture analysis etc. would enable the required specificity to be obtained.

Although many authors have shown strongly altered T1 and T2 relaxation properties for pathology with respect to healthy tissue, the common consensus is that data scatter is a major inhibiting factor for accurate tissue characterisation. This scatter was quantified to some extent by (BOTTOMLEY *et al*, 1984), who reviewed a large database of T1 and T2 relaxation times over a range of frequencies (1-100 MHz) as a function of tissue type, frequency, temperature, *in vivo* versus *in vitro* status, time after excision, and age. The standard deviations of T1 values for tissue groups examined under the same experimental conditions were about 20%, reflecting mainly systematic errors. For the measurement of T2 relaxation times, the systematic errors were about 30%. The two most probable causes of these errors are systematic data collection errors and intrinsic tissue heterogeneity.

In order that the clinical value of tissue characterisation by MRI could be assessed, the Biomedical Engineering Concerted Action Committee of the European Community was activated in 1984. The aims of this concerted action, reported by (PODO, 1988), were for a group of experts from 40 European institutes to discuss the expectation and limitations inherent to tissue characterisation by NMR. The ultimate goal of this work was to minimise systematic errors by establishing guidelines and protocols for ensuring *in vitro* and *in vivo* T1 and T2 measurements obtained from different centres are comparable and reproducible, whatever the equipment used. A small trial of *in vitro* measurements carried out with the imaging protocols developed by the Concerted Action was described by (MATHUR-DE VRE *et al*, 1988). It was shown that in almost all cases, for both T1 and T2, the spread of relaxation values reported from different individual MRI sites was less than for the literature compilation. In a series of system performance assessments devised by the concerted action, (LERSKI *et al*, 1988) found most systems in clinical use not to be optimised for T1 and T2 measurements. In 1988, motivated by the scientific achievements of the first project, a second EC Concerted Action project was activated, (PODO *et al*, 1993). Groups formed by a total of 18 cooperating European institutes were set up for collecting pilot databases of *in vivo* relaxation time measurements on muscle, brain and liver. The clinical MRI equipment was optimised and accurately calibrated using a revised set of performance assessment test objects and protocols, as described by (LERSKI *et al*, 1993a). Under these conditions, (DE CERTAINES *et al*, 1993) carried out a study of T1 and T2 relaxation data from different MRI sites. The results of this work indicated it was possible to obtain comparable relaxation rate values at different sites with a variation of less than 10%, considerably better than that measured by (BOTTOMLEY *et al*, 1984) in his study.

Initial hopes that MRI could characterise tissue based purely on T1 and T2 relaxation times have been somewhat disappointing. In 1989 an additional method of obtaining tissue contrast was introduced by (WOLFF and BALABAN, 1989) called magnetisation transfer contrast (MTC) imaging. The application of a continuous off-resonance RF pulse with an additional RF channel enables the saturation of the motion restricted macromelecular protons. The signal from the unbound water protons is decreased by the amount of saturation transferred by the magnetisation exchange process from the macromelecular protons. Because the signal loss is related to the rate of magnetisation transfer between the two pools of protons, it is dependent on the tissue type. WOLFF and BALABAN demonstrated the quantitative effects of magnetisation transfer with rabbit kidney by constructing parametric image maps of the ratio between proton density images with and without saturation. Improvement in contrast was achieved where large differences in magnetization transfer rate occurred. Parametric image maps were obtained by (ENG *et al*, 1991) for both the magnetisation transfer rate and the observed reduced T1 in rabbit kidney. ENG *et al* found the transfer rate map of the kidney to be consistent with WOLFF and BALABAN's proton density ratio map, but with significantly greater contrast between cortex and medulla. ENG *et al* suggested that both the transfer rate constant and reduced T1 maps should prove useful in the quantitative characterisation of tissue relaxation processes, and in the improvement of tissue contrast in proton MRI. WOLFF and BALABAN's methodology was later improved by (HU *et al*, 1992), enabling the macromelecular protons to be saturated by a train of on-resonance binary pulses.

Cardiac motion degrades the spatial resolution during the long echo T2-weighted sequences required for the detection of ischemic regions, tumours, or transplant rejection. Because MTC has a strong effect on muscle tissue but little or no effect on blood, (BALABAN *et al*, 1991) proposed its use as a method for increasing ventricular-cavity versus myocardium contrast with short echo sequences. Under these conditions the required tissue contrast is achieved with minimal motion artifacts. BALABAN *et al* also suggested MTC an ideal approach for characterisation of cardiac muscle under various pathologic conditions. Myocardium undergoes numerous structural adaptations as it remodels in response to systemic hypertension. A study of the magnetisation transfer in hypertrophic rats carried out by (SCHOLZ *et al*, 1993) found the T1 of unbound water and the transfer rate primarily reflected alterations in tissue water content. SCHOLZ *et al* also found these parameters provided important insights into the sources of variations in relaxation times.

2.3 TEXTURAL TISSUE CHARACTERISATION

Section 2.2 describes the extensive use of parameters and image data collected by different modalities for the purpose of tissue characterisation. This description reports the successes and failures of a large number of authors in the last twenty years. In most cases, the tissue characterisation indices extracted lacked accuracy and specificity. Many authors suggested that better results would be achieved if additional tissue information derived from histology, morphological characteristics, and texture analysis was also considered. Of these three approaches, texture analysis has been identified as probably the most useful technique for MRI tissue characterisation. This is primarily because of its potential sensitivity to pathological detail, and because image texture is used implicitly by clinicians in their appraisal of tissue appearance.

The aim of this research project is to investigate and comment on the efficacy of texture analysis as a source of quantitative tissue features in MRI. The inspiration for this work comes mainly from the successes achieved in non-clinical texture analysis applications described in **chapter 1 section 1.4**, and the promising clinical results obtained in ultrasound, CT, X-ray and MRI described here. The bulk of the clinical application of texture analysis for tissue characterisation, to date, has been carried out in ultrasound.

It was reported by (LERSKI *et al*, 1979) that it was not possible for an experienced observer to classify the state of diffuse liver disease from a visual examination of ultrasound grey-scale B-scans, and suggested more quantitative measurements could be made. LERSKI *et al* achieved 95% accurate differentiation between normal and abnormal liver disease with simple echo amplitude statistics and signal 'texture' parameters derived from the min-max statistics. Some successes were reported with disease classification. This approach was extended with additional textural features derived from the co-occurrence matrix by (LERSKI *et al*, 1981a; LERSKI *et al* ,1981b). The co-occurrence feature set was optimised to remove redundant correlated features. Maximum diagnostic accuracies were achieved using combinations of up to four parameters, the addition of more produced no improvement and in some cases a deterioration. A comprehensive review of this work and contemporary clinical ultrasonic tissue characterisation is provided by (LERSKI, 1982).

The peaks in the frequency spectrum of echoes in a B-mode system were found by (JOYNT

et al, 1979) to give information on the average spacing between tissue scatterers through constructive interference. JOYNT *et al* suggested that this method yielded structural information not easily discernable on the visual display. The human observer performs badly at image discrimination tasks based on second order statistical properties such as periodicity. This method was successfully applied by (JOYNT *et al*, 1980) for the detection of changes in tissue structure associated with disease for *in vivo* liver. This application of the auto-correlation function in ultrasound B-scan images was also demonstrated by (WAGNER *et al*, 1985). The mean spacing of tissue scatterers in the frequency domain was estimated by (SOMMER *et al*, 1981) by applying the auto-correlation function to the back-scattered ultrasonic waveforms of *in vivo* normal and abnormal human livers and spleens. Measures of scatterer cross-section area were also made from the mean amplitude and variance in the amplitude domain. These measurements were found by (SOMMER *et al*, 1982) to be capable of describing the histology of normal and abnormal tissues.

One dimensional texture functions based on the spectral information in the digitised echo train were derived by (CADY *et al*, 1983) to monitor the pathological changes that accompany muscular dystrophy. The results indicated that ultrasound was more sensitive to the changes that occur in the early stages of the disease than X-ray CT. CADY *et al* also found ultrasound tissue texture to be dependent on the angle on incidence of the transducer.

An investigation of *in vivo* ultrasonic characterization of breast disease from backscattered A-mode signals was carried out by (FINETTE *et al*, 1983). Raw A-scans were used as a source of one dimensional texture information. FINETTE *et al* explored the possibility of obtaining characteristic texture features for the classification of malignant, benign, and normal breast tissue. FINETTE *et al* based their classification scheme on texture features derived from first order statistics, power spectrum information, min-max statistics, grey level run length statistics, co-occurrence, correlation functions, and psychophysical information. Some promising results were obtained with these features.

The work of (CRAWFORD *et al*, 1985) provided a study of the texture of digitised echograms of *in vivo* human placentae in smokers and non-smokers, with a method based on the co-occurrence matrix. The tissue texture was described with the angular second moment, contrast, correlation and entropy features. The high contrast speckle pattern due to the coherent tissue scatterers is strongly dependent upon the depth of the scattering volume.

CRAWFORD *et al* failed to demonstrate any differentiation between obviously different placentae with this method due to inadequate depth correction. With a correction for depth, (MORRIS, 1988) was able to detect gross differences between placentae with these texture features. Both CRAWFORD *et al* and MORRIS observed from the texture features that pregnancy induced hypertension and smoking produced opposing changes in the placenta.

For B-mode image evaluation, (RAETH *et al*, 1985) derived texture parameters from first-order grey level distributions, first-order gradient distributions, grey level run length histograms, the grey level co-occurrence matrix, and the power spectrum. Diagnostic classification was carried out on regions in B-mode images with a broad spectrum of liver pathology. RAETH *et al* showed the overall diagnostic accuracy of computerised texture analysis for the classification of diffuse and malignant disease was 96%, compared to 85% for subjective evaluation.

The shape of the probability density function of grey level pixel values for B-scan images was found by (INSANA *et al*, 1986a) to be determined by the specular and diffuse scatterers in the tissue. The relative contributions of these scatterers to the net echo signal constituted a tissue signature. INSANA *et al* reported some success in the detection of abnormal cardiac tissue with this tissue parameter, although he predicted the possibility of errors due to the distribution of scatterer types in the other surrounding soft tissues. It was suggested by (INSANA *et al*, 1986b) that the overall texture tissue signature would be improved by the inclusion of parameters calculated from the first order statistics of the ultrasonic attenuation. Such a scheme was demonstrated for the classification of diseased breast and liver tissue with measures of mean grey level, correlation, entropy, and skewness. The average scatterer spacing was derived from unprocessed RF echoes (average power spectrum) corresponding to selected regions of B-scan images by (INSANA *et al*, 1986a). This parameter was observed to be a sensitive tissue parameter providing diagnostic information not otherwise accessible to observers. This measure was also used by (INSANA *et al*, 1986b) for the detection of diffuse liver disease, where again it proved to be the best discriminative feature.

B-scan ultrasound was analysed by (WAGNER *et al*, 1986) in a feature space based on underlying physical scattering properties. The diffuse and specular backscattering intensities, and the variance of the specular backscattering intensity were derived from the signal power spectrum and employed as a three dimensional feature vector. WAGNER *et al* showed the

statistical ranking of patterns by this method was more sensitive than that of the human observer. The clinical value of this method was also demonstrated for ultrasound of the liver in the detection of a lesion previously identified on a CT scan not visible on the ultrasound scan.

It was possible for (CLOOSTERMANS *et al*, 1986) to discriminated between normal and tumour *in vitro* liver tissue with texture parameters derived from digitised RF *A*-mode echograms. CLOOSTERMANS *et al* used the echogram signal-to-noise ratio and a back-scattering property called relative reflectivity to characterise the echo texture. Both features were found to be correlated with fat and collagen content, and negatively correlated with water content. Although the results displayed useful trends for the histological differences associated with diseased and normal tissue, their value for *in vivo* conditions was not demonstrated.

Experiments were carried out by (NICHOLAS *et al*, 1986) with combinations of texture features from five different methods for the classification of normal liver and spleen from B-scan image data. The texture features were derived from first-order image statistics, spatial echo density (intensity peaks/unit area), gradient measurements, the power spectrum, and co-occurrence matrices.

A diagnostic accuracy of 88% was obtained by (KRATZIK *et al*, 1988) for differentiation between prostatic carcinoma and hypertrophy. Regions of interest from digitised ultrasound images were analysed with a set of five texture measures; textural edgeness, relative extrema density, spread, grey tone run length, and co-occurrence. KRATZIK *et al* suggested that combinations of these features enabled more reliable judgement of echo-patterns of the gland than the human eye.

Computerised texture analysis was used by (SCHUSTER *et al*, 1988) to evaluate conventional B-mode ultrasound images, but the texture measures were calculated locally making them less dependent on tissue inhomogeneities. Parametric texture maps were formed from regions of interest in the image with each pixel equal to the local values of a specific textural property. SCHUSTER *et al* adopted a set of texture features derived from textural edgeness, grey level run length, co-occurrence matrix, and relative extrema density. Subsequent segmentation was carried out in the map by identifying homogeneous areas of pixel values.

Texture analysis was used by (CHANDRASEKARAN *et al*, 1989) for ultrasound myocardial tissue characterisation. The grey level run length and grey level difference texture methods were applied to digitised clinical B-mode images. CHANDRASEKARAN *et al* concluded that with this method he could distinguish between normal, amyloid, and hypertrophic myocardial structures.

Textural analysis was used by (CHUNG-MING *et al*, 1992) to differentiate between normal, hepatoma, and cirrhosis in liver tissue with ultrasound B-scan images. Experimental results demonstrated the performance of conventional texture features to be no better than physical property based features. The texture features were taken from the co-occurrence matrices, the Fourier power spectrum, the grey level difference method, and LAW's energy measures. CHUNG-MING *et al* also demonstrated the successful use of fractal based texture features.

The most commonly reported applications of textural analysis in X-ray are for the detection and classification of pulmonary disease in chest X-rays and for tissue characterisation in X-ray mammography. These methods have been applied to complement diagnosis by visual inspection and quantitative methods based upon measurements of X-ray attenuation.

The researchers (SUTTON and HALL, 1972; CHIEN and FU, 1974; KRUGER *et al*, 1974; KATSURAGAWA *et al*, 1990) all experimented with texture measurement for the automatic classification of pulmonary disease in chest X-rays. SUTTON and HALL employed texture methods based on a variable-distance local-differencing operator, the power spectrum, and directional edge measurements. The directional edge feature provided the best overall classification results. The performance of texture features derived from the power spectrum, however, were disappointing. CHIEN and FU derived a set of texture features from the co-occurrence of grey levels in the pixel neighbourhood and measured their consistency in the characterisation of normal-abnormal class differences in chest X-ray images. In chest X-rays the changes in vascularity caused by the disease process affects the image texture. In CHIEN and FU's experiments, classification failure was due to small amounts of clothing material superimposed on the lung field. Although this method was shown to be potentially useful, the results were not conclusive because only a relatively small number of images were analysed. KRUGER *et al* used texture analysis for the diagnosis of pneumoconiosis from chest radiographs. The first-order statistics of the digitised radiograph were normalised and spatial texture measures derived from the grey level co-occurrence matrix were calculated. The

texture features measured were auto-correlation, moment of inertia, entropy, and homogeneity. KRUGER *et al* also applied a coherent optical approach for fourier analysis. The automated normal-abnormal classification accuracy on a testing basis was no lower than 88%. KATSURAGAWA *et al* demonstrated characterisation of lung texture in digitised chest radiographs from measures derived from the power spectrum. The two features used were the root-mean-square (RMS) variation (texture magnitude), and the first moment of the power spectrum (texture fineness or coarseness). The texture regions were selected from the intercostal spaces in chest. An automated approach to ROI selection was developed for this application by (POWELL *et al*, 1988). The results of KATSURAGAWA *et al* suggested that computerised texture methods may be helpful to radiologists for accurate classification of pneumoconioses.

An investigation into the use of computerised texture analysis for the risk assessment of X-ray mammographs was carried out by (MAGNIN *et al*, 1986). MAGNIN *et al* selected four first-order statistical parameters; dynamic range, variance, skewness and kurtosis. These globally measured parameters were found to give tendencies concerning each risk group, but provided little discriminative information. Similar success was achieved with the entropy, inertia and homogeneity parameters derived from the co-occurrence and the grey level difference texture analysis methods. MAGNIN *et al* suggested that normalisation of the radiographic contrast may increase the reliability of the texture features. She also suggested that assuming that the mammographs contain one global texture rather than several distinct textures was also responsible for inaccuracies. Fractal based texture descriptors were used by (CALDWELL *et al*, 1990) to classify the parenchymal pattern in X-ray mammographs. The results from this texture classification method correlated well with the radiologists. The agreement level between the clinicians and the fractal classifier was 84 percent.

In Computed Tomography (CT), the possibility that tissue structure could be characterised by a texture or pattern measure was first introduced by (PULLAN *et al*, 1978). PULLAN *et al* found that underlying tissue structures in the liver, spleen, and brain significantly influenced the spatial distribution of attenuation values, and that changes in structure associated with disease could be detected by applying the auto-correlation function to the image data. A gradient analysis method was also successfully applied in this study. PULLAN *et al*'s preliminary texture analysis experiments with the liver were based upon measurements from only a small number of patients. The successful results for obtained with the auto-correlation

function by PULLAN *et al* were confirmed on a much larger series of patients by (RITCHINGS *et al*, 1979). RITCHINGS *et al* also detected textural changes in the spleen in the presence of cirrhosis. The extensive use of the auto-correlation function and gradient analysis as second-order statistical tissue characterisation parameters in CT was reported by (DUERINCKX and MACOVSKI, 1979). DUERINCKX and MACOVSKI demonstrated however, that with the dose limitations of clinical CT, these parameters are dominated by artifacts caused by quantum noise. Following on from the work of PULLAN *et al* and RITCHINGS *et al*, (COLEMAN *et al*, 1982) employed the power spectral density (PSD) as a textural measure for selected regions of CT scans of the liver. Similar Fourier domain techniques have been used successfully to classify textures on chest X-rays by (SUTTON and HALL, 1972) and (KRUGER *et al*, 1974). COLEMAN *et al* found the PSD for normal and cirrhotic livers to differ around two specific frequencies, suggested possibly to be due to the difference in the period structure between normal and diseased tissue.

Successes have also been reported for CT tissue characterisation methods that examine image structure and morphological features. The successful use of CT-number spatial *patterns* to characterise pathology in the lungs was reported by (McCULLOUGH and MORIN, 1983). However, specificity was found to be compromised by the sensitivity of patterns to both the reconstruction filter and the site of the region. Mathematical morphological image processing is a sequence of transformations that progressively simplify an image by preserving only the significant part of the information contained in the image. A description of how this method could be used to characterise CT image structure by its geometrical and/or topological properties was provided by (PRETEUX *et al*, 1985) . PRETEUX *et al* also discussed this method as a precursor to segmentation and texture analysis.

The application of texture analysis for tissue characterisation in Magnetic Resonance Imaging (MRI) has not yet become common-place. The researchers (PRENDERGAST *et al*, 1987; JUST *et al*, 1988a; KJAER *et al*, 1991; DE CERTAINES *et al*, 1993) amongst others have recommended that the inclusion of texture measures in a tissue characterisation scheme would improve the specificity achieved with just T1 and T2 relaxation parameters and proton density.

PRENDERGAST *et al* proposed that improved accuracy in the quantitative characterisation of pathology in clinical MRI could be achieved using statistical measures and texture descriptors derived from T1, T2, and proton density maps. In this study an *N*-dimension

feature vector was formed with measurements of energy, entropy, correlation, inertia and local homogeneity from the co-occurrence matrix, and the first four statistical moments. This method was experimentally verified for the differentiation between nine patients with prostatic carcinoma and eight with benign prostatic hypertrophy. PRENDERGAST *et al* suggested that the increased tissue discrimination achieved offered an opportunity for wider clinical application of texture analysis.

A small clinical study of the textural characteristics of MRI images of selected brain tumours undertaken by a concerted action work group was described by (LERSKI *et al*, 1993b). The first and second order statistics of calculated T1, T2 and proton density parameter images were used to describe the texture of different tissue types. The first order statistics originated from the grey level distribution and the second order statistics from gradient distribution and the contrast, angular second moment, entropy and correlation features derived from the co-occurrence matrix. LERSKI *et al* showed that successful discrimination could be achieved between tumour and edema with statistical image texture analysis of MRI images, and suggested a large well defined clinical study was required to verify these results. This study was repeated by (SCHAD *et al*, 1993) using the same texture descriptors as LERSKI *et al*, but with the inclusion of second order image statistics from the grey run length histogram. SCHAD *et al* showed that the T1 and T2 images contained nearly all the information necessary for the successful separation of white matter, grey matter and liquor from pathological lesions in the brain. However, SCHAD *et al* also demonstrated that textural information was necessary for any successful discrimination between oedema and tumour.

2.4 SUMMARY

Section 2.2 of this chapter reviews the use of quantitative tissue characterisation methods based on the physical properties of the modality. In many cases this information has proved to be insufficient for accurate characterisation, resulting in poor specificity. To improve this accuracy, it has been suggested that additional information such as texture analysis should be included. **Section 2.3** provides a review of the use of texture analysis in a range of imaging modalities.

The most comprehensive use of texture analysis in medical imaging has clearly been made in ultrasound. The main reason for this can be put down to it being the first imaging modality

from which digitised data could be routinely extracted. Although X-ray techniques have been around for one hundred years, it is only in the last twenty years that digitisation methods have been widely available to enable the quantitative analysis of X-ray images. In computed tomography, like X-ray, the image contrast relates to tissue density. Because clinicians can normally interpret this information visually, little use of texture analysis has been made in either technique. Although Magnetic Resonance Imaging (MRI) provides excellent soft tissue contrast, the objective identification and subsequent classification of tissue classes from image contrast has proved to be a difficult task for the human observer. The body of the work on MRI tissue characterisation reported in this chapter clearly demonstrates this subjectivity, and its effect on characterisation accuracy. With reference to successes achieved in other modalities, several authors suggest that quantitative methods such as textural analysis will provide the necessary information for more accurate tissue characterisation. Unfortunately, because MRI is a more recent development, there are very few examples of the use of textural analysis in MRI, and even fewer examples where it has been successfully employed.

Almost all the implementations of MRI texture analysis reported in this chapter are limited in some way. Most researchers have been motivated by the successful use of texture in other modalities and have re-applied the methods verbatim, assuming their behaviour in MRI to be comparable. More often than not, this approach has yielded some promising tissue descriptors. In general though, these results do not bring us any closer to quantifying the full potential of texture in MRI tissue characterisation. This is mainly because these experiments are normally carried out on a single pathology, with a few measurements for only a handful of texture parameters. These implementations rarely consider how the features perform under different experimental conditions eg. the effect of ROI size. Listed below are nine experimental details that will affect the performance of a textural tissue characterisation scheme.

- The effect of pixel width variation
- The effect of ROI size variation
- The effect of data normalisation
- The effect of directional sensitivity
- The effect of parameter stability
- The effect of correlation between parameters
- The effect of combining larger numbers of texture parameters
- The effect of texture subject variation

The aim of this project is to assess the clinical efficacy of texture analysis in MRI. This work has been carried out with a large number of texture methods that have been derived from a variety of sources. The behaviour of the texture model, and the influences of each of the variables above on performance has been examined. The object of the exercise has been to identify the best group of texture features, and the conditions under which their performance is optimal. Not only does this project present a far more comprehensive study of the use of texture features in MRI than can be found elsewhere, but also it also provides an insight into the effect of experimental conditions on parameter performance.

The two main groups of texture methods reviewed in **section 2.3** are those that have been used in MRI, and those that have only been used in ultrasound, X-ray and CT. **Table 2.1** provides a list of these methods and the modality in which they have been implemented. Under the heading *Project*, **Table 2.1** also shows the texture methods that have been assessed as part of this project. This list also includes several techniques based on image transforms that have not previously been employed in any modality for tissue characterisation. All the methods listed in the table can be divided into three distinct branches of texture analysis, notably (S) statistical methods, (T) transform methods, and (F) fractal methods. The definition of the textural features included in each of these three groups forms the body of the theory in **chapters 4**, **5** and **6**, respectively.

	Ultrasound	X-ray	CT	MRI	Project
Min-max statistics	YES				
First order statistics	YES	YES		YES	YES (S)
Gradient analysis	YES	YES	YES	YES	YES (S)
Textural edgeness	YES	YES			
Run length	YES			YES	YES (S)
Co-occurrence	YES	YES		YES	YES (S)
Power spectrum	YES	YES	YES		YES (T)
Auto-correlation	YES		YES		YES (T)
Law's energy measures	YES				
Fractal		YES			YES (F)
Psychophysical Info	YES				
Cluster size					YES (S)
Walsh transform					YES (T)
Slant transform					YES (T)

Table 2.1 : Texture methods - (S) statistical, (T) transform and (F) fractal

The main reason for examining texture methods that have already been implemented in MRI is to find a starting point for this project. Repeating the documented successes and failures of the various investigators helps establish the required level of proficiency for new MRI texture methods. Further experiments with these methods also enables the influences of experimental conditions on accuracy and specificity to be assessed. Texture methods used in other modalities may offer the same potential in MRI. The possibility of using these methods, as well as others of non-clinical origin, to complement existing MRI texture techniques has also been explored.

With such a large number of texture features, a great deal of redundancy exists between correlated features. A good classification scheme requires a minimal number of accurate and un-correlated texture indices to spread out the tissue classes in feature space. The final aspect of this work is to identify the most stabile, discriminative, and un-correlated texture features from the methods investigated, and to form classification schemes. Extensive testing of these schemes with large numbers of images with different pathologies will provide strong pointers as to the efficacy of textural analysis in MRI tissue characterisation.

CHAPTER 3 - THE *MAIVIS* IMAGE VISUALISATION SYSTEM

3.1 INTRODUCTION

In order that the goals set out in **chapter one** may be achieved, a flexible image processing environment is required. At the time this project was initiated (October 1989), the Biomedical Systems group at Imperial College had already begun to develop such a system. Because a great deal of computer programming expertise was available *in house*, a commercial alternative was not considered. The first image processing platform to be developed was conceived and implemented on a mainframe graphics terminal, with the programming language C and a basic graphics library (GKS). Although this system enabled sophisticated algorithms to be implemented, the technology of the hardware and the basic software graphics library were limiting factors. In addition to the development of a number of useful image processing tools, some early development of texture algorithms was carried out using this system for this project.

At the end of the first year of this project, a standalone workstation became available with significant improvements in both hardware and system software over the mainframe. It was decided that this hardware should provide the platform for a major software development plan. A powerful image visualisation system would be designed and implemented for use by this project, and all subsequent projects and research programmes associated with the Biomedical Systems group at Imperial College. Because such a system would provide an ideal platform for the development and implementation of this project work, it was decided that a significant amount of the project's time would be dedicated to its conception.

The system that has been developed has been named the '*M*edical *A*pplication *I*mage *V*isualisation and *I*nterpretation *S*ystem' (*MAIVIS*). The basic *MAIVIS* system enables image visualisation and manipulation, and contains a large library of general image processing tools. In addition to the basic image processing tools available, a large library of image texture analysis methods have been specifically developed for this project.

This chapter describes the evolution of *MAIVIS* through the time-span of this project. The description tracks the development from the first prototype that ran on the mainframe with front end graphics terminals, to the current system that comprises a suite of programs that can

be run on any ANSI C and X-Windows compatible workstation.

Since *MAIVIS* has become freely available outside this project to the Biomedical systems group at Imperial College, it has played a key role as a software platform in many other research projects. In particular, *MAIVIS* has been used to develop a network conferencing prototype, and has provided the basic analysis platform for the PhD projects of (FREE, 1994) and (DE WILDE, 1994). It has also become the preferred image analysis platform for the Department of Health Magnetic Resonance evaluation team based at Imperial College (MagNET). As a direct result of its successes within the research group, *MAIVIS*' design features have now become accepted as the standard on which all Image Visualisation and Interpretation Systems (*IVIS*) should be based. The latest *IVIS* to come out of the Biomedical Systems group (developed as part of another project) is called *MIDAS* (Medical Image Display and Analysis System). This system, written in C, was developed to succeed *MAIVIS* with sophisticated programming tools not available at the time of *MAIVIS*' inception. However, the price that has been paid for the software enforced modularity of these tools, in tandem with the rigid program structure adopted, is a loss in design flexibility making it more difficult to build a user friendly application. It is only possible to successfully build user friendly applications with a strict rigid program structure in an object orientated language such as C++. This is because in an object orientated language the main program functions are permitted to behave differently depending on how they are solicited. This kind of control has been successfully achieved with C in *MAIVIS* because its structure deliberately lapses to allow multi-operation program functions with extensive use of global variables. *MAIVIS* also has the distinct advantage over its successors because it was coded in C with the low level X-Windows library (Xlib) instead of a high level programming tool. Coding this way allows more control over the X-Windows components.

Section 3.2 describes the role of *MAIVIS* as an *IVIS* in the context of post image-retrieval in Magnetic Resonance Imaging (MRI), and **Section 3.3** describes the development of *MAIVIS* from an earlier prototype. The remainder of this chapter is dedicated to a description of the technical components of *MAIVIS*. **Section 3.4** gives a brief insight into X-Windows, the powerful software libraries used in the programming language C to achieve the elegant image display and manipulation processes in *MAIVIS*. **Section 3.5** describes the basic structure and function of *MAIVIS*. An account of how the C language program code of an application can be implemented in the *MAIVIS* framework is also provided. The final part of this chapter,

section 3.6, is devoted to the description of the basic image processing tools that have been integrated into the basic *MAIVIS* system, and their function.

3.2 THE ROLE OF *MAIVIS* IN MEDICAL IMAGING

When this project was first initiated, the typical image post-processing power of a Magnetic Resonance Imaging (MRI) system was extremely limited. The system consoles only had a very basic image processing capability; typically to vary such parameters as image contrast and intensity levels. These tasks were often carried out using trackball devices, where the two degrees of freedom represent the number of grey levels ie. the window width, and the intensity level. Other commonly available tools included basic grey level intensity, statistical and geometric measures.

The primary motivation for the development of *MAIVIS* was to serve the requirements of this project. The secondary, and more general, motivation was that *MAIVIS* became a sophisticated image processing platform capable of supplementing the post-processing deficiencies of these early MRI systems. This general use of *MAIVIS* however, is not limited by the tools developed as part of this project. Additional tools may easily be implemented though this basic platform to suit any specific post processing task.

In the timescale of this project, significant technological advances in MRI systems have occurred. In addition to developments in imaging capability, the user interfaces have also become more sophisticated. The original *button press* control consoles and plasma screens are now being superseded by industry standard workstations with keyboard and mouse control, and large graphical displays. The immediate gains of these systems over their predecessors lie in their user friendly operation and speed and ease of use. Although users may experience a small learning curve when moving from a key-command system to a mouse-pointer system, they soon realise its benefits. This progress is analogous to moving from Microsoft DOS to Microsoft WINDOWS on an IBM compatible personal computer.

Significant developments have also occurred in the image post-processing software. Many of the basic image processing tools implemented as part of this project in *MAIVIS* are now widely available on the latest MRI systems. Despite these developments, the *MAIVIS* platform, with its modality- and manufacturer-independent image analysis capability, still has a major

role to play in diagnostic imaging. Although most leading MRI system manufacturers provide extensive post-processing packages, their general functionality is still very limited. *MAIVIS*, however, provides the user with the freedom to develop and implement tools and specific applications independent of manufacturer or modality in an already well defined image processing package.

Because the latest generation of MRI systems are driven by the processing power of industry standard workstations, low cost workstations can easily be networked onto MRI systems as additional independent consoles. In addition to proprietary analysis software, these hardware platforms will also be able to support independent software packages such as *MAIVIS*. This is possible because *MAIVIS* is portable and can function on any UNIX or VMS compatible platform. Although beyond the scope of this project, the resources required to develop the functionality of *MAIVIS* for a specialised clinical application would be minimal.

Because manufacturers are incorporating industry standard workstations in their MRI systems, it will be possible to combine an MRI system with additional independent consoles from different manufacturers and modalities. This is due to the standard hardware available and the imposition of image data standards such as ACR/NEMA - DICOM 3. These workstations can easily be networked together to communicate with existing imaging equipment, and both MRI or CT systems can be evaluated on one workstation. It may also now be possible for image registration between modalities. These developments open up new avenues for diagnostic imaging, and make Picture Archiving and Communication Systems (PACS) a very real proposition.

The ethos behind PACS is to centralise patient information within a hospital or local health authority, and to allow its distribution through a network of dedicated computer terminals. The quantity and technical specification of the terminals would be determined by the clinical practice. The terminals would be strategically positioned around the site to gain maximum benefit; for example, one in each consulting room. The PACS application enables the complete medical records of a patient to be accessed hospital-wide over a computer network. The record would be in the form of computerised text and digitised images from the relevant imaging modalities.

A more technical definition of the components of a PACS system has been offered by

(WIMMER *et al*, 1991). He describes PACS as being comprised of three primary components, namely; application clients, database servers and image servers. The application clients are typically the computer terminals distributed around the site. They provide the interface between the user and the database server, providing the means for retrieval/update of patient information. The database servers represent the least conspicuous component of PACS. They serve as efficient data archiving and retrieval systems, providing the links between the text and digitised images in a patient folder, and managing the large amounts of data that pass in and out of the record database. The image servers are the dedicated displays where the digitised images can be viewed in conjunction with the application client patient text output.

The (WIMMER *et al*, 1991) PACS model can be simplified further by combining the application client and the image server hardware using an industry standard workstation. The workstation is ideal for this role because it is easily networked, manufacturer independent, and standard hardware. Providing adequate workstation software existed, both text and corresponding images from a patient file can be presented simultaneously. Because of the amount of data involved in a hospital-wide record system, the database server must remain a separate PACS component. The software for additional function such as image analysis could also be added to such a system. *MAIVIS* is capable of both image visualisation and text presentation, and application programmability.

An example of the successful integration of a PACS system with a workstation behaving as image server and application client was demonstrated by (WOLFMAN *et al*, 1992). A PACS system was introduced to replace film in the ultrasound section of a clinical radiology department. Three ultrasound units were connected by a fibre optic network via acquisition nodes to a centra data management system (database server), workstation, and optical jukebox (physical data storage).

It was suggested by (HO *et al*, 1991) that a workstation in a multi-modality PACS network must be responsible for the access, display, and analysis of multi-modality digital images with a large variation of formats. He also proposed that the user interface should allow clinicians with minimal or no computer manipulation skills to use complex analysis tools, and recommended that this be achieved with a graphics orientated user interface, with windows and icons. Independently, these aims also governed the design and development of *MAIVIS*.

The software and hardware considerations of such PACS workstations has been described by (BOYCE *et al*, 1991) in terms of the image visualisation packages developed in the Biomedical Systems group at Imperial, with particular reference to *MAIVIS*.

Further comment on the design criteria for the visualisation software in a hospital-wide PACS project is provided by (LIGIER *et al*, 1992). They describe the visualisation and manipulation of images provided by different imaging modalities as one of the most challenging components of PACS. The first consideration they makes, like (BOYCE *et al*, 1991), is that the user interface is portable, ie. workstation independent. This allows users to benefit from the evolving workstation technology without significant re-design of existing software. In addition to a standard set of image-manipulation and processing tools, LIGIER *et al* say that there is a need for such software to be modular such that the platform can easily be expanded and adapted to a variety of specific clinical applications. These considerations form the principal design features of *MAIVIS*, and are described in more detail in **section 3.5**. In the context of this project, the software platform is expanded to accommodate the textural analysis application. A full tenable version of PACS has not yet been implemented, but there are many examples of a smaller scale system known as Mini-PACS, for example SIENET from Siemens provides solutions for workstations, image management, image archiving, and interfaces/gateways.

The potential of *MAIVIS* in this PACS framework is already beginning to be realised through its successor, *MIDAS* (Medical Image Display and Analysis System) (KITNEY and DEWEY, 1993). This system is now being used by all levels of researchers in the Biomedical Systems group, as well as by the International Consortium for Medical Imaging Technology (ICMIT). In a description of the goals set by ICMIT in its collaboration, this system has been described as a *technical demonstrator*, a mock-up of a medical workstation of the near future, (McCRONE, 1993) .

3.3 AN OVERVIEW OF THE SYSTEM DEVELOPMENT

The first medical image visualisation and interpretation system to be developed in the Biomedical Systems group was designed by Drs Janet De Wilde, Sam Free, and Keith Straughan in 1988 (before this project had begun). This venture was undertaken using a VAX 8600 mainframe front-ended with Sigmex graphics terminals running under a VMS operating

system. The software was written in the C programming language with the GKS graphics kernel C libraries linked to the source code to facilitate graphical display. This system (named *SBT*) was written with the broad aim of allowing the display and analysis of medical images. Images were loaded and displayed, and then manipulated using the large number of image processing functions that could be solicited from the menu structure. The system adopted a fixed image window format with a fixed menu window. At the time of its design it was felt that such a regime would be more suitable than floating windows and pull down menus because they clutter the screen working area, especially when dealing with repetitive processing tasks. These criteria are important when developing a user friendly interface.

For ease of development, the *SBT* system C code was split into a number of files that would be linked together after compilation. Each of these program files represented a different aspect of the system function, for example; menu control, screen arrangement, image loading, image display, and region of interest definition etc.

Although state of the art technology in 1988, there were a number of hardware restrictions associated with the VAX 8600 mainframe that made it very difficult for both the code programmer and the system user. As the *SBT* system became more sophisticated, so peak hour usage became increasingly impractical (Mon-Fri 9.00am to 6.00pm). This was due to the system being overloaded with users and the mainframe operating a time slicing regime to provide fair distribution of CPU time. The slow and limited image handling facilities of the GKS graphics kernel also contributed to this effect. These *run-time* speed problems were a direct consequence of the mainframe being a shared resource. The solution would be to transfer the software to an independent hardware platform with a powerful graphical capability, such as a DEC or Unix workstation. The ultimate aim of the Biomedical Systems Group at Imperial is to have developed a universal image visualisation and analysis package, standing on a standard hardware platform, that provides a vehicle for all research work, in both academic and commercial spheres. The decision to transfer the *SBT* image visualisation system to a DEC or Unix workstation was the first significant step in this direction.

The VAX 8600 mainframe ran under a Digital Research VMS operating system. This operating system was considered by many to be inferior in portability and flexibility to Unix, or Ultrix (the Digital Research version). In the early 90's, although well established, Unix was fast becoming the industry standard; most of the latest workstations were running under

Unix or a proprietary form of it. It was therefore considered a high priority to adopt a Unix based hardware platform to ensure compatibility across industry standard workstations. For software compatibility the version of the programming language C also had to be an industry standard. For this project the standard ANSI version of C was used. Fortunately, this version had previously been used for the original VAX 8600 image visualisation package. There was therefore no learning curve associated with C programming in moving the software from the VAX to a workstation.

The first computer programming tasks in this project were carried with the original VAX 8600 image visualisation package. This mainly involved adding fairly simple image processing functions. In addition to providing an introduction to the functionality and overall design features of the existing image processing environment, this exercise enabled several key tasks to be accomplished. Firstly, a significant advancement in the standard of C programming was achieved, secondly, some basic texture algorithms were developed, and finally, an understanding of the practical implications of image processing was gained. Much of this early work also enabled development ideas for an independent workstation based visualisation software to be considered.

The *MAIVIS* system was developed as part of this project as the prototype independent workstation image visualisation system. Based upon the basic design precepts of the Vax 8600 *SBT* system, *MAIVIS* was written on a DEC 3520 workstation with the programming language C and the advanced graphics kernel X-Windows. Although the hardware platform only supported the Digital VMS operating system, and not UNIX, this work was considered by the Research group to be a major step forward.

Subsequent ports of the *MAIVIS* prototype have bee easily made to DEC 5000 workstations running ULTRIX (DEC proprietary UNIX), SUN SPARCstations running UNIX, and UNIX based Silicon Graphics workstations. This work has been carried out as part of this and other research projects in the Biomedical Systems Group. The evolution of *MAIVIS* from the VAX 8600 *SBT* system to the DEC 3520 workstation, and its subsequent ports to the industry standard UNIX workstations has been described by (BOYCE *et al*, 1991). The role of *MAIVIS* in the PACS environment is also covered in this discussion.

Once the *MAIVIS* prototype had been completed, its role in this project was to provide an

environment for the development and implementation of texture analysis algorithms with MRI images. This image processing environment also provided the platform for a number of other MRI related research projects in the group; notably (DE WILDE, 1994; FREE, 1994), and the Department of Health Magnetic Resonance evaluation team based at Imperial College (MagNET).

All the independent workstations in the Biomedical systems are networked together using an Ethernet Local Area Network (LAN). In addition, most of the IBM compatible Personal Computers (PCs) are also networked. This regime makes it possible to run *MAIVIS* on any of the group's workstations with easy passage of image data from one machine to another. With the commercially available product ExceedW for the IBM PC, *MAIVIS* can be run on a workstation remotely at a networked PC terminal. An overview of the Biomedical systems computer network and its resources are shown in **figure 3.1**. This figure shows how the resources of the five main computing areas used by the Biomedical Systems group within Imperial college are networked. In essence this schematic is a PACS prototype; *MAIVIS* can be accessed and run from many different locations (either IBM PC or workstation terminals) with a common data source. The inclusion of the MRI scanner at St Mary's Hospital completes the picture. MRI image data can be retrieved via the scanner's networked independent console and distributed via the network to wherever it may be required.

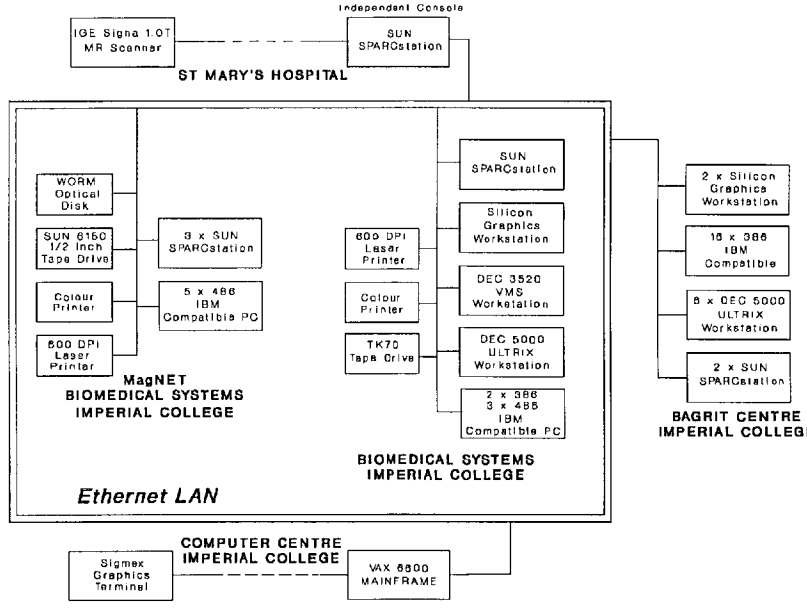


Figure 3.1 : Network overview at Biomedical systems group, Imperial College

3.4 X WINDOWS CONCEPTS

3.4.1 INTRODUCTION

The X Window system is a portable software standard developed at the Massachusetts Institute of Technology's Project Athena. It controls the displays of engineering workstations and provides a standard environment to application software such as page-layout editors and computer-aided-design packages. Applications which use X to operate a workstation display can be easily run on a variety of workstations from a variety of computer vendors. To summarise, the purpose of the X Window System is to provide a *network-transparent* and *vendor-independent* operating environment for workstation software.

The X Window system is a software environment for engineering workstations. It offers a rich and complex environment to the programmer and user of application software. The software environment is defined by a large number of C-language library routines. The *Xlib* library package is provided for the development of applications at the lowest level of programming, enabling complete control over the graphical environment. Alternatively, the less complex high-level X *toolkit* package may be used.

At the time when this project was begun, there was limited access to general X Window literature. The first text to become generally available provided only a description of the *Xlib* library functions, (JONES, 1989). This text formed the basis for the development of the *MAIVIS* software in the X Window environment. Although complex and involving low level programming, the *Xlib* library allowed a large degree of flexibility in *MAIVIS'* design. More recently available is a set of X Windows programming manuals by (NYE and O'REILLY, 1988). This text provides a comprehensive guide to both the *Xlib* and *toolkit* packages. It is now even possible to obtain high level programming languages for designing X Window applications without any prior understanding of either *Xlib* or *toolkit.*

Sections 3.4.2, 3.4.3, and **3.4.4** in this chapter describe three important aspects of X Windows that provide insight into some of the design features of *MAIVIS.* These sections provide a particularly important reference for the description of *MAIVIS'* graphical display, the layout of its screen windows, and the function of its menu structure (all to be found in **section 3.5**).

3.4.2 WINDOW HIERARCHY

The foundation of X Windows is the *base window system*. This system interfaces with the outside world using the X *network protocol*. The *network protocol* is the *only* interface to the *base window system*. *Xlib* provides the means to interface with the *base window system* at a low level. The basic *Xlib* function calls that have been used to initialise the *MAIVIS* X Windows application are described below, aided with the C-language example in **figure 3.2**.

The first X Window System request carried out in any program is to initiate a display connection from the application software to the workstation. **XOpenDisplay** fills the structure *mydisplay* with that information. The *network-node-name* argument contains the workstation network ID. When the argument is null it indicates that the display and the application both relate to the same workstation. The **XDefaultScreen** function returns the workstation's default screen. Finally, **XCreateWindow** creates a new window on mydisplay with the root window as its parent ie. **XDefaultRootWindow**.

```
mydisplay  = XOpenDisplay( "network-node-name" )
myscreen   = XDefaultScreen( mydisplay )
mywindow   = XCreateWindow(mydisplay, XDefaultRootWindow(mydisplay),
             other-parameters )
```

Figure 3.2 : X Windows initialisation (C-language code)

Once this initialisation process has been completed, all the application windows may be created. This is a process of fitting each new window somewhere into the already existing window hierarchy. Each window created must be assigned to an already existing parent window, with the top of the tree being the **XDefaultRootWindow**. There is no limit to the number of assigned children windows to each parent window. As more windows are created, the child window becomes the parent window and the new windows created become its children, and so the hierarchy of windows is set up.

3.4.3 MAPPING WINDOWS

Once a window has been created, it must first be mapped to the screen before it becomes visible. Consequently, any attempt to display either text or image data etc. in an unmapped window will have no effect. If a window is unmapped, making it invisible, then its contents

ie. displayed text, will be lost. Re-mapping such a window makes it visible again, but it will not contain any of the data originally displayed. It is therefore very important to keep a track of mapped and unmapped windows so that visual data is preserved successfully. Unmapped windows may of course be manipulated.

3.4.4 PROGRAM CONTROL AND EVENTS

Once a window has been created and mapped to the screen it can form the basis for the interface between the user and the application. Each window can be configured to provide basic interrupts depending on how the user interacts with it. For example, the window can be defined to send an interrupt whenever the mouse button is clicked over it. This interrupt would notify the computer program and appropriate action would be taken. In X Windows these interrupts are called events. The events are buffered into memory as they occur, and can be read one at a time with **XNextEvent**. Details of this process and various other event management functions are given in (JONES, 1989).

The application therefore can be almost completely defined by the careful construction of a windows hierarchy and the configuration of its windows, and an event handling structure. This is the basic framework around which MAIVIS has been designed.

3.5 THE STRUCTURE OF *MAIVIS*

3.5.1 SYSTEM ORGANISATION

The *MAIVIS* application software consists of the basic application shell, a set of general image processing routines, and a set of user specific image processing routines. The basic shell and the general image processing routines are common to all users. The library of image processing routines is constantly being updated; distribution of new functions is made to users whenever possible. The user specific functions normally constitute personal research work.

The basic system allows the user to load medical images from any digital source (ie. scanning modalities) and then to display them in different sized image windows. There are currently thirteen pre-defined image formats available on the system. Although often a complex process, additional formats can be added to the software whenever required; a significant amount of

image format expertise has been gained as part of this project in providing *MAIVIS* to other users, especially with the Department of Health MRI evaluation centre (MagNET). This is no longer a problem in MRI because the ACR/NEMA standards require MRI manufacturers to conform to a strict image format.

The standard image processing tools are primarily concerned with the image enhancement. In addition, there are tools for edge detection and the measurement of statistics, distances, profiles, and histograms. More details are provided in **section 3.6**.

The user specific functions enable the basic visualisation package to be extended to cover a particular area of work, in effect specifying the application for which the software is to be used. **Figure 3.3** shows the basic application as a box in the centre of the diagram. The file cylinders that point towards the application indicate the potential diversity of modalities as image sources. The specific application is defined by the library of functions used to extend the basic system. The four boxes at the bottom of the figure show typical extensions to the image processing library. For the purpose of this research project a library of textural analysis tools has been developed. *MAIVIS* has also been used by MSc students for project work, MRI image analysis by MagNET, and the PhD work of (DE WILDE, 1994; FREE, 1994).

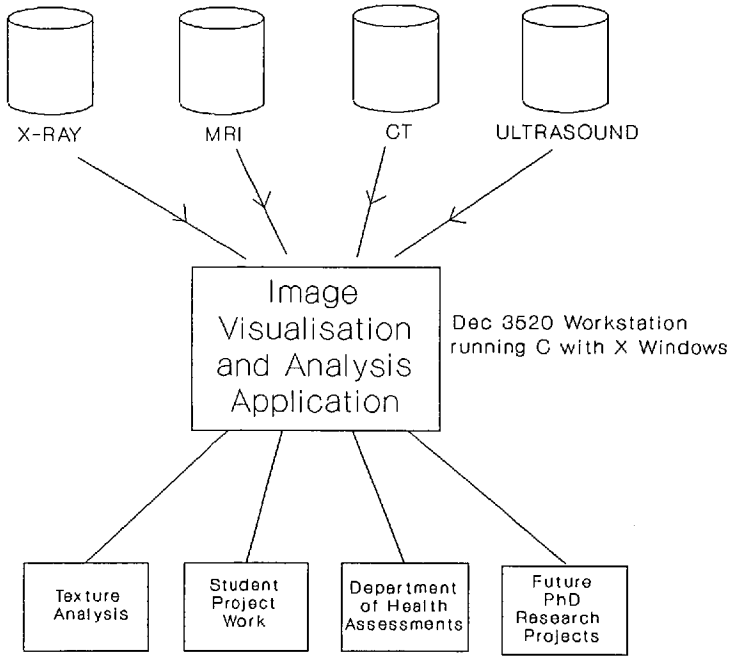


Figure 3.3 : *MAIVIS* applications

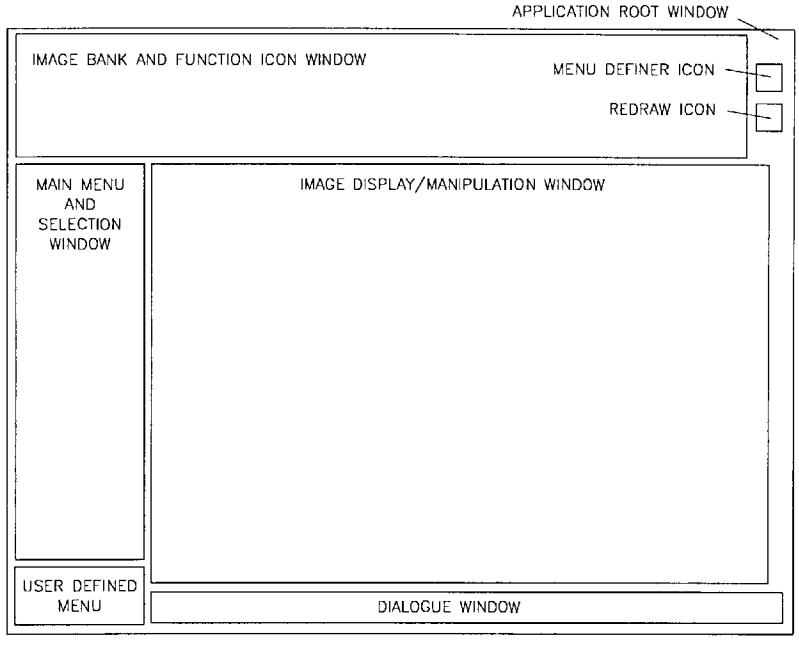


Figure 3.4 : *MAIVIS* main screen windows

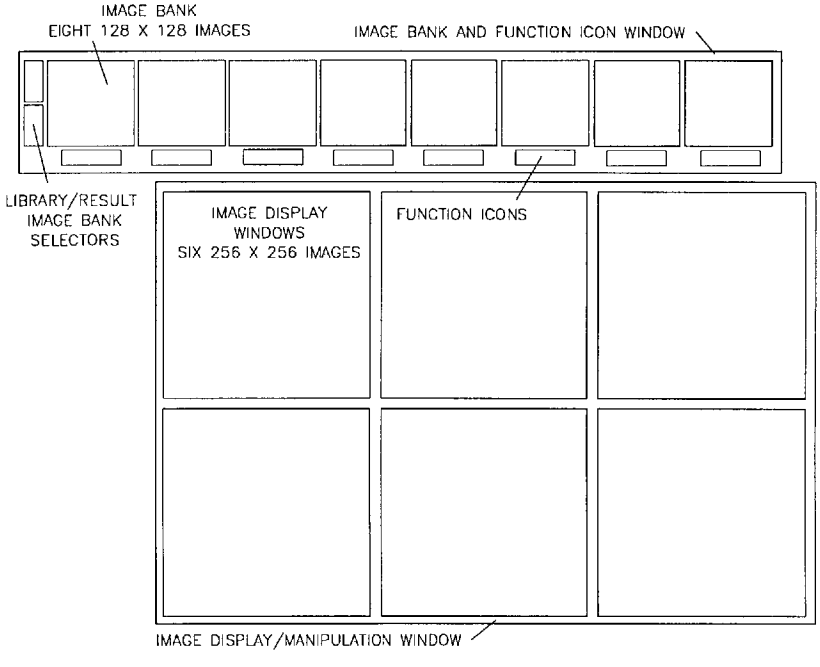


Figure 3.5 : *MAIVIS* image windows

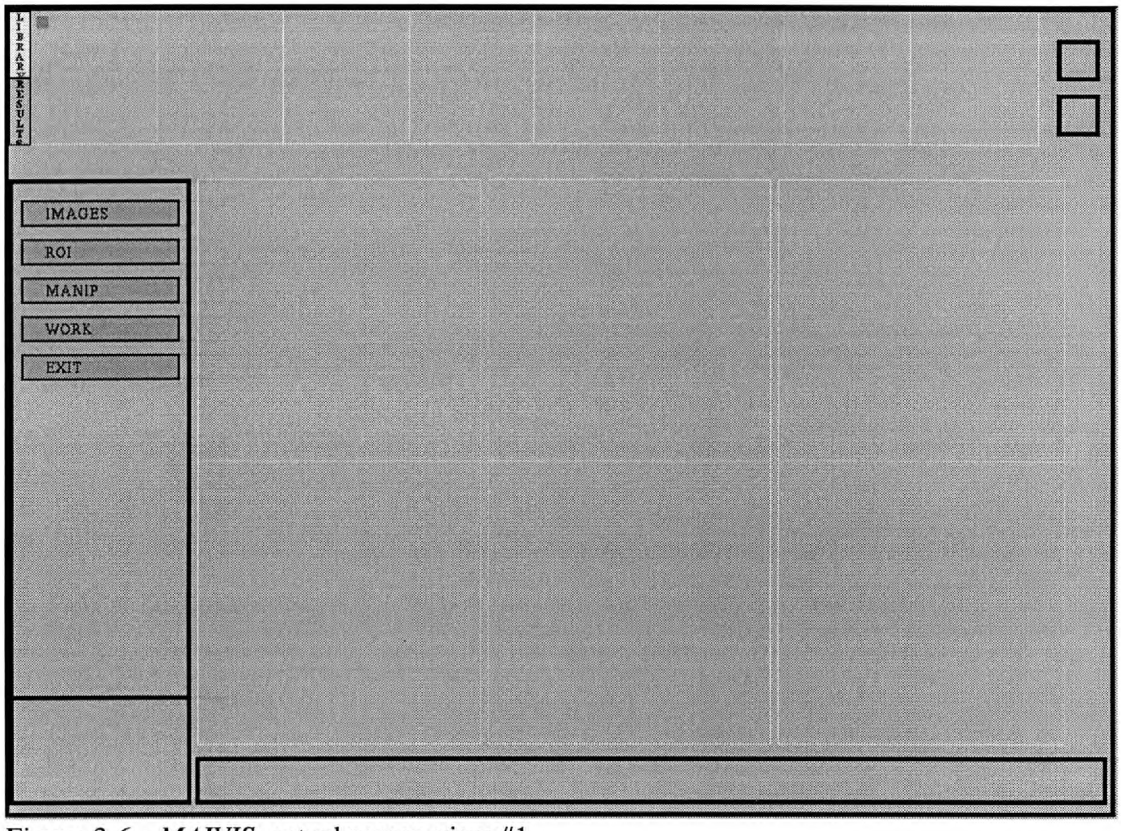


Figure 3.6 : *MAIVIS* actual screen view #1

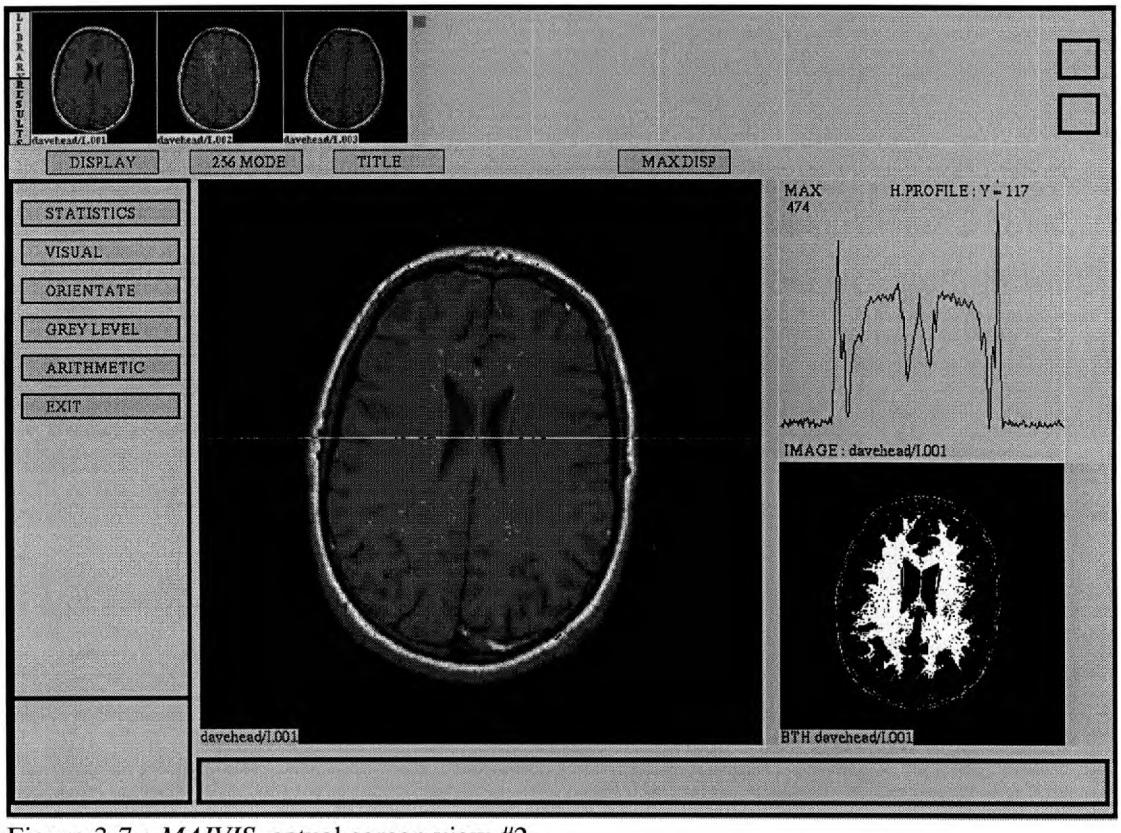


Figure 3.7 : *MAIVIS* actual screen view #2

3.5.2 THE SCREEN LAYOUT

Section 3.4 describes how X Windows can be used to create a hierarchy of windows to interface between the user and the application software. This section gives a brief description of the display windows of *MAIVIS*, how they relate to one another, and how they form the user interface.

The main screen windows shown in **figure 3.4** are all on the same level in the windows hierarchy; a single parent window exists called the *application root window* which has the X Windows **XDefaultRootWindow** as its parent. The *main menu and selection window* is the area of the screen from which the application is controlled. The options of a tree like menu structure are selected from icons in this window. Multiple choice selections are also made from this window; for example a choice of file formats is given when loading an image. **Figures 3.6 & 3.7** give examples of working menus. As well as the fixed menu window, *MAIVIS* has a *user defined menu* window. With the *menu definer icon* this menu can be programmed with up to three popular options from the main menu.

The *dialogue window* is a text input and message window. This window becomes highlighted when either a keyboard input is required or a when the application displays a message. The *redraw icon* is used to refresh the contents of the screen windows after another X Windows application has obscured the display.

The *image bank and function icon window* is another control window. **Figure 3.5** shows the bank of eight 128 x 128 pixel image windows that it contains. There are two pages to the image bank; the *library* bank is for images that have been loaded from file, and the *result* bank is for processed images. The *image bank selectors* shown in the figure are used to toggle between the two banks. **Figure 3.7** shows the *library* bank as indicated by the highlighted selector. Images are compressed where necessary to fit into the bank windows. Each bank of images wraps around; once all eight positions have been filled, images are overwritten. Specified images can be protected so as not to be overwritten. Images are selected for processing by choosing the appropriate window in the bank.

The *image bank and function icon window* also contains a set of eight *function icons*. The latest version of *MAIVIS* only uses four of these icons; two image display functions, one to

change the image display format, and one to display image titles. These are all shown clearly in **figure 3.7**.

The *image display/manipulation window* in **figure 3.4** is the image work space for the application. It operates in two modes; a 256 mode containing six 256 x 256 pixel windows or a 512 mode containing one 512 x 512 pixel window and two 256 x 256 pixel windows. The 256 mode is shown in **figure 3.5** & **3.6** and the 512 mode is shown in **figure 3.7**. The second *function icon* in **figure 3.7** is used to toggle between the modes. Image data is automatically interpolated/decimated appropriately to fit into the display window selected. **Figure 3.7** shows a 512 x 512 MR brain image with its profile plotted in the top right 256 x 256 window. A binary threshold algorithm has been carried out on this image and is displayed in the bottom right 256 x 256 window.

3.5.3 MENU STRUCTURE

A menu based on a tree structure imposes strict control of program flow and is highly organised. Such a structure has been developed for *MAIVIS* based on the idea of a linked list of data structures used in the programming language Pascal. Each node of the tree has forward pointers which point to the other nodes down the tree, the children, and a backward pointer which points to the node up the tree, the parent. When a menu option is chosen, the nodes that it points to are presented as a choice of further options. This process continues until the node of the tree selected does not point to any nodes, ie. it represents the termination of a branch. This node corresponds to a program function, and upon its selection, the appropriate task is carried out. Details of the C-language implementation of this tree menu concept are given in the following text.

Each node of the menu is associated with a defined C structure. The data type of this structure is called *Menu_Type*, and is defined in **figure 3.8**. A data array called *WinMenu*[i] is defined for each node and points to its data structure. This array is initialised in **figure 3.8** for a total of i=*NWINS* menu nodes, where i is the array index. The structure that the data array *WinMenu*[i] points to contains the array indices of the parent and children of the i^{th} node. In the case of a branch termination, where the node has no children, the structure contains information about the program function that must be executed. Information about the node's position amongst siblings in the menu window, and its title are also contained in the structure.

```
struct  Menu_Type    {
                     int      No_Children;
                     int      Parent_No;
                     int      Children[14];
                     int      Screen_Menu_no;
                     char   *MenuName;
                     }   WinMenu[NWINS];
```

Figure 3.8 : Menu data structure

The *No_children* element indicates the number of children, or the number of nodes, to which the current menu node points. This dictates how many options become available once this option has been selected. If children exist then the *Children*[j] element of the current node structure stores the array indices of the node structures of its children. The indices are stored in the order that the option icons occur in the menu window.

If the current menu node corresponds to a program function then the *No_Children* element becomes a function index. This is indicated by the 8th bit being set. The function index is formed from the least significant 7 bits and fed into a lookup table to select one of a possible 127 program functions. If more program functions are required a higher check bit can be used.

The third use for the *No_Children* element is to register the status of the three user definable menu options. These occupy the structures indicated by the array elements i=*NMENUS*, *NMENUS*+1, and *NMENUS*+2, where *NMENUS* is a program constant that allocates the next free structure after the main menu options. For this purpose the 9th bit is set, leaving the least significant 7 bits to hold the menu status. If the status value is zero the menu options are inactive, otherwise the value stored is the array index of the defined function. Again, a higher check bit would be used for a system that required more than 127 program functions.

With this menu concept, the process of moving down the menu tree is simply a case of passing control from parent to child node until a termination node is reached. Moving up the menu tree is achieved just as easily by selecting an *exit* option. These are false nodes that

point back up the tree to the parent of the current node. The *No_Children* element for such a node structure is zero.

The *Parent_No* element in the node structure holds the array index that correspond to the parent node. The parent is the node one level higher up the tree from its children. A negative value indicates that the current node is the top of it menu tree. If *Parent_No* is equal to '-1' the menu option is at the top of the fixed menu. If *Parent_No* is equal to '-2' the menu option is a user definable menu option.

The *Screen_Menu_no* structure element stores the position of the menu option icons when put up in the menu window on the screen with its siblings (1-13). The mapping of a menu option to the screen requires only the position number. The *MenuName* element holds the title for a particular menu option. The title is displayed in the appropriate menu option icon in the menu window.

The menu concept developed for *MAIVIS* is both simple and clever, allowing menu trees to be implemented very quickly and easily. New nodes are also easily defined and linked in with existing menu trees. A detailed example of the function of this menu concept is provided in **appendix 1**.

3.5.4 USER DEFINABLE MENU

A common problem associated with a rigid menu, such as that discussed in **section 3.5.3**, is caused by the fixed position of the program functions within the menu tree. Sometimes a sequence of program functions is required to carry out a particular task. If the program functions required for this task are spread throughout the menu tree, users will be forced to move up and down through the menu structure in order to achieve their goal. This can be laborious and very user unfriendly.

An example of these difficulties can be seen even with a simple process. Consider the tedious task of loading one hundred images and performing a simple operation on each image in turn. If the load function and the processing function are on very different branches of the menu tree, a great deal of time and effort would be expended in the menu manipulation. *MAIVIS* presents the solution to this problem in the form of a separate menu structure that contains

three definable menu options. Once defined, these options, may be selected completely independently of the main menu. These three options can only be defined as program functions, they cannot support a menu sub-tree.

The user definable menu is programmed with a special function called the *menu definer*; solicited from a dedicated icon - **figure 3.4**. An option is added by selecting a menu function from the main menu and then choosing one of three possible positions in the user definable menu window. The new option is completely active once defined and behaves exactly the same as it would have done had it been selected from the main menu. The *menu definer* also allows options to be removed from the definable menu.

3.5.5 PERSONAL MENUS

Once the *MAIVIS* prototype was complete, it was released for use outside this project. It soon became clear that with a large number of users and system programmers, the menu tree would very quickly become large and unmanageable. It was also difficult to keep all users up to date with the changes that were being made to the various copies of *MAIVIS*. All users require the basic *MAIVIS* functionality, but their needs for additional functionality depend on their specific applications. Different groups of users require different sets of functions and procedures. It became apparent that in addition to the basic menu structure, *MAIVIS* should provide a specific user menu tree. Although this concept has not yet been fully implemented in *MAIVIS*, a prototype has been developed.

The user menu tree would be linked into the system menu tree at a predefined node. This idea has already been experimented with for several research projects. Users writing algorithms specific to their own research have been able to define personal menu structures and link them into the system menu at a menu node called *WORK*. Specific menu structure array indices have been reserved for this purpose. The array index of the top node of the user menu tree is a system constant in order to make the link with the main menu compatible. The user menu options are currently defined in the same file as the system menu options. A more flexible solution is for the user menu options to be defined in a separate user menu file. The user menu file and the user program function files then form the basis of each application; the main system files remaining constant between applications.

The link between the user menu tree and main menu tree would be completely transparent. As far as the system user was concerned, the application would appear to have one large menu tree. By programming different user menus, the basic system can be tailored for diverse applications.

3.5.6 RESTRICTIVE MENU OPTIONS

An interesting revision of the fixed menu tree is the introduction of restrictive menu options. This enables menu nodes to be programmed so that their existence depends on the profile of the user, or groups of users, currently operating the *MAIVIS*. The user profile could be entered as an account number when starting the system. A good example of the use of restrictive menu options would be to keep sections of the menu tree that are under development out of general circulation. Users not involved in the development work would not find these options in their menu tree.

Restrictive menu options could also be used to customize *MAIVIS* according to the user profile. A selection of applications such as MR Image Quality Assurance, Breast Mammograph Analysis, and MR Texture Analysis, for example, would all exist as research projects on the same version of *MAIVIS* and share many functions and tools. However, with careful use of restrictive menu options, users would only be aware of their privileged application. In addition to the advantages that may be gained in system management, this feature may also have an important role to play in system security, especially when dealing with sensitive issues such as patient confidentiality, commercial data protection, and the confidentiality of unpublished research.

The concept of restricted menu options has not yet been implemented on any working version of *MAIVIS*. However, the current system could easily be revised by adding the attribute *WinMenu*[i].*Codeword* to the menu data structure, **figure 3.8, section 3.5.3**. The codeword assigned to each of the existing menu nodes would be made up n binary digits, each bit representing a group of users. The bits set in a codeword indicate the groups of users that are entitled to access the corresponding menu option. Users would be assigned to the relevant user group by a special account number, and only the menu options that are permitted for that group would appear in their menu. Non privileged menu options would simply not exist.

3.5.7 RESULT DATA AND RESULT IMAGES

When an image is processed in some way from the functions available to the *MAIVIS* user, the results are either presented in the form of a graph, a set of parameters, or in new image(s) called *result* image(s).

Results presented in the form of parameter lists can be tendered as tables of numbers in either the image display windows, or as external data files. External data files can subsequently be exported to an IBM compatible personal computer (PC). The results from the texture analysis methods developed as part of this project were down-loaded onto a PC and evaluated using a sophisticated spreadsheet - Microsoft EXCEL for windows.

Graphs or graphical interpretations of result data can either be presented in the image display windows directly, or mapped into image structures as *result* images in order that they may be preserved.

By definition, a *result* images is created by some form of mapping process from an existing image from either the *library* or *result* image banks. The new image will inherit its name from its parent or parents (image subtraction for example), and the process involved. If a data image with the title *IMAGE1* is used in a fast Fourier transform (FFT) operation for example, the resulting transformed image will have the title *IMAGE1.FFT*.

The new image will assume the next free position in the *result* bank of images, with the first 18 characters of its name displayed for identification. Because a new image may be the result of many successive functions, it is possible that the title of a *result* image may exceed 18 characters in length. The full title of an image can be displayed in the dialogue window by selecting the *title* function icon, **figures 3.4-3.7**.

Many of the *result* images created on *MAIVIS* may be of use in subsequent analysis sessions. A facility exists to enable these data images to be saved for future recall. Data images can be saved using a variable file format in a specially allocated ULTRIX/UNIX directory.

3.5.8 REGIONS OF INTEREST

In most applications users are only interested in examining specific areas of data images. These areas are commonly referred to as regions of interest (ROI). *MAIVIS* provides the user with a set of functions for the definition of square, rectangular, and circular ROI. Each ROI is defined either from a set of fixed dimensions or with specially written rubber-banding routines. The ROI can the be positioned either with the mouse pointer or with a fixed set of coordinates.

Although not yet been implemented on the latest version of *MAIVIS*, a method for the definition of irregular shaped ROI has also been developed. This method allows the user to select a given number of points around an area of an image, which are subsequently joined up. These points may then be pulled inwards or outwards independently of the rest of the shape enabling the ROI to be moulded to fit over any area of the image.

Once defined, regions of interest can be copied from one image to another; either one by one, or as a whole set. The dimensions and positions of defined ROI may also be found and stored for reference. This information is essential for repeat measurements and registration of features between images. In this project this capability is essential for comparing different texture features from the same region in the same image, and in differently acquired images.

3.6 BASIC IMAGE PROCESSING TOOLS

This section describes the system tools available on the basic *MAIVIS* system. These tools are grouped under seven headings which roughly describe their function. In addition to a set of basic tools enabling image file management and ROI definition, these tools are grouped under the following headings; image statistics, visual manipulation, image orientation, image arithmetic, grey level manipulation, edge detection masks, and spectral tools.

3.6.1 GENERAL TOOLS

Images are loaded into *MAIVIS* by selecting the appropriate file format and then the relevant files from a library of images. Once loaded, images occupy the next available position in the *library* image bank. When this bank is full, loaded images wrap around and start filling from

the first position in the bank, overwriting the images already present. Image protection and un-protection functions enable management of the *library* and *result* image banks to prevent vital images being overwritten.

Some image modalities produce images that are 512 by 512 pixels in size. Because of memory restrictions, *MAIVIS* only allows the user to store one image of this size at a time. In texture analysis, the areas of interest in an image are always small in comparison to the whole image. An image capture facility exists to grab a region of size 256 by 256 pixels from a 512 by 512 image to form a new, more manageable image. A function has also been developed to cater for images that are 1024 by 1024 pixels in size. This tool selects a 512 by 512 region by coordinates only, and loads this as a complete image.

A number of types of regions of interest can easily be created to examine specific image regions. Once created, ROI can be associated with any image to allow region comparison between images. Image ROI can also be magnified, or zoomed, to form an image of size 256 by 256 pixels. This feature has been designed purely for display purposes, therefore no attempt to interpolate the data points has been made.

For the purpose of exporting result data from *MAIVIS* to an IBM PC based analysis package, a pipeline function has been developed. When active, this function sends the results of all subsequent image analysis to an external text file. This file can be read by a PC spreadsheet and then further analysed.

3.6.2 IMAGE STATISTICS

These tools enable the first order statistics of an image to be examined. The maximum, minimum, and first four moments (mean, variance, skewness, and kurtosis) of the pixel values in an image region can be calculated. Individual pixel grey level values can be obtained either by pointing at the pixel, or by stating its coordinates. The distance and angle between two points in an image can also be found.

3.6.3 VISUAL MANIPULATION

With these tools two dimensional horizontal or vertical image profiles can be displayed. A line

is defined on the image by two points and the grey level intensity is plotted against the pixel position on the line. A three dimensional perspective plot of the image can also be displayed.

3.6.4 IMAGE ORIENTATION

These tools enable the orientation of an image to be changed. Image can be rotated through either 45, 90, 180 or 270 degrees. Images can also be flipped either across a vertical or horizontal axis.

3.6.5 IMAGE ARITHMETIC

This set of tools allows image pixel values to be altered using arithmetic functions. Two images of the same size can be either added together or one subtracted from the other. The values of corresponding pixels are either added or subtracted and the result is given to the pixel with the same coordinates in the result image.

Boolean algebra can also be performed between the two images. The bit patterns of the corresponding pixels are either logically *AND*ed, *OR*ed or *XOR*ed to get the desired result. Boolean arithmetic can also be applied on a single image with a constant value supplied by the user.

3.6.6 GREY LEVEL MANIPULATION

This set of tools provides several useful functions for re-mapping image pixel grey level intensities. These methods are generally used for image enhancement.

The invert function transforms images like *photographic negatives* of the originals. The log scale function re-maps images with a narrower distribution of grey level intensities. By reducing the number of available grey levels, the quantisation function is a crude but effective method of grey level segmentation.

Image thresholding is a useful technique for examining an image over specific range of pixel grey levels. This method is used to nullify pixels with grey levels outside a pre-defined range. The range can be defined in one of three ways; above, below, or between user defined values.

The pixels with grey levels within the range are conserved. A binary threshold also exists, whereby these pixels are given a value equal to the maximum possible grey level. The resulting image shows white where the pixels are conserved and black where they have been nullified.

A histogram of the recurrence of each grey level can be plotted. A function also exists to evenly re-distribute this grey level probability density function. A local averaging mask can be used to remove image noise. This function can be applied with a threshold level to ensure that image features are not blurred.

3.6.7 EDGE DETECTION MASKS

The standard set of edge detection tools have been implemented. The set contains Roberts, Sobel, FCline, FCedge and Laplace filters. These tools produce an image where only edge information is given, with varying degrees of success.

3.6.8 SPECTRAL TOOLS

The standard spectral tools available to the user are the fast Fourier transform (FFT), the power spectral density, and the autocorrelation function.

3.7 SUMMARY

A significant amount of project time was dedicated to the development of *MAIVIS*. This effort is totally justifiable because this platform was fundamental for the development, testing, and implementation of the texture analysis tools used for this research project. All the basic functions and tools have been designed with this application in mind. Many of the basic tools described in this chapter, in addition to their defined usage, have been incorporated into more sophisticated functions for texture analysis.

In addition to the specific requirements of this project work, the development of image visualisation platforms is a valid and important area of Magnetic Resonance Imaging (MRI) technology.

CHAPTER 4 - STATISTICAL TEXTURE FEATURES

4.1 INTRODUCTION

The three principal approaches used in image processing to describe texture are statistical, structural, and spectral, (GONZALEZ and WINTZ, 1987). Statistical methods use the statistics of an image to characterise its texture. Structural methods deal with the geometric composition of an image, or its texture primitives; texture primitives are the geometric building blocks of texture. Spectral techniques are based on the properties of the Fourier spectrum of an image. Comparison and performance evaluation of such methods carried out by (WESZKA *et al*, 1976; HARALICK, 1979; CONNERS and HARLOW, 1980) look more favourably at the texture features derived from image statistics.

This chapter looks at a number of established methods that produce texture features from first and second order image statistics. In addition to some suggested alternative approaches to the use of existing techniques, a novel method called 'cluster size' has been developed. Cluster size is a method that extends the concept of 'run length coding' from one to two dimensions.

First order statistics refer to the distribution of grey level intensities in an image region and are presented as a histogram of the frequency of occurrence of each grey level in the image. Although easy to calculate, first order statistics give little texture information because they do not contain spatial information. The statistical moments calculated from this histogram can provide some basic texture descriptors (GONZALEZ and WINTZ, 1987). These texture descriptors are described in **section 4.2**.

Spatial information can be obtained by examining the second and higher order statistics of an image. In order that such information can be extracted properly, it is often required that the first order statistics of a set of images are equalised. If the first order statistics of two images can be equated, then higher order statistical texture descriptors can be used to measure purely higher order differences. This process is often required when comparing texture samples acquired in the same way, but with different average grey level intensities. The first order statistics of an image are normalised either by histogram equalisation or by fixing the image mean and variance to some constant values. The normalisation of the first order statistics of images has been discussed in more detail by (GONZALEZ and WINTZ, 1987) and (LERSKI

and STRAUGHAN, 1987). In order that images from different sources (ie. differently acquired) may be compared on equal terms, more sophisticated normalisation methods are required, beyond the scope of this project.

For the texture methods described in this chapter that investigate second order statistics, we have suggested adopting these normalisation techniques. A comparison will be made of the efficacy of the texture features derived from normalised and un-normalised image data to observe the effects of equalising first order statistics.

The second order statistics of an image are concerned with the distribution and the position of the pixel grey levels in the image. This statistical information is obtained by examining the local relationships of grey levels between image pixels. Because textural information is held in the inter-relationship between neighbouring image pixels, second order statistical techniques are potentially powerful texture descriptors.

Because of the natural variation in the orientation of tissue texture in clinical MRI images, it is difficult to uniquely specify tissue type using a set of the texture descriptors that contain directional information. These parameters are not likely to give the same value if the texture pattern changes its orientation. Many of the texture methods that extract second order statistical information in this chapter are sensitive to orientation changes of a pattern or texture. In these cases we have tried to suggest ways to make modifications the methods to minimise this sensitivity. The negative side of this stipulation is that by loosing the directional specificity of these methods, the amount of textural information that can be acquired is reduced. These methods may also become less effective at differentiating between texture patterns.

In this chapter we discuss five methods for analysing image texture using second order statistical information; run length coding, cluster size, grey difference method, spatial grey level dependence method, and the sum and difference co-occurrence method.

Run length coding is a method that has its origins in early data compression techniques. Its use for texture analysis was suggested by (GALLOWAY, 1974). This method is used to examine the second order statistics of an image by describing each row or column as a collection of runs of pixels of the same grey level. This method is described in **section 4.3**.

The cluster size method is technique that has been developed as part of this project in an attempt to modify run length coding to become a method that is independent of texture orientation. This method is detailed in **section 4.4** with step by step details its algorithm. We suggest that the performance of the resulting texture descriptors should be compared to those from conventional run length analysis.

The grey tone difference method extracts second order statistical information by examining the local differences in grey level intensity within an image. This technique, described in **section 4.5**, uses the statistical moments of the grey level difference histogram as texture measures. Because the grey level differences are measured in a specific direction, this method is sensitive to the orientation of a texture. We have suggested that by summing the grey tone difference histograms over all possible directions the resulting set of texture features will be independent of texture orientation. We suggest that the performance of both the directional and non-directional forms of this method should be evaluated and compared.

The final second order statistical method we have implemented is called the spatial grey level dependence method, or co-occurrence. This method was introduced by (HARALICK *et al*, 1973) and is described in **section 4.6**. The co-occurrence matrix is the normalised 2D histogram of the frequency of the grey levels of co-occurrent image pixels, from which fourteen unique texture features may be extracted. A preliminary study of the behaviour of these features in MRI has been carried out (LERSKI *et al*, 1993b) with algorithms developed as part of this project. Like the grey tone difference method, this method uses a positional operator to identify the co-occurrent pixels. This method is therefore sensitive to the orientation of tissue texture. It was suggested by (GOTLIEB and KREYSZIG, 1990) that the co-occurrence matrices from all directions should be summed to overcome this problem.

Section 4.7 describes the sum and difference co-occurrence method as an alternative to the co-occurrence method of (HARALICK *et al*, 1973). This technique was introduced by (UNSER, 1986) and enables substantial savings in computer time and memory resources for only a small loss in accuracy. This method transforms the conventional second order probability density function for co-occurrent pixels into two un-correlated first order probability density functions. However, this method has not been implemented in this project for two reasons. Firstly because it is an approximation of the full co-occurrence method, and secondly because the computational resources available for this project are unlimited.

4.2 FIRST ORDER STATISTICAL APPROACHES USING MOMENTS

The first order statistics of an image region can be evaluated by examining its grey level histogram, ie. a plot of grey level against the frequency of its occurrence. The statistical moments are calculated from this histogram to provide some basic texture descriptors (GONZALEZ and WINTZ, 1987).

Let z be the random variable denoting discrete grey level, and let L be the number of possible grey levels. Let $h(z)$ be the number of image pixels at the given grey level z. If N is the total number of image pixels (eg. N = 256x256), then the grey level probability distribution can be defined by the expression given in **equation 4.1**.

$$p(z) = \frac{h(z)}{N} \tag{4.1}$$

It can be seen from **equation 4.1** that the probability that a pixel will have a grey level z_i is $p(z_i)$, with i in the range 0 to L. The mean grey level, μ_1, is defined in **equation 4.2**. The subsequent n^{th} statistical moments, μ_n with $n>1$, are defined in **equation 4.3**.

$$Mean = \mu_1 = \sum_{i=1}^{L} z_i p(z_i) \tag{4.2}$$

$$\mu_n(z) = \sum_{i=1}^{L} (z_i - \mu_1)^n p(z_i) \quad n>1 \tag{4.3}$$

$$skewness = (\mu_3/\mu_2)^{3/2} \ , \ kurtosis = (\mu_4/\mu_2)^2$$

The second statistical moment is the variance of the grey level distribution, $\sigma^2(z)$. The variance can be used as a measure of grey level contrast, and can be used to establish descriptors of relative smoothness. One such measure given by GONZALEZ and WINTZ, R, is defined in **equation 4.4**. This descriptor a value $R=0$ for areas of constant intensity, and approaches $R=1$ for large values of variance, indicating a large range of grey level values.

$$R = 1 - \frac{1}{1+\sigma^2(z)} \tag{4.4}$$

Equation 4.4 is a subjective measure of texture as it simply indicates the range of grey level used to construct the texture. It should therefore be used with some care.

The third moment is called the skewness, and represents the degree of asymmetry of the grey level distribution histogram about its mean. The fourth moment is the kurtosis, which represents the flatness of the distribution compared to the Normal distribution. Both measures help describe the shape of the image intensity histogram, and consequently may be useful texture descriptors.

Because the first order statistics of an image contain no spatial information, the texture descriptors are limited. To be useful they require the average illumination of all the texture samples to be the same. Because higher order moments such as skewness and kurtosis describe the shape of the histogram, their values will not be dramatically affected by slight variations in the sample illumination. Significant variation in illumination between texture samples would, however, vary the shape of the grey level histograms dramatically. The histograms would be compressed or elongated along the intensity axis and the column height would vary appropriately. In this case it would become difficult to use the skewness and kurtosis to characterise the image texture.

Although it does not look likely that moments taken from first order statistics will make ideal texture descriptors, a full assessment of these methods will be made.

4.3 RUN LENGTH CODING

Run length coding is a method that has its origins in early image compression techniques. An image row or column can be seen as a series of runs of contiguous pixels with the same value. Run length coding analyses each row or column of the image and replaces it with a description of the length and grey level of the runs of pixels that make it up. This simple technique is an effective method of reducing the amount of space required for data storage.

Methods that analyse first order image statistics compute their texture measures from the grey level histograms only. These texture features carry no information regarding the relative position of pixels with respect to one another. The spatial information contains important textural information and can only be found in the second and higher order image statistics. Run length coding is a second order statistical method that extracts spatial information in addition to first order statistics. Pre-processing methods that normalise the first order statistics of an image could prove useful when comparing textures with very different image statistics.

It was the suggestion of (GALLOWAY, 1974) to use run length coding as a method for extracting image texture information. She defined a grey level run length primitive as a maximal connected collinear set of pixels that all have the same grey level. The runs can be categorized by the grey level, the length of the run, and the direction of the run.

GALLOWAY computed the joint probability of run length and grey level for the vertical, horizontal and the two diagonal directions (0, 45, 90, 135 degrees). For a given direction, the joint probability that there is a run of length j having grey level i is $p(i,j)$. For ease of computation, this function holds the frequency of each event; it has not been normalised by the total number of events as with conventional probability density functions.

GALLOWAY used these probability density functions to compute five functions analogous to those developed by (HARALICK *et al*, 1973) for grey level co-occurrence matrices. GALLOWAY also utilised the set of aerial photographs used by HARALICK *et al* as a source of texture samples. This was a set of samples of six different terrains, namely: lake, scrub, swamp, suburb, railroad and orchard. Using the five texture measures for each of four directions, she illustrated that about 83 percent identification could be made of the six categories of texture.

Because of their complexity, GALLOWAY's five texture measures have been defined in **appendix 2** in **equations A2.1** to **A2.5**, where the number of grey levels is given by N_g, the maximum run length by N_r, and the direction of the run is defined by θ. The significance of these texture descriptors is described below.

Equation A2.1 is the *Short Runs Emphasis*. This function divides each run length value by the length of the run squared. The function is normalised by the total number of runs in the image region. This measure emphasizes the short runs. An abundance of short runs over long runs indicates that the image has a fine texture.

Equation A2.2 is the *Long Runs Emphasis*. This function multiplies each run length value by the length of the run squared. This has the effect of emphasizing the long runs in the image region. The measure is normalised by dividing by the total number of runs in the region. A high value indicates a coarse image texture.

Equation A2.3 is the *Grey Level Non-uniformity*. This function squares the number of run lengths for each grey level. The measure is normalised as in **Equations A2.1** and **A2.2**. Grey level non-uniformity is a measure of how unevenly the runs are distributed over the grey levels. When runs are equally distributed throughout the grey levels, the function takes on its lowest values.

Equation A2.4 is the *Run Length Non-uniformity*. This function squares the number of runs for each length and is normalised in the usual way. This is a measure of how unevenly the runs are distributed over run length. If the runs are equally distributed throughout the lengths, the function will have a low value.

Equation A2.5 is the *Run Percentage*. This measures the fraction of the image in runs. It should have its lowest value for images with the most linear structure.

Image quantisation can be used to reduce the resolution of a texture region. This effectively 'blurs' an image texture and will increases the long run emphasis. The texture will be seen as being more coarse than at the original higher resolution. The image region will be quantised if threshold methods are used, or the number of grey levels used to represent the image are reduced.

Because image quantisation will affect the run length texture parameters, care must be taken with any image pre-processing that is required. These procedures must be standardised for all the data images used in a comparison. This is especially important when an absolute set of parameters are to be extracted, eg. for a classification lookup table.

From the theory discussed, we have some indication of how the long runs emphasis and short runs emphasis measures may be used to quantify and classify texture. A practical implementation of these methods, however, must be made to fully understand them and to evaluate the potential scope of their usage. The grey level non-uniformity, run length non-uniformity, and run percentage measures are very subjective texture descriptors. Their usefulness in the context of this project is not clear from these initial definitions. An experimental approach is needed to assess their significance and to evaluate their potential. The implementation of these run length texture descriptors includes a comparison of the effects of normalisation of the first order statistics and image quantisation.

4.4 CLUSTER SIZE

The run length concept has been adapted to a 'cluster' length idea as part of this project. We have defined the cluster length, or cluster size, as a maximal connected cluster of pixels of the same grey level. The pixels can only be connected in the two orthogonal axis of the image if they are adjacent in either a column or row. The cluster is not defined by a primitive, but by the number of pixels in the cluster.

The motivation for the development of this method is the fact that a cluster is independent of direction. This is a considerable asset in the context of medical imaging, where the tissue texture orientation varies a great deal.

Exactly the same type of image information is extracted using this method as with run length, and the same care must be taken with image pre-processing.

As in the run length method, the joint probability that a cluster of length j having grey level i is $p(i,j)$. This function holds the frequency of each event occurring. The five texture measures used in run length analysis given in **appendix 2** are used again here with the cluster size joint probability density function. We have renamed **equations A2.1** to **A2.5** as *Small Cluster Emphasis*, *Large Cluster Emphasis*, *Grey Level Non-uniformity*, *Cluster Size Non-uniformity*, and *Run Percentage* respectively. Details of how the cluster size method works are provided below.

The original image is shown in **figure 4.1**. For simplicity, this image has very small dimensions $N=8$. The image also has only four grey levels.

The cluster counting algorithm uses a cross mask that relates the centre pixel with its four nearest neighbours on the arms of the cross. The four arms of the cross are defined relative to the centre pixel (x,y) as (x-1,y), (x+1,y), (x,y-1), and (x,y+1). As this mask is moved over each image pixel (x,y) , ($0 \leq x,y < N$), the centre pixel and the of the four nearest neighbours on the arms of the mask are interrogated.

To enable the mask to pass over every single pixel in the image, it must be padded one pixel thick all the way around. This single layer of pixels allows for the reach of the four arms of

the cross. The original image is padded one pixel thick with '-1' as shown in **figure 4.2**.

In addition to the original image, the cluster size method requires an image mask. The pixels in this mask correspond directly to each pixel in the original image. The image mask has its pixel values preset to `-1`, and is padded in the same way as the original image. The image mask is used to identify clusters in the original image, and is shown in **figure 4.3**.

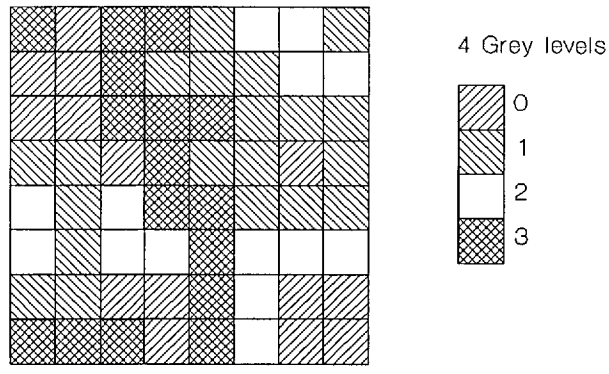


Figure 4.1 : Cluster size method - original image

-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	3	0	3	3	1	2	2	1	-1
-1	0	0	3	1	1	1	2	2	-1
-1	0	0	3	3	3	1	1	1	-1
-1	1	1	0	3	1	1	0	1	-1
-1	2	1	2	3	3	1	1	1	-1
-1	2	1	2	2	3	2	2	2	-1
-1	1	1	0	0	3	2	0	0	-1
-1	3	3	3	0	3	2	0	0	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1

Figure 4.2 : Cluster size method - padded image

-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
-1	-1	-1	-1	-1	-1	-1	-1	-1	-1

Figure 4.3 : Cluster size method - image mask

The clusters are identified using the following *two step* method. The cross shaped window passes with its centre (x,y) along each row of both the original image and the image mask. The four arms of the cross cover the four pixels at the relative coordinates (x,y-1), (x,y+1), (x+1,y) and (x-1,y). The image mask and the original image pixel values are identified by the variables mask[x,y] and image[x,y] respectively.

The image mask is used to hold the cluster identity numbers. These numbers indicate the position of the first sighting of a cluster in the original image. The identity numbers are used to uniquely identify each cluster and to reference the grey level of the cluster in the original image. It is from the numbers in the image mask that the 2D frequency histogram of cluster size and cluster grey level is formed.

STEP 1

Check mask[x,y]; the value of the image mask at the position that corresponds to the pixel (x,y), of the original image.

If the mask[x,y]= -1 (the initial value)

Then ...

If the value of mask[k,m] is not -1, and grey levels image[k,m] and image[x,y] are the same, then set mask[x,y] equal to mask[k,m].

Where (k,m) are the coordinates of the nearest neighbours of (x,y), as defined by the arms of the cross and (k,m) equal either (x,y-1), (x,y+1), (x-1,y) or (x+1,y).

If the value of mask[x,y] is still equal to -1

Then ...

Set the value of the mask (cluster identity number) to correspond to the position in the image, where N is the dimension, or the length of each line of the image.

mask[x,y] = x + (N*y).

Otherwise ...

For each arm of the cross window, image[k,m], where (k,m) equal in turn (x,y-1), (x,y+1), (x-1,y), (x+1,y).

An action must be taken if the image grey levels are equal ie. image[k,m] = image[x,y], and the values of mask[k,m] and mask[x,y] are different, given that mask[k,m] has been defined. The action taken is to look at the two values held by the mask. As the grey levels of these adjacent pixels are the same, although the cluster id's are different, they are from the same cluster. The cluster id. of the most recently found part of this cluster must be changed to the id. of the part found earlier.

STEP 2

Now that the value of mask[x,y] is not -1

Proceed with...

With (k,m) equal to (x,y-1), (x,y+1), (x-1,y) and (x+1,y).

If grey level of the pixels at image[k,m] and image[x,y] are the same, and the image mask, mask[k,m], has the value -1, then copy the image mask value; mask[k,m] = mask[x,y].

The process *STEP 2* smears the cluster identity number (mask value) from the centre to the arms of the cross window if the corresponding adjacent pixels have the same grey level ie. are part of the same cluster. There are no boundary effects with this windowing method because both the image and the mask are padded with '-1's and the pixel values dictate the mask changes.

These two steps are repeated for all pixels in the image (x,y). A flow diagram of this cluster identification process is given in **appendix 3**.

The completed mask contains numbers which identify the pixel where a particular cluster was first discovered in the above process. The number for the pixel (x,y) is given by the expression, $(x + N*y)$, where N is the dimension of the image.

The image mask contains the cluster identity number. This number indicates at which pixel in the image the cluster was first discovered. This is an important reference point because it identifies the grey level of the cluster. The information contained in the cluster mask is used to construct a two dimensional histogram of cluster size and grey level. This in turn is used to construct the two dimensional probability density function needed for the (GALLOWAY, 1974) 'run length' texture descriptors.

A one dimensional array is used to count the number of pixels that occur in each cluster. This array is called *ClusterSize[k]*, where k is the address of the array. This address corresponds to the identity number or pixel position in the image, $k = x + y*N$, where (x,y) are the pixel coordinates and N is the side length of the image. This array is of length $(N*N)$.

As the image mask is scanned pixel by pixel, the array *ClusterSize[k]* is incremented at the address indicated by the value held in each mask position. The address corresponds to the cluster identity number k. The result of this process is a one dimensional histogram with cluster size measured in pixels against identification number. A example plot of this histogram is given in **figure 4.4**. The cluster identity number is given along the conventional x-axis or domain, and the size of the identified cluster is plotted in the conventional y-axis or range. This one dimensional histogram *ClusterSize[k]* is subsequently used to create the 2D histogram of cluster size against grey level.

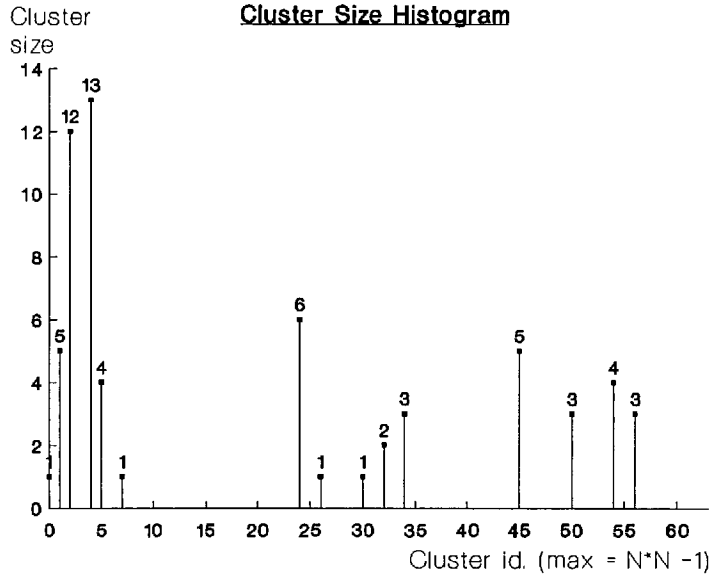


Figure 4.4 : Cluster size method - 1D histogram

It is from the information stored in the one dimensional histogram *ClusterSize[k]* that the two dimensional array of cluster size against grey level is formed. This array is called *Cluster[i,j]*. The coordinates of the array correspond to a cluster of size j and grey level i. The value stored at this address is the number of occurrences of such a cluster in the image.

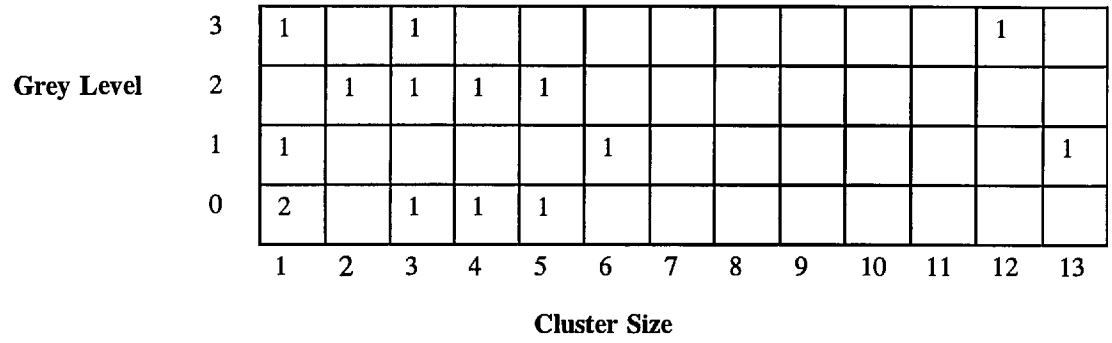
The one dimensional array is scanned through one item at a time. When a cluster is found, the histogram array address k is used to find the starting point of the cluster in the original image. The grey level of the cluster can then be found from this position. The histogram *Cluster[i,j]* is then incremented at the address corresponding to the parameters i=cluster size and j=grey level. The data for the example is shown in **table 4.1**. The values in the two dimensional histogram *Cluster[i,j]* constructed for the example are shown in **figure 4.5**.

The joint probability function of cluster size and grey tone is given by $P(i,j)$, where i and j correspond to the grey level and the cluster size respectively. This joint probability function is given by the normalised 2D histogram of cluster size and grey level. The histogram is normalised by the total number of clusters found in the image, given by N_c. The joint probability function is defined in **equation 4.5**.

Table 4.1: Values of 2D histogram *Cluster[i,j]*

Cluster ID	0	1	2	4	5	7	24	26	30	32	34	45	50	54	56
Grey Tone	3	0	3	1	2	1	1	0	0	2	2	2	0	0	3
Cluster Size	1	5	12	13	4	1	6	1	1	2	3	5	3	4	3

Figure 4.5 : Cluster size method - example 2D histogram



$$P(i,j) = \frac{Cluster[i,j]}{N_c} \qquad (4.5)$$

The number of possible grey levels in the image is given by N_g, and maximum size a cluster by N_{mc}. In the example given the number of grey levels is N_g=4. An image may typically have anything from 7 to 16 bits assigned for the grey levels. However, in practice the image pixel values are likely to be scaled in the range of 5 to 8 bits. The maximum possible size of a cluster is bound by the image size, *N*N*. This value is also bound by the type of image being investigated. It may be necessary to ignore all clusters over a certain fixed maximum size, because they contain no useful texture information and their detection is computer intensive.

We suggest that the (GALLOWAY, 1974) texture measures used for run length in **appendix 2** should be used with the cluster size probability density function to extract second order statistical texture information. The implementation of this method must take into consideration the effects of the normalisation of first order image statistics, and image quantisation.

4.5 GREY TONE DIFFERENCE METHOD

The grey tone difference method is a standard texture technique for examining second order image statistics. This method is based upon the principle that textural information is contained in the local relationship between the grey levels of neighbouring image pixels. This information can be gathered by constructing a histogram of neighbouring pixel grey level, or grey tone, differences. The method described in this section uses the statistical moments calculated from the grey tone difference histogram as texture descriptors. This method originates from descriptions of second order statistical texture methods by (WESZKA et al, 1976) and (CONNERS and HARLOW, 1980).

Let $f(i,j)$ be a digital image quantised to L grey levels (0 to L-1). Let $\delta = (\Delta i, \Delta j)$ be a vector in the (i,j) plane of the image. This vector is used to offset the image before being differenced with itself. The displacement vector is normally taken to be of the order of one pixel in any direction ie. $\Delta i = \pm 1$ & $\Delta j = \pm 1$. The effect of this operation is to compare local changes in grey level intensity in any given direction. The grey level intensity transitions from pixel to pixel are mostly slow or gradual for coarse textures, and more dynamic for fine textures. Because this method examines local changes in grey level, it is appropriate for the extraction of image textural information.

The difference image is defined in **equation 4.6**. The difference in grey level between an image pixel and that of the pixel indicated by the displacement vector is registered in the difference histogram. This assignment is formalised by defining the histogram of the difference image, $h(\lambda)$, with a dirac delta function as in **equation 4.7**, where i,j are the orthogonal image axis ie. image coordinates, and λ is the grey level.

$$\Delta f(i,j) = |f(i+\Delta i, j+\Delta j) - f(i,j)| \tag{4.6}$$

$$h(\lambda) = \sum_{\lambda} \sum_{i} \sum_{j} \delta[\lambda - \Delta f(i,j)] \tag{4.7}$$

The probability density function $P_\delta(\lambda)$ can then be extracted from the histogram by normalising each entry by the total number of entries. This probability density function refers to the probability of the occurrence of a grey level difference λ.

The mean, variance and higher order moments of this distribution can then be found and used as texture measures. A small mean indicates that there are only gradual changes in grey tone, and hence a coarse texture. A large mean indicates a busy image, containing sharp changes in grey tone, and thus a fine texture. The variance gives the range of grey level differences about the mean. A small variance indicates that there is little variation about the difference mean in the histogram. The difference mean can then be taken as a adequate description of texture coarseness. A large value of variance indicates a combination of gradual and sharp changes of grey level in the image. In this situation it is difficult to describe of the overall textural content.

Three more useful texture measures derived from the difference histogram are contrast, angular second moment, and entropy. These three parameters were defined by Weszka (1976) and are shown here in **equations 4.8**, **4.9** and **4.10** respectively.

$$CON \equiv \sum_{i} i^2 P_\delta(i) \tag{4.8}$$

$$ASM \equiv \sum_{i} P_\delta(i)^2 \tag{4.9}$$

$$ENT \equiv -\sum_{i} P_\delta(i) \log_2 P_\delta(i) \tag{4.10}$$

The contrast feature is the second moment of the probability density function P_δ, and can be interpreted in the same way as the histogram variance. The angular second moment is defined to be a measure of the non-uniformity of the difference distribution. Low values indicate that the texture is coarse, and high values indicate that the texture is fine. The entropy measure is large for a uniform difference distribution and small for a very unequal one. This parameter is a measure of distribution uniformity, and can be interpreted texturally in the same way as the angular second moment.

Because displacement vector defines a specific direction in which the image is to be differenced with itself, the grey tone difference method is sensitive to the orientation of image texture. This is of particularly concern to us because the clinical MRI images we are analysing contain texture regions of variable orientation. It will therefore be difficult to uniquely specify a tissue texture. We have suggested using this method in a more general form to overcome

the orientation sensitivity. If grey tone difference histograms taken from all the possible pixel pair orientations are summed, then the measures calculated from the resulting histogram will be independent of orientation. The penalty for this modification is that a certain amount of information will be lost in the process.

The efficacy of the grey tone difference method, or its generalisation, has not yet been established. This evaluation will be carried out as part of this work; with and without the normalisation of the first order image statistics.

4.6 GREY TONE SPATIAL DEPENDENCE MATRICES

Grey tone spatial dependence, or *co-occurrence*, is one of the most popular methods for extracting second order statistical information from image texture. This method is based on the assumption that the texture information in an image is contained in the overall or 'average' spatial relationship which the grey levels in the image pixels have to one another.

A set of grey tone spatial-dependence probability distribution matrices are computed for the an image region. This technique is described by (HARALICK *et al*, 1973), who suggest a set of fourteen texture features which can be extracted from each of these matrices. HARALICK *et al* propose that the texture content information of an image can be adequately specified by the co-occurrence matrix.

The co-occurrence matrix is a 2D histogram measuring the frequency of the grey levels of the co-occurrent pixels in an image region. The matrix has the same number of columns and rows as there are quantised grey levels. The spatial dependency of one pixel on another is defined by a distance d and a direction θ. Each matrix element represents the number of occurrences of two spatial dependent pixels with specific grey levels for an image region. The specific grey levels are indicated by the matrix element positions.

The probability density function used to calculate the texture features is defined by the normalised co-occurrence matrix. This is formed by dividing each element by the total number of pixel pairs interrogated.

Consider a rectangular image $I(x,y)$ with horizontal and vertical resolution N_x and N_y. The

corresponding pixels are quantised to one of N_g grey levels. Let the horizontal spatial domain be $L_x = \{ 1,2, .. ,N_x \}$, the vertical spatial domain be $L_y = \{ 1,2, .. ,N_y \}$, and the set of quantised grey levels be $G = \{1,2, .. ,N_g \}$. The image $I(x,y)$ can be represented by a function that assigns some grey level in the set G at each pixel position (x,y).

The matrix that contains the pairing frequencies, P, is an N_g by N_g matrix. The matrix elements $P(i,j)$ are the number of times two pixels, one of grey level i, and the other of grey level j occur. The pairing of the pixels is all important in creating a matrix that gives maximum textural information.

The use of a position operator is discussed by (GONZALEZ and WINTZ, 1987) ie. the translation that is applied to each pixel in turn to find its dependent pixel. For example, each pixel in the image is paired with the pixel described as "one pixel to the right, and one below". This operator is more formally defined as a vector by (WESZKA *et al*, 1976). The dependent pixel is found by a shift in horizontal and vertical position Δx, Δy. Both (HARALICK *et al*, 1973) and (CONNERS and HARLOW, 1980) use a polar notation, defining the dependent pixel by a distance, d, and an angle θ. The polar notation has been adopted as the most appropriate notation.

As each pixel has eight nearest neighbours, there are therefore eight possible position operators, corresponding to a possible eight co-occurrence matrices. These eight matrices are potentially useful for the detection of image orientation. Details of an image texture can be measured in one of eight directions between adjacent pixels. A knowledge base of textural features can be used as a comparison to detect the orientation of an image.

Given a distance $d = 1$, the possible values for θ are 0°, 45°, 90°, 135°, 180°, 225°, 270°, and 360°, corresponding to the eight nearest neighbours.The work of (CONNERS and HARLOW, 1980) shows that a matrix with $\theta = \phi$ is the transpose of the matrix with $\theta = \phi + \pi$, and makes it clear that the matrices with $\theta \geq \pi$ are therefore superfluous. These four matrices can be constructed by transposing the matrices representing opposite direction dependence vectors.

The generalised construction of the co-occurrence matrices to ignore the distinction between vectors of opposing directions was adopted by (HARALICK, 1979). The purpose of this modification was to reduce the already significant computational requirements of this method.

CONNERS and HARLOW averaged the four asymmetric matrices with their transposes to the same end. The resulting set of four matrices have horizontal, vertical, diagonal from left to right, and diagonal from right to left grey level spatial dependencies. These matrices are symmetric. By halving the number of matrices, only the sum of the contributions of opposing vectors is known. Some information is lost because the individual contribution matrices no longer exist. Because no spatial information is lost, this generalisation does not compromise the textural information.

Two dimensional probability density functions are formed by normalising the co-occurrence matrices. This is achieved by dividing each matrix element by the total number of pixel pairs involved. The normalisation factors are different for each spatial dependency vector. The values for the four matrices are defined in **equations 4.11, 4.12, 4.13** and **4.14**, where $\theta=0^0$ is vertical direction. Subsequent directions are clockwise rotations relative to the vertical .

$$\theta=0^0 \qquad R=2N_y(N_x-d) \qquad d<N_x \tag{4.11}$$

$$\theta=45^0 \qquad R=2(N_y-d)(N_x-d) \qquad d<N_x, N_y \tag{4.12}$$

$$\theta=90^0 \qquad R=2N_x(N_y-d) \qquad d<N_y \tag{4.13}$$

$$\theta=135^0 \qquad R=2(N_y-d)(N_x-d) \qquad d<N_x \tag{4.14}$$

It was suggested by (HARALICK *et al*, 1973) that all the texture information of an image could be extracted using the co-occurrence matrices, and they proposed a set of fourteen texture features to characterise this information. These features, symbolised ***F1***, ***F2***, etc to ***F14***, are defined in **appendix 4**.

Some of these measures relate to specific textural characteristics of the image such as homogeneity, contrast, and the presence of organised structure within the image. Other measures characterise the complexity and nature of grey level transitions which occur in the image. Although it is known that these features contain textural information, it is difficult to identify the specific textural characteristics that are represented by each of these features. However, an attempt to group HARALICK *et al*'s fourteen features according to their nature was made by (GOTLIEB and KREYSZIG, 1990). The four groups (i-iv) are shown on the following page.

i. A group of classifiers that express visual texture characteristics: *second angular moment **F1***, *contrast **F2***, and *correlation **F3***.

ii. A group of classifiers that are based on statistics: *variance **F4***, *inverse difference moment **F5***, *sum average **F6***, *difference average **F6'***, *sum variance **F7***, and *difference variance **F10***.

iii. A group of classifiers that are based on information theory, in particular entropy: *sum entropy **F8***, *entropy **F9***, *difference entropy **F11***.

iv. A group of classifiers that are based on information measures of correlation ***F12***, ***F13***, and *maximal correlation coefficient **F14***.

F1 is a measure of grey level homogeneity in an image region. In an homogenous image region there will be few dominant grey-level transitions; the entries in the co-occurrence matrix will be few, but significant. An image region with random grey levels will have a large number of small entries in the co-occurrence matrix. ***F1*** is calculated by summing the squared grey level transition probability densities. Large grey level transition probabilities such as in homogeneous regions therefore lead to large values of ***F1***.

F2 is a difference moment of the probability density matrix, and is a measure of contrast or the amount of local variation in an image.

F3 is a measure of the grey level linear dependency in an image. A linear structure will have a high correlation feature in the direction it is aligned. Image noise, because it is un-correlated does, however, reduce the correlation values.

Because the features ***F1-F3*** are classifiers that express visual textural characteristics, their significance can be illustrated in image terms. The textural significance of the remaining features from a human perception point of view, however, is not so clear. There are some intuitive expectations of the properties represented by some of these features. For example, one might expect the entropy feature to take higher values for more complex image regions.

A study of the discrimination power of combinations of six selected features was carried out by (GOTLIEB and KREYSZIG, 1990) using 13 (BRODATZ, 1966) texture images. The six

features were chosen in such a way that at least one of them belonged to each of the four groups he defined. GOTLIEB and KREYSZIG proposed that the features chosen were mathematically representative of all fourteen (HARALICK *et al*, 1973) features.

A similar approach to feature evaluation has been adopted by this project using texture regions from MRI images instead of the BRODATZ texture images. A small study is made using all fourteen features. Based on the correlation between the texture measures, a set of features is found that is fully representative of the full set of (HARALICK *et al*, 1973) measures. A more detailed study is then made using combinations of these measures. From this study the textural significance of the HARALICK *et al* features can be established.

The direction sensitivity of the individual spatial dependence matrices can be employed usefully in many applications where the orientation of the image is important. In such applications the sensitivity can be increased by altering the spatial dependency. If the dependent pixel distance, d, is extended, then a larger number of directions, θ, are possible. Because this results in a larger set of co-occurrence matrices, there is greater resolution of rotation detection.

This project is attempting to characterise tissue texture in clinical MRI images. Because of the lack of rigidity and order of many tissues structures in the human body, the orientation of a tissue group may vary across the field of view of an image, or from image to image. In this case it is very important that the texture measures used to classify the textural information are not sensitive to tissue orientation.

A study carried out by (LERSKI *et al*, 1993) applied co-occurrence matrices to a set of (BRODATZ, 1966) textures images and found the texture features to have an appreciable angular dependence. He was also able to differentiate between a set of sample textures. To ensure that the (HARALICK *et al*, 1973) classifiers are invariant under rotation, (GOTLIEB and KREYSZIG, 1990) suggested summing the four spatial dependence. This general model for the co-occurrence matrix has been used in this project. Subsequent calculation of textural features from this matrix removes any directional biases between image regions.

For second order statistical methods like co-occurrence to be useful in texture classification, either the texture features used in classification should be invariant under monotonic grey-level

transformations, or the first order statistics of the image should be normalised. The HARALICK *et al* texture features independent of first order image statistics are the second angular moment, the entropy measures, information measures of correlation, and the maximal correlation coefficient. Normalisation of the first order image statistics ensures that any differences in texture features are due to second order image statistics, or spatial differences.

The size of the co-occurrence matrices is another important practical consideration. Many of the images used in this work have up to 11 bit pixel grey levels. Because the dimensions of the matrices are of the same order, the number of grey levels must be reduced to make the implementation possible. The image regions used in this project are quantised to reduce the number of grey levels in the range of 4-6 bits.

The size of the image region is also important. If it is too small there will not be enough textural information. If it is too large then it will contain objects from several different texture categories. The optimum region size can only be found experimentally.

The co-occurrence matrices used in this project adopt a polar notation with $d=1$ and $\theta = 0°$, $45°$, $90°$, $135°$ with the distinction between opposing vectors ignored. These four matrices are summed to alleviate any directional biases between image regions.

In the initial experiments a full set of (HARALICK *et al*, 1973) texture features is used on a large number of image regions. By measuring the correlation between features, a subset of texture features that are fully representative of the full set is found. Care is taken that, at least one of them belongs to each of the four groups defined by (GOTLIEB and KREYSZIG, 1990).

In the subsequent experiments, parameters such as the region of interest size, the image quantisation and the normalisation scheme are varied, and the subset of texture features are measured. Comparisons are then made to find the optimum parameters.

4.7 SUM AND DIFFERENCE CO-OCCURRENCE MATRICES

The main drawbacks associated with the co-occurrence method are the large computational resources required to obtain the grey-tone spatial-dependence matrices. The number of

operations required to process an image using this method is N^2, compared to $N \log N$ operations for fast Fourier or Hadamard transformations (N is the image dimension). Because the dimensions of the matrices are dependent on the number of grey levels, this method also requires a large amount of computer memory.

Sum and difference co-occurrence was presented by (UNSER, 1986) as substantially quicker alternative to conventional co-occurrence. Because this method transforms the second order probability density function for co-occurrence into two un-correlated first order probability density functions, less computer memory is also required.

Consider the image values I_1 and I_2 of two pixels in a relative position fixed by the displacement vector $(\Delta x,\Delta y)$, **equations 4.15** and **4.16**. The joint probability of these random variables having grey levels i and j forms the basis of the co-occurrence matrices, as shown in **equation 4.17**.

$$I_1 = I(x, y) \tag{4.15}$$

$$I_2 = I(x+\Delta x, y+\Delta y) \tag{4.16}$$

$$P(I_1=i, I_2=j) = P(i, j) \tag{4.17}$$

Because the two random variables I_1 and I_2 represent stationary random processes, they have equal means and variances. The covariance matrix C_I is defined in **equation 4.18**, with the joint expectation in **equation 4.19**. The eigenvalues λ_i and eigenvectors $\hat{u}_i$ of this matrix are given in **equation 4.20**.

$$C_I = \begin{bmatrix} \sigma^2 & \sigma^2 . \rho \\ \sigma^2 . \rho & \sigma^2 \end{bmatrix} \tag{4.18}$$

$$\sigma^2 \; . \; \rho = E\{(I_1-\mu) \; . \; (I_2-\mu)\} \tag{4.19}$$

$$C_I \; . \; \hat{u} = \lambda \; \hat{u} \tag{4.20}$$

Unser says that the eigenvectors of the covariance matrix are the axes of inertia of the joint probability density function (pdf), and that the eigenvalues are the variances along these axes.

He also says that the principal axes of any second order pdf can be defined by the sum and difference of its random variables. Two un-correlated random variables Z_1 and Z_2 are defined in equation 4.21 that correspond to the covariance matrix eigenvectors. The variances of these random variables are equal to the matrix eigenvalues, equation 4.22.

As the two random variables Z_1 and Z_2 are un-correlated, their covariance function is zero. When the random variables are also independent, the joint pdf can be computed from the expression in **equation 4.23**. This expression is always true for Gaussian random variables. Because the images used in this project are not true Gaussian, the last equality will therefore not be properly satisfied. Nevertheless, the product of the sum and difference first order pdfs can still be used as a close approximation of the joint pdf. This is achieved by multiplying the product of the first order pdfs by a normalisation constant to ensures that is unity when summed over all values.

$$\begin{aligned} Z_1 &= (I_1+I_2)/\sqrt{2} \\ Z_2 &= (I_1-I_2)/\sqrt{2} \end{aligned} \tag{4.21}$$

$$\begin{aligned} Var\{Z_1\} &= \lambda_1 = \sigma^2.(1+\rho) \\ Var\{Z_2\} &= \lambda_2 = \sigma^2.(1-\rho) \end{aligned} \tag{4.22}$$

$$P(I_1, I_2) = P(Z_1, Z_2) = P_s(Z_1) \,.\, P_d(Z_2) \tag{4.23}$$

The sum histogram $h_s(i)$ represents the number of occurrences of the random variable $Z_1=i$, where i is the sum of the grey levels of the spatially dependent pixels. The difference histogram $h_d(j)$ represents the number of occurrences of the random variable $Z_2=j$, where j is the difference between the grey levels of the spatially dependent pixels. Both first order probability density functions are then found by normalising the respective histograms by the total number of pixel pair interrogations.

(UNSER, 1986) devised six texture features based on (HARALICK *et al*, 1973)'s co-occurrence parameters from the first order probability density functions that measure the *Mean, Angular second moment, Correlation, Variance, Inverse difference moment, & Entropy.*

Computing sum and difference histograms instead of the conventional co-occurrence matrices has two advantages; the computation is much faster and the amount of memory required for

data storage is much less. The memory and time required for computation is reduced by a factor of N_g /4. The disadvantage with this method is that the approximation of the joint pdf made in equation 4.23 would introduce large errors with the typically non-Gaussian data used in this project. This method has not been used as an alternative to conventional co-occurrence because these potential inaccuracies by far outweigh its computational limitations.

4.8 SUMMARY

The methods described in this chapter derive texture features that can be used to classify tissue by its statistical properties. These methods examine both the first and second order image statistics.

The first order image statistics contain only information about the distribution of grey levels in the image. Although the statistical moments calculated from the image grey level distribution contain no spatial information, they can provide some basic texture descriptors. We have taken these parameters and appraised their efficacy as texture features as part of this project.

The second order image statistics contain spatial information that is useful in the description of image texture. We have discussed several methods that are able to extract such image information in this chapter. In many of these methods the first order image statistics are also extracted. Because first order image statistics refer to the grey level distribution of image pixels, such information may distort the texture features when images of different average illumination with similar texture are compared. It is often necessary therefore, that the first order statistics of image regions be equalised to enable a fair comparison of second order statistics. As part of this project, we have compared the effect of normalised and un-normalised first order statistics on second order texture features.

Because of the natural variation in the orientation of tissue texture in clinical MRI images, it is difficult to uniquely specify texture when the texture features are directionally sensitive. As most of the texture methods described in chapter satisfy this undesirable criteria in their original form, we have strived to generalise them so that they are not sensitive to pattern orientation.

The five second order statistical texture analysis methods described in this chapter are run length coding, cluster size, grey tone difference, co-occurrence, and sum and difference co-occurrence.

Run length coding is a method conventionally used for data compression, but can be used to extract first and second order image statistics with the five texture features suggested by (GALLOWAY, 1974). After initial experimentation with test images, we have examined the performance of these parameters in the classification of a large number of tissue texture samples from clinical MRI images.

The run length concept has been adapted for the novel 'cluster' length idea as part of this project work. We have defined the cluster length, or cluster size, as a maximally connected cluster of pixels of the same grey tone. The pixels can only be connected if they are vertically or horizontally adjacent. The motivation for the development of this method is the fact that the cluster is independent of direction; a considerable asset in the context of the variable orientation of tissue texture in clinical MRI images. GALLOWAY's five texture measures have been redefined in the context of this new method, and their performance as texture descriptors evaluated using test images and samples of tissue texture from clinical MR images.

The grey tone difference method is a standard texture technique that examines second order image statistics. This method, discussed by (WESZKA *et al*, 1976) and (CONNERS and HARLOW, 1980), examines the local differences in grey level intensity. Because the grey level differences are measured in a specific direction, this method is sensitive to the orientation of a texture. This is particularly of concern to us because the orientation of tissue texture in clinical MRI images is extremely variable. As part of this project, we have suggested that the grey tone difference histograms should be found for all directions and summed. The resulting texture measures from this grey level difference histogram will be independent of texture orientation. The efficacy of this method, both directional and non-directional, for the extraction of texture information and subsequent classification of image texture has been evaluated as part of this project.

The spatial grey level dependence or *co-occurrence* method, described by (HARALICK *et al*, 1973), is a technique that extracts second order statistical information from image texture. A 2D histogram is formed that measures the frequency of co-occurrent pixel grey levels. The

probability density function, or co-occurrence matrix, is defined by normalising the histogram by the total number of entries. The 14 texture features described by HARALICK *et al* can then be extracted. Like the grey tone difference method, the co-occurrence method uses a positional operator to identify the co-occurrent pixels, and hence this method is directionally sensitive. In this form, it may prove difficult to uniquely specify a texture sample that varies in orientation. To ensure that the texture features are invariant under any rotation of the texture pattern, (GOTLIEB and KREYSZIG, 1990) suggest that the co-occurrence matrices from all four planes; horizontal, vertical and both diagonals, be summed together. This general model of the co-occurrence has been implemented in this project to extract texture features from test images and a large selection of texture samples from clinical MRI images. The proficiency of these texture measures in the characterisation of tissue texture has then been examined under a variety of normalisation and quantisation schemes.

The main drawbacks with the co-occurrence method described in this chapter occur with the large computational time and memory resources required to obtain the co-occurrence matrices. The sum and difference co-occurrence method is an alternative introduced by (UNSER, 1986) that enables substantial savings in computational time and memory resources at the price of a small loss in accuracy. This method transforms the second probability density function for co-occurrent pixels into two un-correlated first order probability density functions. Although, considerable savings are made in computational resources, we have decided to ignore this innovation and stick with the original form of the co-occurrence method. Our main motivation for this decision comes from our interest in the descriptive and discriminative power of the texture methods. At this juncture we are not so interested in building efficient and fast commercial texture characterisation algorithms, but more in the appraisal of the potential efficacy of the texture features themselves for tissue characterisation. The question of computer resources therefore is not yet a consideration.

CHAPTER 5 - TEXTURE FEATURES FROM TRANSFORMS

5.1 INTRODUCTION

The methods described in this chapter derive textural information by applying existing image transforms. The three image transforms employed are the Fourier transform, the Walsh transform and the Slant transform. The use of such transformations can be seen as both a structural and a spectral approach to texture analysis.

The two dimensional Fourier transform decomposes an image region into its frequency components. The result of this transform is complex and can be expressed in terms of phase and magnitude. The magnitude function, or Fourier spectrum, indicates the relative abundance of each spectral component, while the phase contains edge information. The magnitude function is of particular interest to us because textural information can be inferred from the spectral composition of an image.

The power spectrum is formed from the Fourier transform of an image. This method has been used as a source of texture features with some success. Textural information can be extracted from the power spectrum by taking radial averages. A coarseness measure based on this idea was developed by (WESZKA *et al*, 1976). A coarseness measure was also extracted from the auto-correlation function by (KAISER, 1955). These methods, however, only provide a limited description of texture. As part of this project we have examined the power spectra of a selection of clinical MRI images and derived several textural features in addition to those suggested by KAISER and WESZKA *et al*.

The frequency spectrum can also be useful in the extraction of image textural information from an image (HARALICK, 1979). Although the most obvious method of obtaining spectral information is through Fourier decomposition, the same objective can be achieved using the Walsh transform.

The Walsh transform performs the decomposition of a function with a set of rectangular waveforms rather than the sine-cosine complex exponential waveforms associated with the Fourier transform. It was suggested by (LARSEN and LAI, 1980) that the Walsh transform can be used in its place in some applications because it extracts the same spectral information.

LARSEN and LAI substantiate this claim by showing the performance of a set of statistical features derived from the Fourier and Walsh spectra of sleep EEG signals to be equivalent. The motivation behind LARSEN and LAI's work is the gain in efficiency that can be achieved by replacing the Fourier transform with the Walsh transform in certain applications. The Walsh transform has a vast computational superiority over the Fourier transform. Because the result of the Walsh transform is not complex, unlike the Fourier transform, less computational space is required for its implementation. Very fast and efficient algorithms exits for both the Fourier transform (FFT), (COOLEY and TUKEY, 1965), and the Walsh transform (FWT), (LARSEN, 1976). The fast Walsh transform is about three times faster than an efficient radix-four fast Fourier transform. This is because the algorithms are similar in structure, but the FWT contains no multiplies.

The efficiency of the FWT algorithm has been a key factor in the popularity of the Walsh transform. Walsh functions have been found useful for coding, enhancement and other signal processing tasks. Many of these general techniques have been described by (ANDREWS and PRATT, 1969; ANDREWS and PRATT, 1970). Some interesting work using Walsh functions for pattern recognition and feature extraction has also been carried out; (MELTZER *et al*, 1967; ANDREWS, 1971; ALEXANDRIDIS and KLINGER, 1971b).

It is true to say that both the Fourier and Walsh transform perform spectral decomposition. It is also true to say that textural information can be extracted from the Fourier power spectrum. Based on these premises, we have suggested that it may be possible to extract image textural information by applying the Walsh transform. In an attempt to corroborate this statement, we have taken the set of textural features that we proposed in our investigation of the Fourier power spectrum, and re-defined them in the context of the Walsh transform. Their performance in this role was then examined.

Texture is defined by the distribution and arrangement of elemental shapes called primitives, or the patterns they make. Walsh functions have been used with some success in recognising basic shapes; MELTZER *et al* used the outline of leaves as the subject for his classification scheme. The leaves were identified from histograms of the squared weights at each resolution. ANDREWS digitised handwritten numerical characters and identified them directly from the Walsh transform. ALEXANDRIDIS and KLINGER used a property of the Walsh functions called *axis symmetry* to identify alphanumeric characters.

Because Walsh functions can be used to identify individual shapes, it may be possible to use them to differentiate texture by identifying collections of shapes, such as the groups of primitives that form texture. As part of this project we have examined the possibility of using Walsh functions to differentiate between samples of a variety of different tissue textures in clinical MRI images. However, there are several problems in this application that arise from the nature of clinical MRI images.

We have suggested that the square regions of texture from clinical MRI images introduced to the Walsh transform as input arrays should be selected by hand. We make this recommendation because we have experienced difficulty in finding homogenous regions of tissue texture in MRI images of the body. This is mainly due to the structural complexity of the body, ie. the countless boundaries and sub-boundaries of tissue types. All body tissue has a blood supply; arteries and capillaries bring oxygenated blood to the tissue, and capillaries and veins take waste products away from the tissue. Blood vessels also interrupt the continuity of tissue texture. These flaws in tissue homogeneity will dominate the low frequency and DC components of the transformed region, and distort the subsequent data manipulation and tissue type classification schemes. It is therefore important to select tissue regions with some care.

Another source of difficulty stems from the irregular alignment of the many different layers that make up tissue structures in the human body. Although it may be possible to locate a rectangular region of tissue that, at first glance, appears to contain uniform texture, the texture contained therein is susceptible to geometric transformations across the region due to the nature of the tissue and other associated structures. A good example of this is the concentric orientation of the texture of muscle and fat tissue around the leg. External Body fat will hang with gravity and distort its uniform texture.

For a tissue texture to be uniquely specified, the texture must have a specific orientation and be homogenous across the tissue region being investigated. It practise, these criteria are never fully met, consequently, it is difficult to use tables of reference values for each texture feature in their classification.

The Walsh transform was found not to be positionally invariant by (ALEXANDRIDIS and KLINGER, 1971b). The patterns they examined were not specified uniquely because they were not located at the same coordinates from image to image. ALEXANDRIDIS and

KLINGER overcame this problem by pre-processing the patterns with the Fourier transform and taking the magnitude function as the input to the Walsh function. The Fourier magnitude function is invariant under the translation of the initial pattern. Using this technique, with their orientation restricted to whole units of 90 degrees, ALEXANDRIDIS and KLINGER were able to uniquely specify their patterns.

Because of the subjectivity of the structure of texture, it is a certainty that any method of texture characterisation that is not invariant under translation will not be able to uniquely specify the texture. As part of this project we have used the Walsh transform with the Fourier transform pre-processing step, as suggested by ALEXANDRIDIS and KLINGER, and applied it to a sample of clinical MRI images that contains a variety of tissue textures. Taking care to select wholly uniform texture regions, we have examined the performance of several methods of extracting textural information from this result.

Walsh functions can only hold the values plus and minus one. In the accurate description of natural shapes and patterns using Walsh functions this two level resolution can be limiting. Although Walsh functions provided a good set of basis functions for their earlier work in the specification of biological form, (MELTZER *et al*, 1967) found that they were restricted by the highly discrete form of these functions. MELTZER *et al* were attempting to accurately and uniquely specify the shape of a set of leaves. A more natural set of basis functions, such as a set of leaf shaped functions, would be more useful in this application. This observation is true in many applications that require detail in the specification of natural or irregular form. It is however, by definition, very difficult to devise an orthogonal set of basis functions based on a set of irregular or natural shapes.

A more continuous set of basis functions than those of the Walsh transform are given with the Slant transform. The Slant functions are multi-level functions; the number of levels corresponds to the size of the transform matrix. Although the Slant transform matrix construction is a little more complex, it is implemented in the same way as the Walsh transform. As part of the work carried out in this project, we have replaced Walsh functions with Slant functions and observed the difference in performance between the two sets of texture features.

We discussed earlier that because of the natural variation in the orientation of texture patterns

in tissue in clinical MRI images, it is difficult to uniquely classify Walsh texture features. As part of this project, we have suggested a totally original method of implementing the Walsh and Slant transforms in a texture classification scheme that is insensitive to texture orientation. This method has been developed and tested as part of this project, and will allow unique texture classification from the Walsh or Slant features, regardless of the orientation of the tissue texture being characterised.

5.2 FOURIER POWER SPECTRUM

In the spatial domain coarse textures signify slow and gradual intensity variations across an image. In the same way fine textures signify rapid intensity variations across the image. In the Fourier domain these effects correspond to abundances of low and high frequency components respectively. Because the Fourier transform breaks an image down into its frequency components, assumptions can be made about the nature of the texture by analysing the Fourier spectrum.

The Fourier transform of a real function, such as that of the image data, is generally complex; that is made up of real and imaginary components. It is often convenient to express this complex function in terms of magnitude and phase. The square of the Fourier spectrum or magnitude function, commonly referred to as the power spectrum or the spectral density, is used as a source of texture features. The significance of coarse and fine textures as regard to the distribution of frequency components in the power spectrum are the same as for the Fourier transform.

Consider the discrete data image $f(x,y)$, with (x,y) being the image coordinates. The discrete Fourier transform of this data image is $F(u,v)$, with (u,v) being the coordinates in 2D frequency space. The discrete two dimensional Fourier transform is defined in **equation 5.1** from (GONZALEZ and WINTZ, 1987).

$$F(u,v) = \frac{1}{MN} \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x,y) \exp\left[-j2\pi \left(\frac{ux}{M} + \frac{vy}{N}\right)\right] \qquad (5.1)$$

where $u = 0,1,2, \ldots, M-1$ and $v = 0,1,2, \ldots, N-1$

The Fourier power spectrum is defined by the expression $F(u,v)F(u,v)^*$, where the * denotes

complex conjugation. A simple example of the use of the power spectrum in extracting texture information from image data is provided by (WESZKA *et al*, 1976) who suggested that the texture coarseness can be measured by taking averages of the power spectrum over ring-shaped regions centred at the origin. **Equation 5.2** shows this coarseness measure with ($0 \leq u,v \leq n$-1) and r_1 and r_2 respectively being the inner and outer radii of the ring. This radial average function is often used to estimate the one dimensional Fourier spectrum or power spectrum of a two dimensional function such as an image.

$$\phi_{r_1 r_2} \equiv \sum_{r_1^2 \leq u^2+v^2 \leq r_2^2} |F(u,v)|^2 \qquad (5.2)$$

Texture can be described in terms of the spatial size of the tonal primitives in an image (HARALICK, 1979). There is a relationship between coarse textures and large tonal primitives, and between finer textures and tonal primitives of smaller size. The tonal primitive in this model is the grey tone, or grey level.

The auto-correlation function can be used to analyze these primitives. HARALICK states that if the tonal primitives are relatively large then the auto-correlation will drop off slowly with distance. If the tonal primitives are small then the auto-correlation will drop off quickly with distance. This can be explained in terms of the primitive size. A large tonal primitive will be correlated across a larger area of the image than a smaller tonal primitive. Therefore, from the definition of the auto-correlation function, the function drops off more slowly.

The periodicity of the auto-correlation function also affects the drop off of the function. In a totally un-correlated image there are no periodicities. The auto-correlation function of such an image would comprise of a very narrow spike at the origin. As periodicities are introduced into the image, such as the effect of repeating the same tonal primitives across the image, the auto-correlation function widens and drops off more slowly.

Thus two aspects effect the shape of the auto-correlation function; the size of the tonal primitives and the periodicities in the image. The size is inversely related, and the periodicity is directly related to the drop off.

The auto-correlation function $p(x,y)$ is defined by the normalised inverse Fourier transform

of the power spectrum, **equation 5.3**. The normalisation factors are defined in **equations 5.4** and **5.5**, where $f(x,y)$ denotes the image pixel values at the coordinates (x,y). The image is bounded by $(0 \leq x \leq L_x)$ and $(0 \leq y \leq L_y)$, where $L_x = L_y = N\text{-}1$ and N is the image dimension.

$$p(x,y) = \frac{1}{f_{avg} \cdot N_{x,y}} \mathcal{F}^{-1} \{|F(u,v)|^2\} \tag{5.3}$$

$$f_{avg} = \frac{1}{(N-1)^2} \sum_{x,y}^{N-1} f^2(x,y) \tag{5.4}$$

$$N_{x,y} = (N-1)^2 - (N-1)(|x|+|y|) + |x||y| \tag{5.5}$$

In practise, the implementation of the discrete Fourier transform on digital computers proves to be extremely expensive in terms of the computer's resources, ie. memory and processor time. The process requires something proportional to M^2 operations in order to compute the finite Fourier transform of a series of M (complex) data points. The number of data points present in an image of dimension N is $M=N^2$.

An algorithm for the computation of Fourier coefficients which requires much less computational effort was reported by (COOLEY and TUKEY, 1965). This method is widely known as the *fast Fourier transform*. The fast Fourier transform (FFT) is a method for efficiently computing the discrete Fourier transform. The economy of this method yields a reduction in the computational requirements from M^2 to $M \log_2 M$ operations.

The long and interesting history of the fast Fourier transform has been described in a paper by (COOLEY *et al*, 1967). In this paper, the contributions of many investigators are described and placed in historical perspective. The complete derivation of this algorithm, with examples to demonstrate the concepts involved, is presented in a paper by (COCHRAN *et al*, 1967).

For practical purposes, the fast Fourier transform algorithm incorporated in some of the texture techniques presented in this project has been taken from a library of standard computational tools, (PRESS *et al*, 1991).

An experiment using the auto-correlation function as the basis for a simple texture descriptor was carried out by (KAIZER, 1955) and further discussed by (HARALICK, 1979). KAIZER

made the assumption that the two dimensional auto-correlation function was circularly symmetric. KAIZER's texture descriptor was based on computing the auto-correlation function and finding the radial distance d, such that $p(d) = 1/e$. The experiment KAIZER carried out compared the ranking of seven images by human eye, on a scale from fine detail to coarse detail, with his measure $p(d)$. There was a 99 percent correlation between the results from the human subjects and his auto-correlation measure. From these results KAIZER established that his subjects were measuring the same textural features as his function.

Upon further investigation of the 1 percent error in classification, he found that a relatively flat background (low frequency) can be interpreted as either a fine or coarse textured area. A fundamental of texture is that it cannot be analyzed without a reference frame being stated or implied. For any smooth grey tone surface there exists a scale such that when the surface is examined, it has no texture. As the resolution is increased, it takes on a fine texture, and then a coarse texture.

A simple analogy of this phenomena is to take the view of a corn field from an aeroplane. At a great height, the field will have no detail and no texture. As the plane comes closer to the ground more detail can be seen and the field takes on a fine texture. In close proximity we begin to reach the limit of resolution and the texture begins to become coarse.

The texture models of (WESZKA *et al*, 1976) and (KAIZER, 1955) both measure the degree of coarseness of a texture. Both methods classify textures on a scale from coarse to fine. These texture descriptors are however one dimensional, and the scope of their usage in more general texture classification is limited. It is clear from this work, however, that the power spectrum contains some valuable texture information.

Some time has been dedicated as part of this project to examining the power spectra of a selection of image samples that contain a variety of image texture. The texture samples we have used are from clinical MRI images. In the first instance we have proposed a set of potentially useful texture features that can be extracted from the power spectrum. By comparing the separation of feature space for different texture regions, the proficiency of these feature as texture descriptors can be assessed.

The proposed texture features are the spectrum energy, the cutoff frequency, the exponent of

frequency which relates each power spectrum component to the frequency value, and the spectra first order statistical moments. The spectrum energy is found from the sum of the power components. The exponent of frequency is a constant that describes the power spectrum. This value is the power to which the frequency is raised for each power component. This constant can be found from the gradient of the *log-log* plot of the spectrum against frequency. The four moments taken from the first order statistics of the power spectrum are the first four, namely; the mean, the variance, the skewness, and the kurtosis.

It may prove to be that these texture features are not useful in texture discrimination when used alone as classification criteria. In combination with other measures, however, a precise classification scheme can be constructed.

5.3 HADAMARD AND WALSH TRANSFORMS

A large class of image processing operations are linear; an object image field is formed from linear combinations of pixels of an input image field. The generalised linear operator is defined by (PRATT, 1978) in **equation 5.6** where the N_1 by N_2 element input image array is described as $F(n_1,n_2)$ and the M_1 by M_2 element output image array by $P(m_1,m_2)$. The operator kernel $O(n_1,n_2;m_1,m_2)$ defines the weighting constants for the image elements. In general this operator is a function of both input and output image coordinates.

$$P(m_1,m_2) = \sum_{n_1=1}^{N_1} \sum_{n_2=1}^{N_2} F(n_1,n_2)\, O(n_1,n_2;m_1,m_2) \qquad (5.6)$$

The Hadamard transform is a unitary transform. The unitary transformation is a specific type of linear transformation in which the basic linear operation of **equation 5.6** is exactly invertible and the operator kernel satisfies certain orthogonality conditions. The mathematical requirements for the kernels to be orthonormal are also given by PRATT.

A good description of the Hadamard transformation is provided by (PRATT, 1978; CLARKE, 1985; GONZALEZ and WINTZ, 1987; SCHALKOFF, 1989). The Hadamard transform is based upon the Hadamard matrix which is a square array of plus and minus ones, whose rows and columns are orthogonal. A normalised Hadamard matrix, $\boldsymbol{H}$, satisfies the condition in **equation 5.7**.

$$HH^T = I \tag{5.7}$$

The smallest size orthonormal Hadamard matrix is a *2* by *2* matrix, **equation 5.8**. The construction of Hadamard matrices of higher order is achieved using a matrix relating the desired matrix order *N* to the matrix order *N/2* as shown in **equation 5.9**.

$$H_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{5.8}$$

$$H_N = \frac{1}{\sqrt{2}} \begin{bmatrix} H_{N/2} & H_{N/2} \\ H_{N/2} & -H_{N/2} \end{bmatrix} \tag{5.9}$$

The Hadamard transform is performed by employing the Hadamard matrix, ***H***, in the form shown in **equation 5.10**, where [*f*] is the image function matrix, and [*F*] is the resulting transformed image matrix.

$$[F] = H[f]H \tag{5.10}$$

An alternative to performing the Hadamard transformation through matrix algebra is to express the Hadamard matrix in the form of an operator kernel, as represented by $O(n_1,n_2;m_1,m_2)$ in the generalised linear operator equation, **equation 5.6.** The two dimensional operator kernel in **equation 5.11** is defined by (GONZALEZ and WINTZ, 1987), where the summation in the exponent is carried out in modula 2 arithmetic.

$$O(n_1,n_2;m_1,m_2) = \frac{1}{N}(-1)^{\sum_{i=0}^{n-1} [b_i(n_1)b_i(m_1)+b_i(n_2)b_i(m_2)]} \tag{5.11}$$

where $b_k(z)$ is the *k*th bit in the binary representation of *z*, and *n* comes from $N = 2^n$. For example, if $n=3$ and $z=6$ (110 in binary), we have $b_0(z)=0$, $b_1(z)=1$, and $b_2(z)=1$.

The image function matrix [*f*] can be written in terms of *basis function matrices*. A function can be decomposed to a set of weighted basis functions. This concept is described in detail by (SCHALKOFF, 1989). The basis function matrices have the same dimensions as the image

function matrix. The relationship between the basis functions and the image function is given in **equation 5.12** where E_{mn} is the basis function and f_{mn} its weighting factor. The number of basis functions used to represent an N by N image is N^2.

$$[f] = \sum_{m=1}^{N} \sum_{n=1}^{N} f_{nm} E_{nm} \tag{5.12}$$

The Fourier transform of an image function represents a translation from two dimensional spatial or temporal space to two dimensional frequency space. The result of the transform is a frequency image matrix. This image is complex and contains the phase and magnitude information associated with the frequency components present in the original image function. Clues to the frequency components present in the original image can be seen from the distribution of grey level intensity throughout the image function. Large grey level intensity gradients within the image correspond to high frequency information, while slow or gradual changes in intensity across the image represent low frequency information.

The Fourier transformation can be thought of as a process that identifies the weighting factors, f_{mn}, of the fundamental image matrices or basis functions, E_{mn}, of the image. The resulting matrix contains the set of basis function weights. The image can be reconstructed from the superposition of the basis functions weighted by these coefficients.

The Fourier basis functions are a set of complex exponentials that may be decomposed into Sine and Cosine components. Each Fourier basis function represents a specific frequency component. The actual frequency depends on the sampling rate in spatial/temporal space. Because the basis functions are arranged in order of increasing frequency, the scale in the Fourier domain represents an increase in frequency from the origin..

Like the Fourier transform, the Hadamard transform has basis functions. The Hadamard basis functions are a set of rectangular functions with amplitude of plus and minus one. The conventional definition of frequency does not apply because the zero-crossings of the functions are not equally spaced along the axis of the independent variable, eg. time. A generalisation of the concept of frequency called *sequency* has been designed by (HARMUTH, 1968) to deal with this situation. This is taken to be the number of sign changes along a column of the Hadamard matrix, or the number of zero crossings in a basis vector.

The one dimensional Hadamard basis functions for a function of size N=8 data points are shown with their corresponding sequency from top to bottom of **figure 5.1**. These functions can only hold two values: plus one and minus one. The sequency is taken from the number of zero crossings between these values. It can be seen from the figure that the eight basis functions are not in any order of sequency. The first basis function has a sequency of 0. This function would contain only a DC frequency component in Fourier space. The second basis function has the highest sequency. The shape of this function indicates that it would correspond to a high frequency component in Fourier space.

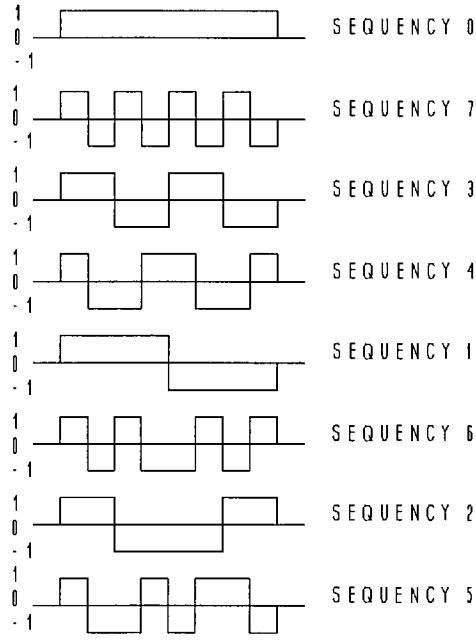


Figure 5.1 : 1D Hadamard basis functions

When the one dimensional Hadamard transformation is applied to a function a set of coefficients are produced. These values represent the contribution that each of the 8 basis functions makes to the original function. This function can be reconstructed by summing each of the basis functions weighted by the corresponding coefficient value.

A more useful form of the Hadamard transformation is called the Walsh transform, or the Walsh-Hadamard transform as it is known in some image processing literature. The Walsh transform is formed in the same way as the Hadamard transform, but with the rows and columns arranged such that the basis functions are in order of increasing sequency. This

formulation is analogous to the Fourier transform, where frequency also increases as a function of increasing reference coordinates. The results of applying the Walsh transformation can therefore be interpreted in a similar way to those of the Fourier transform. This analogy is discussed by (PRATT, 1978; CLARKE, 1985; GONZALEZ and WINTZ, 1987; SCHALKOFF, 1989). The Walsh matrix merely performs the decomposition of a function by a set of rectangular waveforms rather than the sine-cosine waveforms associated with the Fourier transform. GONZALEZ and WINTZ point out that because the Walsh transform is real, it requires less computer storage for a given problem than for the Fourier transform which is generally complex.

These properties of the Walsh transform provided the motivation for (LARSEN and LAI, 1980) to use the discrete orthogonal set of Walsh functions as a basis of processing EEG signals. Features selected from the Walsh spectrum of sleep EEG data are compared to corresponding Fourier features. LARSEN and LAI show that the Walsh spectral features classify the data in much the same way as the Fourier spectral features. This is considered as sufficient justification for usage of Walsh spectral features in place of Fourier spectral features, enabling one to take advantage of the vast computational superiority of the fast Walsh transform over the fast Fourier transform. Details of the fast Walsh transform algorithm are provided by (GONZALEZ and WINTZ, 1987).

Like the Hadamard, the Walsh transformation can be applied using the generalised linear operator formula in **equation 5.6**. The operator kernel is very similar to that of the Hadamard transform, **equation 5.11**. The two dimensional operator kernel for the Walsh transform is defined in **equation 5.13** from GONZALEZ and WINTZ, where $b_k(z)$ is the kth bit in the binary representation of z, and n comes from the expression $N = 2^n$. The function $p_k(z)$ is defined in **equation (5.14)**.

$$O(n_1, n_2; m_1, m_2) = \frac{1}{N}(-1)^{\sum_{i=0}^{n-1} [b_i(n_1)p_i(m_1) + b_i(n_2)p_i(m_2)]} \tag{5.13}$$

$$\begin{aligned} p_0(z) &= b_{k-1}(z) \\ p_1(z) &= b_{k-1}(z) + b_{k-2}(z) \\ p_2(z) &= b_{k-2}(z) + b_{k-3}(z) \\ &\vdots \\ p_{k-1}(z) &= b_1(z) + b_0(z) \end{aligned} \tag{5.14}$$

Because the operator kernel is dependent only on the coordinates of the input and output image arrays, the kernel can be calculated in advance of the transformation process. The operator formula in **equation 5.6** therefore becomes simplified to the product of two matrices.

The one dimensional Walsh basis functions are shown, for N=8 data points, starting from the top, in **figure 5.2**. Note that the basis functions are in order of increasing sequency.

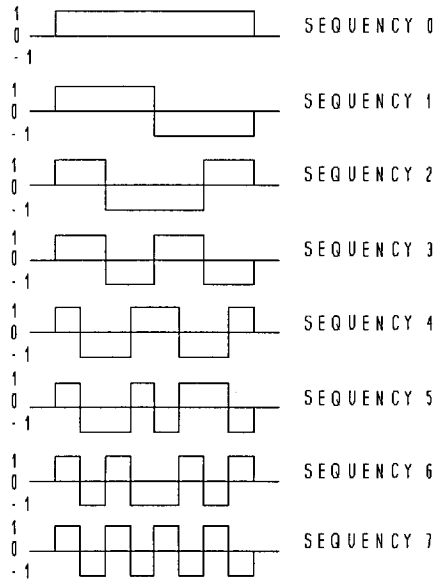


Figure 5.2 : 1D Walsh basis functions

Because of its recursive structural properties, the Walsh matrix can only be a square matrix of dimension N, where $N=2^n$ and n is an integer such that N = 2, 4, 8, 16, 32 etc. Because the input image array must be the same size as the operator matrix, these dimension restrictions also apply to the input image array. The MRI images used for analysis in this project originate in two sizes; 256 by 256 pixels and 512 by 512 pixels. The *MAIVIS* image visualisation and analysis software platform facilitates the selection of smaller regions of interest from library images, **chapter 3 section 3.5.8**. The software can be used to select a region of a predefined size and shape from anywhere within an image. Selected image regions can be used as input arrays for the Walsh transformation.

The two dimensional Walsh basis functions for a matrix size N=4 are shown in **figure 5.3**. Like the one dimensional basis functions, these functions can only hold two values; plus one

and minus one. These values are represented respectively by the black and white areas in the sixteen functions. The zero sequency, or DC, basis function is in the top left corner of the figure; the black square. The one dimensional function sequency patterns can be seen in the top row and left column functions. The sequency of the other functions is defined by a combination of the vertical and horizontal sequency components. The two dimensional transformation yields a result image that has the same dimensions as the input array and the operator matrix. The input image array can be reconstructed by summing each basis function weighted by the value found in the corresponding row or column of this result image array.

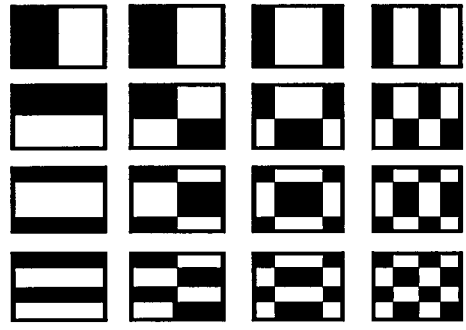


Figure 5.3 : 2D Walsh basis functions

Two dimensional Walsh functions have been used by (MELTZER *et al*, 1967; MELTZER and SEARLE, 1968b) in the numerical specification of biological form. The development of an organ or organism is characterised by a succession of shapes each of which may involve distinctive metabolic states or morphological patterns. MELTZER *et al* used the shape of leaves to demonstrate his technique. Problems associated with the rigidity of the Walsh functions were also examined by (MELTZER, 1968a). The Walsh derived functions are most suitable as an initial set of basis functions for the specification of biological form because of their simplicity. Their characteristic checkerboard patterns, however, are not the most natural to use for the analysis of familiar kinds of biological objects. Meltzer suggests that it might be possible to construct a different orthonormal set of basis patterns which would be more 'leaf-like' in shape. A more natural set exists in the form of the Slant transform, **section 5.4**. The Slant transform basis functions are continuous functions with N levels of luminescence, where N is the dimension of the transform matrix, in contrast to the discrete two level Walsh basis functions.

In the search for useful image characterisation parameters that can be derived from the Walsh transform, (ALEXANDRIDIS, 1971a) introduced a property of Walsh functions called *axis symmetry*. Consider a 1D Walsh basis function written in the form, $x = (x_N x_{N-1} \ldots x_1 x_0)$, where x_i are the N data entries. The data values at these points are either plus or minus one. A basis function with $N=2^n$ data points has n axes of dyadic translation. **Figure 5.4** shows the three axes a_2, a_1 and a_0 defined for the basis functions with $N=2^2$ and $N=2^3$ entries. A function possesses axis symmetry when it is invariant under dyadic translation about any of these defined axes.

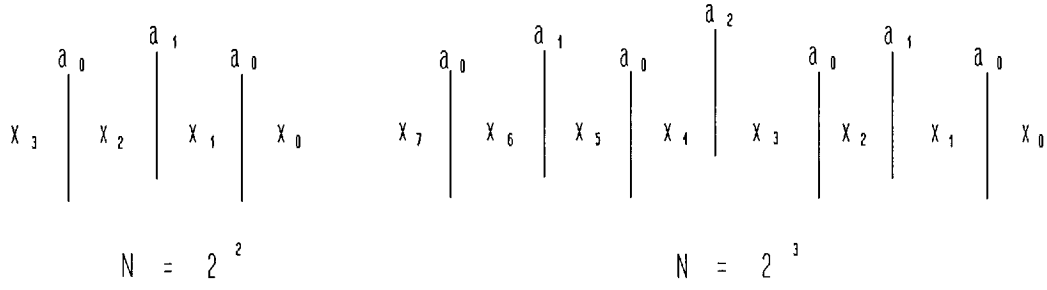


Figure 5.4 : Axes of dyadic translations

$$A^{(n)} = a_{n-1} a_{n-2} \ldots a_1 a_0 \tag{5.15}$$

The axis symmetry of a Walsh function is given by the binary expression in **equation 5.15**, where the axis parameter, a_i, has the value of 1 when the function is invariant under translation about that axis. The axis parameters are zero in all other cases. The axis symmetry for two dimensional Walsh basis functions are calculated from the Boolean AND operation between the one dimensional axis symmetry values of the corresponding row and column.

The use of Walsh functions in the design of a feature extraction algorithm was described by (ALEXANDRIDIS and KLINGER, 1971b). The *axis-symmetry* property of the Walsh functions is used to decompose the geometric patterns. A histogram is then formed from the axis symmetry information associated with a pattern. These histograms contain implicit information about symmetries, periodicities and discontinuities present in a figure. Because the Walsh transformation is not positionally invariant, the sequency spectrum does not specify the pattern uniquely. ALEXANDRIDIS and KLINGER overcome this problem by performing a pre-processing normalisation on the input pattern through Fourier transformation. Although

complex FFT is not invariant to horizontal and/or vertical translations of the pattern, its magnitude is invariant to such translations. This method results in great reduction in dimensionality in feature space, which leads to a computationally simpler classification task. ALEXANDRIDIS and KLINGER used this technique successfully to identify alphabetic characters as long as the patterns differed only in translation or a 90 degree rotation.

This project deals with the extraction of textural information from clinical MRI images. Because of the structural complexity of the body, it is very difficult to find homogenous regions of tissue texture. Either a region of texture is punctuated with other structures, or the tissue itself changes because of structural constraints. For the texture of a tissue region to be uniquely specified, the texture must have a specific orientation and be homogenous across the region. We have applied the Walsh transform to a selection of tissue regions with a variety of different texture types.

It has been shown by (HARALICK, 1979) that textural information can be extracted from the frequency spectrum of a region of texture in an image. The Walsh transform performs the decomposition with a set of rectangular waveforms rather than the sine-cosine complex exponential waveforms associated with the Fourier transform. Because the Walsh transform can be used as an alternative to the Fourier transform for such spectral decomposition, (LARSEN and LAI, 1980) and (GONZALEZ and WINTZ, 1987), it is reasonable to assume that the Walsh transform can be used to extract textural information from image functions. As part of this project, we have investigated the performance of the texture features originally suggested for the Fourier spectrum in **section 5.2** in the context of the Walsh transform, namely; the spectrum energy, the cut-off sequency, the exponent of sequency, and the spectrum first order statistical moments. The Walsh spectrum is calculated as the square of the Walsh transformed matrix, cf. Fourier power spectrum. This method has been applied with and without the Fourier transform pre-processing stage suggested by (ALEXANDRIDIS and KLINGER, 1971b), and a comparison made in the performance of the two sets of texture features. The Fourier transform pre-processing step renders the process invariant under translation.

We have also taken the result of the Walsh transform and calculated the axis symmetry and subsequently the axis symmetry histograms for a large number of tissue regions. Again, a comparison of the performance of this method has been made, with and without the Fourier

transform pre-processing.

The Walsh transform in its present form, however, is not the most ideal source of textural information. There are several serious shortfalls that make it difficult for this method to be used to uniquely specify texture.

Because texture is defined by the relationship between its primitives or fundamental shapes, any method that is sensitive to translation, such as the Walsh transform, is not able to uniquely specify a single texture that is arbitrarily sampled from region of tissue. This problem can be overcome by taking the Fourier magnitude of the tissue texture region as an input to the Walsh transform, because the Fourier magnitude function is invariant under translation (ALEXANDRIDIS and KLINGER, 1971b).

The problems encountered in finding homogenous regions of tissue texture can only be overcome by careful selection of the regions.

It has been suggested by (MELTZER *et al*, 1967) that the Walsh basis functions are too rigid for the specification of natural form such as the image data we are dealing with. A more natural set of basis functions are given by the Slant transform, described in **section 5.4**.

The variation in orientation of a texture within seemingly homogenous tissue regions, and from region to region, make it very difficult to uniquely specify tissue texture. Careful selection of tissue regions is the only approach open to us to minimise the subjectivity with this form of the Walsh transform. As part of this project, we have developed an alternative form of the Walsh and Slant transform that proposes to be invariant under rotation, and hence will permit unique specification of tissue texture. This method is described in **section 5.5**.

5.4 SLANT TRANSFORM

In **section 5.3** we define the Walsh transform and discuss how it may be employed for the purpose of the extraction of textural information from clinical MRI images. The motivation for this application of the Walsh transform comes from the work of (LARSEN and LAI, 1980; MELTZER *et al*, 1967; ALEXANDRIDIS and KLINGER, 1971b). LARSEN and LAI show that spectral features selected from the Walsh transform are analogous to those from the

Fourier transform. The consequence of LARSEN and LAI's work is that textural features inferred from spectral information can be derived from the Walsh transform. MELTZER *et al* found Walsh functions to be wholly suitable in the numerical specification of biological form; he used a set of 'leaf-like' shapes as the subject of his experiments. Because texture is defined by the spatial relationship between small patterns called primitives that repeat across the image, we suggest that Walsh functions may be used to characterise tissue texture. MELTZER *et al*, however, found the highly discrete form of the Walsh basis functions to be too approximate for the analysis of such variable form as biological structure, and suggested that a more natural form of the Walsh transform would provide a more precise specification of biological form. In the characterisation of texture from clinical MRI images, such an alternative to the Walsh transform would also prove to be an asset.

The Slant transform is similar both in structure and properties to the Walsh transform, but possesses a more continuous, ie. less discrete, set of basis functions. The Slant basis functions have N levels of luminescence, where the Walsh basis functions have only two. The number of levels, N, is derived from the dimension of the Slant matrix. In the decomposition of image texture, the Slant transform potentially offers a greater degree of specification than the Walsh transform. This section describes the Slant transform as an alternative to the Walsh transform for the purpose of characterising tissue texture.

The Slant transform belongs to a class of image processing operations that are linear. The Slant transform is another realisation of the general linear operator transformation that is given in **equation 5.6**. The Slant operator kernel, like the Walsh kernel, is invertible and orthonormal. The Slant matrix satisfies the orthogonality condition $SS^T=I$, where S is the Slant matrix and I is the identity matrix. The definition of the Slant transform is described in this section from (PRATT, 1978) and (SCHALKOFF, 1989).

For a vector of length $N=2$, the Slant transform matrix S_N is identical to the Hadamard transform matrix of order two. This transform matrix is given in **equation 5.16**.

$$S_2 = H_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{5.16}$$

$$S_N = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 & \vdots & & \vdots & 1 & 0 & \vdots & \\ & & \vdots & [0] & \vdots & & & \vdots & [0] \\ a_N & b_N & \vdots & & \vdots & -a_N & b_N & \vdots & \\ \cdots & \cdots & \cdots & \vdots & \cdots & \vdots & \cdots & \cdots & \vdots & \cdots \\ & [0] & & \vdots & I_{(N/2)-2} & \vdots & & [0] & \vdots & I_{(N/2)-2} \\ \cdots & \cdots & \cdots & \vdots & \cdots & \vdots & \cdots & \cdots & \vdots & \cdots \\ 0 & 1 & \vdots & & \vdots & 0 & -1 & \vdots & \\ & & \vdots & [0] & \vdots & & & \vdots & [0] \\ -b_N & a_N & \vdots & & \vdots & b_N & a_N & \vdots & \\ \cdots & \cdots & \cdots & \vdots & \cdots & \vdots & \cdots & \cdots & \vdots & \cdots \\ & [0] & & \vdots & I_{(N/2)-2} & \vdots & & [0] & \vdots & -I_{(N/2)-2} \end{bmatrix} \begin{bmatrix} S_{N/2} & & [0] \\ & & \\ & & \\ \cdots & \cdots & \cdots \\ [0] & & S_{N/2} \\ & & \end{bmatrix} \tag{5.17}$$

The transform matrices of higher order are defined recursively. The Slant matrix of order N is given in terms of the Slant matrix of order $N/2$ by the recursive relation in **equation 5.17**. In this equation $S_{N/2}$ denotes the Slant transform of order $N/2$, and I_N represents the N by N identity matrix. The scalar values a_N and b_N are determined for $N \geq 2$ by the **equations 5.18** and **5.19**.

$$a_{2N} = \left[\frac{3N^2}{4N^2-1} \right]^{\frac{1}{2}} \tag{5.18}$$

$$b_{2N} = \left[\frac{N^2-1}{4N^2-1} \right]^{\frac{1}{2}} \tag{5.19}$$

The 2D Slant transform may be implemented with the Slant matrix via the separable transform given in **equation 5.20**. Because the Slant matrix is not dependent on the data being transformed, the transform is separable. This property, also seen in the Walsh transform, allows the matrix to be calculated in advance of any data processing. Note the similarity between the 2D Slant transform shown in **equation 5.20** and the 2D Walsh transform, with the Walsh matrix H, in **equation 5.10**. Because of the nature of the matrix algebra required to perform the transformation, the square Slant matrix and the square image function [f] must have the same dimensions. It also follows that the dimensions of the result image function [F] will also be the same.

$$[F] = S[f]S^T \tag{5.20}$$

The one dimensional Slant basis functions for a function of N=8 data points are shown in order from top to bottom of **figure 5.5**. Note from the number of zero crossings that the basis functions are in order of increasing sequency. This property is consistent with that defined for the Walsh transform in **section 5.3**. Each function varies in the horizontal direction with constant slope with N discrete levels, where N is given by the dimension of the Slant transform matrix.

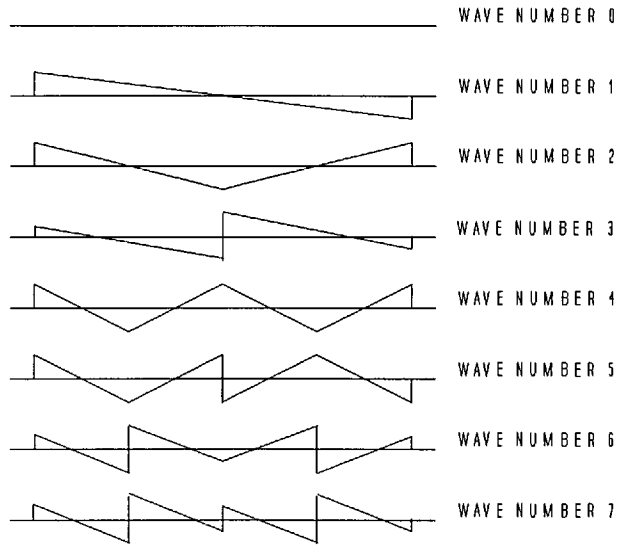


Figure 5.5 : 1D Slant basis functions

When the one dimensional Slant transform is applied to a function a set of coefficients are produced. These values represent the contribution that each of the basis functions makes to the original function. The input function may then be reconstructed by summing each of the basis functions weighted by the corresponding coefficient value.

The two dimensional Slant basis functions are defined by a vertical and horizontal combination of one dimensional basis functions. Because the one dimensional functions vary horizontally with a constant slope, the two dimensional functions have ramp like characteristics with the

same number of discrete levels as their one dimensional counterparts. The two dimensional basis functions are best visualised using a three dimensional representation such as that employed in **figure 5.6**. This figure shows a selection of two dimensional Slant transform basis functions for $N=8$; (a) S_{12} (b) S_{22} (c) S_{88} (d) S_{81}. The sequency of these functions can be deduced from one dimensional vertical and horizontal sequency components. The two dimensional transformation yields a result image with the same dimensions as the input array and the operator matrix containing basis function weighting. The input image can be reconstructed by summing each basis function weighted by the value found in the corresponding row or column of the result image.

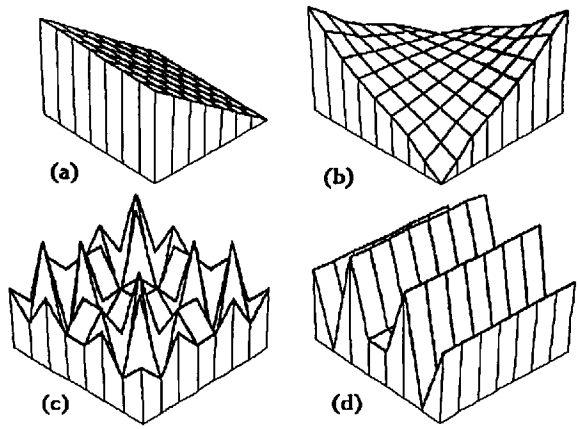


Figure 5.6 : 2D Slant basis functions

Like the Walsh transform, the Slant transform has computational advantages over the Fourier transform. A fast computational algorithm implementation exists for the Slant transform (FST), (PRATT *et al*, 1974). This algorithm also requires less computational space than the Fourier transform because, unlike the Fourier transform, the Slant transform is not complex. The fast computational algorithm implementation and high energy compaction properties of the Slant transform, like the Walsh transform, endear it to applications such as image compression, (PRATT, 1978). It has been shown by (LARSEN and LAI, 1980) that the Walsh transform can be used to extract the same spectral information as the Fourier transform, with the added bonus of computation superiority. With the subject of LARSEN and LAI's work as a key motivation, **section 5.3** describes how textural features can be inferred from the spectral information derived from the Walsh transform. Because the Slant transform is similar in both structure and function to the Walsh transform, we suggest that it may be a source of textural information in the same way.

As part of this project, we have taken the set of textural features that we proposed for the Fourier spectrum in **section 5.2**, and re-defined them in the context of the Slant transform. The parameters we have chosen are the spectrum energy, the cut-off sequency, the exponent of sequency, and the spectrum first order statistical moments. The Slant spectrum is calculated as the square of the Slant transformed matrix, cf. Fourier power spectrum.

Like the Walsh transform, the Slant transform is not positionally invariant. The Slant sequency spectrum will not be able to specify texture uniquely because the relative positions of the texture primitives from one sample of texture to another are likely to be different. The effects of pattern translation between samples of the same texture can be eliminated by performing the Fourier transform pre-processing stage suggested by (ALEXANDRIDIS and KLINGER, 1971b), and used in **section 5.3** for the Walsh transform. The magnitude function of the Fourier transform of a texture region is taken as the input to the Slant transform. This enables the process to be invariant under texture translation.

The Slant sequency spectrum has been analysed with and without the Fourier transform pre-processing stage, and a comparison made in tissue type characterisation performance of the two sets of suggested texture features from a variety of tissue regions.

A useful image characterisation parameter derived from the Walsh transform is the *axis symmetry*. This figure, introduced by (ALEXANDRIDIS, 1971a), is a measure of the symmetries within the basis functions. Although the Slant basis functions are in order of increasing sequency, it is not possible to associate the axis symmetry property with the result of the transform because of the continuous nature of the basis functions.

Because of the nature of the complex structures that make up the body, it is very difficult to identify uniquely homogeneous regions of tissue texture within clinical MRI images. Careful selection of the subject tissue regions by hand is the only approach open to us to minimise this source of error. In many cases subtle changes in orientation deny the texture from being uniquely specified. As part of this project, we have developed an alternative form of the Walsh and Slant transform that proposes to be invariant under rotation, and hence will favour the unique specification of tissue texture. This method is described in **section 5.5**.

5.5 TEXTURE ANALYSIS USING NON-DIRECTIONAL FORMS OF WALSH AND SLANT TRANSFORMS

The orientation of texture patterns in the tissue in clinical MRI images varies grossly as a result of the complex structure of the human body. The techniques used in **section 5.3** and **5.4**, the Walsh and Slant transformations, decompose an image by its geometrical properties. It is therefore difficult to use these methods to uniquely classify clinical MRI texture features. As part of this project we have suggested a totally original method of implementing the Walsh and Slant transforms in a texture classification scheme that proposes to be insensitive to texture orientation. This method has been developed and tested as part of this project, and will allow unique texture classification from the Walsh or Slant features, regardless of the orientation of the tissue texture being characterised.

When the Walsh or Slant transformation is applied to an image matrix, the resulting two dimensional matrix contains the coefficients of the transform basis functions. The initial image can be subsequently reconstructed by the superposition of the basis functions weighted by the coefficient matrix. Because these techniques are discrete and break down an image into a set of basic structural primitives, they are very sensitive to orientation. An image will be decomposed differently if its orientation has changed. It is not always possible to uniquely classify texture with a such a measure.

A solution to this problem is to develop a method that is able to registrate the image texture regions prior to the texture analysis. Image regions could be forced into a particular orientation as a rule. This method would require a complex decision making process and a large knowledge base of the textures involved. This is an unrealistic approach to the problem.

Image texture is a property of the relationship between the neighbouring pixels in an image. Let us therefore make the assumption that all the texture information is contained in the relationship between a pixel and its eight nearest neighbours, for all pixels across the image. Given this assumption, there are only eight possible orientations that a texture may adopt; namely, through 0, 45, 90, 135, 180, 225, 270 and 315 degrees. In order to eliminate the orientation sensitivity of our Walsh and Slant texture measures, we suggest the following pre-transformation step. Eight copies of the texture region being examined should be made; each copy adopting one of the eight possible orientations of the original image region. The eight

matrices should be used as inputs to the Walsh or Slant transforms. The eight transform result matrices should then be averaged to form a single matrix that is normalised with respect to texture orientation. This matrix can then be analysed as a conventional Walsh or Slant result matrix, and used as the basis for texture classification.

The advantage of this normalisation process is that given the texture features derived from the transform methods are good enough, it will be possible to uniquely specify a texture type despite changes in orientation across the region.

There are, however, several shortfalls with this method. Because we have assumed that textural information is contained within the relationship between a pixel and its eight nearest neighbouring pixels, we must consider the interpolation errors occurring from orientation changes not considered in our set of eight 45 degree increment rotations. Also we are discretising the texture by restricting the relationship between a pixel and its more distant neighbours. Two other potential difficulties arise because we are calculating the transformation of the texture region at different orientations, and averaging the result. The first is the amplification of image structural features in the result of the transformation. To reduce this effect we must chose the texture regions more carefully to avoid the presence of tissue boundaries, arteries etc. The second potential difficulty stems from the simple fact that by averaging the result of eight image transforms we will be dramatically reducing the information content of the result. This may reduce the resolving power of the texture methods and limit their ability to uniquely characterise texture.

We have talked about how textural information is held in the relationship between a pixel and its neighbours. The human observer is unable to contemplate these inter-pixel, or micro, relationships. The human observer witnesses the overall, or macro, effect of all these micro relationships across the whole image region. Many of the statistical texture analysis methods described in **chapter 4** extract textural information from the micro relationships between pixels. The Walsh and Slant transforms, however, examine the structural composition of an image. This equates to the macro effect of the inter-pixel relationships. Therefore, the orientation normalisation of the micro relationships guarantees a continuity of the macro or structural texture. This is the motivation and the justification for this method. This process may appear to be speculative, and therefore, will require some experimentation to confirm its validity.

In order for the image data from the eight image region orientations to be consistent, the image region must be bounded. Consider a square image region that contains the figure '*Y*', **figure 5.7a**. If the region is rotated through 45 degrees about its centre, then the data from the corners will overlap the boundaries set by the region's square shape, and data will be lost. If we clip the image data from the corners, a hexagonal region within the square image region is formed, as bounded by the dotted line in the figure. This hexagonal region of image data within the square region can then be rotated about its centre through the orientations shown in **figure 5.7b**. The data being processed from all orientations will therefore be consistent. Because we will only be processing this hexagonal sub-region of the texture region, this must be taken into account when selecting the square texture region to interrogate.

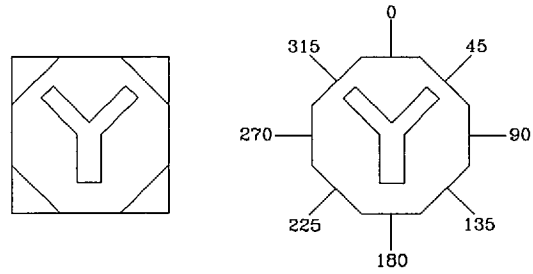


Figure 5.7 : Clipping image

The image region is stored in the computer as a two dimensional array or matrix. In order to enable the rotation of this square region, it must copied into a much larger matrix to give the corners clearance. **Figure 5.8** shows this large matrix with the square image region in the centre. The square with a dotted outline shows the position of the image region when rotated through 45 degrees about its centre, and indicates the clearance required for the corners. The large matrix is padded with zeroes around the side of the central image region. The corners of the image region are zeroed as bounded by the lines from the rotated image in the figure. The only non-zero values within the large matrix are those in the hexagon image region common to all orientations. This region is identified in **figure 5.8** by the area in common between the two image regions centred in the large matrix.

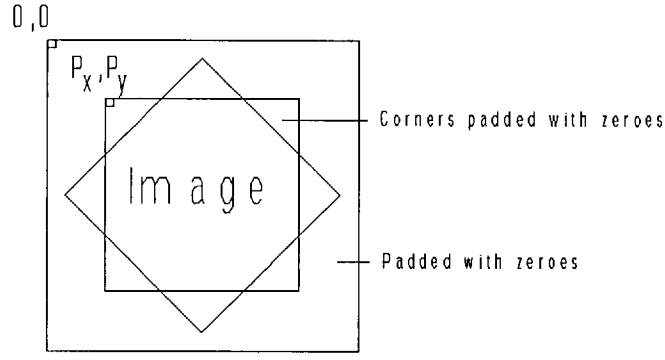


Figure 5.8 : Image padding

The position of the square image region within the large matrix is defined using a cartesian coordinate system. The zero axis (0,0) is defined as the top left corner of the large matrix and the relative position of the image region is given as (P_x,P_y), **figure 5.8**. These coordinates are defined in **equation 5.21**, where s is the side length of the image region. The dimensions of the image region corners are defined by the isosceles triangle of side lengths; x, x, and $x*\sqrt{2}$, where x is given in **equation 5.22**. These values are used as matrix indices and are rounded up to the next integer to avoid computational errors.

$$P_x = P_y = \frac{s}{2(1+\sqrt{2})} \tag{5.21}$$

$$x = \frac{s}{(2+\sqrt{2})} \tag{5.22}$$

The rotation transformations used to acquire the eight orientations of the image region $I(x,y)$ are given in **equations 5.23** to **5.30**, where $t=1/\sqrt{2}$, and s is the image side length. The image is rotated clockwise about its centre within the large matrix by 0, 45, 90, 135, 180, 225, 270, and 315 degrees, and the resulting transformed image regions are I_0, I_{45}, I_{90}, I_{135}, I_{180}, I_{225}, I_{270} and I_{315} respectively. The positions of the pixels in the large matrix are found using the offset coordinates (P_x,P_y).

$$0 \text{ degrees} \quad I_0(x,y) = I(x,y) \tag{5.23}$$

$$45 \text{ degrees} \quad I_{45}(x,y) = I(k,m)$$
$$k = t*(x-y) + s/2$$
$$m = t*(x+y) + s*(1/2-1/t) \tag{5.24}$$

90 *degrees* $I_{90}(x,y) = I(y, s-1-x)$ (5.25)

135 *degrees* $I_{135}(x,y) = I_{45}(y, s-1-x)$ (5.26)

180 *degrees* $I_{180}(x,y) = I(s-1-x, s-1-y)$ (5.27)

225 *degrees* $I_{225}(x,y) = I_{45}(s-1-x, s-1-y)$ (5.28)

270 *degrees* $I_{270}(x,y) = I(s-1-y, x)$ (5.29)

315 *degrees* $I_{315}(x,y) = I_{45}(s-1-y, x,)$ (5.30)

For the four compass point orientations (0, 90, 180 and 270 degrees) the rotation transformation is a simple translation in any of the two orthogonal axis. These transformations are precise processes. For the remaining four orientations (45, 135, 225, 315 degrees) the rotation transformation is a more complicated procedure. The transformed coordinates are calculated using the arithmetical expressions given in **equation 5.24**. Because the new coordinates of each pixel are non-integer, they must be approximated to the nearest whole integer. Another consequence of the non-integer transformed pixel coordinates is that many of the pixels in the new image are un-assigned. An interpolation process is required to fill the missing data points in the rotated image region.

The interpolation process we have used is the most basic; the missing values are estimated from the arithmetic mean of the pixel values of their neighbours. The rotation transformation process uses a mask to indicate which pixels have been assigned in the newly orientated image. After the new image has been created, this mask is examined and the missing values identified. Each missing value is found from the arithmetic average of the values from the eight nearest pixels. If the mask indicates that any of these eight neighbouring pixels have not been assigned, they are simply excluded from the estimation.

Once the eight differently orientated versions of the original image region have been acquired, then either the Walsh or Slant transformation can be applied to them in turn. These image regions are square matrices with a hexagon of image data bordered with zeroes. The resulting image functions from the Walsh and Slant transforms are $[F]_H$ and $[F]_S$, as given in **equations 5.10**, **5.20**, respectively.

The eight result functions are then summed to give a single result matrix. Because each input

function consists of a hexagon of image data bordered with zeroes to form a square, and the Walsh and Slant transforms are sensitive to structure, the result matrix will be dominated by boundary information. An image mask the same size as the processed image regions must therefore be created with the same hexagon bounded by zeroes. The constant value held by the pixels inside the hexagon is chosen arbitrarily to create a step in pixel value across the boundary. This image mask is transformed in the same way as the eight differently orientated image regions. Because the hexagon boundary in the image regions is more significant than any texture detail, it will contribute the large values to the result matrix. The result matrix from the image mask will contain the same boundary information, although the values may be scaled up or down. The unwanted boundary information in the image region result matrix is removed by subtracting the scaled image mask result matrix. Because the largest values in the result matrix are due to the boundary information, the mask result matrix is scaled to the maximum value held in the image result matrix.

The final result of the process described in this chapter is a matrix not unlike that one would expect from a straight forward application of the Walsh or Slant transform on an image region. We can therefore use some of the interpretative methods described in **section 5.3** and **5.4** to extract some textural features from the result matrix. In these sections we have examined textural features extracted from the sequency spectrum from both the Walsh and Slant transforms. We have also investigated the axis symmetry property and subsequently the axis symmetry histograms that may be calculated from the results of the Walsh transform.

As with the standard application of the Walsh and Slant transforms, this process is sensitive to translation of patterns or texture. Applying the Fourier transform to image data prior to implementing the Walsh or Slant transformas as suggested by (ALEXANDRIDIS and KLINGER, 1971b) ensures that the result is invariant under translation. This approach can be incorporated in the process described in this chapter.

As part of this project, we have suggested three speculative approach to the analysis of the result matrix. The first method is to calculate the first order statistical moments of the result matrix. The statistics of the matrix may hold key textural information. The mean, standard deviation, skewness and kurtosis measures will be found, and their efficacy as textural features will be appraised. The second and third methods extract from the result matrix the sets of textural features that are normally associated with the analysis of run length and co-

occurrence matrices. We have suggested the use of these sets of texture features because there are certain structural similarities between the run length, co-occurrence and Walsh/Slant result matrices. Although the axis represent different properties, the values in all three matrices have positional significance.

The run length result matrix, described in **chapter 4 section 4.3**, is a joint probability function of run length and grey level. The Walsh/Slant result matrix contains unbounded negative and positive values that represent weights of the basis functions. This matrix must be converted to an un-normalised probability density function ie. a function that records the frequency of events. The unsigned values of the Walsh/Slant result matrix are taken in order that this method can be used. The run length textural measures are given in **appendix 2**; **equation A2.1** short run emphasis, **equation A2.2** long runs emphasis, **equation A2.3** grey level non-uniformity, **equation A2.4** run length non-uniformity, and **equation A2.5** run percentage. The run length matrix has the run length axis along the rows and the grey level axis running down each column. We have interpreted the five run length textural features in the context of the multidirectional Walsh/Slant transforms as followings; **A2.1** a measure of low frequency contribution in both the x & y directions, **A2.2** a measure of high frequency contribution in both the x & y directions, **A2.3** sequency non-uniformity - this feature measures the distribution of sequency components; a uniform distribution gives a low value, **A2.4** the same as **A2.3**, and **A2.5** indeterminate structural measure.

The co-occurrence matrix, described in **chapter 4 section 4.6**, measures the frequency of the grey levels of the co-occurrent pixels. Like the run length matrix, the co-occurrence matrix is a probability density function, and so the Walsh/Slant matrix must be normalised in the same way as described above. In **chapter 4** we discuss the extraction of the 14 textural features suggested by (HARALICK *et al*, 1973). We propose to extract textural information from the Walsh/Slant result matrix by applying the same method of analysis. The 14 HARALICK *et al* texture measures ***F1-F14*** can be read about in **chapter 4 section 4.6**, and their definition can be found in **appendix 4**. It is very difficult to reason a description of each of these features in the context of the multi-directional Walsh/Slant method from a purely theoretical stance. We suggest that the interpretation of these features be left to a later stage when this method has been fully tested on a large selection of texture regions.

The theoretical justification for these three methods is somewhat obscure. The motivation for

employing these methods is purely experimental; the efficacy of the suggested texture features can only be measured by practical implementation.

5.6 SUMMARY

Three image transforms have been employed in the texture characterisation processes described in this chapter; the Fourier, the Walsh and the Slant transform. All the methods described are sensitive to structural detail. If the regions of texture chosen for analysis contain large structures such as tissue boundaries, arteries etc, then the structural effects will dominant. The selection of homogenous regions of texture for analysis also enables more accurate texture specification.

The Fourier transform is used in its own right to provide the frequency spectrum of regions of texture. The power spectrum was used by (HARALICK, 1979) to calculate the auto-correlation function as a texture measure. It was suggested by (WESZKA *et al*, 1976) that textural information can be extracted by taking averages of the power spectrum over ring-shaped regions centred at the origin. Both these measures describe the texture only in terms of its coarseness. We have suggested a set of parameters that should be extracted from the power spectrum to give us more dynamic indices to characterise texture.

The Walsh transform has been described by (LARSEN and LAI, 1980) as a valid alternative to the Fourier transform for extracting spectral information in some applications because its behaviour is similar and it is computationally faster than the Fourier transform. The set of Fourier spectral parameters suggested as texture descriptors have been defined for use in the context of the Walsh spectrum. The *axis symmetry* property of the Walsh functions has been described by (ALEXANDRIDIS and KLINGER, 1971b), and how it can be used to decompose the geometric patterns to form a histogram. These histograms contain implicit information about symmetries, periodicities and discontinuities present in a texture region. We have used this property to differentiate between different textures. One shortfall of the Walsh transform is that it is sensitive to any translation within the image region being scrutinised. This is particularly important in the analysis of texture because the texture primitives are very small and a texture region can only be arbitrarily sampled from the image. ALEXANDRIDIS and KLINGER suggested that the texture region is first Fourier transformed, and then the magnitude function taken as the input to the Walsh transform. The process renders the process

invariant under translation. We have therefore examined the performance of the Walsh transform with and without this pre-processing stage.

The motivation for using the Slant transform comes from the comments made by (MELTZER *et al*, 1967) when using the Walsh transform to specify biological form. It was suggested that the Walsh basis functions are too rigid for the specification of natural form, and a more 'lifelike' set should be used. Because the Slant basis functions are less discrete than the Walsh basis functions, they are more suited for the analysis of clinical MRI images. We have assumed that (LARSEN and LAI, 1980)'s suggestion for the Walsh transform also applies to the Slant transform. The Fourier spectral parameters suggested as texture descriptors have therefore also been defined for use in the context of the Slant spectrum. Because the Slant basis functions are continuous, the *axis symmetry* property is not valid for the Slant transform. Like the Walsh transform, the Slant transform is sensitive to any translation within the image being examined. We have therefore measured the performance of the Slant transform with and without this Fourier pre-processing stage.

Because the Walsh and Slant transformations decompose images by their geometrical properties, the variation in texture orientation inherent in clinical MRI images makes it difficult to uniquely specify the texture type. We have suggested a method of implementing the Walsh and Slant transforms in a scheme that proposes to be insensitive to texture orientation. In addition to the methods of analysis already discussed for the Walsh and Slant transformations, we have suggested three more for this approach. The first method calculates the statistical moments from the resulting transformed image. The second and third methods take the result image and analyse it as if it were a run length or a co-occurrence matrix, respectively. The performance of these three speculative techniques can only be measured through experimentation.

CHAPTER 6 - FRACTAL TEXTURAL FEATURES

6.1 INTRODUCTION

In natural scenes the perceived texture varies according to scale. As one moves closer to an object more detail becomes apparent, and the texture appears to change. It seems that our perception texture is a function of proximity. A good analogy of this phenomena is the view of a ploughed field from an airplane. From a great distance, the field takes on a lined texture due to the visible furrows. As the plane gets nearer, the *field of view* becomes effectively smaller, and more detail is visible. In the limit of this analogy, the *field of view* is restricted to within one furrow and the texture is mottled, due to the soil.

From our example of the ploughed field, it is possible to begin to form a clear picture of the concept of visual texture. A texture pattern can be regarded as being composed of basic elements or primitives (HARALICK, 1979). HARALICK describes these texture primitives as small shapes or patterns that repeat regularly according to a set of placement rules. When examined more closely, ie. at greater magnification, it becomes evident that the texture primitives may also contain their own texture. The texture primitive sub-textures can subsequently be defined in terms of their own set of primitives etc. This definition of texture clearly fits in with our example of the ploughed field. This model of texture is also adopted by (LIPKIN and ROSENFELD, 1970), who suggested that texture can be described using a three components model; the element, the sub-pattern, and the placement rule. The texture consists of small sub-patterns which are repeated according to defined placements rules. The sub-patterns themselves may be comprised of structured elements or, more generally, of sub-sub-patterns, which may themselves be made up of still smaller elements etc.

The concept of fractal geometry was first introduced by (MANDELBROT, 1977). He originally suggested fractal geometry as a method of mathematically modelling natural phenomena that would be described as *chaotic* systems such as long term weather forecasting and economics. In the most general terms, fractals are defined by a property called self-similarity, (FEDER, 1988; LIU, 1992). This term arises from the property of fractals objects to be similar at different scales. If the pattern of a fractal signal or surface is examined closely, it is found to be comprised of smaller sub-patterns that obey the same structural properties as the overall pattern. These sub-patterns are subsequently found to be comprised

of a number of sub-sub-patterns that again obey the same structural rules, but at a further smaller scale etc. The underlying structure is repeated through all scales of the fractal surface, and is independent of the scale of observation.

A good example of this self similarity property using the lung as the fractal object was given by (BARTLETT, 1991). The lung has a branching structure which repeats over and over as the trachea divides into bronchi and those divide into yet smaller branches, and so on to the bronchioles. The basic structure of a single tube dividing into two smaller tubes is repeated over length scales ranging from the size of the trachea to the size of terminal bronchioles. Fractals, therefore, do not have a single length scale, but rather have structure on multiple scales of length, (GOLDBERGER, 1992).

If we apply this self similarity property to the (LIPKIN and ROSENFELD, 1970) model for texture structure, then we can define a texture with fractal properties. The texture would be comprised of texture primitives that are small patterns that repeat regularly according to a placement rule. If the texture was fractal then the primitive patterns, when examined at a larger scale, would resemble the overall texture they form. In other words, the structure of the texture pattern and the pattern contained in the texture primitives would appear the same when viewed at the same scale. This texture is fractal because the observed texture will always be the same, regardless of the viewing scale; ie. the texture is invariant under scale.

Ideal fractals objects are those which obey the self similarity criterion strictly. A formal mathematical definition for fractals objects is provided by (PENTLAND, 1984) and (KUBE and PENTLAND, 1988). Although there are few examples of ideal fractal geometry in nature, there are many examples of objects or surfaces that possess fractal properties. An ideal fractal object possesses self-similarity at all scales. For a non-ideal fractal there are always discrepancies between the observed object at different scales. It is equally difficult to find a natural texture that is completely fractal. PENTLAND states that natural texture exhibits only a limited range of scale for self similarity. Image texture becomes less fractal as there is more variation in the observed structure at different scales. The clinical MRI images that have been used for analysis as part of this project present texture samples that are clearly not fully fractal. However, it is clear that these images will at least exhibit limited fractal properties by virtue of the texture patterns. If there is a method by which the degree of fractal behaviour of image texture can be measured, then it may form the basis for a useful texture

characterisation tool.

A widespread interest in fractal geometry and its potential applications has developed in the wake of the introduction of the concept of fractal geometry by (MANDELBROT, 1977; MANDELBROT, 1982). Although most applications of fractals involve their use in modelling natural phenomena, it is interesting to note that fractal patterns such as those created using *Mandelbrot* and *Julia* sets have become very popular as visual works of art. In the cinema such fractal images have even been used for effect in film sequences. Reference is made to two sci-fi films that use these methods by (LINNETT and CLARKE, 1990); "Star Trek II: The Wrath of Khan" and "The Last Starfighter".

It is suggested by (PENTLAND, 1984) that representational schemes such as Plato's notion of ideal forms (e.g. spheres, cylinders, and cubes) are inadequate to describe natural objects. He proposed that a method of shape representation capable of succinctly describing the surface of natural objects should be sought. A method of obtaining such a description from raw data should also be determined. PENTLAND decided that fractal functions provided such a model of natural surfaces for two reasons; because many basic physical processes produce fractal surfaces, but even more importantly because fractals *look* like natural surfaces. The defining characteristic of a fractal is that it has a *fractal dimension*. This is a measure of a surface that corresponds quite closely to our intuitive notion of roughness. This measure loosely describes the degree of fractal behaviour of a surface and therefore may be of use to us in texture classification. The full definition of the *fractal dimension* is provided in **section 6.2**.

PENTLAND found that the *fractal dimension* is appropriate in describing the perceptual smoothness of a surface, and proposed to use it for texture segmentation. It wa suggested by (MEDIONI and YASUMOTO, 1984) that because the *fractal dimension* suffers from drawbacks associated with any single feature measurement space, it is not sufficient to uniquely categorise texture; it describes one aspect of a texture and therefore can only separate texture which differ enough in roughness. He implies that it should be used in conjunction with other texture descriptors so that feature space can be sufficiently partitioned according to the texture classes being examined.

Further justification for the use of the fractal model as a source of textural features is given by (ALLINSON and LAWSON, 1990), who examined the fractal properties of two sets of

image data. The first set were digitised images of natural textures from (BRODATZ, 1966). The second set were artificially generated fractal surfaces; using the recursive mid-point displacement method given by (DODD, 1987) and (BARNSLEY, 1988). The motivation for ALLINSON and LAWSON's work was the desire to perform accurate texture characterisation using fractal techniques. They found their artificial textures to be fully fractal as expected. They also found that the natural textures exhibited fractal properties. Although the range of scale of self similarity in the natural textures was limited, ALLINSON and LAWSON concluded that fractal techniques were a valid alternative to statistical and structural methods for the description of natural texture. These conclusions re-enforce the notion that fractal methods can be useful in the description of image texture and hence tissue characterisation in clinical MRI images.

Another example of fractal methods being used in the description of texture is provided by (LINNETT and CLARKE, 1990), who used fractal methods in the modelling and analysis of texture in studies of the seabed. Real data arrives in the form of digitised sonar or video data. LINNETT and CLARKE uses the *fractal dimension* measure to model the texture of the seabed. They also address the concerns of (MEDIONI and YASUMOTO, 1984) about single parameter feature space by enhancing his model with additional parameters such as bathometry data (height).

Based on the conclusions of much of the work described in this introduction, it is reasonable to assume that the fractal properties of image texture can be used to differentiate between images containing different textures. As part of the work undertaken for this project we will be applying a variety of fractal methods to clinical MRI images to extract textural information. The clinical images used in this study have been identified as containing natural texture patterns. Because the MRI data acquisition method imposes a large number of noise artifacts on the images, we can only speculate as to the degree of fractal behaviour that these images display. Although the image texture exhibits a limited range of scale of self similarity, fractal measures such as the *fractal dimension* may prove useful in discriminating between different texture regions within the images. The concept of *fractal dimension* is introduced and described in **section 6.2**.

As part of this project, three different methods of calculating the *fractal dimension* of an image region have been investigated, namely the *box* method, the *blanket* method and the

power spectrum method; described in **sections 6.3**, **6.4**, and **6.5** respectively. The *box* method (CALDWELL *et al*, 1990), and the *blanket* method (PELEG *et al*, 1984), are presented in this chapter as they appear in the literature. Most of the practical considerations for these methods have been met in the original text. Work has, however, been carried out in the development of the *box* and *blanket* method algorithms and their implementation in the *MAIVIS* image visualisation and analysis application developed as part of this project, **chapter three**. The *power spectrum* method has been developed fully as part of this project from the mathematical theory of (PENTLAND, 1984). This required some postulation as to the correct relationship between the power spectrum of an image function and its *fractal dimension*. The complicated practical implications of this method were overcome, and an algorithm was developed and implemented. The analysis of the result data from all three methods required some careful interpretation. It is the intention of this study to examine all three methods of obtaining the *fractal dimension*, and to measure their potential efficacy as parameters for the characterisation of clinical MRI tissue texture.

Section 6.6 describes an interesting application of fractals suggested by (KELLER *et al*, 1987) for the identification of outlines of shapes. We have recognised this method as being potentially useful for the automatic selection of tissue regions for texture analysis. Such a method would serve as an ideal precursor to the tissue texture characterisation process. The implications of such a method are discussed in the context of our study of clinical MRI images.

6.2 THE FRACTAL DIMENSION

The classic example used by (MANDELBROT, 1982) to illustrate the concept of fractal geometry is the measurement of the coastline of Norway. In his book MANDELBROT dedicates a whole chapter, called 'How long is the coast of Britain', to the measurement of various coastlines and boundaries. The example computes a measure of fractal behaviour from the relationship between the measured length of the coastline and the length of the measuring ruler, for different length rulers.

If it were possible to measure the length of a coastline with a ruler of length ε, then the observed length $L(\varepsilon)$ would be the number of rulers needed to be put end to end to cover the distance, multiplied by the ruler length. As the size of the ruler is reduced, so the accuracy

at which the coastline can be measured increases. This observed length of the coastline will therefore approach the actual coastline length L_0 as the ruler length tends to zero. The effect of varying the ruler length on the measurement accuracy can be illustrated by considering the approximations that occur in the measurement of the length of a curve with a straight ruler; as the ruler becomes smaller in length it can trace the contours of the curve more accurately.

MANDELBROT defines the relationship between the measured length of the coastline and the ruler length in **equation 6.1**. In this equation the parameter L_0 is the actual coastline length, and D_T is the topological dimension; in the case of line length measurements the topological dimension is one. The parameter D is a fractional measure of dimension that gives an indication of fractal behaviour. It was explained by (FORTIN *et al*, 1992) that this measure of dimension offers a way to numerically quantify the property of self similarity. This measure is called the *fractal dimension*, and is often referred to in the literature as the Hausdorf-Besicovitch dimension.

$$L(\epsilon) = L_0 \epsilon^{D_T - D} \tag{6.1}$$

A simple definition of fractals in terms of the *fractal dimension* is provided by (FEDER, 1988). Non-fractal objects such as ordinary curves or surfaces, ie. objects with Euclidian geometry, have a *fractal dimension* equal to the topological dimension. Fractal objects, such as the line in MANDELBROT's coastline example, have a non-integer *fractal dimension* that strictly exceeds the topological dimension, ie. $D \geq D_T$. The term fractal originates from the non-integer nature of the *fractional dimension* of a fractal object.

If we take the example of the branching tracheo-bronchial tree given by (BARTLETT, 1991), a fractal like structure has a dimension between 2 and 3 since it converts a volume of gas in the trachea ($D=3$) into something approaching a surface area ($D=2$) in the alveoli. BARTLETT describes the *fractal dimension* as giving a measure of how space-filling and object is. In other words, if an object exists within a given volume of space, a more convoluted object will take up more of that space. Similarly, a more convoluted object would have a higher *fractal dimension.*

The measurement of the area of a fractal surface is an extension to the coastline example. The surface area is found by counting the number of square tiles of side length ε that are required

to cover the whole area, and multiplying that number by the area of a one tile (ε^2). The relationship (MANDELBROT, 1982) defines between the measured surface area A(ε) and the tile side length ε is given in **equation 6.2**, where A_0 is the actual surface area. As in **equation 6.1**, D_T is the topological dimension (for a surface D_T=2), and the parameter D is the *fractal dimension* of the surface

$$A(\epsilon) = A_0 \epsilon^{D_T - D} \tag{6.2}$$

An interpretation of the *fractal dimension* of a surface relevant to textural analysis is provided by (FEDER, 1988). FEDER sees the *fractal dimension* as a measure that describes the coarseness of texture, and was able to reason his argument by examining the rate of change of the measured surface area A(ε) with the size of the tile ε, and its effect on the *fractal dimension* in **equation 6.2**. A smooth surface like a slowly undulating curve would be described as having a coarse texture. For such a surface the rate of change of measured area with tile size would be slow, and hence a low *fractal dimension*. In the limit where the *fractal dimension* equals the topological dimension, $D=D_T$, the surface is flat and Euclidean. For a surface that contains a 'busy' or fine texture, the rate of change of surface area with tile size would be larger, and the subsequent value of *fractal dimension* would be greater.

These findings can be summarised by describing the *fractal dimension* as a measure that defines how jagged a surface is; a large *fractal dimension* indicates fine texture, and a low *fractal dimension* indicates coarse texture.

Chapter 4 describes methods that characterise texture types using the second order image statistics. Statistical texture descriptors are the most well known and best documented class of texture characterisation parameters. Further justification for the use of the *fractal dimension* as a texture parameter is provided by the findings of (JARDINE and WHITWORTH, 1992), who found that the second order statistics of a fractal texture are completely defined by its fractional (*fractal*) dimension.

Fractal properties are sometimes described in terms of the self similarity factor, H, instead of the *fractal dimension*. The relationship between these two measures, given by (DODD, 1987), is provided in **equation 6.3**. The self similarity factor can hold values in the range of

0 to 1; for a smooth surface with a low *fractal dimension* this factor will have a high value, and for a jagged surface with a high *fractal dimension* this factor will have a low value.

$$D=D_T+1-H \tag{6.3}$$

The expression in **equation 6.2** relating the measured surface area for a particular tile size to the actual surface can be used to estimate the *fractal dimension* of a surface. When several measurements of the surface area are made with different tile sizes, in addition to the equation constants (topological dimension and actual surface area), the value of *fractal dimension* should be consistent between measurements. **Equation 6.2** can be simplified by taking natural *logs* from each side of the equation. The resulting expression, given in **equation 6.4,** is the equation of a line with a gradient equal to the difference between the topological dimension and the *fractal dimension* (D_T-D).

$$Log\ A(\epsilon) = K + (D_T-D)\ Log\ \epsilon \tag{6.4}$$

The *fractal dimension* of surface can therefore be estimated by plotting the natural *log* of the surface area measured against the natural *log* of the tile size for all measurements, and obtaining the gradient of the line that best fits the plotted points. The *fractal dimension* is subsequently found by subtracting the line gradient from the topological dimension.

The images we are examining as part of this project have two dimensional intensity functions with fixed dimensions (256 by 256 pixels). Each image can be though of as a surface, where the *xy*-plane of the surface is represented by the image matrix, and the *z*-component of the surface is the intensity value.

If a square tile is placed over any region of an image, then the pixel intensity measured at that position is an average of all the pixels that the tile covers. In order that we can measure the surface area of an image with different sized tiles, we must artificially create a set of successively lower resolution images, where the resolution in terms of pixels is equivalent to the tile size. This decimation can be achieved by successively averaging the pixel values from groups of four neighbouring pixels to form larger pixels (by a factor of two). The surface area measurements made from each image in this set are therefore analogous to the measurements

made on a fractal surface with different size tiles.

6.3 THE FRACTAL DIMENSION AND THE BOX METHOD

As described in **section 6.2**, the *fractal dimension* of a surface can be found from the gradient of the line fitted to the points on the *log-log* plot of its measured with different sized tiles. The expression that relates the gradient of this line with the surface fractal dimension is given in **equation 6.4**.

The surface of an image can be visualised by the 3D projection of the grey level intensity values out of the matrix plane. In order that we can estimate the *fractal dimension* of an image we need to be able to make several measurements of its surface area with different tile sizes. This process may achieved by means of the *box* method. The box method was used for surface area measurements by (CALDWELL *et al*, 1990) to such end for a set of digitised mammograms. CALDWELL *et al* employed this method because of its simplicity both in understanding and practical implementation. For similar reasons, we have developed a *box* method algorithm based on this scheme, and implemented it for the estimation of the *fractal dimension* of clinical MRI images.

The *box* method relies on the analogy that the surface of an image is a collection of adjacent skyscrapers. Each skyscraper has a square 'roof' of side length ε. The skyscraper 'roof' is defined by the resolution at which the image surface is being measured ie. the minimum pixel size as constrained by the tile size ε. Pixels equal to the tile size are formed by averaging the original image pixels that the tile covers when placed on the image. The height of each column or skyscraper is defined by the corresponding pixel grey level intensity.

The surface area for a particular tile size is found from the sum of the area of the 'roofs' and the area of the exposed sides of the columns. The measured area is given by **equation 6.5**, where the grey level intensity function $I(x,y)$ gives the height of the column at the pixel position (x,y).

$$A(\epsilon) = \sum_{x,y} \epsilon^2 + \sum_{x,y} \epsilon \{abs[I(x,y) - I(x+1,y)]\} + \sum_{x,y} \epsilon \{abs[I(x,y) - I(x,y+1)]\} \tag{6.5}$$

The basis for the calculation of the *fractal dimension* of a surface is the measurement of its area with different size tiles, **section 6.2.** CALDWELL *et al* estimated the *fractal dimension* by measuring the surface area of his images A(ε) firstly with a tile size equal to 1 pixel, and subsequently with the tile side ε equal to 2, 3, 4, 5 and 6 pixels in length. He was able to find the *fractal dimension* of his images by calculating the gradient of the line that passes through the corresponding points on the *log-log* plot of tile size and measured surface area.

In his critique of the *fractal dimension*, (BARTLETT, 1991) comments as to the scope of usage of the *box* method. BARTLETT states that because there is no way of ensuring a fixed relationship between the pixel size and the pixel *thickness* from image to image, then a unique *fractal dimension* value cannot be calculated using this method. Pixel *thickness* refers to the number of the image grey levels, and in the analogy refers to the height of the skyscrapers. The pixel *thickness* depends on the parameters of the image acquisition process; eg. the average illumination and the depth of digitisation ie. the number of grey levels. Any variation in the ratio of surface 'height' to the pixel size between images will bias any comparison of their calculated *dimensions*.

The pixel *thickness* is the third dimension of the surface and must be normalised with respect to the pixel size before the *box* method can be used to generate unique classification parameters. We recommend that images for comparison be pre-normalised to n-bits, and N by N matrix size. The scope of use of the *fractal dimension* will therefore be limited to the each set of normalised images; it will not be possible to make fair comparisons of this texture parameter between differently normalised data sets.

Another source of error described by (BARTLETT, 1991) is the crude approximation of the surface provided by the rectangular columns. Therefore, there is some doubt as to how accurate the *dimensions* can be calculated. The aim of this work is to establish whether the *box* method provides competent indices for the classification of the texture of a selection of tissue types in clinical MRI images. In our favour, BARTLETT reports that in the experiments with mammograms carried out by (CALDWELL *et al*, 1990), the pathological changes in *dimension* are much larger than the errors introduced by the surface approximation.

In practice, the observed *log-log* plot of measured surface area versus measuring tile size is not always linear. The nonlinear part of the plot is due to the limitations of image data and

occurs as the image approaches the limit of decimation, where the surface is flat, ie. large tile size. At this point the surface area is only equal to the sum of the area of the 'roofs'. The plot becomes curved and tails off in this limit. The gradient should therefore only be found for the linear part of the plot. A more constructive use of this shortcoming is to measure the gradient at N specific points along the plot and record a fractal signature ie. an N-dimensional feature vector. The use of a fractal signature as a measure of texture was suggested by (PELEG *et al*, 1984) in conjunction with the *blanket* technique, **section 6.4**. This vector can subsequently be used in texture classification.

We have implemented the (CALDWELL *et al*, 1990) *box* method and made measurements of image surface area with tile sizes in the range $\varepsilon = 1$ to 6. The clinical MRI images have been normalised with respect to pixel depth. Regions of image texture of a fixed size have then been examined with this method and comparisons made. The *log-log* plot of surface area versus tile size has been analysed in one of two ways. The gradient of the linear part of the line that intersects the points on the plot is measured using a standard linear regression technique, (PRESS *et al*, 1990). This gradient is then used to estimate the overall *fractal dimension* of the image region. Alternatively, a fractal signature can be found for the image region by calculating the localised gradient along the plotted line. This signature can subsequently be used as a characteristic *fractal dimension* feature vector. In theory, both methods can be used to extract potentially useful image texture information. The efficacy of these methods as texture descriptors will be found by experimentation as part of this project.

6.4 THE FRACTAL DIMENSION AND THE BLANKET METHOD

If the *fractal dimension* of a surface can be calculated using more than one method, then the likelihood of accurate texture classification based on this parameter will increase because of the higher dimension feature vectors available. An alternative to the *box* method for the measurement of surface area, and subsequently the calculation of *fractal dimension*, is provided by the *blanket* method.

The *blanket* method was based by (PELEG *et al*, 1984) on one of (MANDELBROT, 1982)'s methods for the measurement of curve length. MANDELBROT commonly used a coastline to illustrate his curve measurement ideas. All the points on the curve with distances to the coastline of no more than ε were considered; these points formed a strip of width 2ε. The

suggested length $L(\varepsilon)$ of the coast can therefore be found from the area of the strip divided by 2ε.

PELEG *et al* employed this method of measuring image surface area in his *blanket* method because the surface extension is computationally efficient. In the extension from curve to surface, all points in three dimensional space at a distance ε from the surface were considered covering the surface with a 'blanket' of thickness 2ε. The covering blanket is defined by its upper surface u_ε and its lower surface b_ε. The surface area is given by the volume occupied by the blanket divided by the thickness 2ε. The initial upper and lower blanket surfaces are defined in **equation 6.6**, where $I(x,y)$ is the image grey level function. The subsequent upper and lower blanket surfaces for radial distance $\varepsilon = 1, 2, 3$, etc., are defined in **equation 6.7** and **6.8**.

$$u_0(x,y) = b_0(x,y) = I(x,y) \tag{6.6}$$

$$u_\epsilon = \max\left\{u_{\epsilon-1}(x,y)+1, \max_{|(m,n)-(x,y)|\le 1} u_{\epsilon-1}(m,n)\right\} \tag{6.7}$$

$$b_\epsilon = \min\left\{b_{\epsilon-1}(x,y)-1, \min_{|(m,n)-(x,y)|\le 1} b_{\epsilon-1}(m,n)\right\} \tag{6.8}$$

The volume of the blanket is the volume of space that is enclosed by its upper and lower surfaces, and is defined for radial distance ε in **equation 6.9**. PELEG *et al* initially defined the surface area as the measured volume V_ε divided by the thickness of the blanket 2ε. Because the surface area depends on the small changes that occur from scale ε-1 to scale ε, PELEG *et al* suggested that the surface area be found from the difference of the measured volume at different scales. The surface area for radial distance ε is therefore estimated by subtracting the volume $V_{\varepsilon-1}$ from V_ε. The result is divided by 2 to account for both the upper and lower surfaces, as in **equation 6.10**.

$$V_\epsilon = \sum_{x,y} \left(u_\epsilon(x,y) - b_\epsilon(x,y) \right) \tag{6.9}$$

$$A(\epsilon) = \frac{V_\epsilon - V_{\epsilon-1}}{2} \tag{6.10}$$

Like the *box* method in the previous section, the *fractal dimension* can be found from the gradient of a *log-log* plot of the measured surface area against the radial distance (or scale) as in **equation 6.4**. The gradient of the line can be found using a standard linear regression technique. Again, like the *box* method, the observed plot has a gradient that varies along the line. In preference to estimating the overall gradient of the plot, and subsequently a single *fractal dimension*, (PELEG *et al*, 1984) suggested the use of a fractal signature take full advantage of such variations. The fluctuations in line gradient are present because clinical MRI images do not represent true fractals surfaces, and the degree of self similarity from one scale to another is highly variable. However, we propose that images containing the same tissue texture will have similar trends of self similarity between different scales, and hence similar *log-log* plots when analysed. The fractal signature of an image is derived from the set of local gradient values found between each point and its two neighbours along the plot. The *fractal dimension* values are found from the gradient values as described in **equation 6.4**.

The fractal signature can subsequently be used as a feature vector in a simple Euclidean classification scheme. The tissue texture of one image region can be compared with that of another from the vector difference between their fractal signatures, $D(i,j)$. Tissue type characterisation can therefore be achieved by finding the nearest vector in feature space from a set of vectors that represent known tissue type classes. The difference vector is defined in **equation 6.11**, where S_i and S_j represent the fractal signatures of the two image texture regions being compared. The weighting factor $log[(\varepsilon + 1/2)/(\varepsilon - 1/2)]$ compensates for the unequal spacing of the points on the *log-log*.

$$D(i,j) = \sum_{\epsilon} \{S_i(\epsilon) - S_j(\epsilon)\}^2 \log\left(\frac{\epsilon+\frac{1}{2}}{\epsilon-\frac{1}{2}}\right) \qquad (6.11)$$

Because the surface area is calculated from the difference in volume measurements at different scales, this method is not sensitive to variations in pixel size and pixel thickness as pointed out by (BARTLETT, 1991) in the case of the *box* method in **section 6.3**. However, to avoid computational errors such as rounding etc. in the calculations, we recommend that all the image regions being compared with one another be pre-normalised to n-bits, and N by N matrix size.

A source of error described by (BARTLETT, 1991) with this method, which also presents itself in the *box* method, is the crude approximation made in the discrete representation of the surface being analysed. The *blanket* method calculates the surface area from the volume enclosed by two artificial surfaces. There is therefore some doubt as to how reliable the *fractal dimension* parameters calculated using this method will be in the classification, and subsequent characterisation of clinical MRI image tissue types from their texture. The aim of this part of the project is to establish the efficacy of the *blanket* method as a source of image texture characterisation parameters.

6.5 THE FRACTAL DIMENSION AND THE POWER SPECTRAL DENSITY

6.5.1 RELATIONSHIP BETWEEN PSD AND FRACTAL DIMENSION

The problem of representing natural shapes such as mountains, trees and clouds etc. and computing their description from image data was addressed by (PENTLAND, 1984) by relating natural surfaces to their images. PENTLAND found the 3D fractal model provided satisfactory characterisation of these 3D surfaces. Furthermore, he found this characterisation process to be stable over transformations of scale and linear transforms of intensity. PENTLAND also found the *fractal dimension* of these functions could be successfully measured from the two dimensional Fourier power spectrum $P_H(f,\theta)$ of the surface intensity function. We have adopted this approach as our third method for the calculation of the *fractal dimension* of tissue texture. These three methods when applied to our study data will provide us with a comprehensive profile of the fractal characteristics of the tissue samples from which we can effect tissue type characterisation.

PENTLAND described the frequency attenuation in the Fourier power spectrum as the key to the fractal properties of an image region. According to PENTLAND, the power spectrum of an ideal fractal function is proportional to the frequency raised to the negative power of a constant β, **equation 6.12**. This constant is a measure of the attenuation, or tail off, of the power spectrum as frequency increases, and because it is also related to the self-similarity factor, H, the *fractal dimension* of the surface can be found from this value. The relationship between the attenuation coefficient and the self-similarity factor is defined in **equation 6.13**. The fractal dimension can subsequently be found by subtracting the self-similarity factor from the topological dimension (2 for a surface) and adding $+1$, as in **equation 6.3**. We have

expressed the *fractal dimension* of a surface in terms of this attenuation factor in **equation 6.14** for our own use in this study, where D is the *fractal dimension* and D_T is the topological dimension (2) .

$$P_H(f,\theta) \propto f^{-\beta} \tag{6.12}$$

$$\beta = 2H+1 \tag{6.13}$$

$$D = D_T - \frac{3-\beta}{2} \tag{6.14}$$

The Fourier power spectrum and the *fractal dimension* can both be interpreted in terms of similar image features. Traditionally we associate the amount of high frequency content in an image with the detail present. This can be demonstrated by observing the loss of detail and the blurring that occurs when an image has high frequency information removed using low pass filtering methods. PENTLAND describes how the *fractal dimension* of a surface corresponds quite closely to the observer's intuitive notion of roughness. Similarly, (MEDIONI and YASUMOTO, 1984) suggest that the *fractal dimension* defines how jagged a surface is. We have used this close association between the power spectrum and the *fractal dimension* to explain the significance of the magnitude of the Fourier attenuation coefficient in the estimation of image *fractal dimension.*

A quick tail off of the power spectrum means a large value of the attenuation coefficient β, which would indicate that the surface has a limited range of high frequency components ie. the surface is smooth and lacks detail. Such a surface would have a large self-similarity factor and consequently a low *fractal dimension.*

A slow tail off of the power spectrum means a small value of β which would indicate that the surface has a large range of frequency components ie. the surface is possibly rough and jagged with a significant amount of detail. Such a surface would have a low self-similarity factor and consequently a high *fractal dimension.*

A complete description of the practical implementation of this method could not be found in the literature. As part of this project, we have taken the relationships detailed in **equations 6.12** and **6.14** and addressed the practical problems that have posed themselves in their

realisation. The rest of this section sets out to describe the method developed as part of this project that extracts the image *fractal dimension* from the Fourier power spectrum.

Equation 6.12 states that the power spectral density is proportional to the frequency raised to the negative power of the attenuation coefficient β. If natural *logs* are taken of both sides of this relationship a linear form of this equation can be found, **equation 6.15**. This suggests that if we make a *log-log* plot of the power spectral density against frequency it should, in theory, be a straight line and the *fractal dimension* can be found from its gradient. We have developed a recursive least squares regressive technique method to estimate the gradient ie. the attenuation coefficient β.

$$\log\{P(f)\} = K - \beta \log(f) \tag{6.15}$$

The fast Fourier transform (FFT) is the most usual method of obtaining the power spectrum because of its economy in computer effort. For practical purposes, the FFT algorithm incorporated in this method has been taken from a library of standard computational tools.

The power spectrum of a 2D function such as an image is itself a 2D function. In order that we can estimate the gradient ie. the attenuation coefficient β of the line defined in **equation 6.15**, we must average the 2D function radially to obtain a 1D power spectral density function. There are two good reasons for collapsing the power spectral density function. The first is that it makes it easier to estimate the function decay with frequency. The second reason is that the structural elements of the FFT are averaged out, so that the method is not orientation sensitive. This is important in the context of clinical MRI images, where the texture samples we wish to characterise may vary in orientation.

Because the clinical MRI images we are experimenting with do not represent ideal fractal functions, we expect a deviation from the theoretical ideal. Evidence of non-ideal fractal behaviour can be seen in the non-linear parts of the *log-log* plot of power spectral density against frequency. To optimise the result, we have used a recursive least squares regressive technique to estimate the gradient of the 'line'. The general time-series regressive method from which the technique is developed is introduced by (YOUNG, 1984). A method of weighting the regression that depends on the number of values contributing to the rotational average has been incorporated, as suggested by (BARTLETT, 1991).

Another potential source of error is the matrix size of the tissue regions being examined. The number of points on the *log-log* plot is exactly half the matrix dimension. If the regions are too small then there will be too few points on the plot to make a reliable measurement of line gradient.

6.5.2 MANIPULATING THE FREQUENCY INFORMATION

If the 2D fast Fourier transform (FFT) of an intensity function $I(x,y)$ is $F(u,v)$, then the Fourier power spectrum $P(u,v)$ formed is equal to $|F(u,v)|^2$, where $0 \leq (x, y, u, v) \leq N$, and N is the dimension of the image matrix. Because of the periodicity properties of the Fourier transform (GONZALEZ and WINTZ, 1987), the Fourier power spectrum matrix is split up into four quadrants (+u,+v), (-u,+v), (+v,-u), and (-u,-v). The four quadrants are organised in the power spectrum matrix in such a way so that their origins are at each corner, and the high frequency components meet at a reflective boundary at the centre. This arrangement can be seen in **figure 6.1**, where the four quadrants are labelled *q0*, *q1*, *q2* and *q3*. If each quadrant is translated diagonally then one would be looking at a conventional 2D Fourier power spectrum with the origin at the centre.

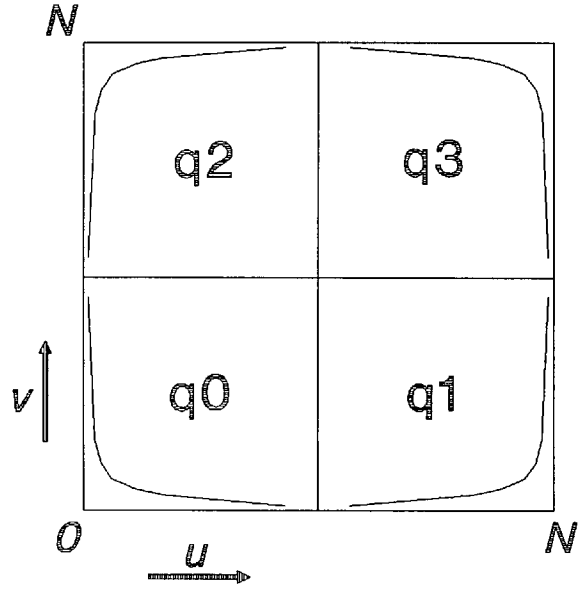


Figure 6.1 : Fourier transformed image quadrants

In order that we can plot a profile of power spectral density against frequency, we must average the 2D power spectrum radially to obtain a 1D function. This can be achieved by first

stacking and summing the four 2D frequency quadrants in the correct orientation, and then radially averaging this single quadrant. The three quadrants (*q1, q2, q3*) backed up at half a period (*N*/2) to the image origin quadrant (*q0*) must be flipped back on themselves to overlap, so that all four quadrants can be summed. The transformations required to translate the quadrants $P_{qi}(u,v)$ are defined in **equations 6.16** to **6.19**, with $0 \leq (u,v) \leq N/2$ and $i=\{0,1,2,3\}$. The four quadrants are then summed and averaged to form a single quadrant $P_A(u,v)$ as defined in **equation 6.20**. This result is subsequently collapsed radially to form a the 1D function of power spectral density.

$$P_{q0}(u,v) = P(u,v) \tag{6.16}$$

$$P_{q1}(u,v) = P(N-u,v) \tag{6.17}$$

$$P_{q2}(u,v) = P(u,N-v) \tag{6.18}$$

$$P_{q3}(u,v) = P(N-u,N-v) \tag{6.19}$$

$$P_A(u,v) = \frac{1}{4}\sum_{i=0}^{3} P_{qi}(u,v) \tag{6.20}$$

The method for radially collapsing the single quadrant which has been adopted for this project calculates the radial average $P(r)$ from those points that are said to be at the distance r from the origin of the 2D function. This radial function is defined in **equation 6.21**, where N_{points} is the total number of points found at each radial distance.

$$P(r) = \frac{1}{N_{points}}\sum_{u,v} P_A(u,v) \;\Big|_{(r-1)^2<u^2+v^2\leq r^2} \tag{6.21}$$

Figure 6.2 shows the radial positions of the pixels in the 2D quadrant function; the numbers in the boxes indicate the radial distance of each pixel from the origin.

This radial scheme has been chosen because it is probably the simplest. No attempt has been made to interpolate the exact pixel values at each radial distance because the aim of this study is primarily feasibility, and such methods would over complicate the process.

The 1D power spectrum $P(f)$ is equivalent to the radial function $P(r)$, and therefore can be

found as a result of the method of manipulation of the 2D Fourier power spectrum described in this section. In order that the fractal dimension of an image region can be found from its frequency information, the decay or attenuation coefficient β of the 1D power spectrum must be found and substituted into **equation 6.14**.

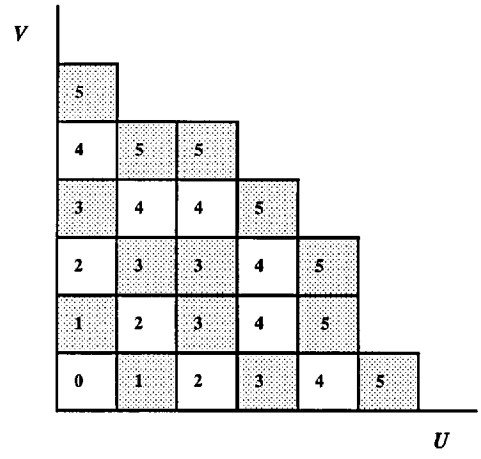


Figure 6.2 : Frequency values at radius $R = \text{sqrt}(U^2+V^2)$

6.5.3 CALCULATING THE FRACTAL DIMENSION

Although discussed as a source of fractal information, the method for estimating the *fractal dimension* from the power spectrum has not actually been described in any literature we have found. In response to this lack of practical details, we have developed a recursive least squares regression method to estimate the *fractal dimension* from the power spectrum. **Equation 6.15** presents a linear relationship between the *log* of the power spectrum value at a specific frequency, the *log* of that frequency, and the overall spectrum frequency attenuation coefficient β relating them both. This coefficient is the constant from which the *fractal dimension* is calculated, **equation 6.14**. **Equation 6.23** is the simplification of the linear expression in **equation 6.15**, using the assignments made in **equation 6.22**. We have transformed **equation 6.15** into a discrete form for two reasons. The first is that in our experiments we are dealing with a discrete source of data points, and the second is that it makes the algebra required in our formulation of the least squares regression method more straight forward.

$$\begin{aligned} y_i &= \log\{P(f)\}\,|_{i=f} \\ m_i &= \log(f)\,|_{i=f} \end{aligned} \tag{6.22}$$

$$y_i = K - \beta m_i \tag{6.23}$$

(YOUNG, 1984) presents the solution for the general problem in the estimation of the set of n unknown parameters a_j; $j=1$ to n, which appear in a linear regression relationship of the form given in **equation 6.24**, where the x_j; $j=1$ to n are *exactly known*, linearly independent variables. The observation y of x_0 is contaminated by noise e_y, as in **equation 6.25**. The minimization of the least squares criterion function for k samples ie. the sum of the squared noise for all samples, **equation 6.26**, requires that all the partial derivatives of J_2 with respect to each of the parameter estimates $\bar{a}_j$ should be set simultaneously to zero. Such a procedure yields a set of n linear, simultaneous algebraic equations which can be solved for the parameter estimates at the k^{th} instant.

$$x_0 = a_1x_1 + a_2x_2 + \ldots + a_nx_n \tag{6.24}$$

$$y = x_0 + e_y \tag{6.25}$$

$$J_2 \triangleq \sum_{i=1}^{k} e_{yi}^2 = \sum_{i=1}^{k} \left[\sum_{j=1}^{n} x_{ji}\bar{a}_j - y_i\right]^2 \tag{6.26}$$

The relationship between the power spectrum, the frequency, and the attenuation coefficient given in **equation 6.23** is an example of a linear regression relationship with the two unknown parameters, namely K and β. **Equation 6.27** shows how we have substituted this expression into x_0 with the known parameters $x_1=1$ and $x_2=m_i$, and the unknown parameters $a_1=K$ and $a_2=-\beta$, and the observation of x_0, namely y.

$$\begin{aligned} x_0 &= a_1x_1 + a_2x_2 = \mathbf{K}.1 - \beta.m_i \\ y &= y_i \end{aligned} \tag{6.27}$$

We have introduced a weighting factor into **equation 6.23** to take into account the number of values used to calculate each radial average of the power spectrum, as suggested by (BARTLETT, 1991). This modification is given in **equation 6.28**, where N_i is the number of values used to calculate each radial average, and N_{TOT} is the total number of values used

to calculate all the radial averages.

$$\frac{N_i}{N_{TOT}} y_i = K\frac{N_i}{N_{TOT}} - \beta\frac{N_i}{N_{TOT}} m_i \tag{6.28}$$

The least squares criterion function J_2 for our application we have given in **equation 6.29**, where the number of samples, $N/2$, corresponds to the number of data points on the 1D power spectrum. The unknown parameters in this equation, K and β, are written with a *bar* above them, $\overline{K}$ and $\overline{\beta}$, to denote that they are estimates. The method suggested by (YOUNG, 1984) requires that both the partial derivatives of this function with respect to the parameter estimates be set to zero. The partial derivatives with respect to $\overline{K}$ and $\overline{\beta}$ have been derived in **equations 6.30** and **6.31** respectively. These two simultaneous algebraic equations can then be solved for the parameter estimates at the $N/2$ instant.

$$J_2 \triangleq \frac{1}{N_{TOT}^2} \sum_{i=1}^{N/2} N_i^2 \left[(\overline{K} - \overline{\beta} m_i) - y_i \right]^2 = \sum_{i=1}^{N/2} e_{yi}^2 \tag{6.29}$$

$$\frac{\partial J_2}{\partial \overline{K}} \triangleq \overline{K} \sum N_i^2 - \overline{\beta} \sum N_i^2 m_i - \sum N_i^2 y_i = 0 \tag{6.30}$$

$$\frac{\partial J_2}{\partial \overline{\beta}} \triangleq -\overline{K} \sum N_i^2 m_i + \overline{\beta} \sum N_i^2 m_i^2 + \sum N_i^2 m_i y_i = 0 \tag{6.31}$$

We have found the solutions for the two consistent equations given **equations 6.30** and **6.31** The estimates for the unknown parameters K and β are given in **equations 6.32** and **6.33**. These parameters are found at the $N/2^{th}$ data point instant, ie. the summations in the equations are made for all data points, $i=1$ to $N/2$.

$$\overline{K} = \frac{\sum N_i^2 m_i \sum N_i^2 m_i y_i - \sum N_i^2 y_i \sum N_i^2 m_i^2}{\sum N_i^2 m_i \sum N_i^2 m_i - \sum N_i^2 \sum N_i^2 m_i^2} \tag{6.32}$$

$$\overline{\beta} = \frac{\sum N_i^2 \sum N_i^2 m_i y_i - \sum N_i^2 y_i \sum N_i^2 m_i}{\sum N_i^2 m_i \sum N_i^2 m_i - \sum N_i^2 \sum N_i^2 m_i^2} \tag{6.33}$$

If **equation 6.15** is re-examined, one finds it to be the equation of a line, with K being the x-value when the line intersects the y-axis and β the gradient. Because the line this equation

represents is the 1D power spectrum, the constant K is the *log* of the DC spectrum component, and the gradient β is a parameter used to relate the spectrum values with their frequencies, **equation 6.12**. Both the parameters estimated with this regression method are therefore very useful in the analysis of the power spectrum. We are however particularly interested in the estimate of the gradient parameter $\bar{\beta}$ for our calculation of the *fractal dimension*. The *fractal dimension* is found by substituting the line gradient into **equation 6.14**.

The process described in this section can be summarised as follows; A 2D texture region is selected from a clinical MRI image. The Fourier transform is applied to this function and its 2D power spectrum is subsequently found. This 2D function is manipulated and then collapsed to a 1D power spectrum function. A regression technique is then applied to the power spectrum to find the relationship between the spectrum values and the frequency at which they occur. This relationship subsequently provides us with the *fractal dimension* of the image region being analysed. The algorithm for this process has been fully developed as part of this project. To avoid repetition, the fast Fourier transform (FFT) routine has been taken from a popular library source, (PRESS *et al*, 1991).

6.6 THE FRACTAL DIMENSION AND AUTOMATED REGION SELECTION

In this section we discuss the potential usefulness of an interesting application of fractals that has been suggested by (KELLER *et al*, 1987) for the characterisation of objects from the *fractal dimension* of their outlines. The study carried out by KELLER *et al* uses fractal techniques to distinguish silhouettes of trees from the silhouettes of mountains. The method was applied to tree lines and mountain lines from a large number of photographs. The *fractal dimensions* of these lines were measured and the lines appropriately classified. Considerable success in classification accuracy was achieved by KELLER *et al* with this method.

We suggest that such a method has potential in the context of clinical MRI images as a process for the selection of image regions for tissue characterisation. This method could be used in conjunction with simple image segmentation techniques to identify the areas of a subject image that contain interesting pathologies; such that require more specific texture analytic processes for their characterisation.

For example, significantly different tissue regions can usually be identified using segmentation

methods based on differences in first order image statistics. With the aid of edge detection techniques and methods that re-enforce these edges into meaningful sets of object boundaries, these different tissue regions can be identified. An introduction to both edge detection, and edge linking and boundary detection is provided by (GONZALEZ and WINTZ, 1987).

Once different tissue regions have been identified, it should be possible to characterise the general tissue type (ie. liver, lung, etc.) from the *fractal dimension* calculated from the outline of the tissue region. From this general tissue type classification it will be possible to automatically select a region with a particular tissue type for more specific analysis of the pathologies that it contains. For example, this method may be used to initially identify the liver by the fractal properties of its outline. A more specific tissue characterisation approach may then be adopted to identify pathologies within the liver.

Although the three methods of calculating the *fractal dimension* described in this chapter are defined for the analysis of 2D surfaces, any of these methods may easily be adapted for the analysis of 1D outlines. However, we do not propose to suggest which of the three methods would prove most proficient in this application without some experimentation.

We have suggested that simple segmentation techniques can be used to initially identify tissue types with significantly different first order statistics, such as different soft tissue structures. More powerful segmentation techniques, outside the bounds of this project, could potentially identify all variations in tissue type, thus recognising the presence of pathologies in healthy tissue. Subsequent analysis of the outlines of all segmented tissue regions would provide a method of completely characterising pathology.

This method has not been implemented as part of this project because it would require considerably more research resources than are available bearing in mind, at present it is merely a conjecture initiated by the work of (KELLER *et al*, 1987).

6.7 SUMMARY

This chapter considers the notion that image texture has fractal properties. The basis for this premise is discussed in detail in the introductory section. The *fractal dimension* is a measure that describes how jagged a surface is; a large value indicates fine texture and a low value

indicates coarse texture. Based on this assumption, we have investigated three methods of calculating a fractal descriptor called the *fractal dimension* to describe the texture of image regions.

The three methods discussed for calculating the *fractal dimension* are called the *box* method, the *blanket* method, and the *power spectrum* method. These methods are able to characterise tissue texture by visualising image regions as surfaces, ie. 3D projections of the grey level intensity functions out of the matrix plane. The final section in this chapter describes a method that has been suggested as part of this project for characterising tissue type from the fractal properties of the tissue outline.

The *box* method relies on the analogy that the image surface is a collection of adjacent skyscrapers made up of roofs or tiles (xy-plane), and vertical rectangles (z-plane). The area of the surface is found at different resolutions by varying the size of the roof tiles. The *fractal dimension* is then found from the gradient of a *log-log* plot of surface area against tile size. Care must be taken with this method to ensure that the ratio between surface *height* and tile size is fixed between measurements from different image regions to guarantee an even comparison. We suggest that this is achieved by pre-normalising the image data with respect to grey levels and matrix size.

The *blanket* method employs a curve measuring algorithm developed by (MANDELBROT, 1982) to measure the surface area of the image region being analysed. Like the box method, this method calculates the *fractal dimension* from the gradient of the *log-log* plot of surface area against resolution. Because the surface area is calculated from the difference in volume measurements at different scales, this method is not sensitive to variations in the ratio between pixel size and pixel thickness. However, we recommend that image regions being compared be normalised with respect to grey levels and matrix size to avoid computational errors.

Because the *log-log* plots from both the *box* and *blanket* methods contain non-linear parts, care must be taken to measure the line gradient from the linear part only. An alterative method suggested by (PELEG *et al*, 1984) was to measure the gradient over a three data point window that moves along the plot. The fractal signature that is calculated from the resulting gradient function can then be used as a unique feature vector for characterisation.

The third method of calculating the *fractal dimension* of an image region derives from its relationship with the power spectral density function. The *dimension* is related to the attenuation coefficient of frequency that defines the power spectrum. A method of extracting this information has been realised as part of this work.

An innovative application of fractal geometry to MRI is discussed in **section 6.6**. The primary aims of this method are to characterise segmented large tissue structures (ie. liver, lung etc.) by the *fractal dimension* of their outlines. This enables automatic selection of structures for specific tissue analysis eg pathology characterisation using texture analysis in the liver. With advanced segmentation techniques outside the bounds of this work, *all* tissue types may be identified and completely characterised by examining the fractal properties of their outlines. This method has not yet been implemented because its success depends on the efficacy of the fractal methods described in this chapter, the results of which will only be known in the final stages of this project.

Although the three methods of calculating the *fractal dimension* described in this chapter are, in theory, measuring the same parameter, it is clear from the different approach taken by each method that their values will be different. Therefore, there is no definitive *fractal dimension*. It is a relative figure; one tissue will have a higher value than another, indicating that one tissue has more jagged texture than another. Ideally there will be correlation between the different methods, but that cannot be guaranteed. We recommend that the *fractal dimension* from all three methods be used in a 3D feature vector from which tissue can be characterised. In order such tissue characterisation can be carried out, a large library of tissue types and their fractal dimensions for the three methods needs to be built up.

CHAPTER 7 - EXPERIMENTAL DETAILS

7.1 INTRODUCTION

This chapter provides some of the important experimental details that enabled the texture analysis methods discussed in the previous three chapters to be examined. The first section is dedicated to a description of the data collection process using the *MAIVIS* image visualisation package. This account includes details of how texture features are extracted from image Region Of Interest (ROI), and a summary of the various image pre-processing options available such as normalisation. A more detailed account of the data normalisation methods employed as part of this project is given in **section 7.3**.

Once the texture features have been extracted from the image data using *MAIVIS*, this data is exported to a PC platform where the discriminative performance of each method is assessed. Details of this process and the extensive use of spreadsheets are discussed in **section 7.4**.

The final section in this chapter provides information about the MRI image data sets and texture samples used in the studies carried out as part of this project. This section includes a description of each data set and a sample image with the relevant texture classes clearly indicated.

7.2 DATA COLLECTION

There are three texture options on the main *MAIVIS* menu structure, a button for each of the methods investigated; statistical, transform and fractal. Under each texture menu option there are two functions; one to edit the parameters that control the feature extraction process and one to extract the texture features. The edit function uses one of the image windows to display the current setting for each parameter and the dialogue window for changes. This function also enables the user to select the texture features to be extracted. Apart from the user selecting the image to process, the feature extraction function can be completely guided by the control parameters.

The three most influential control parameters are the normalisation mode, the number of grey levels that the image data is represented by (depth), and the size of the ROI from which the texture feature is extracted. There are three normalisation modes that can be chosen; (0) none, (1) histogram equalisation and (2) statistical moments. The mechanics of methods (1) and (2) are described in the next section of this chapter. The permitted range of grey level depth is 4 to 8 bits. The software automatically fixes the grey level depth after the image data has been normalised.

Any number up to thirteen ROI can be defined on the image being investigated. The size and position of the ROI can either be defined numerically or by freehand. The ROI size control parameter determines which ROI can be investigated in the selected image. There are some cases where this parameter is limited because of the nature of a particular texture feature. Another parameter controls how the image ROI are chosen; either individually, or all those of appropriate size. The later option enables large amounts of data to be collected with the minimum of intervention.

The output of the texture analysis process is directed by the file status control parameter. This data can either be directed to a text window or to a text file, in which case the filename parameter must be set. The latter option enables the texture feature data to be exported for analysis. **Figure 7.1** below gives an example extract of a typical data file. A header precedes the texture features extracted from each ROI in each image. In this example the data is extracted from ROI number '0' of size '12x12' from image filename (including UNIX path) '198_559_3/I.035'. Two fractal ('Frac') features have been calculated from 8-bit data ('Bits:8') without any normalisation ('Norm:0'). The comment 'Blnk-pts:7' in the header refers to the number_of_blanket_points control parameter.

Figure 7.1 : Extraction from an exported texture feature file

**** Image:198_559_3/I.035 ROI:0 Matrix:12x12 ****	
**** Features:Frac Blnk-pts:7 Norm:0 Bits:8 ****	
Fractal tile FD	2.35737
Fractal blanket FD	3.08742
**** Image:198_559_3/I.035 ROI:1 Matrix:12x12 ****	
**** Features:Frac Blnk-pts:7 Norm:0 Bits:8 ****	
Fractal tile FD	2.34421
Fractal blanket FD	3.06115

Many of the texture analysis functions produce graphical output in addition to the texture feature value. The graph control parameter provides the user with the option of either using the default output windows or to specifically designate output windows. This parameter can also be set to disable graphical output.

In addition to the generic control parameters, all three texture analysis methods require specific parameters for their function. An example of this is the number_of_blanket_points parameter ('Blnk-pts:7') shown in the header in **figure 7.1**.

7.3 DATA NORMALISATION

Normalisation is the process by which the first order statistics of image data are standardised. This process can be used to eliminate the differences in average image intensity between image samples. It was suggested by (LERSKI and STRAUGHAN, 1987) that a study of the effects of image data should be performed. The two methods of data normalisation they suggested have been employed in the texture analysis studies carried out as part of this project; (1) histogram equalisation and (2) statistical moments. In order that images from different sources (ie. differently acquired) may be compared on equal terms, more sophisticated normalisation methods are required, beyond the scope of this project.

The histogram of image intensity value gives rise to the probability density function, from which the probability of occurrence of each grey level can be found. The histogram equalisation method works on the principle that the first order image statistics can be standardised by evenly redistributing the histogram so that the probability of occurrence of each grey level is the same. The problem with this method comes from the practical difficulties associated with splitting the histogram entry, or bin, for each grey level. Information is needed about the neighbourhood of the corresponding pixels so that a decision can be made as to which pixels should have their grey level changed. The easiest solution to this problem is to move the bins and not to split them up. The bins should be moved in such a way as to evenly distribute the histogram ie. the probability density function. This is a standard fix for histogram equalisation (GONZALEZ and WINTZ, 1987) employed by many applications. The algorithm used to implement this method was written for *MAIVIS* by (FREE, 1994). An example of histogram equalisation is given in **figures 7.2** and **7.3**. Both figures are shown to the same scale, with grey level on the x-axis and frequency on the y-axis.

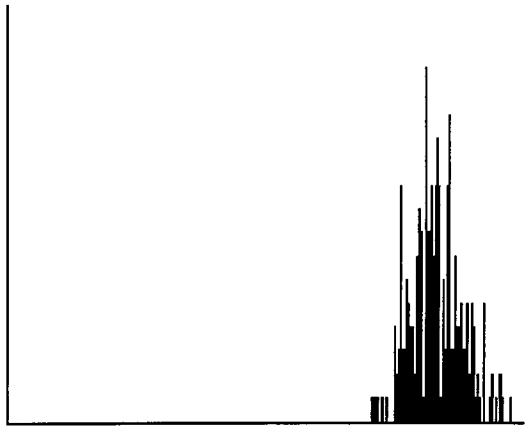


Figure 7.2 : Histogram from ROI

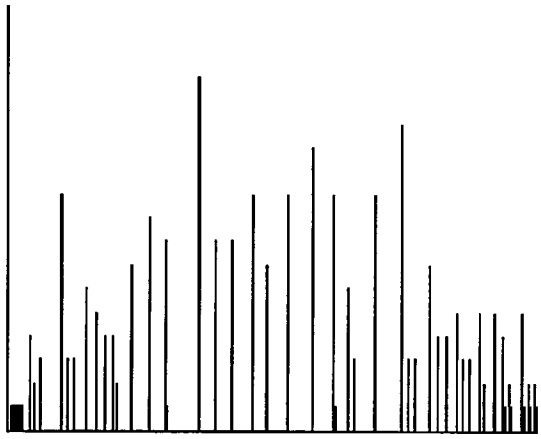


Figure 7.3 : Equalised histogram from ROI

The second method of data normalisation employed in the texture analysis study works by fixing the first two statistical moments (mean and standard deviation) of the image data to some constant values. Although this method is basic in principle, there are many practical problems associated with its implementation. An algorithm to perform this tailored normalisation has been developed as part of this project.

Equation 7.1 shows how each image pixel is normalised, where F(x,y) is the pixel in the original image with mean and standard deviation (μ, σ), and N(x,y) is the pixel in the normalised image with fixed mean and standard deviation (μ_{const}, σ_{const}). The fixed mean is chosen to be in the middle of the range of grey levels ie. equal to the number of grey levels divided by two, **equation 7.2**. In a gaussian distribution, all the values are contained within $\pm$ 2 standard deviations either side of the mean. To accommodate points that are outside this range of values, the fixed standard deviation is chosen so that there are $\pm$ 3 standard deviations either side of the fixed mean, **equation 7.3**.

$$N(x,y) = \frac{\sigma_{const}}{\sigma}(F(x,y) - \mu) - \mu_{const} \quad (7.1)$$

$$\mu_{const} = \frac{grey_levels}{2} \quad (7.2)$$

$$\sigma_{const} = \frac{grey_levels}{6} \quad (7.3)$$

During testing of this method, several of the pixels in the normalised image were found either to be negative or to overflow beyond the maximum permitted value. These pixels correspond to points that are at either end of the original distribution, a large distance from the mean. The solution to this problem is either to make the fixed standard deviation very much smaller than the whole range of grey levels, or to shift the points corresponding to these pixels along the normalised image distribution back in range ie. closer towards the mean. The former solution must be disregarded because of the image averaging effects that become greater as the relative size of the standard deviation is reduced.

Two methods of shifting the points along the normalised image distribution have been investigated. The first method assigns these 'out of range' points to equal either the minimum or maximum of the distribution, before re-normalising the data. This process is repeated until all the points are within range. This method was found to be very unstable and did not reliably converge to a solution. A more successful solution was to assign the neighbourhood average to each image pixel corresponding to an 'out of range' point on the distribution, before re-normalising the data. This process is also repeated until all the points are within range.

7.4 DATA ANALYSIS AND PERFORMANCE EVALUATION

The texture feature data extracted using *MAIVIS* for each study is first exported from the workstation to a Personal Computer (PC) platform before the discriminative performance analysis is carried out. The PC was preferred to the workstation for this purpose for purely practical reasons. All the analysis was carried out using Microsoft EXCEL spreadsheets.

The exported texture feature files contain delimited data such that when they are imported into EXCEL, the headers, feature titles and feature values are correctly separated into the cells of the spreadsheet. After the data has been imported into the spreadsheet, it is mapped into tables from which discriminative performance can be measured eg. tables with texture features down the columns and ROI for different tissue classes along the rows. A large diversity of spreadsheets have been developed for this purpose.

Three methods of measuring the discriminative performance of texture features are considered in this section: the Student's t-test and the Fisher linear classifier (both standard methods), and the DB performance measure (developed as part of this project).

The Students t-test is commonly used to establish the discriminant performance of texture features. This function determines the probability that two data samples have come from two parent populations with the same mean. The data samples correspond to the features calculated for the two classes of texture. The result of the Student's t-test is a number between zero and one. A high value indicates that the features calculated for the two classes of texture are very similar, and therefore the feature is a poor discriminator. A low value, however, indicates the feature is able to discriminate well between the two texture classes. The typical range of this measure (p) was demonstrated by (PRENDERGAST *et al.*, 1987) for MRI tissue characterisation of benign hypertrophy and carcinoma in the prostate. The two tissue classes were not separable with texture features that tested $p>0.05$, but were significantly separated with those that tested $p<0.001$.

The Student t-test measure has been calculated using the Microsoft Excel function library. The possibility of using this measure in the performance overview of the texture features examined as part of this project was explored briefly. This prospect was not considered because a considerable amount of difficulty was experienced in identifying the threshold of satisfactory performance. Values from the Student's t-test are, however, quoted in the discussion that takes place in **chapter 11** to authenticate the results obtained with the DB performance measure.

The DB performance measure was proposed as part of this project as a much simpler alternative to methods such as the Student's t-test. This method has proved to be a good indicator of discriminative behaviour, and was used in the overview of texture feature behaviour presented in **chapters 8** to **10**.

The conditional probability distribution P_i associated with each texture class i is defined in **equation 7.4**, where $f(x)$ is the texture feature value, and ω_i the class. The mean μ_i and standard deviation σ_i associated with this distribution are defined in **equations 7.5** and **7.6**. The DB performance measure is defined in **equation 7.7**, where σ_1 an σ_2 are the standard deviations of the two distributions, and $\sigma_{1,2}$ is the standard deviation of the joint distribution for both classes.

$$P_i = P(f(x)|\omega_i) \qquad (7.4)$$

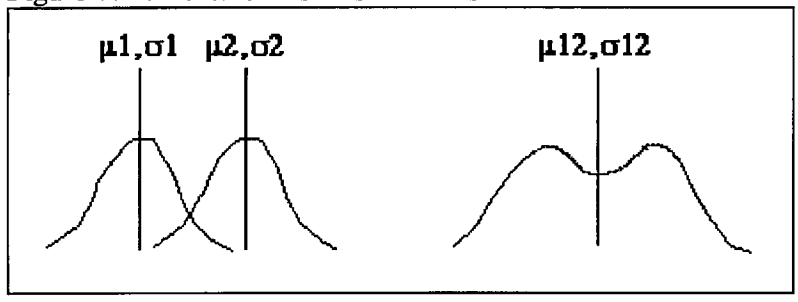
$$\mu_i = E[f(x)|\omega_i] \qquad (7.5)$$

$$\sigma_i^2 = E[(f(x) - \mu_i)^2 | \omega_i] \qquad (7.6)$$

$$DB_{1,2} = \frac{\sigma_{1,2}}{\sigma_1 + \sigma_2} \qquad (7.7)$$

The behaviour of the DB measure is illustrated in the description of the probability distributions in **figure 7.4**. The left side of the figure shows the two texture class probability distributions, and the right side shows the joint probability distribution of the two classes. If the standard deviations of the two distributions remain fixed and the means become further apart, the ratio of the joint distribution standard deviation to the sum of the separate class standard deviations becomes greater, and therefore better DB performance is noted. If the means of the two distributions remain fixed and the standard deviations become smaller, then again this ratio becomes greater, and better performance is measured.

Figure 7.4 : Texture class distributions



The feature performance studies described in **chapters 8** to **11** are carried out on image data containing three tissue classes. The three_class DB performance measure used in these studies is defined in **equation 7.8**, where σ_1, σ_2 and σ_3 are the standard deviations of the separate class distributions and $\sigma_{1,2,3}$ is the standard deviation of the joint distribution of all three.

$$DB_{1,2,3} = \frac{\sigma_{1,2,3}}{\sigma_1 + \sigma_2 + \sigma_3} \qquad (7.8)$$

This three_class DB measure is minimised when the tissue classes appear the same to the texture feature. If all three tissue class distributions come from the same parent population, then the standard deviations of the individual distributions and also the joint distribution will be the same. The theoretical minimum value of the DB measure is therefore 1/3. DB values approaching this limit were seen in the performance experiments reported in

chapters 8 to **10**. A DB measure threshold of 0.5 was also identified as a result of these experiments, indicating the limit of satisfactory discriminative performance.

A similar measure to the DB measure is the Fisher linear classifier (THERRIEN, 1987). This parameter is based upon the difference of the probability distributions of the two texture classes. The Fisher classifier is defined as the mean squared over the variance of this new distribution, as shown in **equation 7.9**. This method has not been implemented as part of the performance overview, but has been explored in the final comparison presented in **chapter 11.**

$$F = \frac{\mu_{diff}^2}{\sigma_{diff}^2} = \frac{(\mu_1 - \mu_2)^2}{\sigma_1^2 + \sigma_2^2} \qquad (7.9)$$

In addition to studies of the discriminant behaviour of the texture features, a great deal of information about the behaviour of features can be found from studies of the correlation between them. Comprehensive correlation studies are carried out between all the features in each of the three groups examined as part of this study. High correlation between features means that one or many of the features in the group are redundant and can be eliminated. The correlation coefficient for any two texture values is calculated are using the Microsoft Excel function library. The correlation is measured between the values obtained from all tissue classes for both texture features.

7.5 IMAGE DATA SETS

The final section in this chapter provides details of the MRI image data sets and texture samples used in the studies carried out as part of this project. **Table 7.1** provides a description of the seven sets of clinical data used, including acquisition details and the number of slices examined. The tissue classes defined in each of these data sets are identified in **table 7.2** along with the figure number of the corresponding image example. The numbered boxes found on the images in **figures 7.5(a)** to **7.6(c)** mark the tissue classes investigated. Texture features are calculated from carefully placed ROI in these tissue regions in the discriminant performance studies carried out as part of this project. Extreme

care is taken that the ROI are consistently placed in these regions both within each image and through any series of slices examined. On average 26 ROI are defined in each tissue sample investigated.

Table 7.1 : Clinical data sets (256x256 images, 12 bit)

Data set	Plane	Sequence, TR/TE (ms)	Slice_width, Gap (mm)	FOV (mm)	Slices	Image type
Thigh	TRA	SE, 2800/30	8, 2.5	240	3	Proton density
Abdomen	TRA	SE, 600/12	8, 4	400	3	T1-weighted
Leg104	COR	SE, 700/12	5, 2	360	2	T1-weighted
Leg192	SAG	SE, 600/12	5, 2	360	2	T1-weighted
Brain1	TRA	SE, 600/20	4, 1	250	3	T1-weighted
Brain2	TRA	SE, 600/20	4, 1	250	3	T1-weighted
Brain3	TRA	SE, 600/20	4, 1	250	3	T1-weighted
Key: TRA = Transverse, SAG = Sagittal, COR = Coronal, SE = Spin Echo						

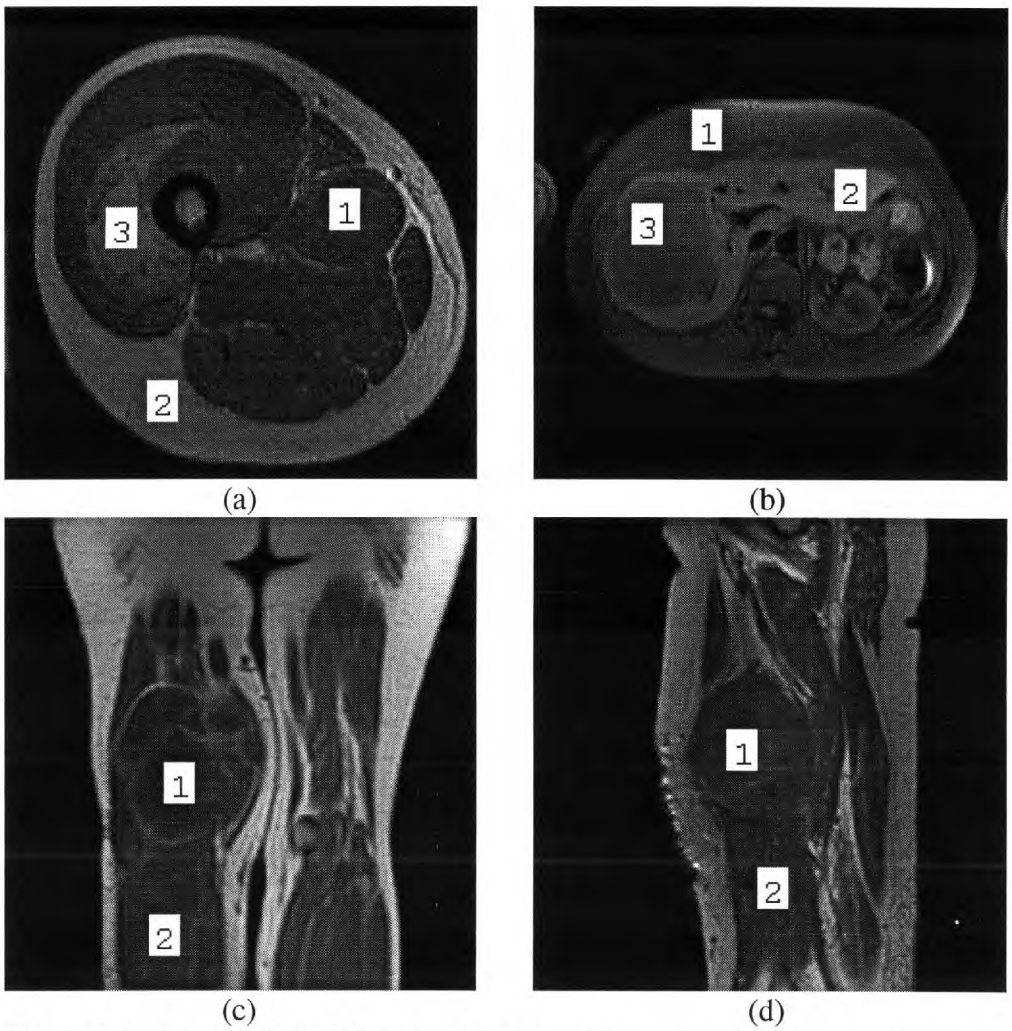
Table 7.2 : Clinical tissue classes

Data set	Tissue 1	Tissue 2	Tissue 3	Example
Thigh	Muscle	Fat	Tumour	Figure 7.5(a)
Abdomen	Fat	Liver	Tumour	Figure 7.5(b)
Leg104	Tumour	Muscle	-	Figure 7.5(c)
Leg192	Tumour	Muscle	-	Figure 7.5(d)
Brain1	Grey matter	-	-	Figure 7.6(a)
Brain2	Grey matter	-	-	Figure 7.6(b)
Brain3	Grey matter	-	-	Figure 7.6(c)
It has not been necessary to clinically identify any of the tumours				

In addition to the clinical image data sets, several texture sample images have been used in the studies carried out as part of this project. Although these images have mainly been used in the development of texture feature algorithms, they have also proved to be very valuable test data. A description of this set of four images is given in **table 7.3**, along with the figure number. These images were digitised from the (BRODATZ, 1966) book of photographs of natural textures.

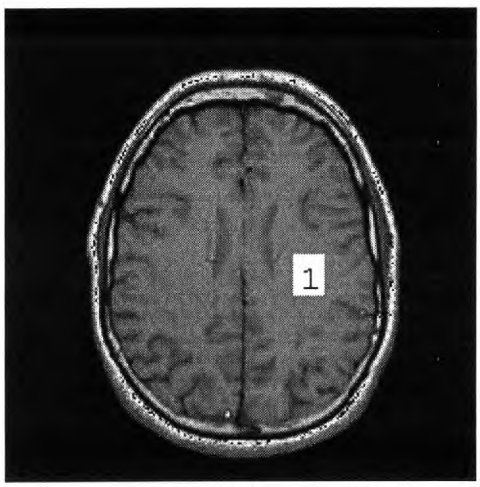
Table 7.3 : Non-clinical data sets (256x256 images, 8-bit digitised)

Data set	Slices	Example
Cloud	1	Figure 7.7(a)
Granite	1	Figure 7.7(b)
Wood	1	Figure 7.7(c)
Water	1	Figure 7.7(d)

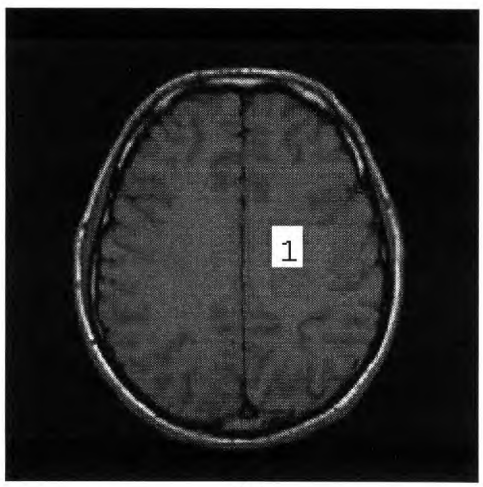


(a) (b) (c) (d)

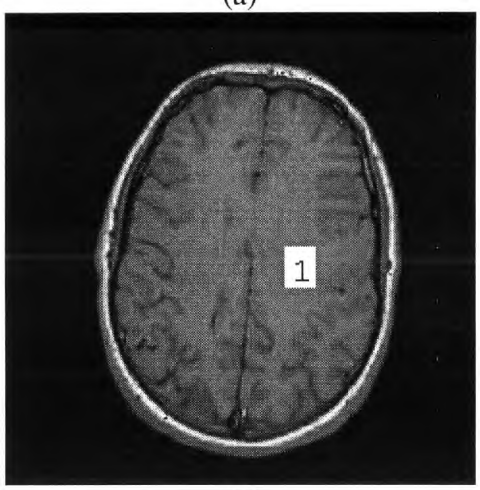
Figure 7.5: General clinical data sets: (a) Thigh (b) Abdomen (c) Leg104 (d) Leg192



(a)



(b)



(d)

Figure 7.6: Brain clinical data sets: (a) Brain1 (b) Brain2 (c) Brain3

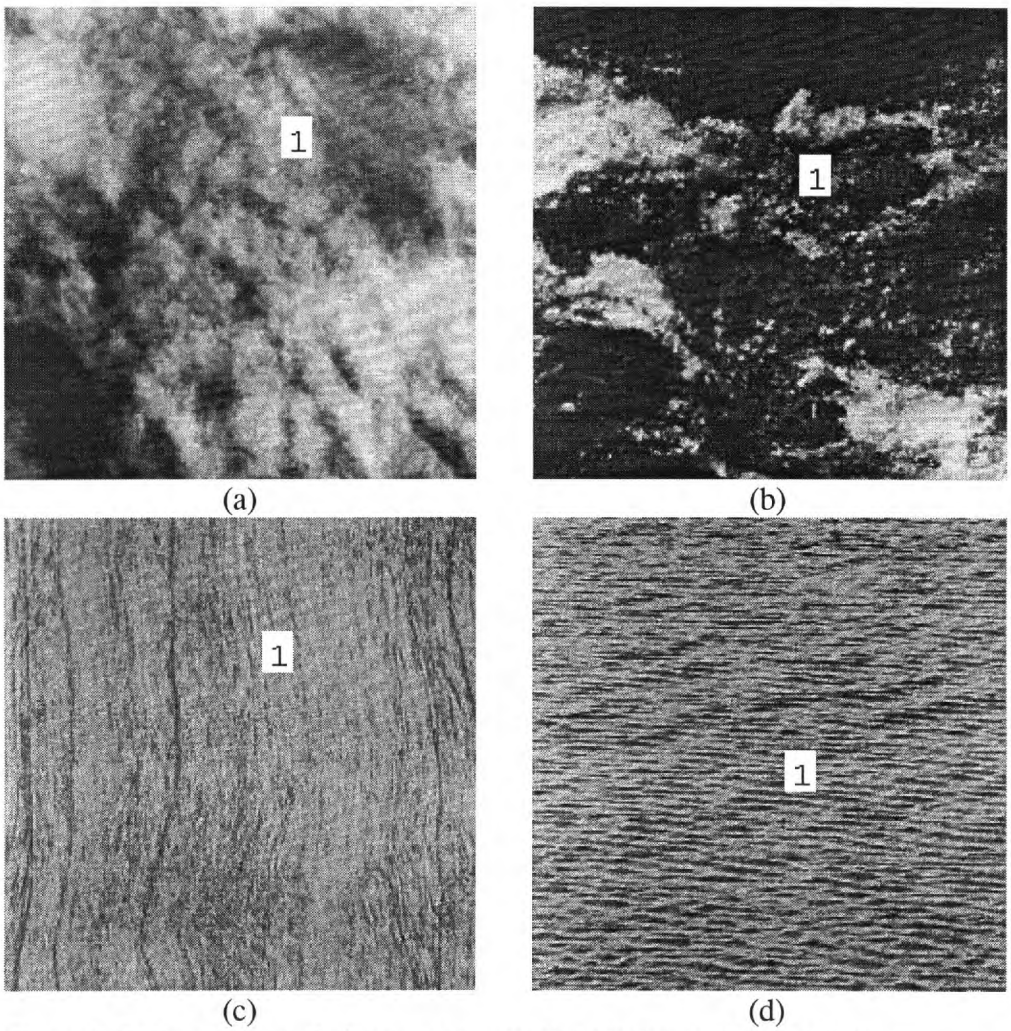


Figure 7.7: Texture sample images sets: (a) Cloud (b) Granite (c) Wood (d) Water

CHAPTER 8 - RESULTS: STATISTICAL TEXTURE FEATURES

8.1 INTRODUCTION

This chapter presents a study of the four approaches to extracting statistical texture information described in **chapter 4**; first order statistics, run length analysis, grey tone difference and co-occurrence. A summary of the texture features derived from each of these methods is given in **table 8.1**. This table provides a reference between the features and the keys used in the graphs and tables presented throughout this chapter.

The results presented in this chapter are clearly sub-divided into six main sections. The first four sections (**sections 8.2** to **8.5**) are devoted to the performance evaluation of each of the four groups of statistical texture features. Each of these four sections contains a performance overview, a study of correlation, and a description of the optimised feature set. The fifth section of this chapter (**section 8.6**) provides a comparison (including correlation) between the optimised features sets derived for each of the four methods, while the sixth section (**section 8.7**) gives examples of their use for discrimination between tissue classes in clinical MRI images.

The discriminative performance tests are carried out on several transverse MRI images of the thigh (see 'thigh' image data in **chapter 7**). These images contain three distinct classes of tissue; (1) muscle, (2) fat and (3) tumour. The discriminative behaviour of each feature is quantified using the DB performance measure defined in **chapter 7** (based on the relationship between the statistics of tissue classes). Unless stated, the DB performance measure is calculated for the discrimination between all three tissue classes (1-3).

The regions of interest (ROI) in the image data are standardised by fixing the number of grey level bits or by implementing some normalisation scheme. This process is carried out before any texture information is extracted. The pre-processing options are indicated by 'B' for the number of grey level bits (in the range of 4-8), and 'N' for the normalisation scheme; (1) none, (2) histogram equalisation, and (3) statistical moments. A detailed description of normalisation schemes (2) and (3) are given in **chapter 7**.

The results of the performance studies for each set of features attempt to demonstrate the effect of data pre-processing on the discriminative behaviour of each texture feature. Consideration is

also given to the influence of ROI size. The purpose of this study is to identify the most robust and best performing features, and the conditions under which they perform optimally.

Table 8.1 : Statistical texture features

Statistical method	Texture features	Key
First order statistics	Mean	M1
	Standard deviation	M2
	Skewness	M3
	Kurtosis	M4
	Coarseness	M5
Run length analysis	Vertical RL: short run emphasis	VSRE
	Vertical RL: long run emphasis	VLRE
	Vertical RL: grey level non-uniformity	VGLN
	Vertical RL: run length non-uniformity	VRLN
	Vertical RL: run percentage	VRPC
	Horizontal RL: short run emphasis	HSRE
	Horizontal RL: long run emphasis	HLRE
	Horizontal RL: grey level non-uniformity	HGLN
	Horizontal RL: run length non-uniformity	HRLN
	Horizontal RL: run percentage	HRPC
	Cluster size: short run emphasis	CSRE
	Cluster size: long run emphasis	CLRE
	Cluster size: grey level non-uniformity	CGLN
	Cluster size: run length non-uniformity	CRLN
	Cluster size: run percentage	CRPC
Grey tone difference (multi-directional)	Contrast	GT1
	Second angular moment	GT2
	Entropy	GT3
Co-occurrence (multi-directional)	Second angular moment	F1
	Contrast	F2
	Correlation	F3
	Variance	F4
	Inverse difference moment	F5
	Sum average	F6
	Difference average	F6′
	Sum variance	F7
	Sum entropy	F8
	Entropy	F9
	Difference variance	F10
	Difference entropy	F11
	Information measure	F12
	Information measure	F13

The correlation study in each section enables the relative behaviour of the statistical texture features to be investigated. Similar behaviour would be expected from highly correlated features. When two features are highly correlated, therefore, one feature can be discarded without any

significant loss of information. Large numbers of texture features can be reduced to more manageable subsets by examining the correlation between each and every feature, and eliminating redundant features.

The tissue characterisation examples in the final section of this chapter (**section 8.7**) provide an illustration of the efficacy of the optimised statistical texture features with a number of sets of clinical images containing a variety of different tissue types.

8.2 FIRST ORDER STATISTICAL TEXTURE FEATURES

8.2.1 DISCRIMINATIVE PERFORMANCE OVERVIEW

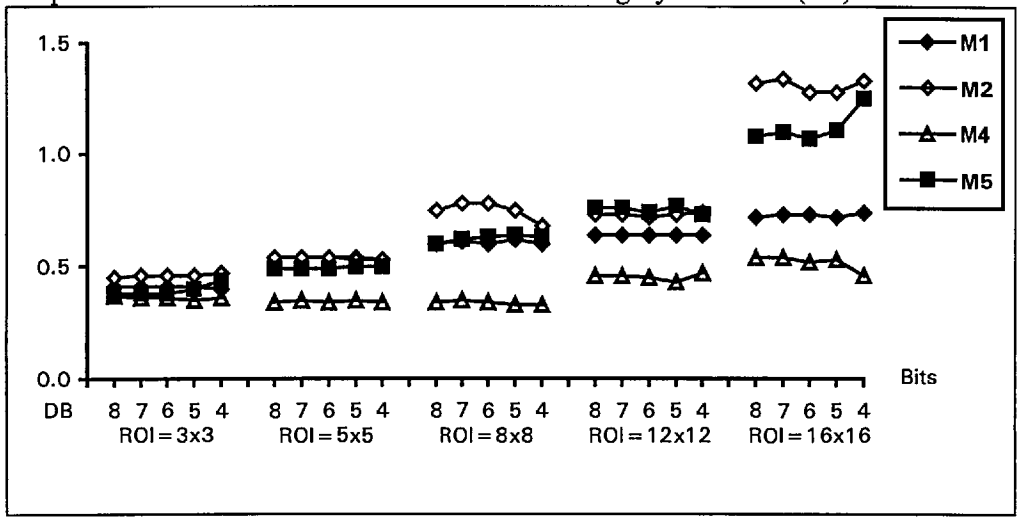
The plots in **graphs 8.1** to **8.4** indicate the discriminative performance trends of the first order statistical texture features extracted from ROI under different pre-processing conditions. For each point on the graph, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (ROI size and either bits or normalisation scheme). The 'Thigh' MRI image data used in this overview contains three distinct classes of tissue. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions (muscle, fat & tumour) on several of images.

Graphs 8.1 and **8.2** show the variation of the DB performance measure with the number of grey level bits (4-8) assigned to the image data by the pre-processing for the five parameters investigated. This behaviour is examined for a series of different ROI sizes. The data used in these graphs was acquired without normalisation (N0).

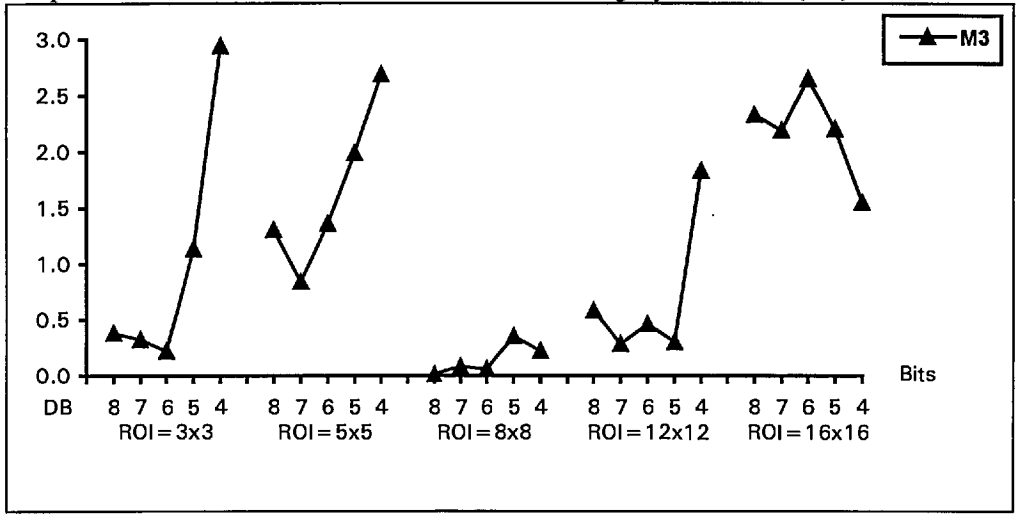
The overall performance of statistical texture features M1, M2, M4 and M5 (see **table 8.1**) without normalisation is illustrated in **graph 8.1**. The most dramatic influence on performance is ROI size. Optimum performance was achieved with ROI of 16x16 for all four features. The effect of varying the number of data bits proved not to be a significant factor in discriminative performance. The performance of feature M3 (see **table 8.1**) is presented separately in **graph 8.2** to draw attention to its less predictable behaviour. Significant improvements in the performance of this feature are seen with a reduction in the

number of data-bits for ROI sizes 3x3, 5x5 and 12x12. This behaviour, however, is not demonstrated with ROI sizes 8x8 and 16x16.

Graph 8.1 : Behaviour of first order features with grey level bits (N0) #1



Graph 8.2 : Behaviour of first order features with grey level bits (N0) #2

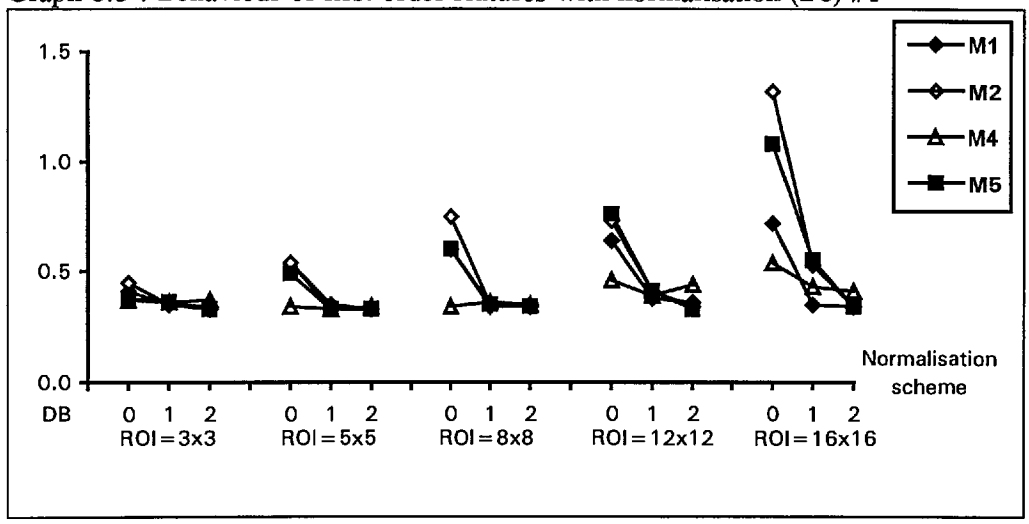


Graphs 8.3 and **8.4** show the variation of the DB performance measure with normalisation scheme (0-2) for the five parameters investigated. This behaviour is also examined for a series of different ROI sizes. The values used in these graphs were acquired with 8-bit image data (B8).

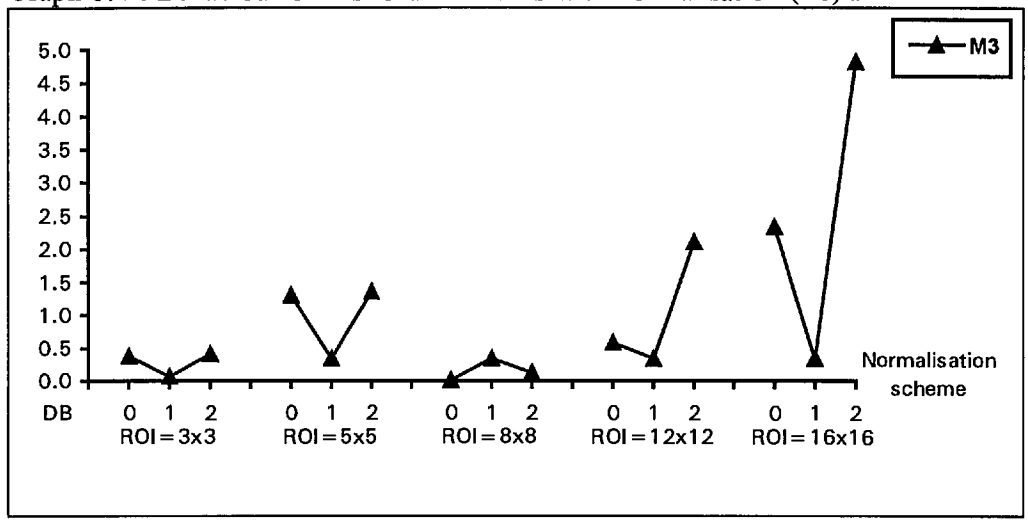
Both normalisation scheme (1) and (2) standardise the first order image statistics by fixing the distribution mean and standard deviation to arbitrary values. Statistical texture features M1 (mean) and M2 (standard deviation) are therefore not useful when normalisation has

first been applied to the image data (**graph 8.3**). Feature M5 is equally unusable under these circumstances because it is a function of standard deviation. The behaviour of kurtosis (M4) is less affected by the normalisation process because it only causes subtle changes in distribution shape.

Graph 8.3 : Behaviour of first order features with normalisation (B8) #1



Graph 8.4 : Behaviour of first order features with normalisation (B8) #2



The effect of data normalisation on the skewness (M3) depends on the scheme adopted as shown in **graph 8.4**. The effect of applying normalisation scheme (1) (a histogram equalisation method) is generally detrimental to the DB performance measure. This effect can be explained by the distortion of the symmetry of the distribution about its mean caused by the non-linear re-distribution of histogram elements. Normalisation scheme (2) shifts and scales the distribution to an arbitrary mean and variance. This process should not

affect the skewness of the distribution. Equal or better performance is achieved with this method over no normalisation.

8.2.2 CORRELATION STUDY

In order to establish the relationship between the five first order statistical texture features, a study of the correlation between them has been carried out. Redundant less well performing features can be eliminated from our features set using the results of this study.

The correlation between each pair of texture features is measured over 78 ROI evenly distributed among the three tissue classes on 'Thigh' image data. This is the same source of data used for the discriminative performance evaluation in the overview in **section 8.2.1**. The texture feature are collected from ROI of 16x16 with 8-bit data and without normalisation. These conditions are optimum for all features with the exception of M3. The performance of M3 can be further improved by applying normalisation scheme (2). The correlation coefficient is calculated using Microsoft EXCEL (see **chapter 7** for details), varying between 0 (no correlation) and 1 (high correlation). A negative sign implies negative correlation.

Table 8.2 : First order feature correlation

	M1	M2	M3	M4	M5
M1	-	-0.81	-0.40	0.21	-0.82
M2	-0.81	-	0.17	-0.08	0.99
M3	-0.40	0.17	-	-0.84	0.16
M4	0.21	-0.08	-0.84	-	-0.05
M5	-0.82	0.99	0.16	-0.05	-

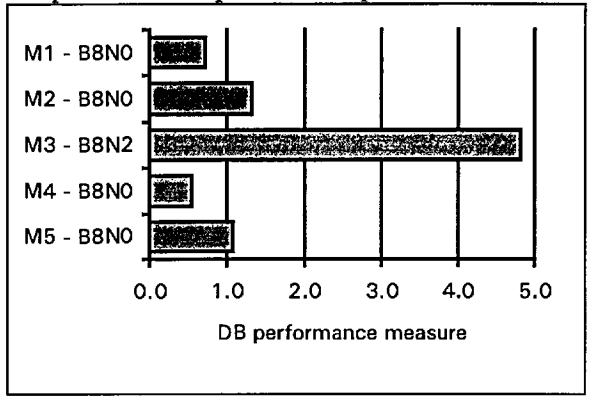
Features M2 (standard deviation) and M5 (coarseness) are highly correlated because M5 is derived from the standard deviation. Feature M1 (mean) is well correlated with both M2 and M5. Very little correlation exists between M2 and both M3 and M4, while the correlation between the latter pair is quite high. Based on the correlation between each of these five features, the feature set could be reduced to just two (M2 and M3) with minimal loss of texture information.

8.2.3 OPTIMISED FEATURE SET

The number of data bits was not found to be a significant factor in the discriminant performance of any of the five first order statistical texture features, and is arbitrarily taken to be 8. With the exception of feature M3, no improvement in performance was noted with the use of normalisation. Performance was found to improve with ROI size for all features. A comparison of the relative performance under optimal conditions for each feature is presented in **graph 8.5**.

High correlation was found between the following pairs of features; M2 & M5, M1 & M2, M1 & M5, and M3 & M4. **Graph 8.5** shows that M2 performs better than M1 and M5, and that the performance of M3 is several factors better than that of M4. This result enables the three least well performing features to be eliminated. To summarise, the performance and correlation studies indicate that the un-correlated features M2 (standard deviation) and M3 (skewness) are the most useful first order statistical texture features.

Graph 8.5 : Comparison of optimised first order statistical features



8.3 RUN LENGTH ANALYSIS TEXTURE FEATURES

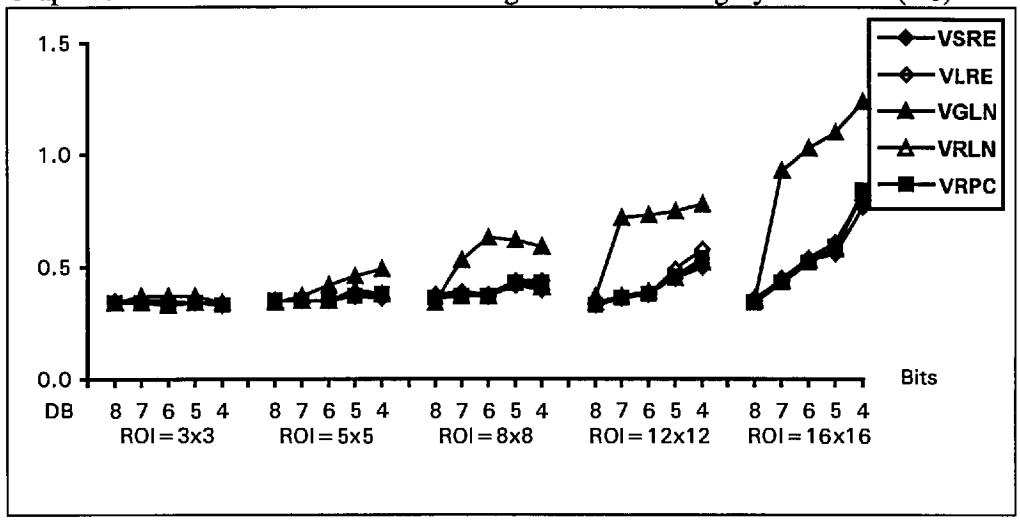
8.3.1 DISCRIMINATIVE PERFORMANCE OVERVIEW

The plots in **graphs 8.6** to **8.11** indicate the discriminative performance trends of the run length analysis texture features extracted from ROI under different pre-processing conditions. For each point on the graph, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (ROI size and either bits or normalisation scheme). The 'Thigh' MRI image data used in this overview contains three

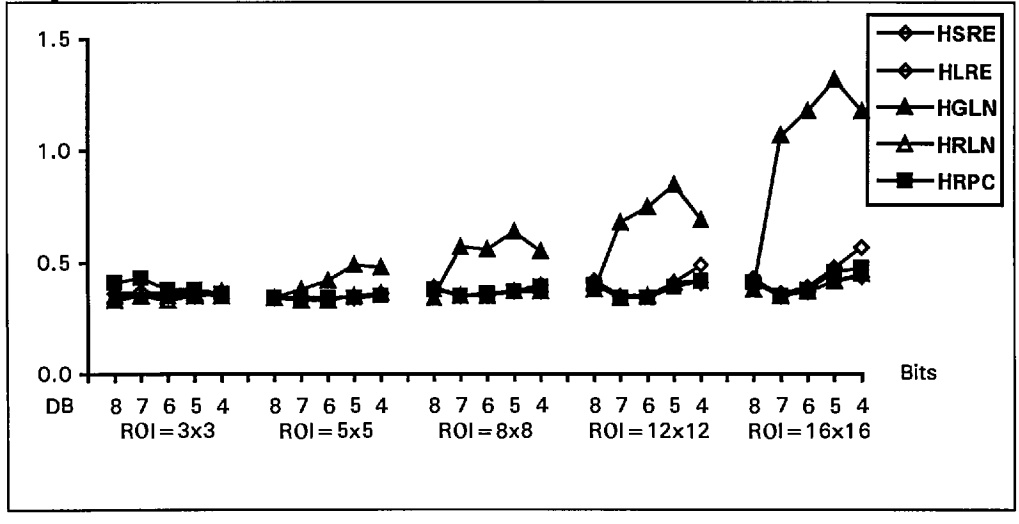
distinct classes of tissue. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions (muscle, fat & tumour) on several of images.

Graphs 8.6 to **8.8** show the overall performance of the vertical run length, horizontal run length and cluster size texture features (see **table 8.1**) without normalisation (N0). All three graphs show the significant influences of ROI size and number of data bits (4-8) on the DB measure of discriminative performance. Optimum performance for all three run length methods was generally achieved with ROI of size 16x16 and with 5-bit data.

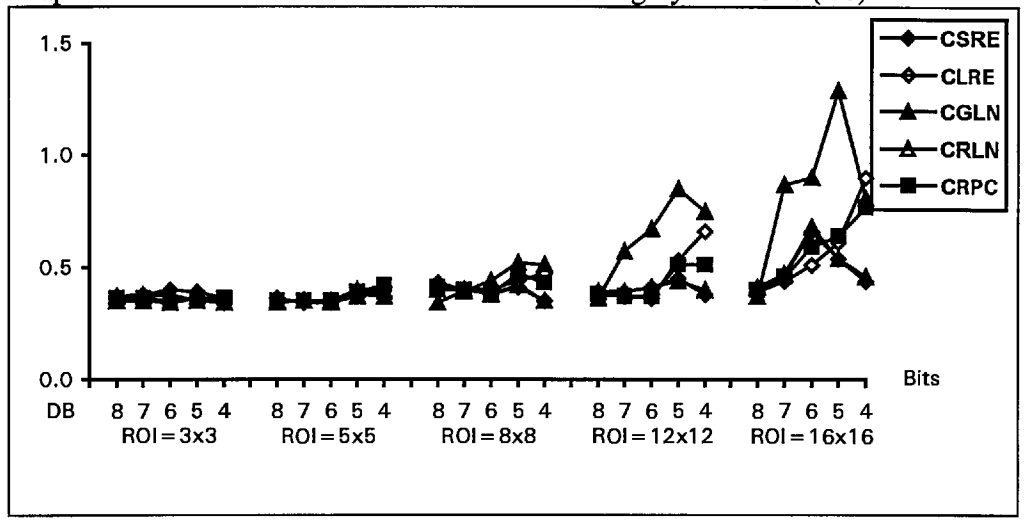
Graph 8.6 : Behaviour of vertical run length features with grey level bits (N0)



Graph 8.7 : Behaviour of horizontal run length features with grey level bits (N0)

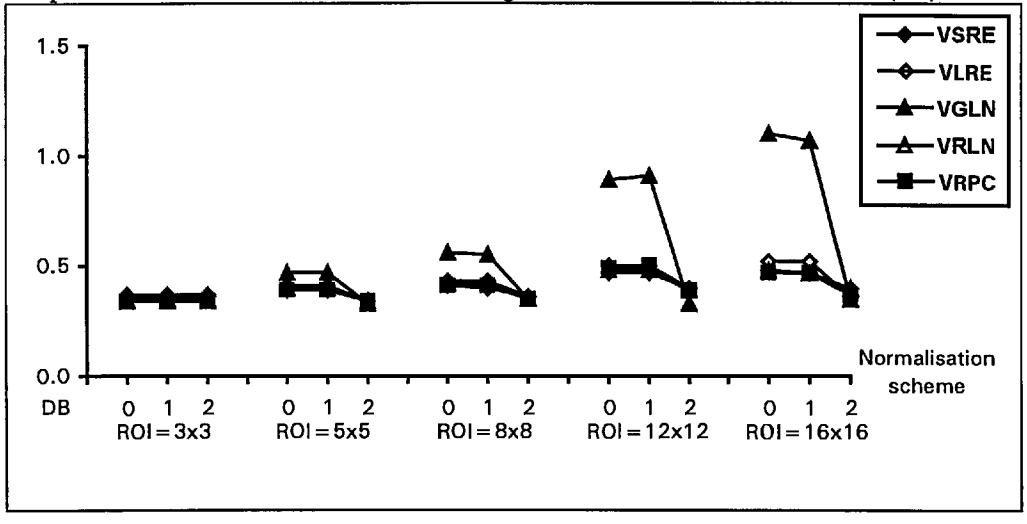


Graph 8.8 : Behaviour of cluster size features with grey level bits (N0)

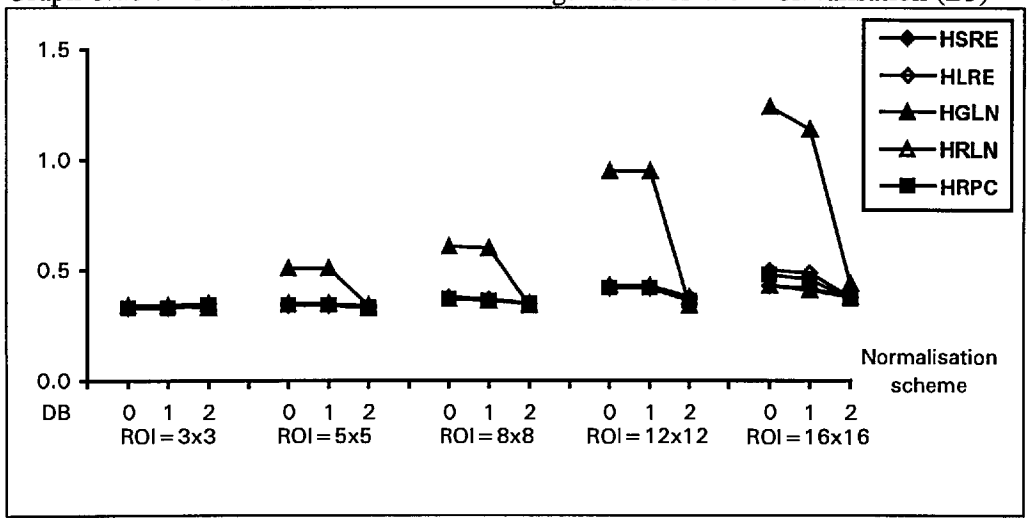


Graphs 8.9 to **8.11** show the variation of the DB performance measure with normalisation scheme (0-2) for each set of run length texture features. This behaviour is also examined over a range of different ROI sizes. The graph values were acquired with 5-bit image data (B5), the optimum bit rate established from **graphs 8.6** to **8.8**. Equal or marginally worse performance was generally achieved with normalisation scheme (1) over no normalisation at all. The effect of normalisation scheme (2) is in all cases, however, significantly detrimental to the performance of the texture features.

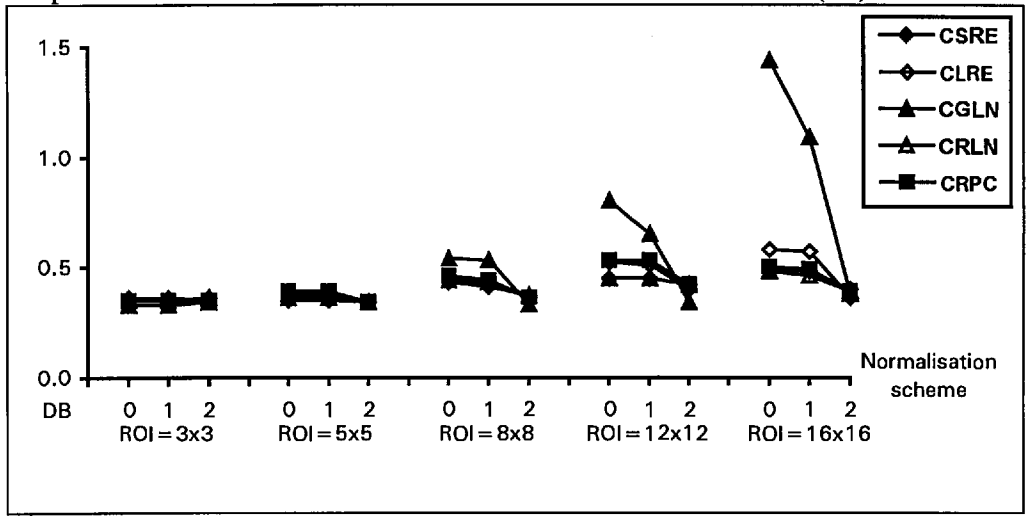
Graph 8.9 : Behaviour of vertical run length features with normalisation (B5)



Graph 8.10 : Behaviour of horizontal run length features with normalisation (B5)



Graph 8.11 : Behaviour of cluster size features with normalisation (B5)



8.3.2 CORRELATION STUDY

In the performance overview described in the previous section, the grey level non-uniformity feature (VGLN, HGLN & CGLN) proved consistently to the best performing feature for all three sets of run length features. The focus of the correlation study is therefore limited to these three texture features.

The correlation between each pair of texture features is measured over 78 ROI evenly distributed among the three tissue classes on 'Thigh' image data. This is the same source of data used for the discriminative performance evaluation in the previous section. The texture feature are collected from ROI of 16x16 with 5-bit data and without normalisation. The

correlation coefficient is calculated using Microsoft EXCEL (see **chapter 7** for details), varying between 0 (no correlation) and 1 (high correlation). A negative sign implies negative correlation.

All three run length features were found to be highly correlated. The correlation coefficients between these feature are given in **table 8.3**. This result indicates that despite the different directions in which the run length distributions are acquired, the three texture features are measuring the same property.

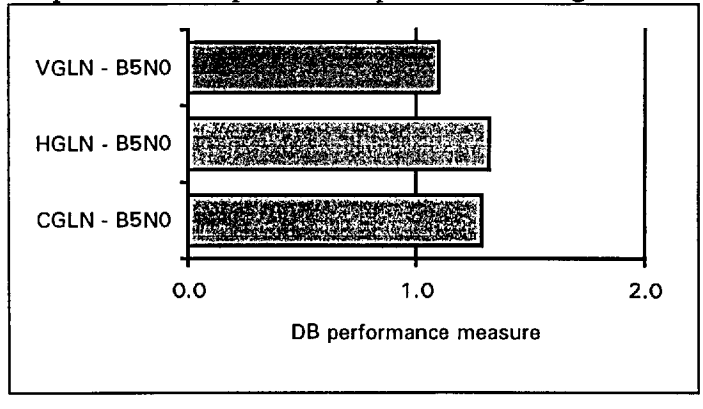
Table 8.3 : Run length feature correlation

	VGLN	HGLN	CGLN
VGLN	-	0.98	0.95
HGLN	0.98	-	0.96
CGLN	0.95	0.96	-

8.3.3 OPTIMISED FEATURE SET

The optimum conditions for acquiring run length texture features (VGLN, HGLN & CGLN) are from ROI of 16x16 with 5-bit data and without normalisation. A comparison of the optimum performance of these highly correlated features, presented in **graph 8.12**, shows little disparity between their discriminative abilities. Texture feature HGLN is arbitrarily chosen as being representative of these three features.

Graph 8.12 : Comparison of optimised run length features



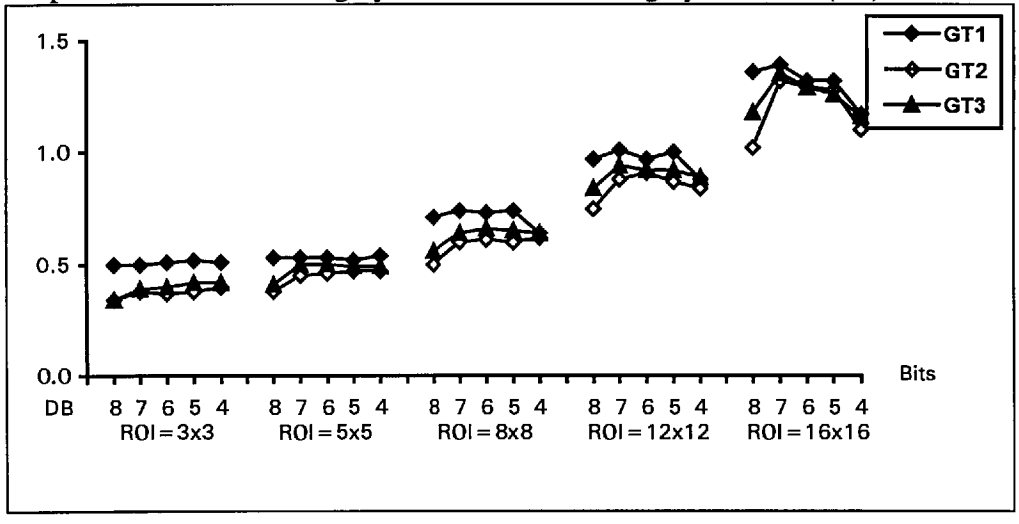
8.4 GREY TONE DIFFERENCE TEXTURE FEATURES

8.4.1 DISCRIMINATIVE PERFORMANCE OVERVIEW

The plots in **graphs 8.13** to **8.15** indicate the discriminative performance trends of grey tone difference texture features extracted from ROI under different pre-processing conditions. For each point on the graph, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (ROI size and either bits or normalisation scheme). The 'Thigh' MRI image data used in this overview contains three distinct classes of tissue. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions (muscle, fat & tumour) on several of images.

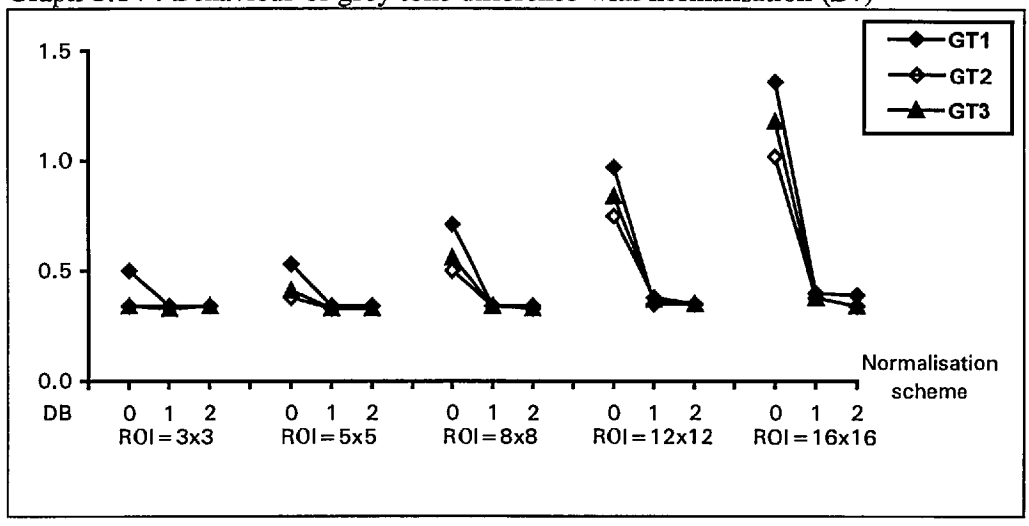
Graph 8.13 shows the variation of the DB performance measure with number of grey level bits (4-8) for the three texture features extracted from a non-directional implementation of the grey tone difference method (see **table 8.1**). This behaviour is also examined over a range of different ROI sizes. The graph data was acquired without normalisation (N0). The overall discriminative ability of these features proved to be significantly influenced by ROI size and number of data bits. Optimum performance was generally achieved with ROI of size 16x16 and 7-bit data.

Graph 8.13 : Behaviour of grey tone difference with grey level bits (N0)

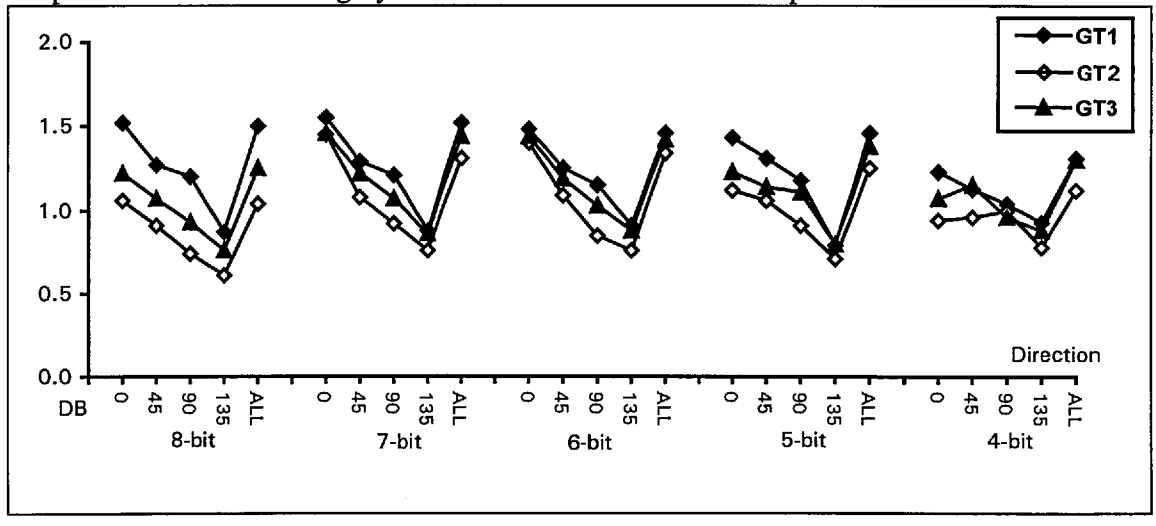


Graphs 8.14 shows the variation of the DB performance measure with normalisation scheme (0-2) for each non-directional grey tone difference texture feature. This behaviour is also examined for a series of different ROI sizes. The graph values were acquired with 7-bit image data (B7), as prescribed by the results obtained from **graph 8.13**. The effect of data normalisation on all three texture features is illustrated in **graph 8.14**. In all cases normalisation was found to be detrimental to the discriminative performance.

Graph 8.14 : Behaviour of grey tone difference with normalisation (B7)



Graph 8.15 : Behaviour of grey tone difference with direction dependence



Graph 8.15 shows the variation of the DB performance measure with the direction in which the grey tone difference displacement vector (one pixel) is applied. The non-directional grey tone difference method (indicated by Dirn=ALL on the graph) incorporates the information from all four directions. The graph values were acquired

without normalisation (N0), with a ROI of 16x16, and for a range of grey level bits (4-8). The overall behaviour of the texture features from both the directionally sensitive and non-directional implementations, however, are comparable. The most important result from this study is that there are no apparent losses of information associated with using the non-directional implementation of the grey tone difference method.

8.4.2 CORRELATION STUDY

The correlation between each pair of the grey tone difference texture features is measured over 78 ROI evenly distributed among the three tissue classes on thigh image data. This is the same source of data used for the discriminative performance evaluation in the previous section. The texture feature are collected from ROI of 16x16 with 7-bit data and without normalisation. The correlation coefficient is calculated using Microsoft EXCEL (see **chapter 7** for details), varying between 0 (no correlation) and 1 (high correlation). A negative sign implies negative correlation.

Table 8.4 presents the correlation coefficients between the three grey tone difference texture features. All three features proved to be highly correlated. This result indicates that the three features are measuring the same property.

Table 8.4 : Grey tone difference feature correlation

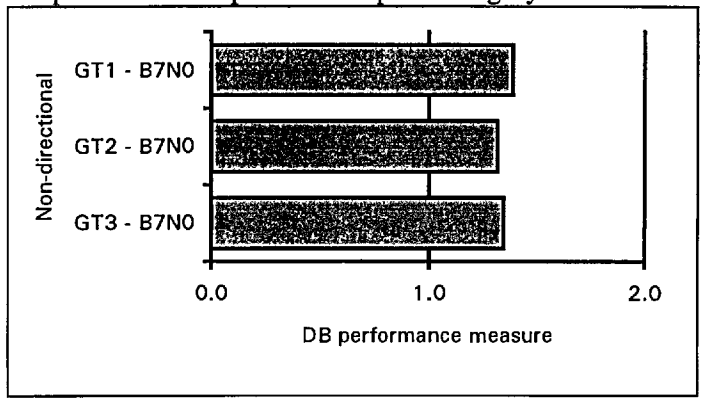
	GT1	GT2	GT3
GT1	-	-0.98	0.99
GT2	-0.98	-	-0.99
GT3	0.99	-0.99	-

8.4.3 OPTIMISED FEATURE SET

Texture analysis methods that are not sensitive to subject orientation are clearly more stable for applications such as clinical MRI tissue characterisation. The performance of the non-directional implementation of the grey tone difference method has been shown to be at least as good as that of the directional method. Optimum performances for all three texture features were achieved with ROI size 16x16, 7-bit data, and without data normalisation. A comparison of the DB performance of these highly correlated features under optimum conditions, presented in **graph 8.16**, shows little disparity between their discriminative

abilities. Texture feature GT1 is arbitrarily chosen as being representative of these three features.

Graph 8.16 : Comparison of optimised grey tone difference features)



8.5 CO-OCCURRENCE TEXTURE FEATURES

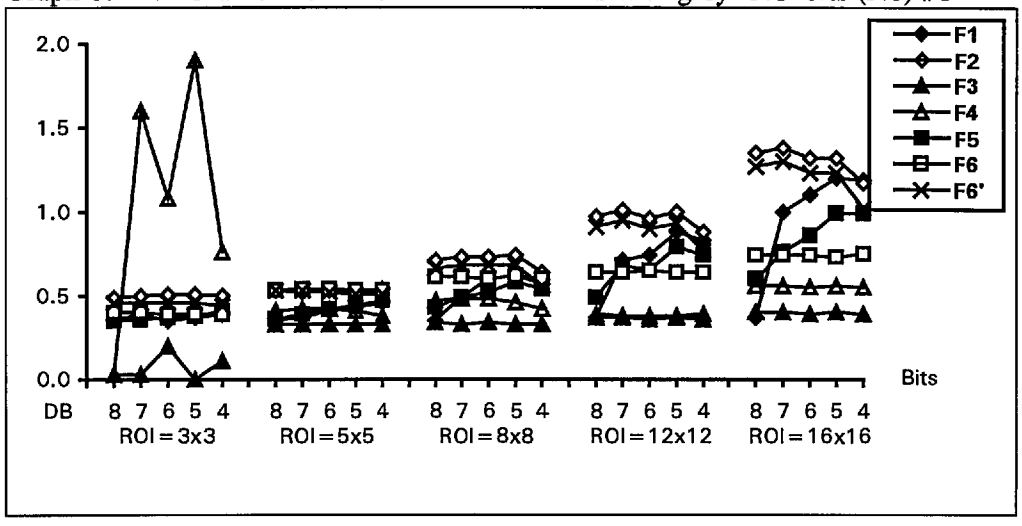
8.5.1 DISCRIMINATIVE PERFORMANCE OVERVIEW

The **graphs 8.17** to **8.20** indicate the discriminative performance trends of co-occurrence texture features extracted from ROI under different pre-processing conditions. For each point on the graph, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (ROI size and either bits or normalisation scheme). The 'Thigh' MRI image data used in this overview contains three distinct classes of tissue. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions (muscle, fat & tumour) on several of images.

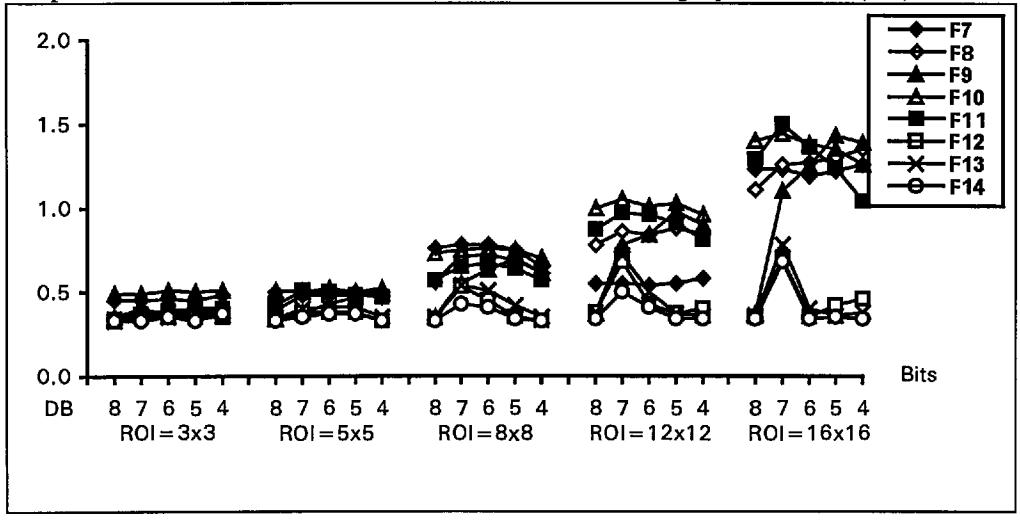
Graphs 8.17 and **8.18** show the variation of the DB performance measure with the number of grey level bits (4-8) assigned to the image data by the pre-processing. The texture features are extracted from a non-directional implementation of the co-occurrence method. This behaviour is examined over a range of different ROI sizes. The graph data was acquired without normalisation (N0).

Graphs 8.19 and **8.20** show the variation of the DB performance measure with normalisation scheme (0-2) for each texture feature. This behaviour is also examined over a range of different ROI sizes. The graph values were acquired with 8-bit image data (B8).

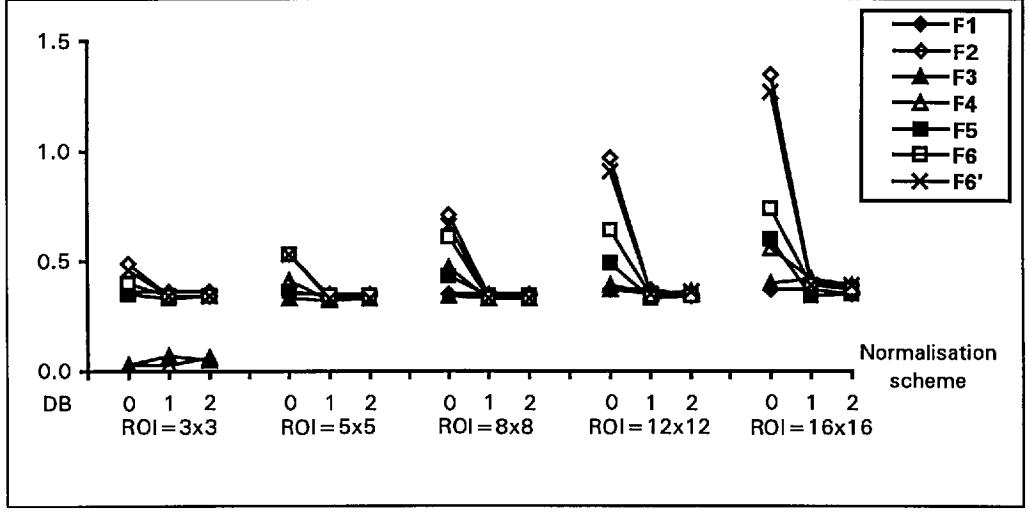
Graph 8.17 : Behaviour of co-occurrence features with grey level bits (N0) #1



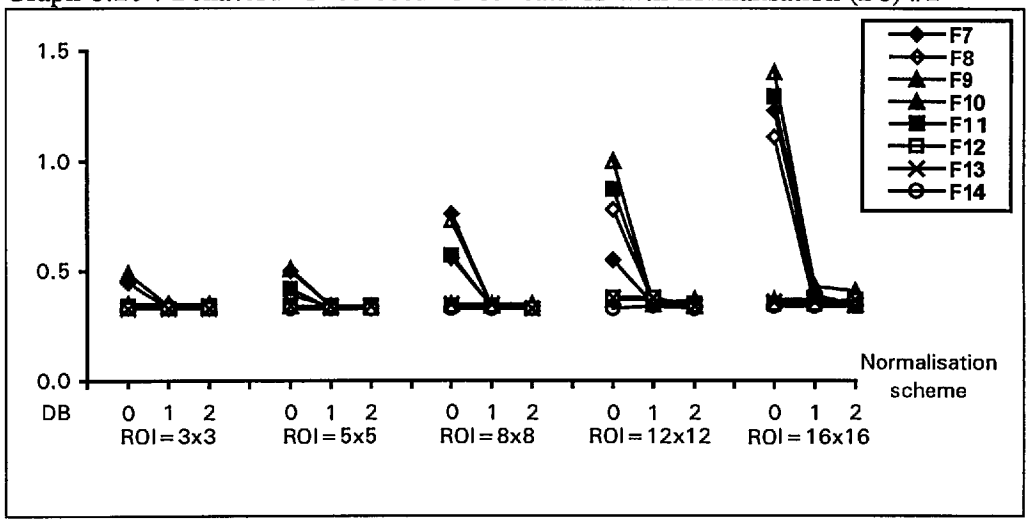
Graph 8.18 : Behaviour of co-occurrence features with grey level bits (N0) #2



Graph 8.19 : Behaviour of co-occurrence features with normalisation (B8) #1

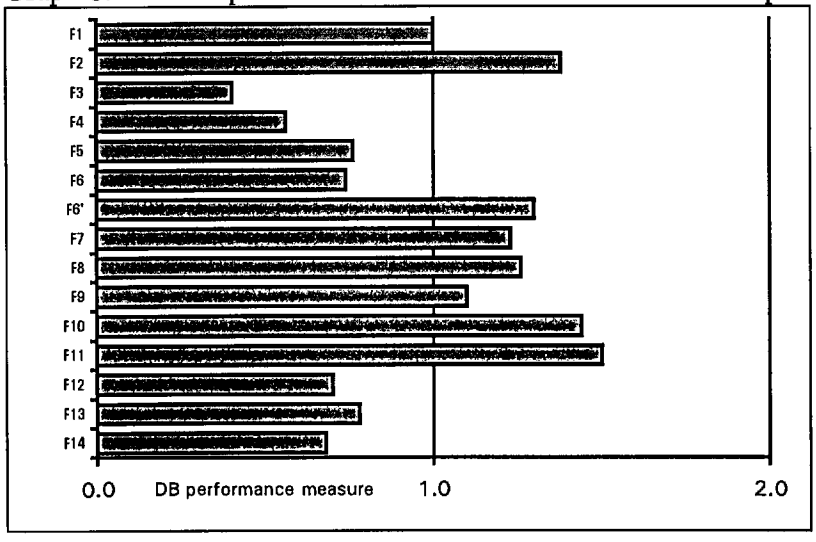


Graph 8.20 : Behaviour of co-occurrence features with normalisation (B8) #2



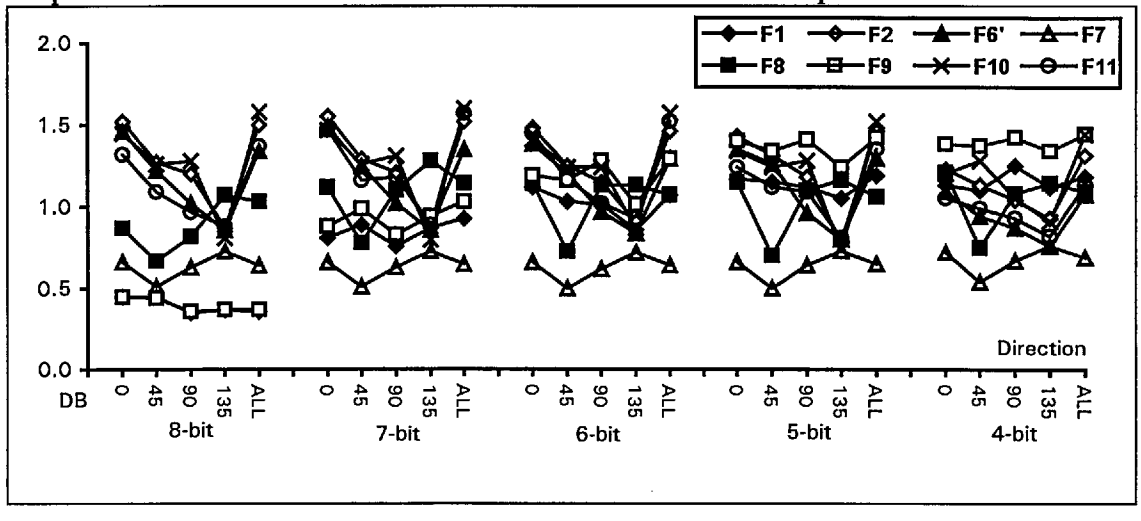
The overall discriminative ability of the co-occurrence texture feature set is significantly influenced by ROI size. Optimum performance was generally achieved with ROI of 16x16 and 7-bit data. In all cases the effect of data normalisation was found to be detrimental to the discriminative performance of the texture features. The co-occurrence method provides a large number of texture features (total of 15) that extend a range of discriminative ability. In order that the less useful features could be eliminated from the feature set, the performance of each feature under optimum conditions was examined. **Graph 8.21** provides a comparison of the DB performance measure for each feature acquired from ROI of size 16x16 with 7-bit data and without normalisation. All texture features with a DB performance measure less than 1.0 (arbitrarily chosen cut-off) were eliminated from the proposed feature set ie. F3, F4, F5, F6, F12, F13 and F14.

Graph 8.21 : Comparison of all co-occurrence features under optimum conditions



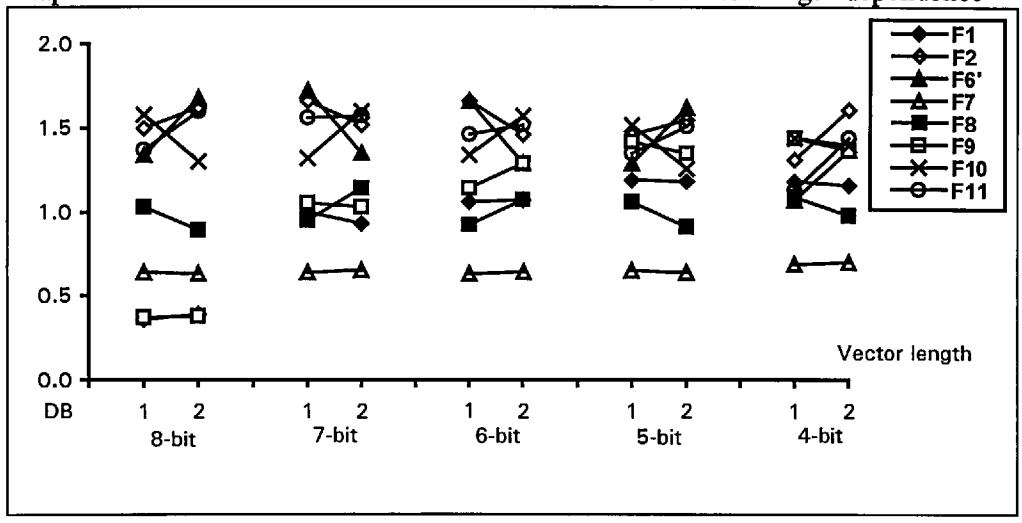
The discriminant performance overview summarised in **graph 8.21** was carried out using the non-directional implementation of the co-occurrence method. This method relies on information averaged from directionally sensitive implementations of the co-occurrence matrix. In theory, the non-directional implementation must be favoured in such applications as clinical MRI because of the naturally occurring variations in tissue orientation ie. texture. The loss of texture information resulting from such averaging is the subject of the study presented in **Graph 8.22**. This graph shows the variation of the DB performance measure with the direction in which the co-occurrence displacement vector (one pixel) is applied. The direction in which this vector is applied is indicated by Dirn=0, 45, 90 & 135 degrees under the x-axis of graph. The label 'Dirn=ALL' refers to the data from the non-directional co-occurrence method. This study was carried out on the reduced texture feature set previously identified. The graph values were acquired from ROI of size 16x16 without normalisation, and for a range of grey level bits (4-8). The overall behaviour of the texture features from both the directionally sensitive and non-directional implementations are comparable.

Graph 8.22 : Behaviour of co-occurrence features with direction dependence



Graph 8.23 presents a comparison of the discriminative performance of the reduced set of co-occurrence texture features for different displacement vector distances. This brief study only examined feature behaviour for vector sizes of one or two pixels, and was carried out with the non-directional implementation of the co-occurrence method. The graph values were acquired from ROI of size 16x16 without normalisation, and for a range of grey level bits (4-8). The overall behaviour of the texture features with different length displacement vectors was comparable.

Graph 8.23 : Behaviour of co-occurrence features with vector length dependence



8.5.2 CORRELATION STUDY

In **section 8.5.1** the large number of co-occurrence texture features was reduced to seven on the basis of discriminative performance. A study of the correlation between these features enables redundancy to be identified, and the feature set to be further reduced.

The correlation between each pair of texture features is measured over 78 ROI evenly distributed among the three tissue classes on 'Thigh' image data. This is the same source of data used for the discriminative performance evaluation in the previous section. The texture feature are collected from ROI of 16x16 with 7-bit data and without normalisation. The correlation coefficient is calculated using Microsoft EXCEL (see **chapter 7** for details), varying between 0 (no correlation) and 1 (high correlation). A negative sign implies negative correlation.

Table 8.5 : Co-occurrence feature correlation

	F1	F2	F6'	F7	F8	F9	F10	F11
F1	-	-0.94	-0.96	-0.91	-0.93	-0.98	-0.93	-0.95
F2	-0.94	-	1.00	0.97	0.97	0.98	1.00	0.99
F6'	-0.96	1.00	-	0.96	0.96	0.98	0.99	1.00
F7	-0.91	0.97	0.96	-	0.99	0.95	0.97	0.96
F8	-0.93	0.97	0.96	0.99	-	0.96	0.97	0.96
F9	-0.98	0.98	0.98	0.95	0.96	-	0.97	0.98
F10	-0.93	1.00	0.99	0.97	0.97	0.97	-	0.99
F11	-0.95	0.99	1.00	0.96	0.96	0.98	0.99	-

Table 8.5 presents the correlation coefficients between the co-occurrence texture features acquired from the non-directional implementation. All seven features were highly correlated ie. they have correlation coefficients greater that 0.9. It is possible, however, to place the most highly correlated features into four separate groups, as shown in **table 8.6**.

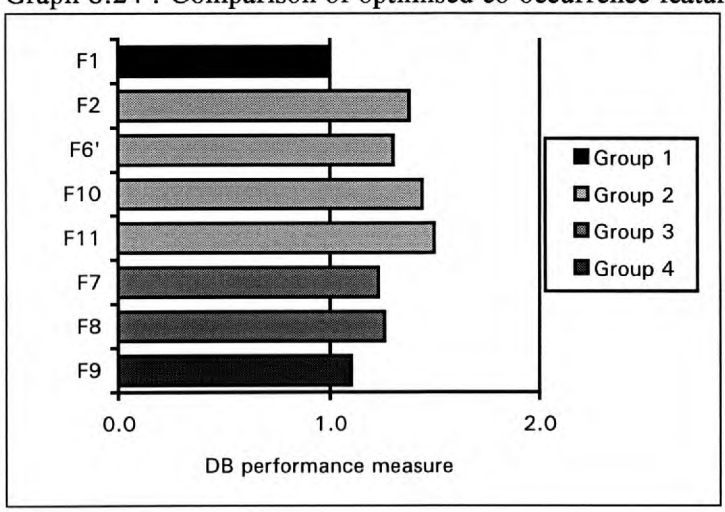
Table 8.6 : Grouping of most correlated co-occurrence texture features

Group	**Texture features**
Group 1	F1 (Second angular moment - homogeneity)
Group 2	F2 (Contrast) F6' (Difference average) F10 (Difference variance) F11 (Difference entropy)
Group 3	F7 (Sum variance) F8 (Sum entropy)
Group 4	F9 (Entropy)

8.5.3 OPTIMISED FEATURE SET

Texture analysis methods that are not sensitive to subject orientation are more stable for applications such as clinical MRI tissue characterisation. The performance of the non-directional co-occurrence method was found to be comparable with the directionally sensitive implementations. No noticeable change in discriminative performance of this method was seen by varying the displacement vector from one to two pixels.

Graph 8.24 : Comparison of optimised co-occurrence features

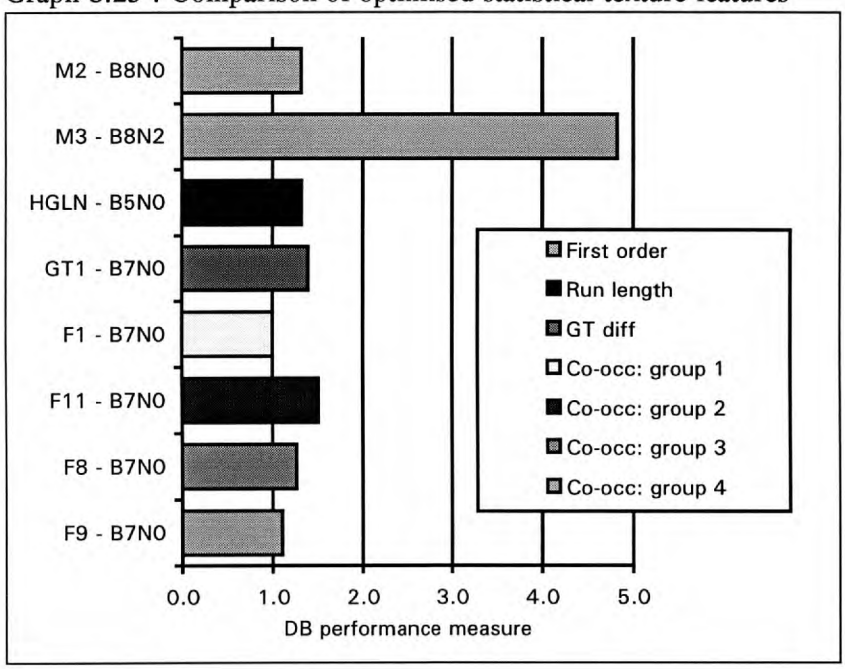


A comparison of the best performing texture features extracted using the non-directional implementation of the co-occurrence method with a displacement vector of length one pixel is presented in **graph 8.24**. The texture features are extracted from ROI of size 16x16 containing 7 bit image data (B7) and without normalisation. The most highly correlated features identified in the previous section are grouped together. The best performing features in each group are F1 (group 1), F11 (group 2), F8 (group 3) and F9 (group 4).

8.6 COMPARISON OF OPTIMAL STATISTICAL TEXTURE FEATURES

This section provides a comparison of the optimum texture features identified for all four methods examined in the previous sections of this chapter; first order statistics, run length analysis, grey tone difference, and co-occurrence. **Graph 8.25** compares the DB measure of discriminative performance for the optimum texture features identified for each of these methods. All eight features (details given in **table 8.1** in **section 8.1**) are extracted under the normalisation and data-bit conditions indicated on the graph, and from ROI of size 16x16. Further details of how the graph values were obtained can be found in the respective performance evaluation sections of this chapter, where the data were originally presented.

Graph 8.25 : Comparison of optimised statistical texture features



The skewness (M3) first order statistical standard deviation texture feature demonstrated a discriminative ability quantified by the DB measure to be a factor of four better than the other seven features. The overall discriminative performance of all other features were of the same order.

More information about the relationship between features derived from different methods can be found from the study of correlation between them presented in **table 8.7**. This table shows a high correlation between all statistical texture features, with the exception of feature M3. From this result, it is a reasonable assumption to make that little additional information is provided by including all the highly correlated features (M2, HGLN, GT1, F1, F8, F9 & F11) with M3 in a feature set. The subset of highly correlated features should be reduced on the basis of discriminant performance. Examples of the discriminant ability of such reduced feature sets with a variety of MRI tissue subjects are given in the following section.

Table 8.7 : Correlation between best performing statistical texture features

	M2	M3	HGLN	GT1	F1	F8	F9	F11
M2	-	0.38	-0.98	0.98	-0.92	0.99	0.96	0.98
M3	0.38	-	-0.45	0.39	-0.56	0.44	0.50	0.39
HGLN	-0.98	-0.45	-	-0.97	0.96	-0.98	-0.98	-0.98
GT1	0.98	0.39	-0.97	-	-0.94	0.97	0.98	0.99
F1	-0.92	-0.56	0.96	-0.94	-	-0.93	-0.98	-0.95
F8	0.99	0.44	-0.98	0.97	-0.93	-	0.96	0.96
F9	0.96	0.50	-0.98	0.98	-0.98	0.96	-	0.98
F11	0.98	0.39	-0.98	0.99	-0.95	0.96	0.98	-

8.7 EXAMPLES OF TISSUE CHARACTERISATION

This section provides examples of the application of the optimised statistical texture features for discrimination between tissue classes in a variety of MRI images. The eight optimised texture features (as described in the previous section) are listed in **table 8.8** with their corresponding graph keys. **Table 8.9** provides a list of the MRI image data sets investigated and the tissue classes they contain. The degree to which the texture features characterise the tissue classes in the case of each data set is illustrated in **graphs 8.26** to **8.45**. A summary of the conditions under which the data for each of these graphs was acquired is provided in **table 8.10**.

Table 8.8 : Statistical texture features - graph key

Statistical texture feature	Key
First order statistics: standard deviation	M2
First order statistics: skewness	M3
Horizontal run length: grey level non-uniformity	HGLN
Grey tone difference: contrast	GT1
Co-occurrence: Second angular moment	F1
Co-occurrence: Sum entropy	F8
Co-occurrence: Entropy	F9
Co-occurrence: Difference entropy	F11

Table 8.9 : Description of image data

Data sample	Tissue classes
Thigh	Muscle, Tumour, Fat
Abdomen	Fat, Liver, Tumour
Leg104	Muscle, Tumour
Leg192	Muscle, Tumour
Brains	Brain1, Brain2, Brain3

Table 8.10 : Examples of tissue segmentation with statistical texture features

Graph	Data sample	ROI size	X-axis feature	Y-axis feature
8.26	Thigh	16x16	M2 - B8N0	M3 - B8N2
8.27		16x16	HGLN - B5N0	GT1 - B7N0
8.28		16x16	F1 - B7N0	F8 - B7N0
8.29		16x16	F9 - B7N0	F11 - B7N0
8.30	Abdomen	12x12	M2 - B8N0	M3 - B8N2
8.31		12x12	HGLN - B5N0	M3 - B8N2
8.32		12x12	F1 - B7N0	F8 - B7N0
8.33		12x12	F9 - B7N0	F11 - B7N0
8.34	Leg104	16x16	M2 - B8N0	M3 - B8N2
8.35		16x16	HGLN - B5N0	M3 - B8N2
8.36		16x16	F1 - B7N0	F8 - B7N0
8.37		16x16	F9 - B7N0	F11 - B7N0
8.38	Leg192	16x16	M2 - B8N0	M3 - B8N2
8.39		16x16	HGLN - B5N0	M3 - B8N2
8.40		16x16	F1 - B7N0	F8 - B7N0
8.41		16x16	F9 - B7N0	F11 - B7N0
8.42	Brains	16x16	M2 - B8N0	M3 - B8N2
8.43		16x16	HGLN - B5N0	M3 - B8N2
8.44		16x16	F1 - B7N0	F8 - B7N0
8.45		16x16	F9 - B7N0	F11 - B7N0

Graphs 8.26 to **8.29** show good examples of successful segmentation between the crucial muscle tissue and muscle tumour tissue classes in the 'Thigh' image data set. It was not, however, always possible to clearly differentiate between tumour and fat tissue in these examples. The clustering of the data points for each of the three tissue classes was very good.

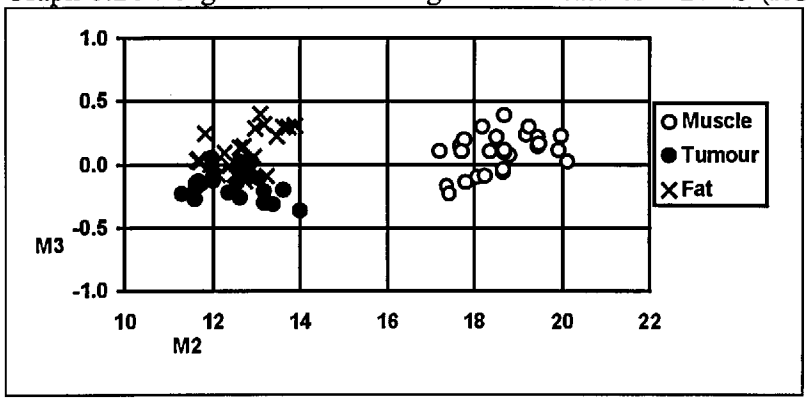
Graphs 8.30 to **8.33** show poor discrimination between the crucial liver tissue and liver tumour tissue classes from the 'Abdomen' data set. There is evidence of clustering of same tissue class data points, but there is a considerable amount of overlap between these groups. In all graphs it appears that there are two liver tissue clusters; the main cluster and a satellite cluster of four points amongst the fat cluster. This second smaller cluster could be attributed to liver regions with fat striations.

Graphs 8.34 to **8.37** show the first of two cases of segmentation of muscle tumour from muscle tissue in the leg (data set 'Leg104'). Better clustering was achieved for the tumour tissue group. The variability of texture parameters for muscle tissue is probably due to the in-homogeneity across the ROI selected from the image, as a result of the fibrous nature of the tissue.

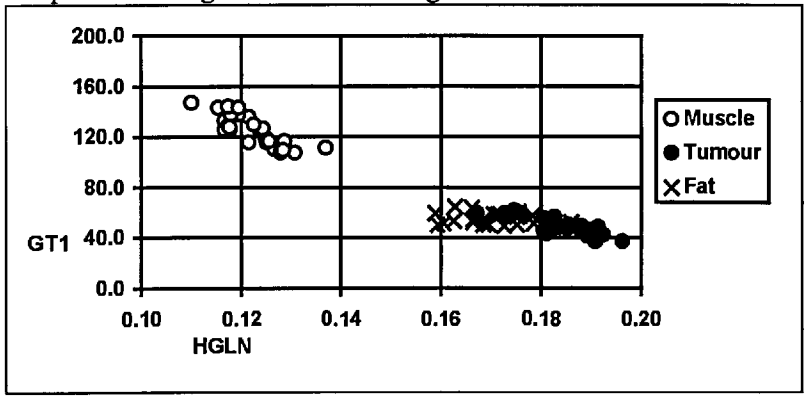
Graphs 8.38 to **8.41** show the second case of segmentation of muscle tumour from muscle tissue in the leg (data set 'Leg192'). In three out of the four graphs, the two tissue classes are clearly separated in feature space. Features M2 and F11 appear to perform very well in this example.

Graphs 8.42 to **8.45** show the results for texture features extracted from grey matter in a series of mid-brain slices in three healthy volunteers; 'Brain' image data. Feature stability can be assessed by examining the variability of features both in each subject class and in all three classes combined. For good stability, one is looking for minimal variability firstly in each class and then in all classes. There is no evidence to suggest that the texture features examined in **graphs 8.42** to **8.45** are stable parameters. The data points corresponding to the three brain tissue classes vary significantly in feature space, and the values occur over a large range.

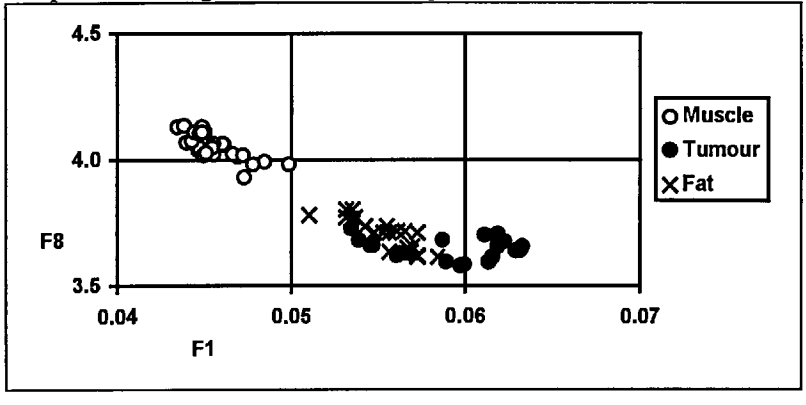
Graph 8.26 : Segmentation of 'Thigh' data - features M2vM3 (ROI 16x16)



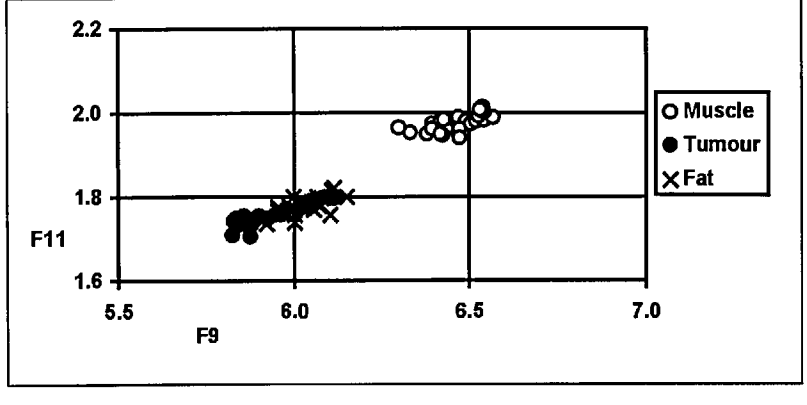
Graph 8.27 : Segmentation of 'Thigh' data - features HGLNvGT1 (ROI 16x16)



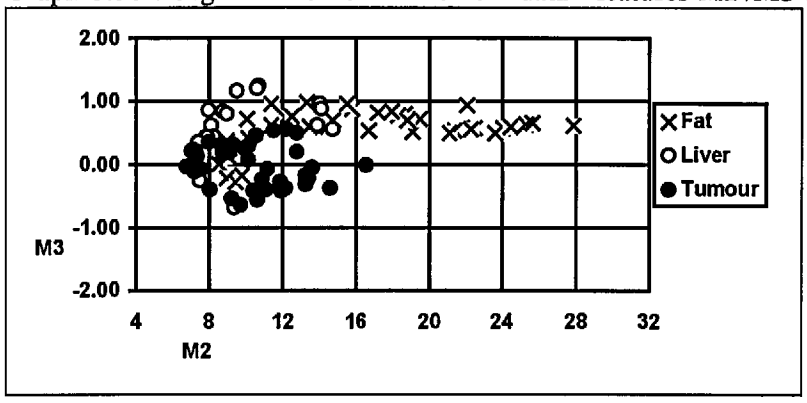
Graph 8.28 : Segmentation of 'Thigh' data - features F1vF8 (ROI 16x16)



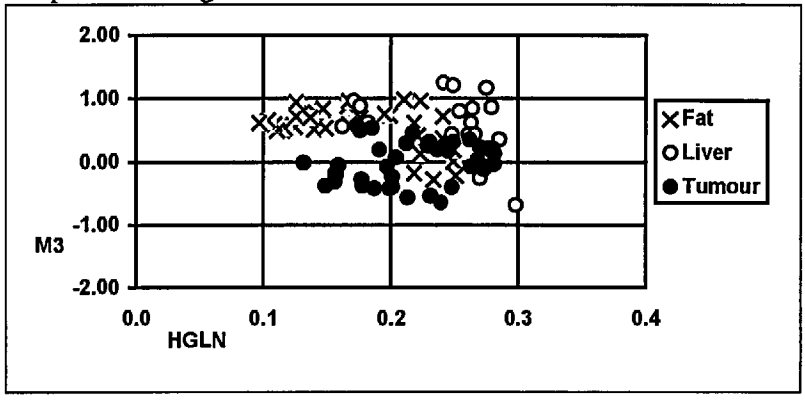
Graph 8.29 : Segmentation of 'Thigh' data - features F9vF11 (ROI 16x16)



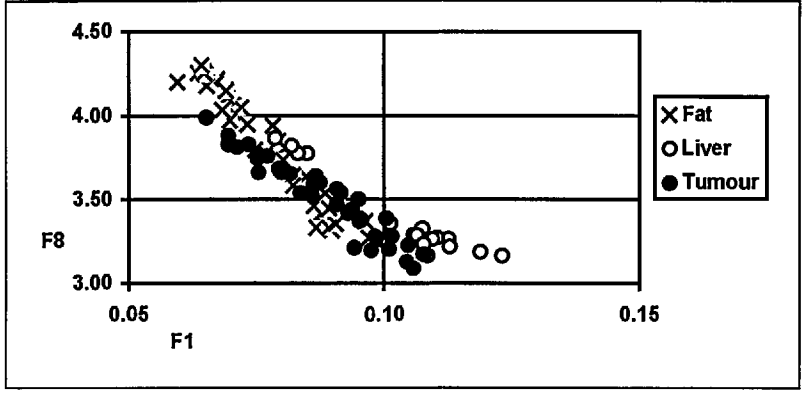
Graph 8.30 : Segmentation of 'Abdomen' data - features M2vM3 (ROI 12x12)



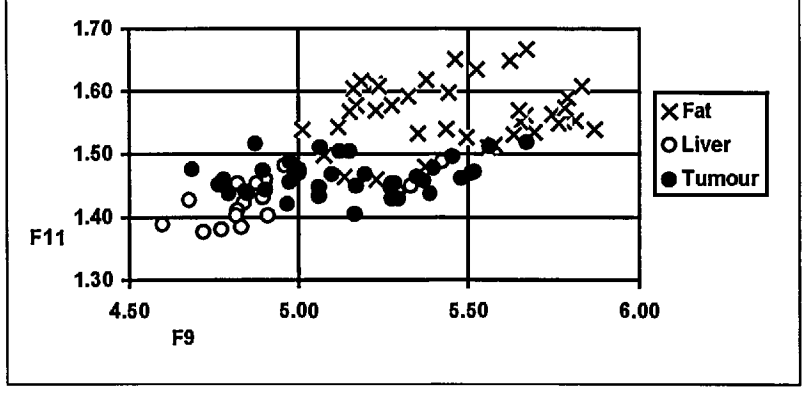
Graph 8.31 : Segmentation of 'Abdomen' data - features M3vHGLN (ROI 12x12)



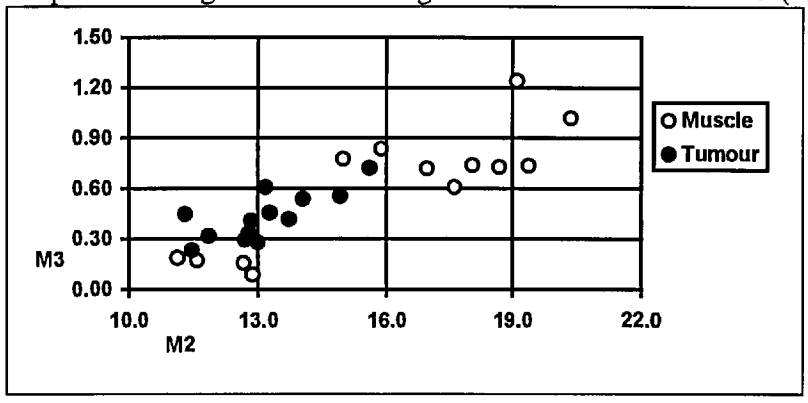
Graph 8.32 : Segmentation of 'Abdomen' data - features F1vF8 (ROI 12x12)



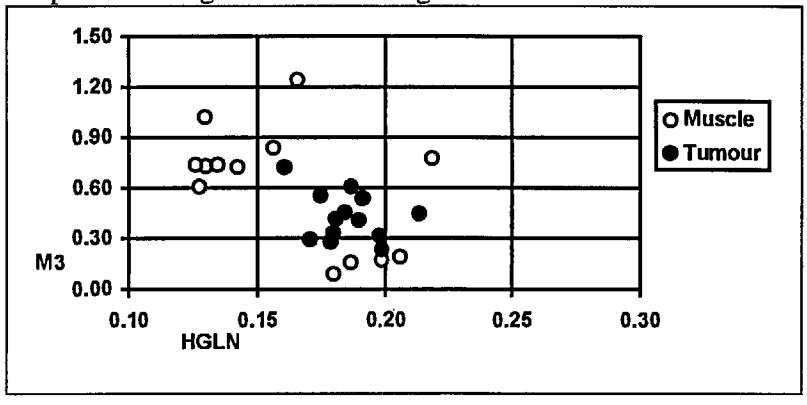
Graph 8.33 : Segmentation of 'Abdomen' data - features F9vF11 (ROI 12x12)



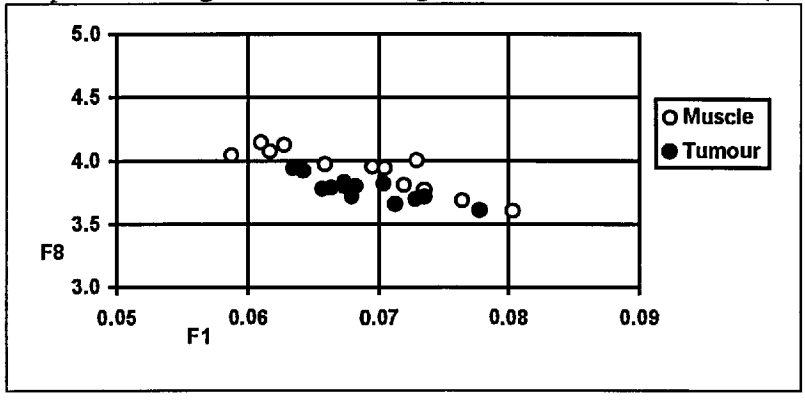
Graph 8.34 : Segmentation of 'Leg104' data - features M2vM3 (ROI 16x16)



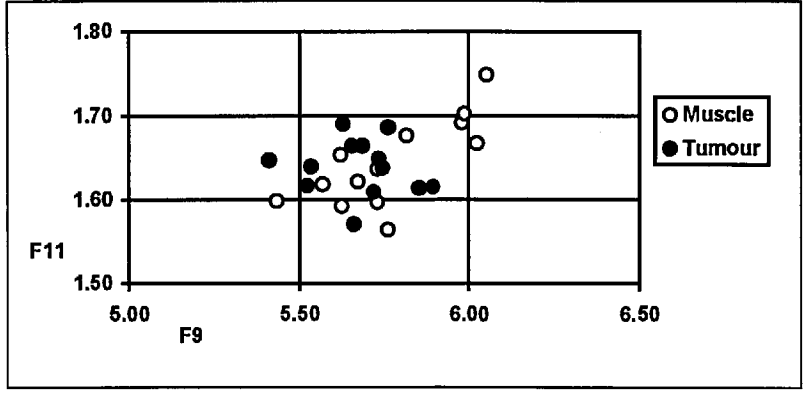
Graph 8.35 : Segmentation of 'Leg104' data - features M3vHGLN (ROI 16x16)



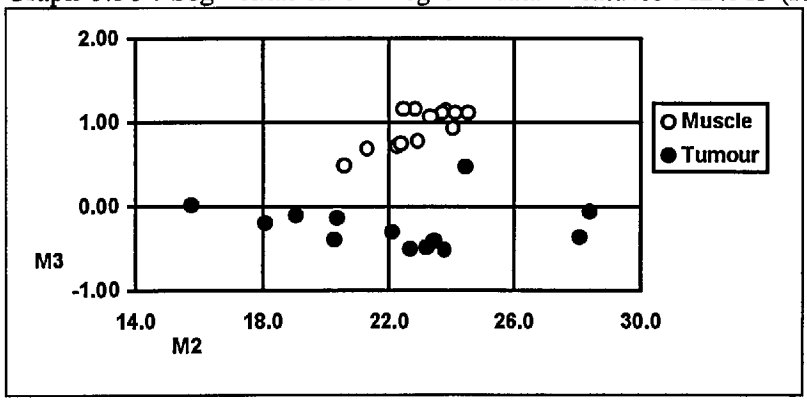
Graph 8.36 : Segmentation of 'Leg104' data - features F1vF8 (ROI 16x16)



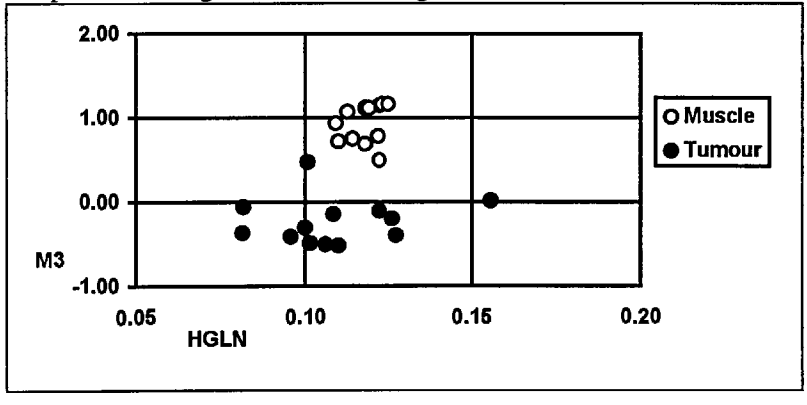
Graph 8.37 : Segmentation of 'Leg104' data - features F9vF11 (ROI 16x16)



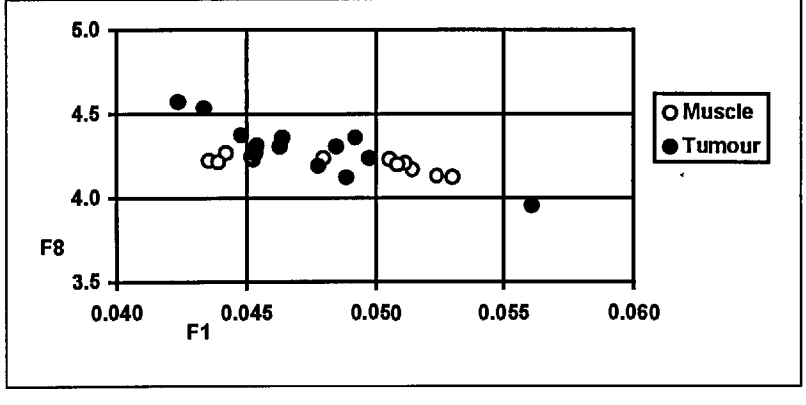
Graph 8.38 : Segmentation of 'Leg192' data - features M2vM3 (ROI 16x16)



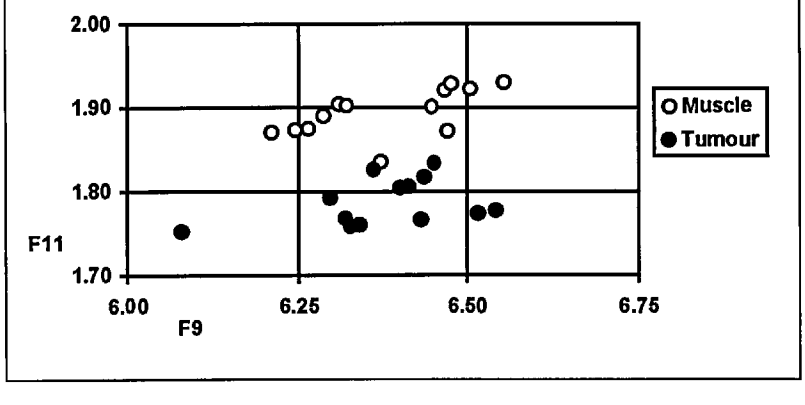
Graph 8.39 : Segmentation of 'Leg192' data - features M3vHGLN (ROI 16x16)



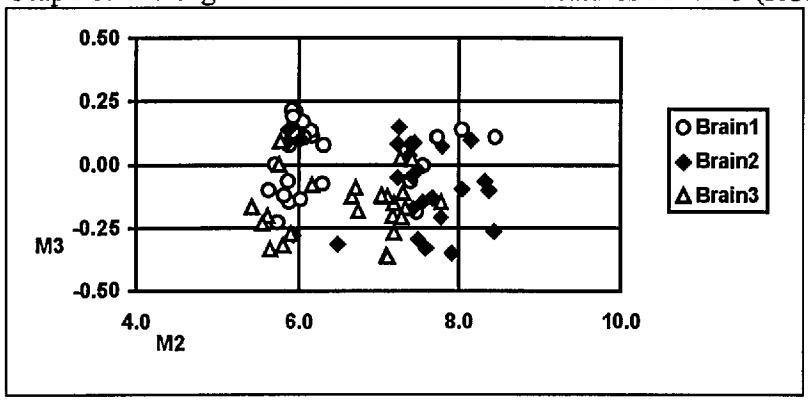
Graph 8.40 : Segmentation of 'Leg192' data - features F1vF8 (ROI 16x16)



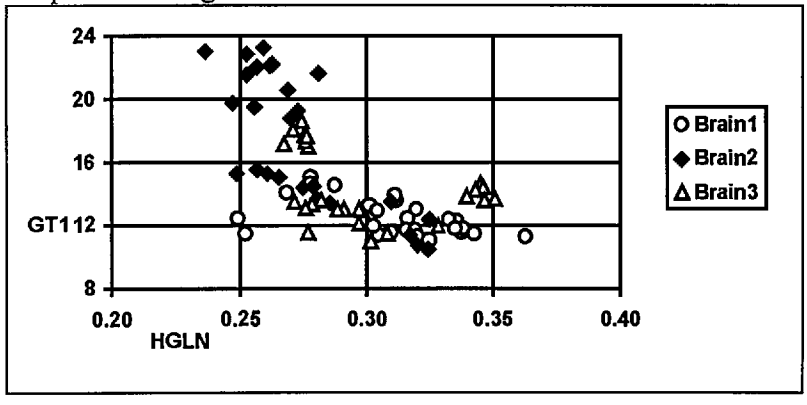
Graph 8.41 : Segmentation of 'Leg192' data - features F9vF11 (ROI 16x16)



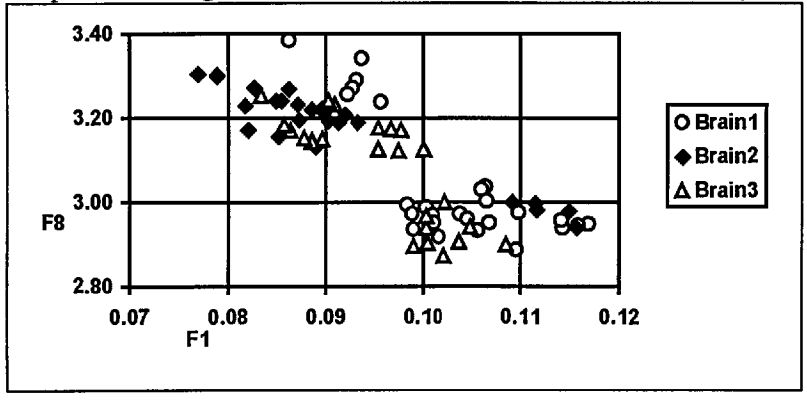
Graph 8.42 : Segmentation of 'Brain' data - features M2vM3 (ROI 16x16)



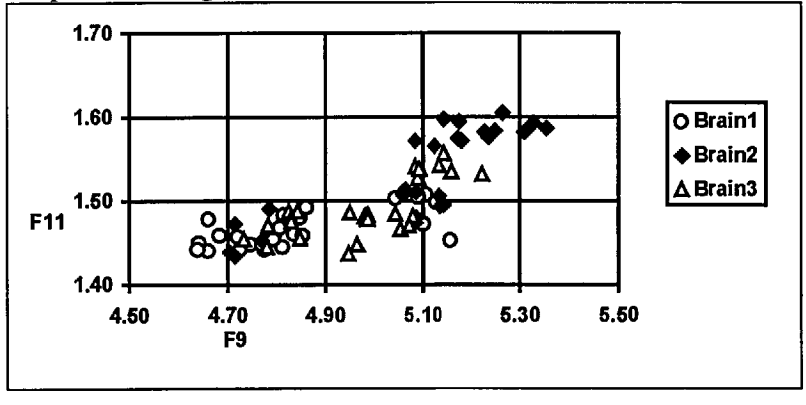
Graph 8.43 : Segmentation of 'Brain' data - features HGLNvGT1 (ROI 16x16)



Graph 8.44 : Segmentation of 'Brain' data - features F1vF8 (ROI 16x16)



Graph 8.45 : Segmentation of 'Brain' data - features F9vF11 (ROI 16x16)



CHAPTER 9 - RESULTS: TRANSFORM TEXTURE FEATURES

9.1 INTRODUCTION

This chapter presents the results of a study of three methods of deriving textural information by applying image transform functions to image data. These methods, described in **chapter 5**, extract image texture information following the application of the Fourier transform and directional and non-directional implementations of the Walsh and Slant transforms.

Each of the three transform methods is dealt with in a separate section of this chapter (**sections 9.2** to **9.4**). Each section begins with an overview of the discriminative performance of the transform texture features, and the optimum conditions under which they can be extracted. Also presented is a study of the correlation between the texture features, a description of the optimised feature set, and examples of the use of this feature set for discrimination between a selection of tissue classes.

The discriminative performance tests are carried out on several transverse MRI images of the thigh (see 'Thigh' image data in **chapter 7**). These images contain three distinct classes of tissue; (1) muscle, (2) fat and (3) tumour. The discriminative behaviour of each feature is quantified using the DB performance measure defined in **chapter 7** (based on the relationship between the statistics of tissue classes). Unless stated, the DB performance measure is calculated for the discrimination between all three tissue classes (1-3).

The regions of interest (ROI) in the image data are standardised before any texture information is extracted by fixing the number of grey level bits or by implementing some normalisation scheme. The pre-processing options are indicated by 'B' for the number of grey level bits (in the range of 4-8), and 'N' for the normalisation scheme; (1) none, (2) histogram equalisation, and (3) statistical moments. A detailed description of normalisation schemes (2) and (3) are given in **chapter 7**.

The results of the performance overview attempt to demonstrate the effect of data pre-processing on the discriminative behaviour of each texture feature. Consideration is also given to the influence of region of interest (ROI) size. The purpose of this study is to identify the most robust and best performing features, and the conditions under which they perform optimally.

The correlation study in each section enables the relative behaviour of the transform texture features to be investigated. Similar behaviour would be expected from highly correlated features. When two features are highly correlated, therefore, one feature can be discarded without any significant loss of information. Large numbers of texture features can be reduced to more manageable subsets by examining the correlation between each and every feature, and eliminating redundant features.

The tissue characterisation examples provided for each transform method demonstrate the efficacy of the corresponding optimised texture features set with a number of sets of clinical images containing a variety of different tissue types.

9.2 POWER SPECTRUM TEXTURE FEATURES

9.2.1 INTRODUCTION TO POWER SPECTRUM TEXTURE FEATURES

The Fourier transform decomposes an image region to a spectrum indicating the abundance of each frequency component. Because texture can be described loosely in terms like fine or coarse, the composition of the frequency spectrum clearly provides textural information. Although the auto-correlation function has successfully been used to derive texture coarseness measures, this method has been overlooked because of the one dimensional description of texture it provides ie. coarse or fine. It is clear, however, that the power spectrum contains some valuable texture information. For this reason it was proposed that the energy and statistics of the radially collapsed 1D image spectra be examined as suitable texture features. A list of the texture features derived from power spectrum that are to be investigated are given below in **table 9.1**. This table provides a reference between the features and the keys used in the graphs and tables presented throughout this section.

Table 9.1 : Power spectrum texture features

Transform method	Texture features	Key
Fourier transform	Energy	FT1
(Power spectrum)	Cut-off	FT2
	Mean	FT3
	Standard deviation	FT4
	Skewness	FT5
	Kurtosis	FT6

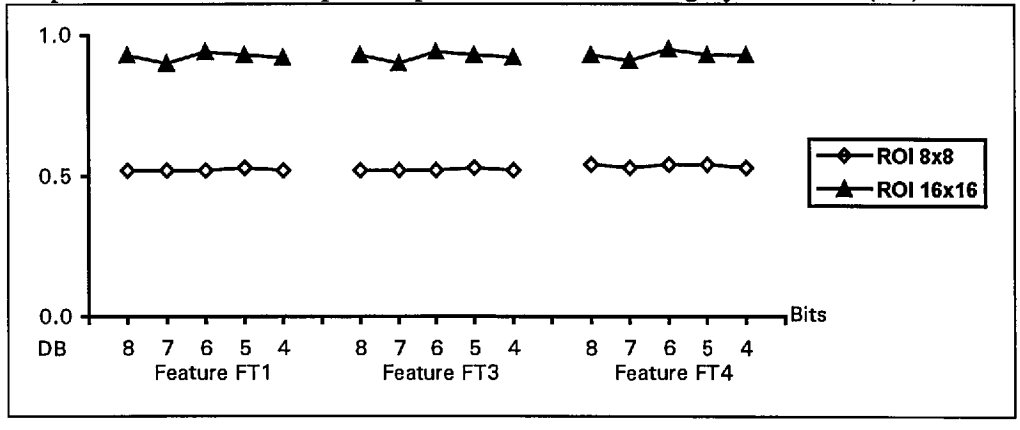
9.2.2 PERFORMANCE OVERVIEW

The plots in **graphs 9.1** to **9.4** indicate the discriminative performance trends of the power spectrum texture features extracted from ROI under different pre-processing conditions. For each point on the graphs, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (either bits or normalisation scheme). The graph legend labels each series of data points plotted with the ROI size (8x8 or 16x16). The 'Thigh' MRI image data used in this overview contains three distinct classes of tissue. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions (muscle, fat & tumour) on several images.

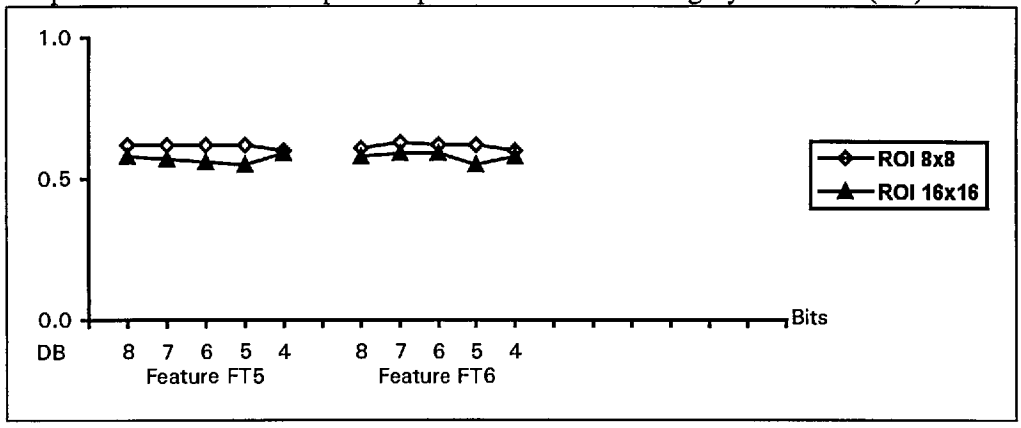
Graphs 9.1 and **9.2** show the variation of the DB performance measures for the power spectrum features with the number of grey level bits (4-8) assigned to the image data by the pre-processing. The data used in these graphs was acquired without normalisation (N0) and with both 8x8 and 16x16 ROI sizes. Features FT1 (energy), FT3 (mean) and FT4 (standard deviation) are presented in **graph 9.1**, and features FT5 (skewness) and FT6 (kurtosis) are presented in **graph 9.2**. Feature FT2 (cut-off) was found not to be influenced at all by image texture. This feature proves only to be a property of the number of data points on the power spectrum, and therefore can be eliminated from the study at this stage.

The effect of number of grey level bits does not appear to affect the discriminant performance of the power spectrum texture features. In the case of features FT1, FT3 and FT4, better DB performance measures were attained for 16x16 than for 8x8 ROI size. For features FT5 and FT6 the effect of ROI size was less significant, in fact, marginally better DB performance was achieved with the smaller ROI.

Graph 9.1 : Behaviour of power spectrum features with grey level bits (N0) #1

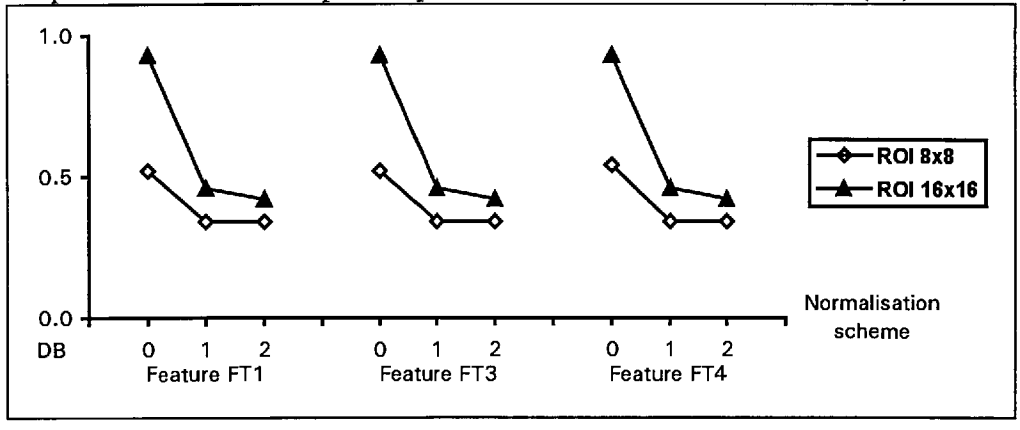


Graph 9.2 : Behaviour of power spectrum features with grey level bits (N0) #2

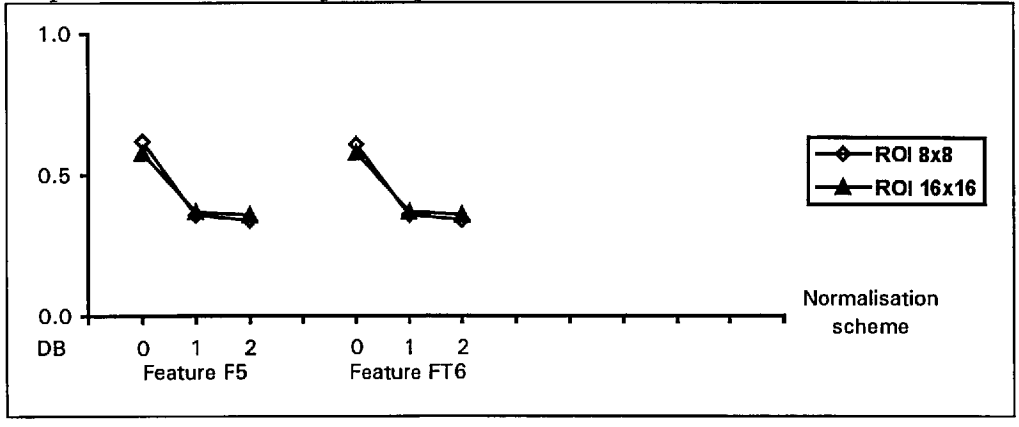


Graphs 9.3 and **9.4** show the variation of the DB performance measures for the power spectrum features with normalisation scheme (0-2). The data used in these graphs was acquired with 8-bit data and with both 8x8 and 16x16 ROI sizes. Both graphs indicate that normalisation is detrimental to the discriminant performance of all the power spectrum texture features.

Graph 9.3 : Behaviour of power spectrum features with normalisation (B8) #1



Graph 9.4 : Behaviour of power spectrum features with normalisation (B8) #2



9.2.3 CORRELATION STUDY

The correlation between each pair of texture features is measured over 78 ROI evenly distributed among the three tissue classes on 'Thigh' image data. This is the same source of data used for the discriminative performance evaluation in the **section 9.2.2**. The texture features are collected from ROI of sizes 8x8 and 16x16, with 8-bit data and without normalisation. The correlation coefficient is calculated using Microsoft EXCEL (see **chapter 7** for details), varying between 0 (no correlation) and 1 (high correlation). A negative sign implies negative correlation.

Tables 9.2 and **9.3** below show that the power spectrum texture features can be separated into two groups of highly correlated features; group 1: FT1, FT3 & FT4 and group 2: FT5 & FT6. These groups correspond well with the discriminant behaviour shown in **graphs 9.1** and **9.2** in the previous section. On the basis of the comparative performance between the members in each group, the feature set examined in this section can be reduced to just two features, one from each group.

Table 9.2 : Power spectrum feature correlation (B8N0 ROI 8x8)

	FT1	FT3	FT4	FT5	FT6
FT1	-	1.00	1.00	0.56	0.56
FT3	1.00	-	1.00	0.56	0.56
FT4	1.00	1.00	-	0.58	0.57
FT5	0.56	0.56	0.58	-	1.00
FT6	0.56	0.56	0.57	1.00	-

Table 9.3 : Power spectrum feature correlation (B8N0 ROI 16x16)

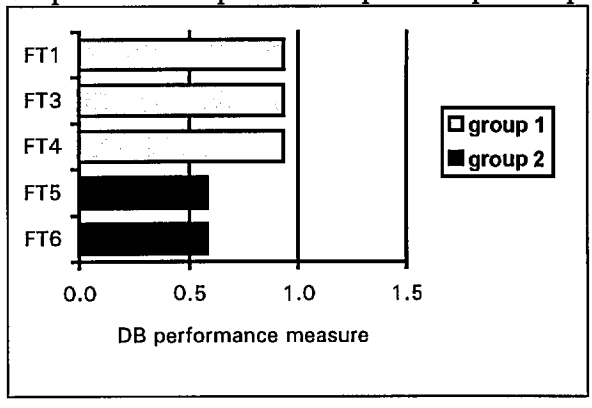
	FT1	FT3	FT4	FT5	FT6
FT1	-	1.00	1.00	0.75	0.76
FT3	1.00	-	1.00	0.75	0.76
FT4	1.00	1.00	-	0.76	0.76
FT5	0.75	0.75	0.76	-	0.99
FT6	0.76	0.76	0.76	0.99	-

9.2.4 OPTIMISED FEATURE SET

Optimum discriminant performance was achieved for all the power spectrum texture features investigated with 8-bit data and without data normalisation. Better performances were generally achieved from ROI of size 16x16 rather than 8x8. This result is consistent

with the improved specificity of the power spectrum that is expected with larger ROI. The number of data points on the power spectrum is equal to half the side length of the square ROI for the Fourier transform algorithm used for this project. A comparison of the DB performance of the power spectrum texture features acquired under optimum conditions, presented in **graph 9.5**, shows little disparity between the features in each of the two groups of highly correlated features (indicated in the graph legend). Features FT1 and FT5 are arbitrarily chosen as being representative of the whole feature set.

Graph 9.5 : Comparison of optimised power spectrum features



9.2.5 EXAMPLES OF TISSUE CHARACTERISATION

This section provides examples of the application of the optimised texture features derived from the power spectrum for discrimination between tissue classes in a variety of MRI images. The two optimised texture features (as described in the previous section) are listed in **table 9.4** with their corresponding graph keys. **Table 9.5** provides a list of the MRI image data sets investigated and the tissue classes they contain. The extent to which the texture features characterise the tissue classes in the case of each data set is illustrated in **graphs 9.6** to **9.11**. A summary of the conditions under which the data for each of these graphs was acquired is provided in **table 9.6**.

Table 9.4 : Power spectrum texture features - graph key

Power spectrum texture feature	Key
Energy	FT1
Skewness	FT5

Table 9.5 : Description of image data

Data set	Tissue classes
Thigh	Muscle, Tumour, Fat
Abdomen	Fat, Liver, Tumour
Leg104	Muscle, Tumour
Leg192	Muscle, Tumour
Brains	Brain1, Brain2, Brain3

Table 9.6 : Examples of tissue segmentation with power spectrum texture features

Graph	Data sample	ROI size	X-axis feature	Y-axis feature
9.6	Thigh	8x8	FT1 - B8N0	FT5 - B8N0
9.7		16x16	FT1 - B8N0	FT5 - B8N0
9.8	Abdomen	16x16	FT1 - B8N0	FT5 - B8N0
9.9	Leg104	16x16	FT1 - B8N0	FT5 - B8N0
9.11	Leg192	16x16	FT1 - B8N0	FT5 - B8N0
9.12	Brains	16x16	FT1 - B8N0	FT5 - B8N0

Graphs 9.6 and **9.7** show good examples of successful segmentation between the crucial muscle tissue and muscle tumour tissue classes in the 'Thigh' image data set. The muscle tumour and fat tissue classes were not separated in these examples, as indicated by the almost complete overlap of their data points in feature space. Better performance, indicated by tighter clustering of same class data points, was demonstrated when the texture features were acquired with the larger ROI size of 16x16.

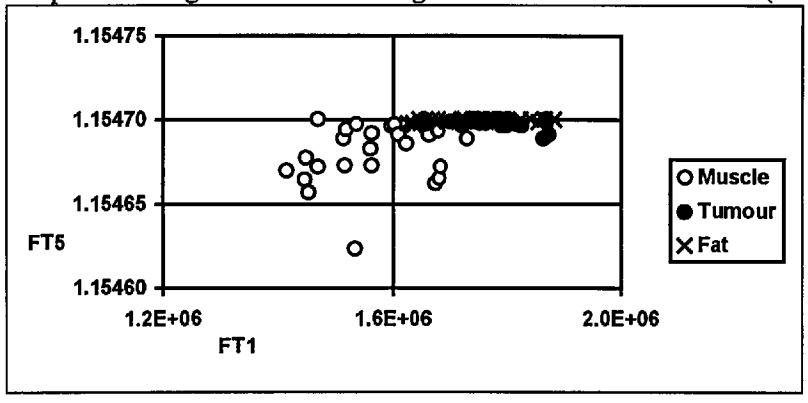
Graphs 9.8 shows the segmentation possible between liver tissue and tumour tissue within the liver from the 'Abdomen' data set. This example demonstrates that these two tissue classes can be clearly differentiated between using power spectrum texture features.

Graphs 9.9 and **9.10** show two cases of the segmentation of muscle tumour from muscle tissue in the leg. The 'Leg192' example in **graph 9.10** clearly shows two different tissue classes in feature space. The boundary in feature space between the two tissue classes is less well defined for the 'Leg104' data set in **graph 9.9**, although there is some grouping of the data points corresponding to the two classes.

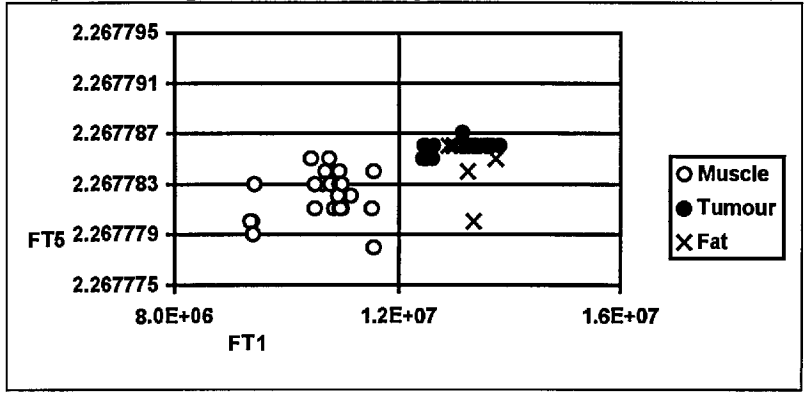
Graph 9.11 shows the results for texture features extracted from grey matter in a series of mid-brain slices in three healthy volunteers; 'Brain' image data. Feature stability can be assessed by examining the variability of features both in each subject class and in all three classes combined.

For good stability, one is looking for minimal variability firstly in each class and then in all classes. High stability is demonstrated by feature FT5, beyond the precision that the feature can be calculated. The variability of feature FT1 appears to be equivalent in all classes.

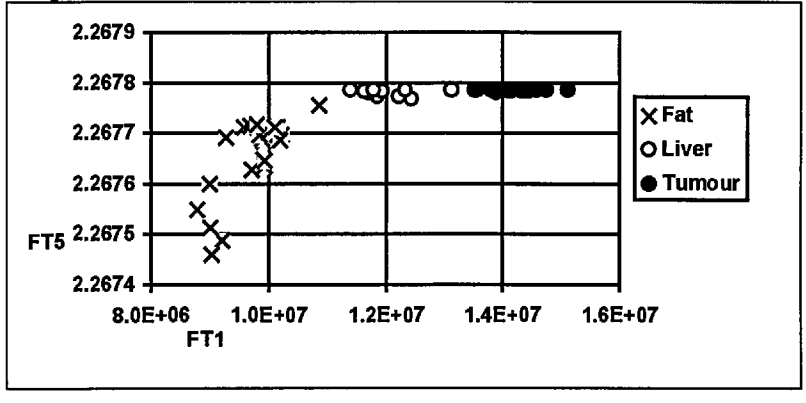
Graph 9.6 : Segmentation of 'Thigh' data - features FT1vFT5 (ROI 8x8)



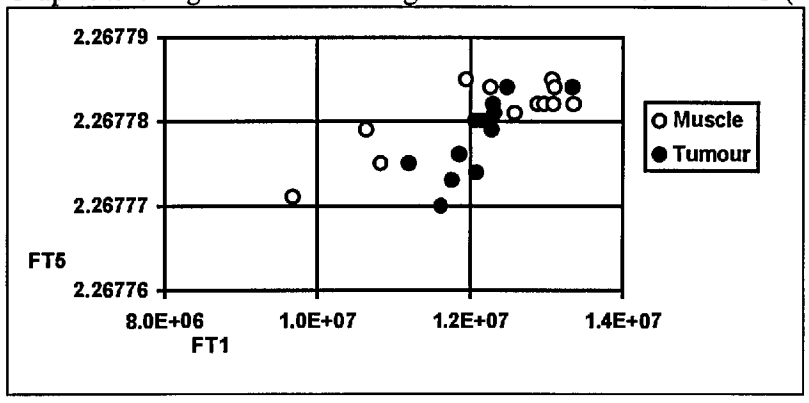
Graph 9.7 : Segmentation of 'Thigh' data - features FT1vFT5 (ROI 16x16)



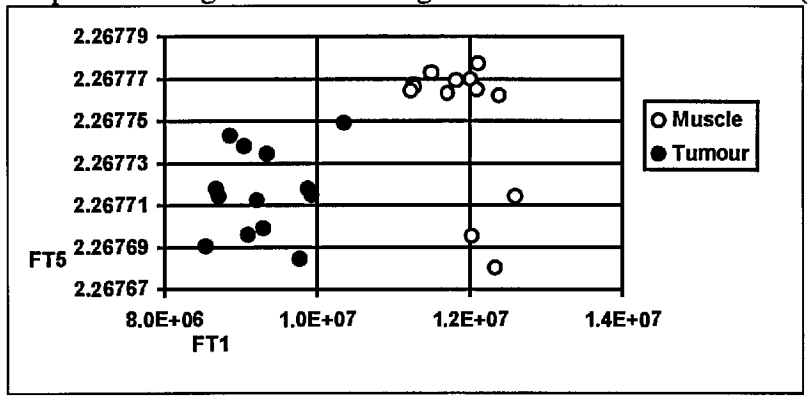
Graph 9.8 : Segmentation of 'Abdomen' data - features FT1vFT5 (ROI 16x16)



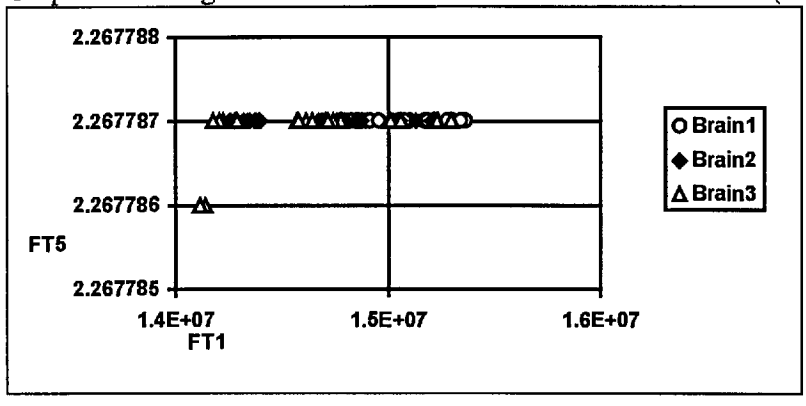
Graph 9.9 : Segmentation of 'Leg104' data - features FT1vFT5 (ROI 16x16)



Graph 9.10 : Segmentation of 'Leg192' data - features FT1vFT5 (ROI 16x16)



Graph 9.11 : Segmentation of 'Brain' data - features FT1vFT5 (ROI 16x16)



9.3 WALSH AND SLANT TEXTURE FEATURES

9.3.1 INTRODUCTION TO WALSH AND SLANT TEXTURE FEATURES

The results of the previous section indicate that the Fourier transform can be used to extract textural information. Because the Walsh and Slant transforms perform similar spectral decomposition to the Fourier transform, but with different sets of basic waveform functions, one can assume that they can be equally well used to extract textural information. This section presents the results obtained with this unique use of the Walsh and Slant transforms for texture analysis. Details of the implementation of these transforms are given in **sections 5.3** and **5.4** in **chapter 5**, respectively.

The results of the transform process can be thought of as an *image,* where position indicates the waveform components and intensity their abundance. Assuming that the texture will have a more pronounced characteristic pattern in the transform *image*, because it comprises of the component waveform functions of the texture, reliable texture descriptors can be obtained by further processing the *image*. A large number texture indices have been obtained by applying four methods of analysis to the transformed texture samples; co-occurrence analysis (**chapter 4 section 4.6**), vertical, horizontal and cluster run length analysis (**chapter 4 sections 4.3** and **4.4**), Fourier analysis (**chapter 5 section 5.2**), and axial symmetry analysis (**chapter 5 section 5.3**). Co-occurrence and run length analysis are applied directly to the transformed data. Fourier analysis is applied to the 1D collapsed transform. Axial symmetry analysis, only possible on the result of the Walsh transform, is also applied directly to the transformed data. **Tables 9.7** and **9.8** summarise the respective post-transform analysis methods and the texture features extracted for the Walsh and Slant transforms. These tables provide a reference between the features and the keys used in the graphs and tables presented throughout this section.

Like the Fourier transform, both Walsh and Slant transforms have dominant DC components. Because these components are common to all image subjects, they must be zeroed to prevent them from dominating the subtler discriminating features of the texture. Both the Walsh and Slant transforms contain signed values, corresponding to a positive or negative contribution of the relevant basis function. In order to perform the four analysis methods described above, the magnitude of the transform data is taken for analysis.

Table 9.7 : Texture features derived from Walsh transformed data

Methods	Texture features	Key
Vertical run length	Short runs emphasis	WVRL1
	Long runs emphasis	WVRL2
	Grey level non-uniformity	WVRL3
	Run length non-uniformity	WVRL4
	Run percentage	WVRL5
Horizontal run length	Short runs emphasis	WHRL1
	Long runs emphasis	WHRL2
	Grey level non-uniformity	WHRL3
	Run length non-uniformity	WHRL4
	Run percentage	WHRL5
Cluster size	Small cluster emphasis	WCS1
	Large cluster emphasis	WCS2
	Grey level non-uniformity	WCS3
	Cluster size non-uniformity	WCS4
	Cluster percentage	WCS5
Fourier transform	Energy	WFT1
	Fractal dimension	WFT2
Axial symmetry	Mean	WAS1
	Standard deviation	WAS2
	Kurtosis	WAS3
	Skewness	WAS4
Co-occurrence	Haralick features F1-F14	WF1-WF14

Table 9.8 : Texture features derived from Slant transformed data

Methods	Texture features	Key
Vertical run length	Short runs emphasis	SVRL1
	Long runs emphasis	SVRL2
	Grey level non-uniformity	SVRL3
	Run length non-uniformity	SVRL4
	Run percentage	SVRL5
Horizontal run length	Short runs emphasis	SHRL1
	Long runs emphasis	SHRL2
	Grey level non-uniformity	SHRL3
	Run length non-uniformity	SHRL4
	Run percentage	SHRL5
Cluster size	Small cluster emphasis	SCS1
	Large cluster emphasis	SCS2
	Grey level non-uniformity	SCS3
	Cluster size non-uniformity	SCS4
	Cluster percentage	SCS5
Fourier transform	Energy	SFT1
	Fractal dimension	SFT2
Co-occurrence	Haralick features F1-F14	SF1-SF14

In this experiment, both transforms have been applied with and without Fourier pre-processing, a process shown to render the process invariant under translation in the conventional pattern recognition application of these methods. Because the Walsh and Slant transforms are very sensitive to image structure, extreme care must be used to ensure that image texture regions are homogenous and contain no additional features eg. tissue boundaries or blood vessels.

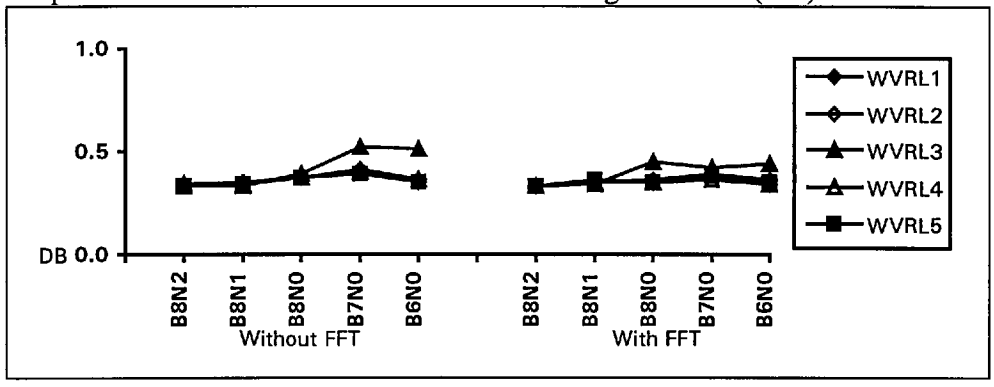
9.3.2 PERFORMANCE OVERVIEW

The aim of this overview is examine the performance trends of the texture features extracted from Walsh and Slant transformations under different pre-processing conditions. The best performing features identified as part of this study are examined in more detail in subsequent sections. The plots in **graphs 9.12** to **9.23** illustrate the discriminative behaviour of these texture features. For each point on the graphs the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (bits and normalisation scheme). The graph legend indicates the texture feature represented by each series of data points. The 'Thigh' MR images used in this overview contain three distinct classes of tissue. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions (muscle, fat & tumour) on several images.

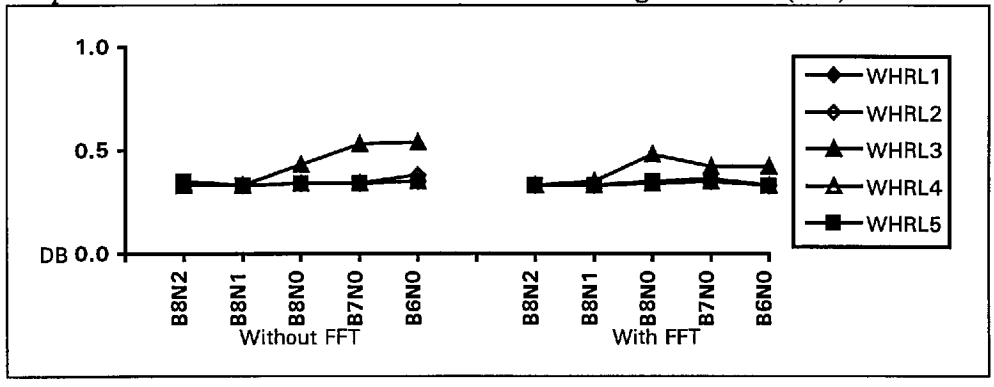
Graphs 9.12 to **9.17** show variation of the DB performance measures for the Walsh transform texture features grouped by method as summarised in **table 9.7**. **Graphs 9.18** to **9.23** show variation of the DB performance measures for the Slant transform texture features grouped by method as summarised in **table 9.8**. All twelve graphs show feature behaviour with number of grey level bits (6-8), normalisation scheme (0-2), and with the use of Fourier transform pre-processing. The data used in these graphs was acquired with 8x8 ROI size.

A large number of texture features are considered in this overview. The purpose of this study is to assess the discriminant ability of each feature and to consider whether it would provide useful texture information. Only those features that demonstrate a satisfactory DB performance measure ($DB > 0.5$) have been considered for closer examination. These features have been identified from **graphs 9.12** to **9.23**.

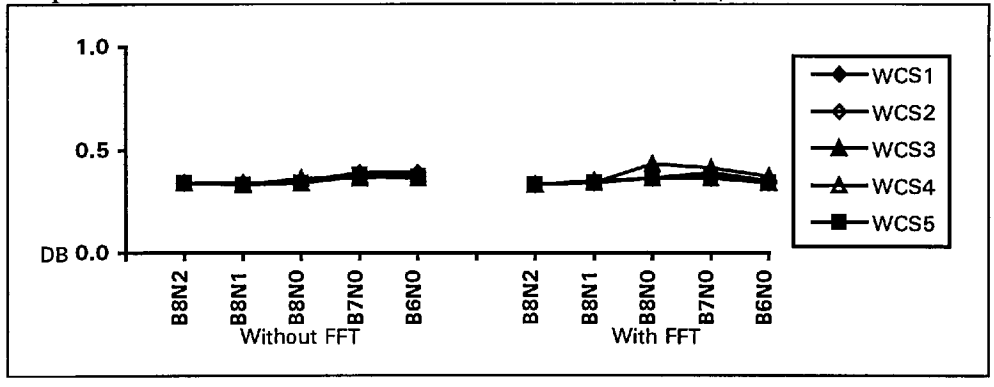
Graph 9.12 : Behaviour of Walsh vertical run length features (8x8)



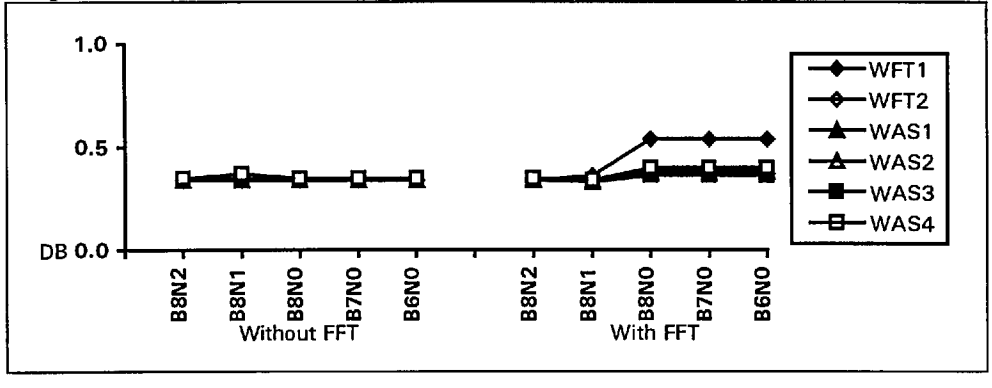
Graph 9.13 : Behaviour of Walsh horizontal run length features (8x8)



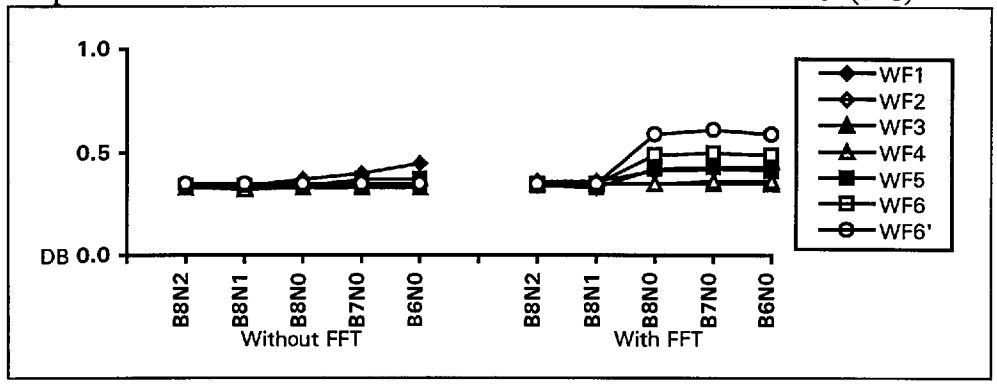
Graph 9.14 : Behaviour of Walsh cluster size features (8x8)



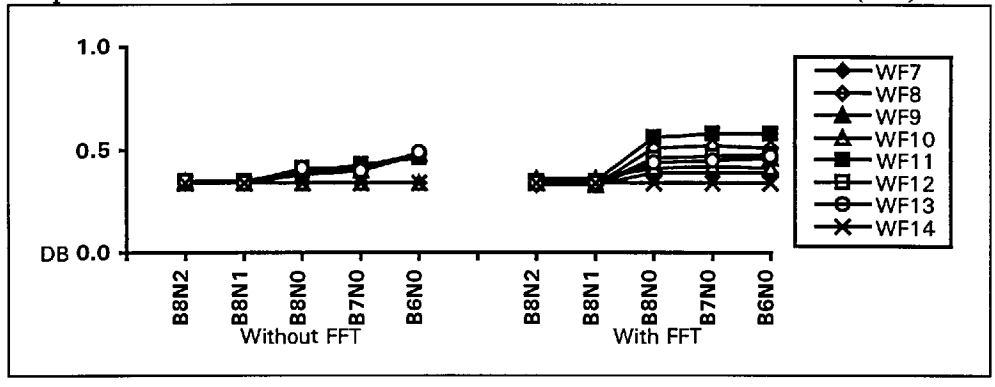
Graph 9.15 : Behaviour of Walsh Fourier and axial symmetry features (8x8)



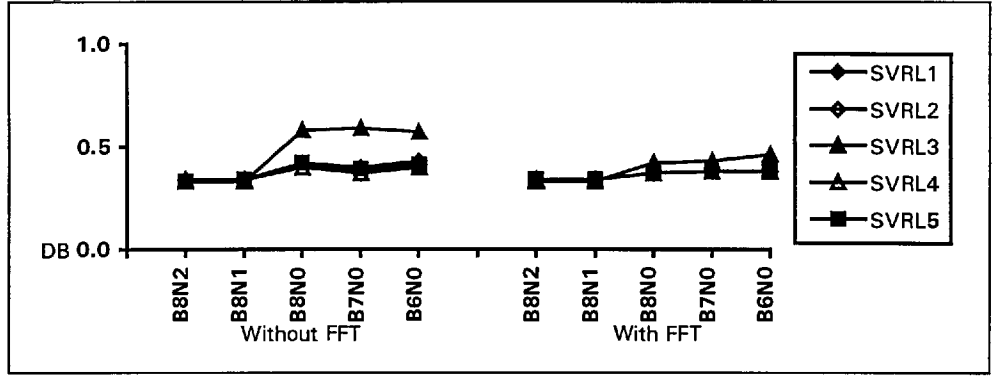
Graph 9.16 : Behaviour of Walsh co-occurrence features WF1-WF6' (8x8)



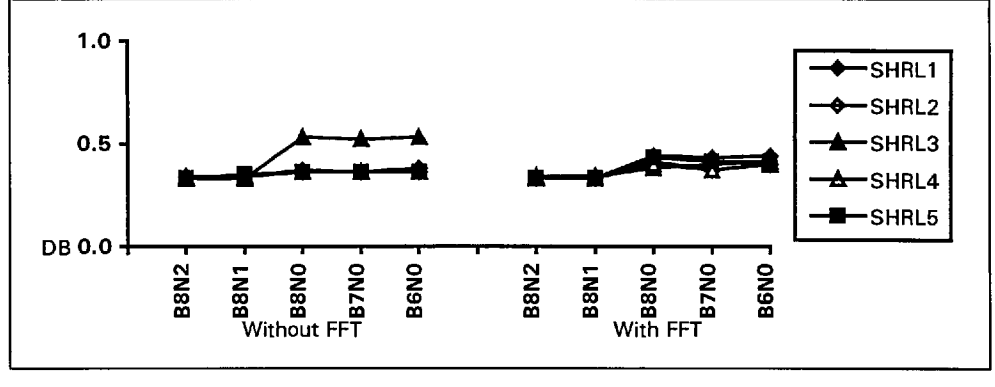
Graph 9.17 : Behaviour of Walsh co-occurrence features WF7-WF14 (8x8)



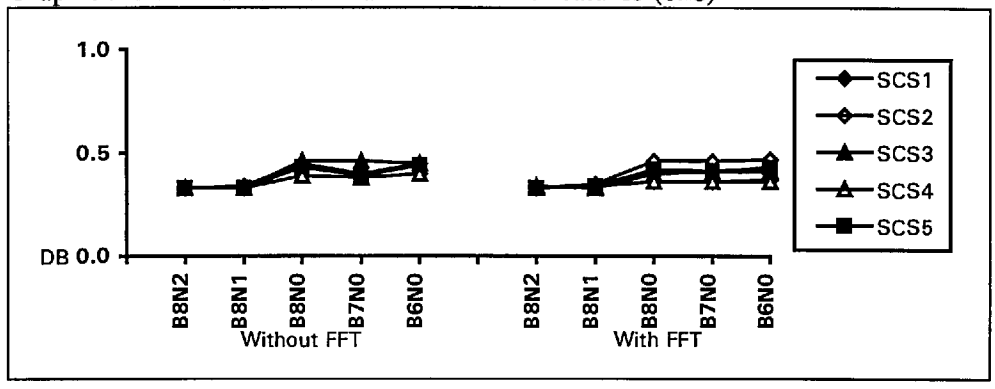
Graph 9.18 : Behaviour of Slant vertical run length features (8x8)



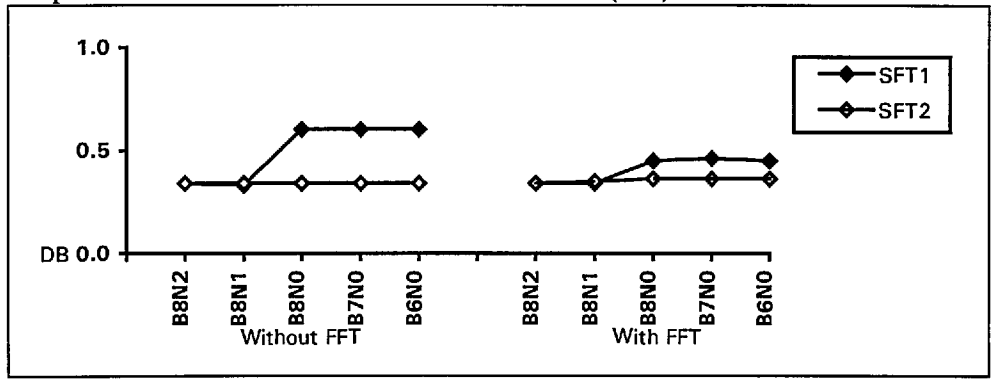
Graph 9.19 : Behaviour of Slant horizontal run length features (8x8)



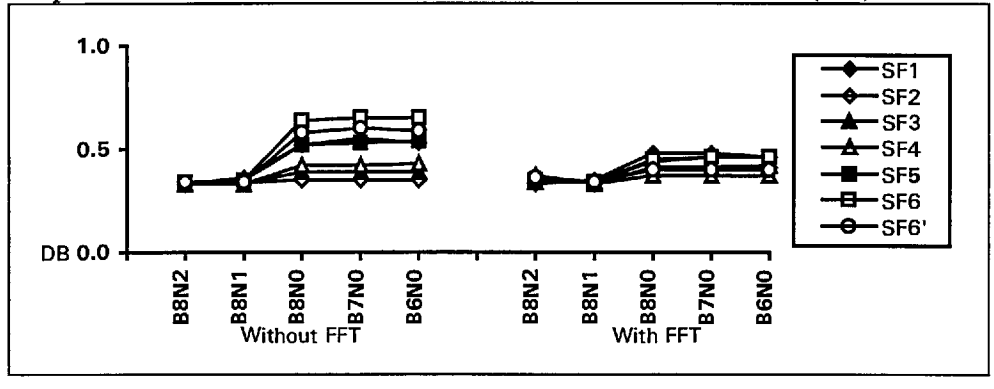
Graph 9.20 : Behaviour of Slant cluster size features (8x8)



Graph 9.21 : Behaviour of Slant Fourier features (8x8)



Graph 9.22 : Behaviour of Slant co-occurrence features SF1-SF6' (8x8)



Graph 9.23 : Behaviour of Slant co-occurrence features SF7-SF14 (8x8)

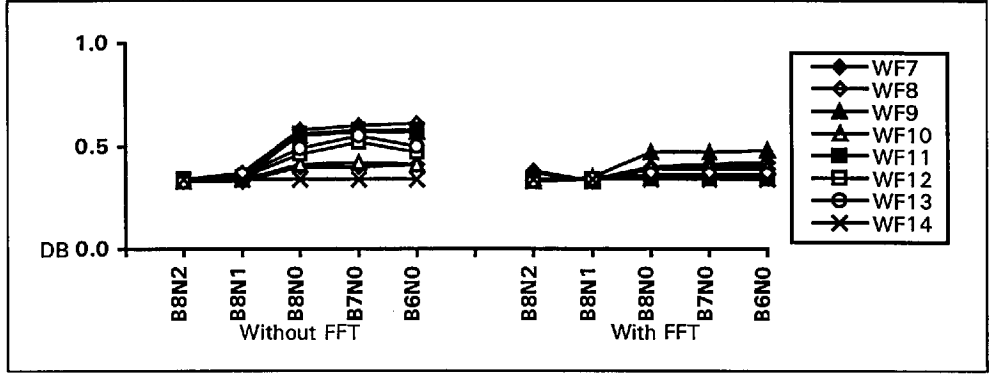
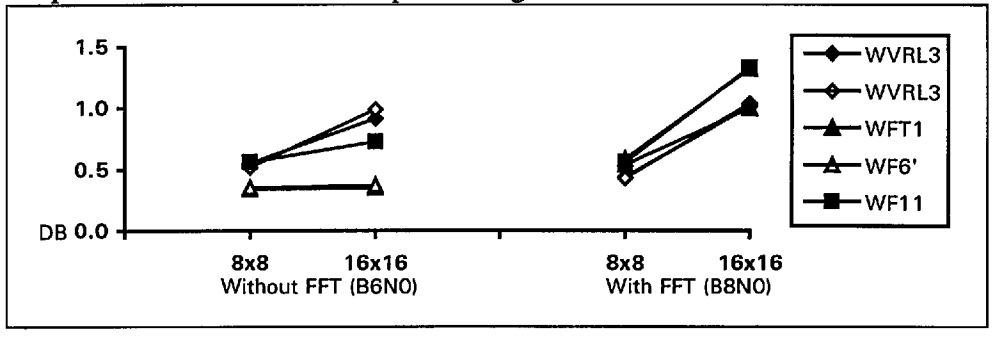


Table 9.9 provides a short-list of the texture features derived from Walsh and Slant transformations that demonstrate a satisfactory level of performance. Normalisation was found to be detrimental to discriminative performance in all cases. Optimum performances were achieved with 8-bit data for Fourier pre-processing, and with 6-bit data without Fourier pre-processing. **Graphs 9.24** and **9.25** summarise the performances of the selected Walsh and Slant texture features, while also demonstrating the influence of ROI size. Better discriminative performances are generally achieved for ROI of size 16x16 rather than 8x8.

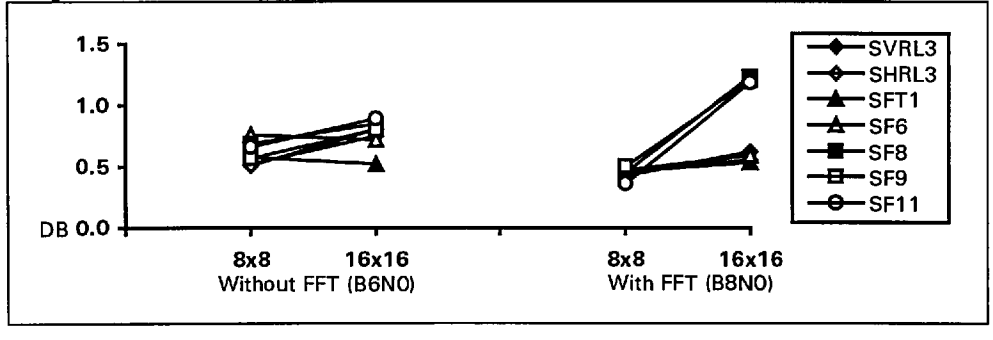
Table 9.9 : Best performing Walsh and Slant features

Transform	Feature	Processing method
Walsh	WVRL3	Run length (vertical)
	WHRL3	Run length (horizontal)
	WFT1	Fourier
	WF6'	Co-occurrence
	WF11	Co-occurrence
Slant	SVRL3	Run length (vertical)
	SHRL3	Run length (horizontal)
	SFT1	Fourier
	SF6	Co-occurrence
	SF8	Co-occurrence
	SF9	Co-occurrence
	SF11	Co-occurrence

Graph 9.24 : Behaviour of best performing Walsh features



Graph 9.25 : Behaviour of best performing Slant features



9.3.3 CORRELATION STUDY

A large number of texture features derived from both the Walsh and Slant transform have been identified in the previous section as providing valuable texture information. Any large group of features derived from one method is likely to contain redundant information. Redundant features can be eliminated by studying the correlation between the features for each transform. More objective tissue characterisation can be achieved with the smaller features sets that result from such studies.

The correlation between each pair of texture features is measured over 78 ROI evenly distributed among the three tissue classes on 'Thigh' image data. This is the same source of data used for the discriminative performance evaluation in **section 9.3.2**. The texture features are collected from ROI of size 8x8 and 16x16 with (8-bit data) and without (6-bit data) Fourier pre-processing, and without data normalisation. The correlation coefficient is calculated in Microsoft EXCEL from the function libraries (see **chapter 7** for details). The coefficient varies between 0 (no correlation) and 1 (high correlation), and a negative sign indicates that the two features are inversely correlated.

Tables 9.10 to **9.13** present the correlation coefficients between the five Walsh texture features identified in **table 9.9** for different ROI size, with and without Fourier (FFT) pre-processing. These values enable elimination of redundant Walsh texture features.

Table 9.10 : Walsh correlation (8x8)

	WVRL3	WHRL3	WFT1	WF6'	WF11
WVRL3	-	0.90	-0.19	-0.11	-0.41
WHRL3	0.90	-	-0.19	-0.12	-0.45
WFT1	-0.19	-0.19	-	0.72	0.48
WF6'	-0.11	-0.12	0.72	-	0.66
WF11	-0.41	-0.45	0.48	0.66	-

Table 9.11 : Walsh correlation (16x16)

	WVRL3	WHRL3	WFT1	WF6'	WF11
WVRL3	-	0.98	-0.78	-0.89	-0.91
WHRL3	0.98	-	-0.77	-0.87	-0.89
WFT1	-0.78	-0.77	-	0.86	0.83
WF6'	-0.89	-0.87	0.86	-	0.98
WF11	-0.91	-0.89	0.83	0.98	-

Table 9.12 : Walsh correlation (8x8 FFT)

	WVRL3	WHRL3	WFT1	WF6'	WF11
WVRL3	-	0.88	0.55	-0.83	-0.87
WHRL3	0.88	-	0.57	-0.82	-0.85
WFT1	0.55	0.57	-	-0.66	-0.69
WF6'	-0.83	-0.82	-0.66	-	0.98
WF11	-0.87	-0.85	-0.69	0.98	-

Table 9.13 : Walsh correlation (16x16 FFT)

	WVRL3	WHRL3	WFT1	WF6'	WF11
WVRL3	-	0.98	0.86	-0.96	-0.97
WHRL3	0.98	-	0.85	-0.97	-0.97
WFT1	0.86	0.85	-	-0.91	-0.92
WF6'	-0.96	-0.97	-0.91	-	0.99
WF11	-0.97	-0.97	-0.92	0.99	-

The optimum conditions for extracting Walsh texture features without Fourier pre-processing is with 6-bit image data and without normalisation. The correlation coefficients between Walsh features under these conditions are provided in **tables 9.10** and **9.11**. Features WVRL3 and WHRL3 are both highly correlated, due to the inherent symmetries of the Walsh transform matrix. WHRL3 is arbitrarily eliminated because both features correlate the same with other features and their discriminative performances are almost identical. There is some correlation between features WF6' and WFT1, but because they perform generally worse than other features they can also be eliminated.

The optimum conditions for extracting Walsh texture features with Fourier pre-processing is with 8-bit image data and without normalisation. The correlation coefficients between Walsh features under these conditions are provided in **tables 9.12** and **9.13**. Both features WVRL3 and WHRL3 are highly correlated and are eliminated because of their poor performance. Although WF6' and WF11 are highly correlated and of comparable performance, WF11 is arbitrarily eliminated.

Tables 9.14 to **9.17** present the coefficients between the seven Slant texture features identified in **table 9.9** for different ROI size, with and without Fourier (FFT) pre-processing. These values enable elimination of redundant Slant texture features.

Table 9.14 : Slant correlation (8x8)

	SVRL3	SHRL3	SFT1	SF6	SF8	SF9	SF11
SVRL3	-	0.94	-0.78	-0.85	-0.90	-0.95	-0.80
SHRL3	0.94	-	-0.77	-0.84	-0.88	-0.95	-0.76
SFT1	-0.78	-0.77	-	0.93	0.88	0.81	0.86
SF6	-0.85	-0.84	0.93	-	0.95	0.90	0.92
SF8	-0.90	-0.88	0.88	0.95	-	0.95	0.93
SF9	-0.95	-0.95	0.81	0.90	0.95	-	0.85
SF11	-0.80	-0.76	0.86	0.92	0.93	0.85	-

Table 9.15 : Slant correlation (16x16)

	SVRL3	SHRL3	SFT1	SF6	SF8	SF9	SF11
SVRL3	-	0.79	-0.78	0.15	-0.89	-0.91	-0.87
SHRL3	0.79	-	-0.74	-0.19	-0.88	-0.83	-0.81
SFT1	-0.78	-0.74	-	0.13	0.84	0.78	0.87
SF6	0.15	-0.19	0.13	-	0.16	-0.02	0.08
SF8	-0.89	-0.88	0.84	0.16	-	0.96	0.94
SF9	-0.91	-0.83	0.78	-0.02	0.96	-	0.93
SF11	-0.87	-0.81	0.87	0.08	0.94	0.93	-

Table 9.16 : Slant correlation (8x8 FFT)

	SVRL3	SHRL3	SFT1	SF6	SF8	SF9	SF11
SVRL3	-	0.60	-0.59	-0.60	-0.73	-0.81	-0.47
SHRL3	0.60	-	-0.54	-0.54	-0.74	-0.78	-0.45
SFT1	-0.59	-0.54	-	0.98	0.70	0.71	0.49
SF6	-0.60	-0.54	0.98	-	0.72	0.73	0.46
SF8	-0.73	-0.74	0.70	0.72	-	0.92	0.64
SF9	-0.81	-0.78	0.71	0.73	0.92	-	0.58
SF11	-0.47	-0.45	0.49	0.46	0.64	0.58	-

Table 9.17 : Slant correlation (16x16 FFT)

	SVRL3	SHRL3	SFT1	SF6	SF8	SF9	SF11
SVRL3	-	0.38	-0.50	-0.50	-0.71	-0.72	-0.73
SHRL3	0.38	-	-0.53	-0.60	-0.52	-0.54	-0.46
SFT1	-0.50	-0.53	-	0.98	0.71	0.71	0.71
SF6	-0.50	-0.60	0.98	-	0.77	0.78	0.76
SF8	-0.71	-0.52	0.71	0.77	-	1.00	0.99
SF9	-0.72	-0.54	0.71	0.78	1.00	-	0.99
SF11	-0.73	-0.46	0.71	0.76	0.99	0.99	-

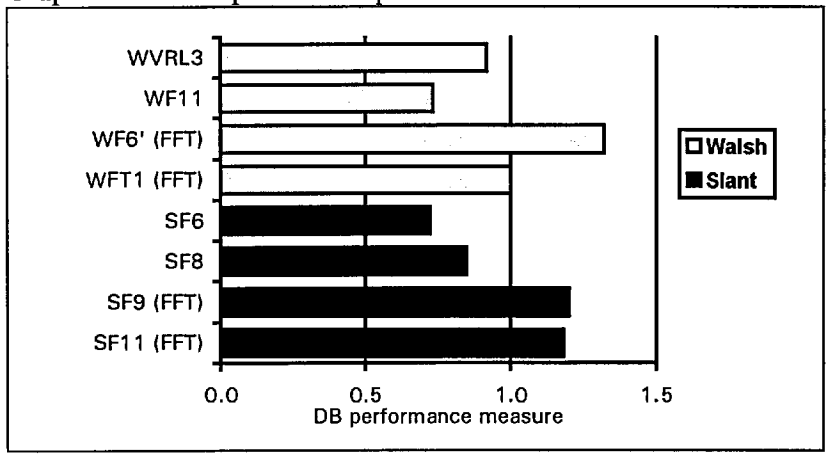
The optimum conditions for extracting Slant texture features without Fourier pre-processing is with 6-bit image data and without normalisation. The correlation coefficients between Slant features under these conditions are provided in **tables 9.14** and **9.15**. Features SVRL3 and SHRL3 are highly correlated, due to the inherent symmetries of the Slant transform matrix. The three highly correlated co-occurrence features SF8, SF9 and SF11 and the Fourier feature SFT1 all correlate well with these run length features. All but one of the features in this group can therefore be eliminated. Feature SF8 is retained because it demonstrated marginally better discriminative performance. SF6 is largely un-correlated with any of the other Slant features and demonstrates excellent performance by comparison.

The optimum conditions for extracting Slant texture features with Fourier pre-processing is with 8-bit image data and without normalisation. The correlation coefficients between Slant features under these conditions are provided in **tables 9.16** and **9.17**. Comparable performance was demonstrated by all seven features for ROI size 8x8. At larger ROI, however, the features could be divided into two distinct groupings by performance. The four features in the lower performance group are eliminated (SVRL3, SHRL3, SFT1 and SF6). The three remaining features are highly correlated (SF8, SF9 and SF11). Feature SF8 is arbitrarily eliminated from this scheme, leaving features SF9 and SF11 for tissue characterisation.

9.3.4 OPTIMISED FEATURE SET

Optimum discriminant performance was achieved for all Walsh and Slant features without data normalisation, and with 6-bit and 8-bit data respectively for implementations without and with Fourier pre-processing. Better performances were achieved from ROI of size 16x16 rather than 8x8. This result is consistent with the improved specificity of the transform matrix that is expected with larger ROI ie. more data points. A comparison of the DB performance of the two sets of each Walsh and Slant texture features acquired under optimum conditions are presented in **graph 9.26**, where (FFT) indicates the features acquired with Fourier pre-processing. This graph demonstrates the benefit of FFT pre-processing on performance. The following section provides examples of the application of these four feature sets for tissue characterisation.

Graph 9.26 : Comparison of optimised Walsh and Slant features



9.3.5 EXAMPLES OF TISSUE CHARACTERISATION

This section provides examples of the application of the optimised texture features derived from the Walsh and Slant transforms for discrimination between tissue classes in a variety of MRI images. Reference information about the optimised texture features is provided in **table 9.9** in the previous section of this chapter. **Table 9.18** provides a list of the MRI image data sets investigated and the tissue classes they contain. The tissue characterisation achieved in the case of each data set is illustrated in **graphs 9.27** to **9.46**. A summary of the information presented in each graph is provided in **table 9.19** on the following page. All graph values are acquired from ROI of size 16x16.

Table 9.18 : Description of image data

Data set	Tissue classes
Thigh	Muscle, Tumour, Fat
Abdomen	Fat, Liver, Tumour
Leg104	Muscle, Tumour
Leg192	Muscle, Tumour
Brains	Brain1, Brain2, Brain3

Table 9.19 : Examples of tissue segmentation with Walsh and Slant texture features

Graph	Data sample	Transform	X-axis feature	Y-axis feature	Conditions
9.27	Thigh	Walsh	WVRL3	WF11	B6N0 - No FFT
9.28		Walsh	WF6	WFT1	B8N0 - FFT
9.29		Slant	SF6	SF8	B6N0 - No FFT
9.30		Slant	SF9	SF11	B8N0 - FFT
9.31	Abdomen	Walsh	WVRL3	WF11	B6N0 - No FFT
9.32		Walsh	WF6	WFT1	B8N0 - FFT
9.33		Slant	SF6	SF8	B6N0 - No FFT
9.34		Slant	SF9	SF11	B8N0 - FFT
9.35	Leg104	Walsh	WVRL3	WF11	B6N0 - No FFT
9.36		Walsh	WF6	WFT1	B8N0 - FFT
9.37		Slant	SF6	SF8	B6N0 - No FFT
9.38		Slant	SF9	SF11	B8N0 - FFT
9.39	Leg192	Walsh	WVRL3	WF11	B6N0 - No FFT
9.40		Walsh	WF6	WFT1	B8N0 - FFT
9.41		Slant	SF6	SF8	B6N0 - No FFT
9.42		Slant	SF9	SF11	B8N0 - FFT
9.43	Brains	Walsh	WVRL3	WF11	B6N0 - No FFT
9.44		Walsh	WF6	WFT1	B8N0 - FFT
9.45		Slant	SF6	SF8	B6N0 - No FFT
9.46		Slant	SF9	SF11	B8N0 - FFT

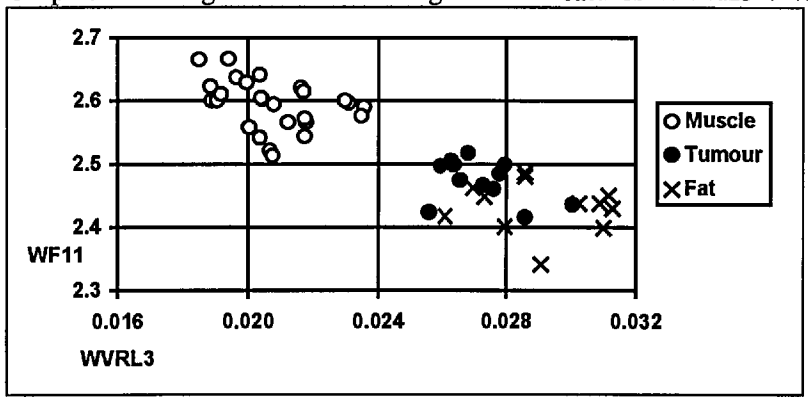
Graphs 9.27 to **9.30** show good examples of successful segmentation between the crucial muscle tissue and muscle tumour tissue classes in the 'Thigh' image data set. Despite some overlap between the muscle tumour and fat tissue class data points, all three tissue classes were tightly clustered.

Graphs 9.31 to **9.34** show the segmentation possible between liver tissue and tumour tissue within the liver from the 'Abdomen' data set. Promising separation between these two crucial tissue classes was demonstrated with the two sets of Walsh transform texture features in **graphs 9.31** and **9.32**. The boundary between these two tissue classes in feature space for the two sets of Slant texture features in **graphs 9.33** and **9.34** is, however, less well defined.

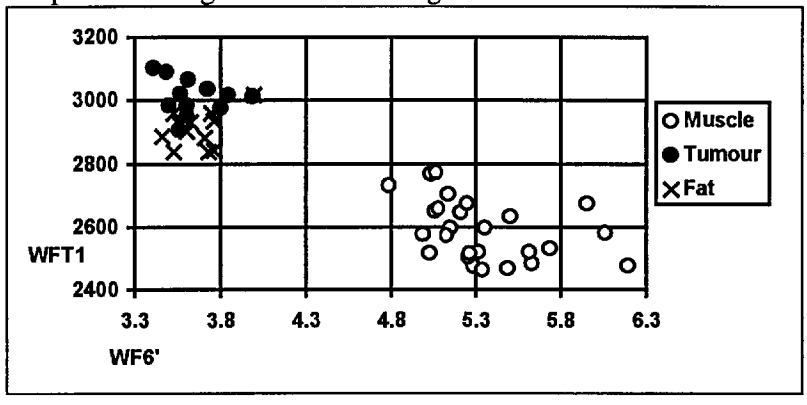
Graphs 9.35 to **9.42** show two cases of the segmentation of muscle tumour from muscle tissue in the leg; 'Leg104' and 'Leg192' image data sets. The feature space boundaries between the muscle and tumour tissue classes in the 'Leg104' examples given in **graphs 9.35** to **9.38** are not clearly defined. There is, however, evidence of clustering of same class data points indicating a degree of tissue segmentation. In the case of the 'Leg192' examples presented in **graphs 9.39** to **9.42**, feature space is better segmented into the two tissue classes. The overall performance of the two sets of both Walsh and Slant features are comparable for this application.

Graphs 9.43 to **9.46** show the results for Walsh and Slant texture features extracted from grey matter in a series of mid-brain slices in three healthy volunteers; 'Brain' image data. Feature stability can be assessed by examining the variability of data points both in each subject class and in all three classes combined. For good stability, one is looking for minimal variability firstly in each class and then in all classes. According to these subjective criteria, the Walsh features in **graphs 9.43** and **9.44** appear to be highly stable. The larger variation seen between the data points for the Slant features in **graphs 9.45** and **9.46** indicates that they are less stable.

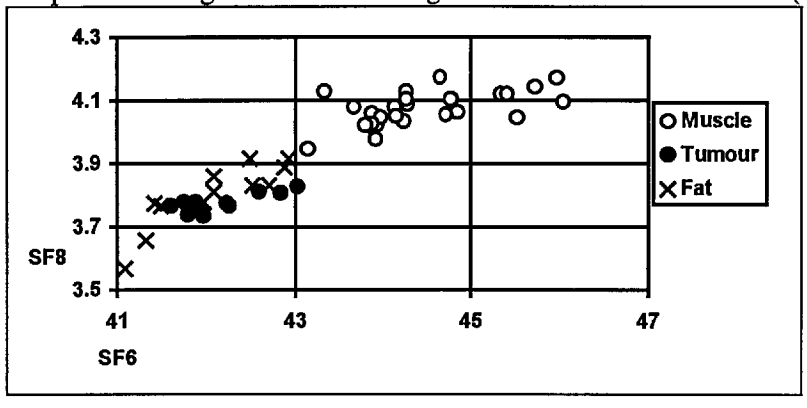
Graph 9.27 : Segmentation of 'Thigh' data - features WVRL3 v WF11 (No FFT)



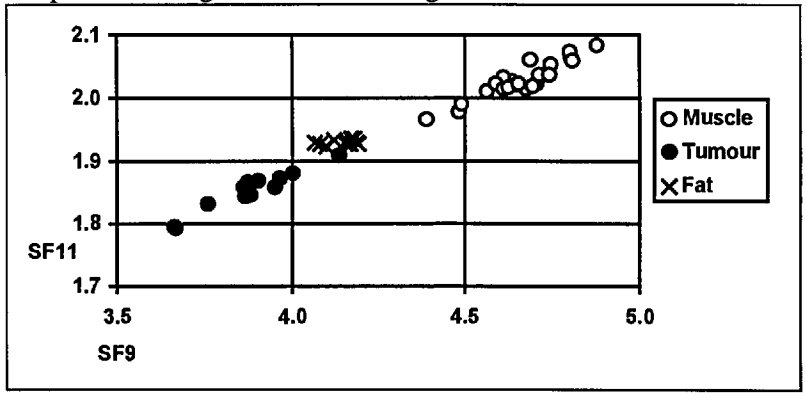
Graph 9.28 : Segmentation of 'Thigh' data - features WF6' v WFT1 (FFT)



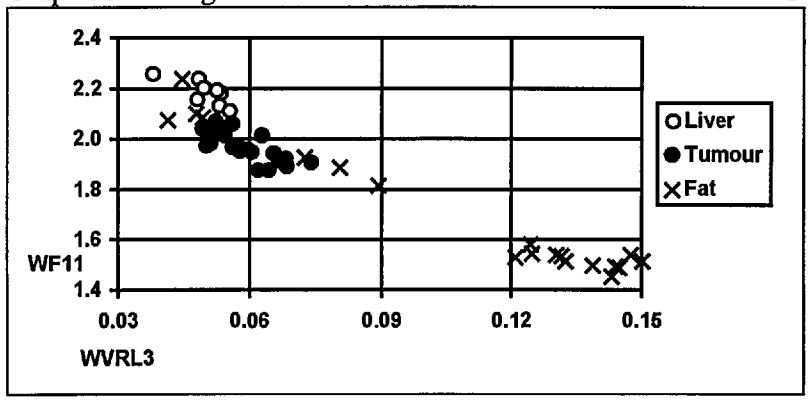
Graph 9.29 : Segmentation of 'Thigh' data - features SF6 v SF8 (No FFT)



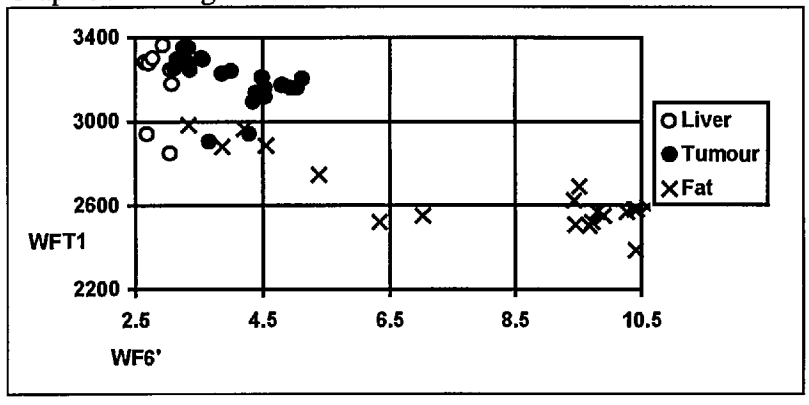
Graph 9.30 : Segmentation of 'Thigh' data - features SF9 v SF11 (FFT)



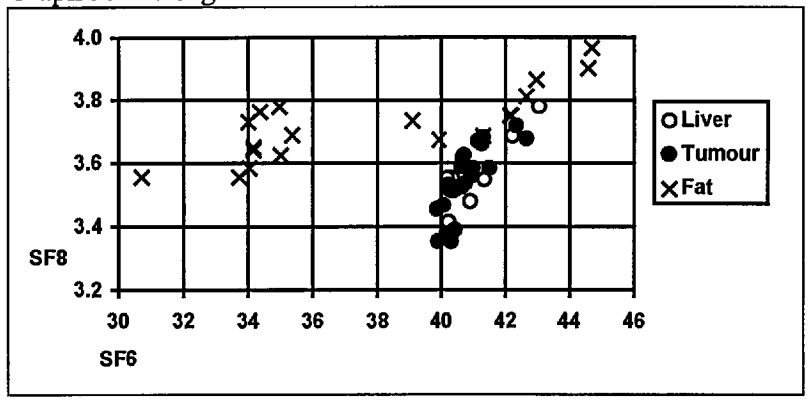
Graph 9.31 : Segmentation of 'Abdomen' data - features WVRL3 v WF11 (No FFT)



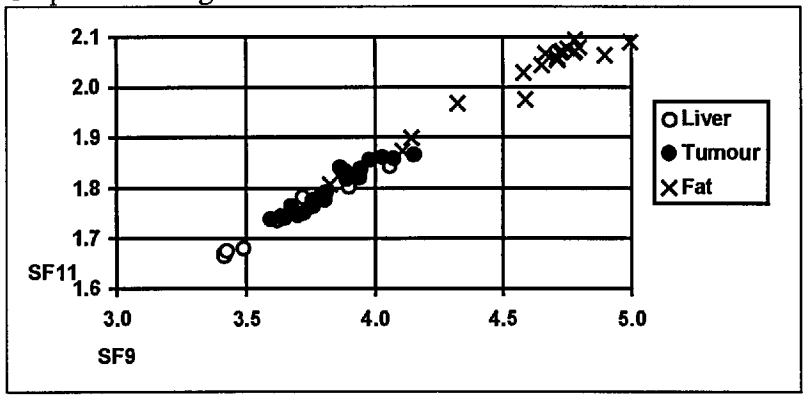
Graph 9.32 : Segmentation of 'Abdomen' data - features WF6' v WFT1 (FFT)



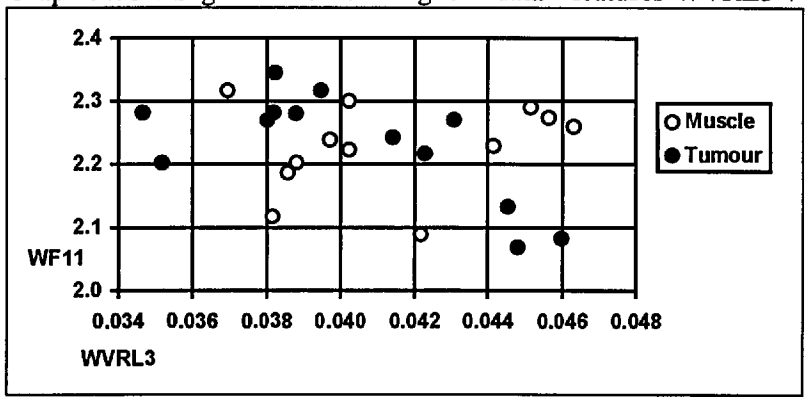
Graph 9.33 : Segmentation of 'Abdomen' data - features SF6 v SF8 (No FFT)



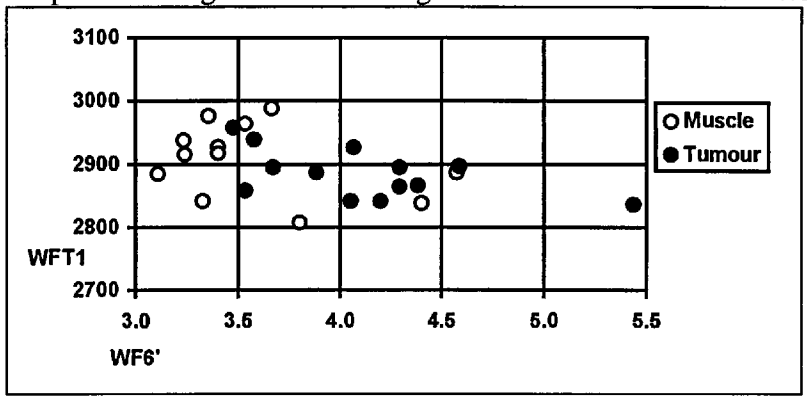
Graph 9.34 : Segmentation of 'Abdomen' data - features SF9 v SF11 (FFT)



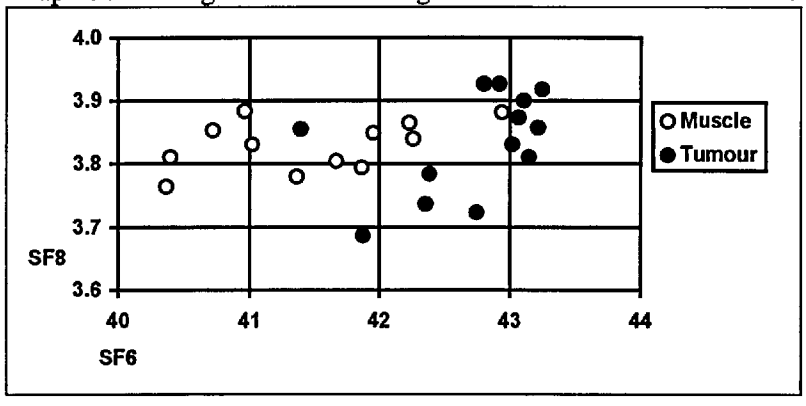
Graph 9.35 : Segmentation of 'Leg104' data - features WVRL3 v WF11 (No FFT)



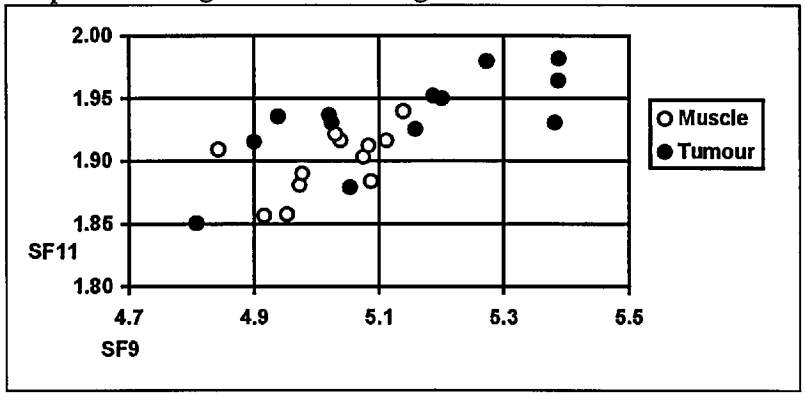
Graph 9.36 : Segmentation of 'Leg104' data - features WF6' v WFT1 (FFT)



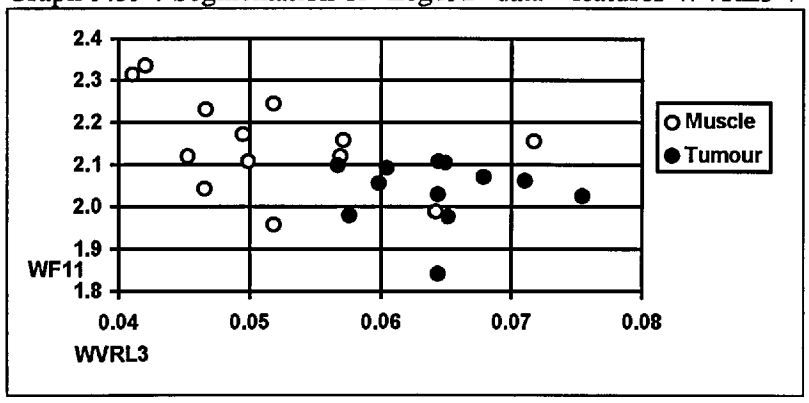
Graph 9.37 : Segmentation of 'Leg104' data - features SF6 v SF8 (No FFT)



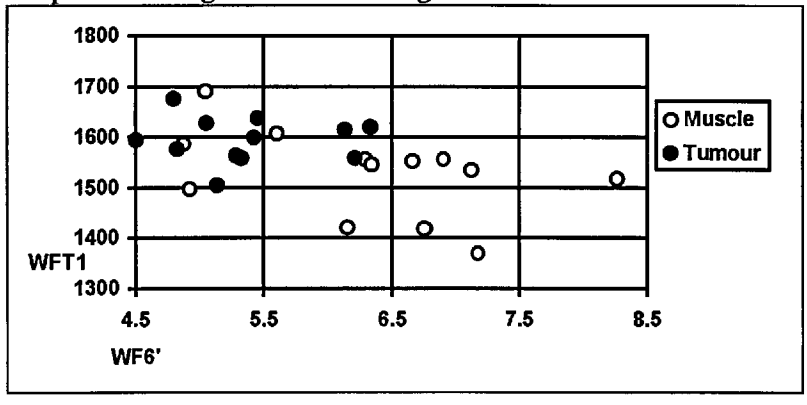
Graph 9.38 : Segmentation of 'Leg104' data - features SF9 v SF11 (FFT)



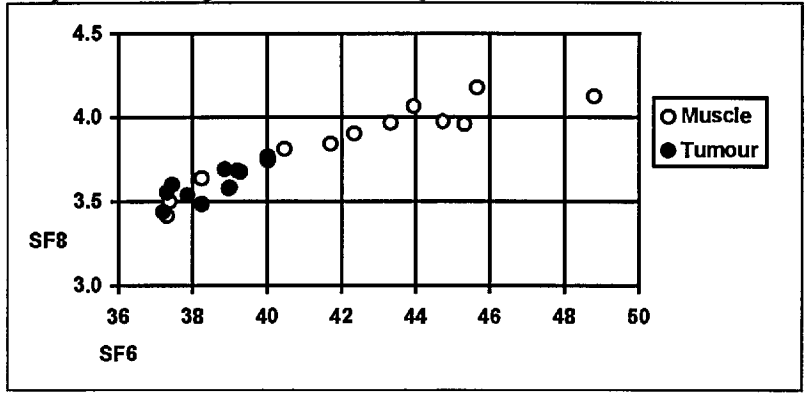
Graph 9.39 : Segmentation of 'Leg192' data - features WVRL3 v WF11 (No FFT)



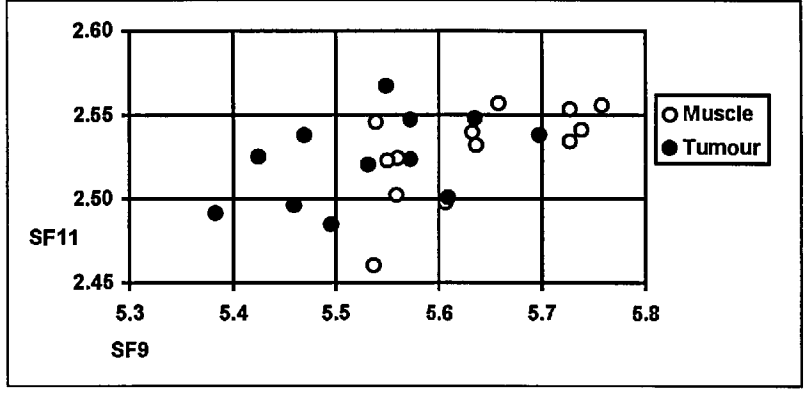
Graph 9.40 : Segmentation of 'Leg192' data - features WF6' v WFT1 (FFT)



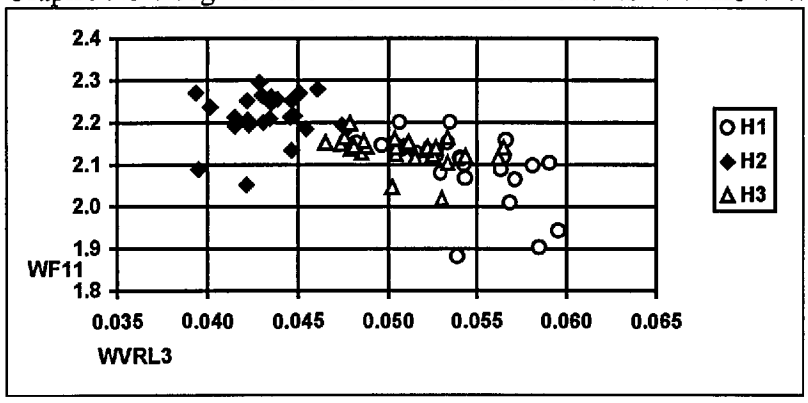
Graph 9.41 : Segmentation of 'Leg192' data - features SF6 v SF8 (No FFT)



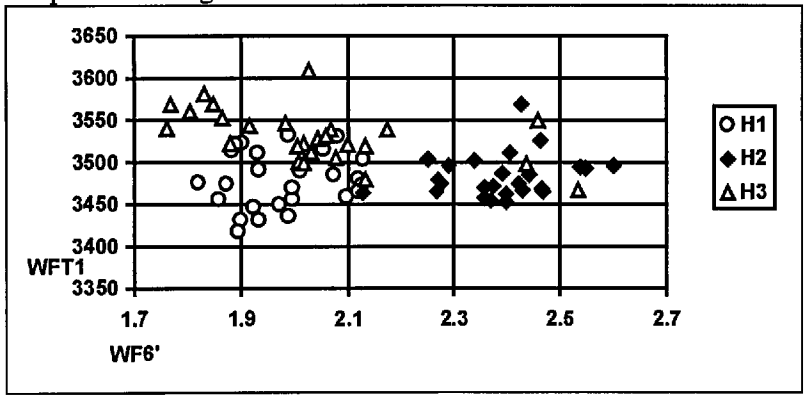
Graph 9.42 : Segmentation of 'Leg192' data - features SF9 v SF11 (FFT)



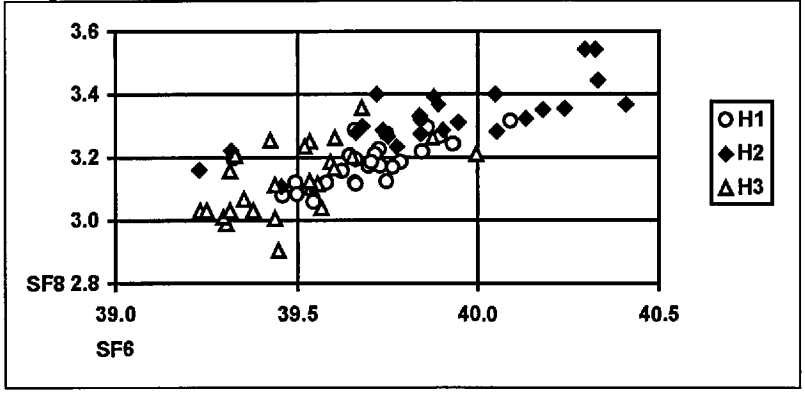
Graph 9.43 : Segmentation of 'Brain' data - features WVRL3 v WF11 (No FFT)



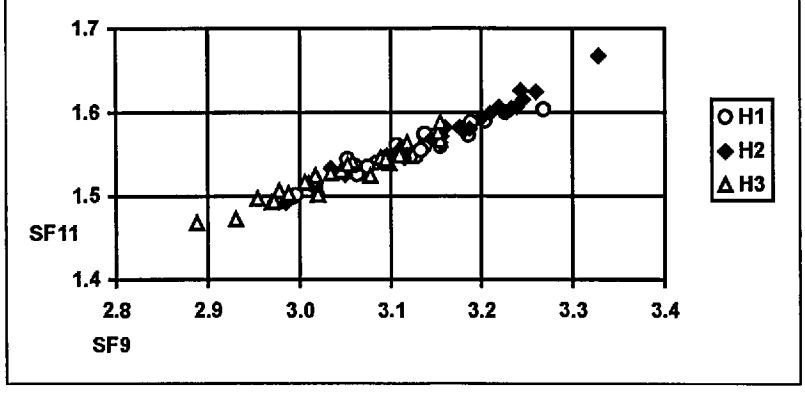
Graph 9.44 : Segmentation of 'Brain' data - features WF6' v WFT1 (FFT)



Graph 9.45 : Segmentation of 'Brain' data - features SF6 v SF8 (No FFT)



Graph 9.46 : Segmentation of 'Brain' data - features SF9 v SF11 (FFT)



9.4 NON-DIRECTIONAL WALSH & SLANT TEXTURE FEATURES

9.4.1 INTRODUCTION TO NON-DIRECTIONAL WALSH AND SLANT

The results of the **section 9.3** demonstrate that proficient texture features can be derived from both Walsh and Slant transforms in a similar method as that applied to the Fourier transform in **section 9.2**. The standard implementations of the Walsh and Slant transforms decompose an image by its geometrical properties. However, because of the significant variation of tissue orientation in clinical MRI images, these methods are subject to inconsistencies. A totally original method of implementing the Walsh and Slant transform in a texture classification scheme that proposes to be insensitive to texture orientation has been suggested as part of this project. This method is described in **section 5.5** in **chapter 5**. This section presents a study of the performance of the texture features derived from these non-directional implementations of the Walsh and Slant transforms.

Like the standard Walsh and Slant implementation described in the previous section, a large number texture indices have been obtained by applying four methods of analysis to the transformed texture samples; co-occurrence analysis (**chapter 4 section 4.6**), vertical, horizontal and cluster run length analysis (**chapter 4 sections 4.3** and **4.4**), Fourier analysis (**chapter 5 section 5.2**), and axial symmetry analysis (**chapter 5 section 5.3**). Co-occurrence and run length analysis are applied directly to the transformed data. Fourier analysis is applied to the 1D collapsed transform. Axial symmetry analysis, only possible on the result of the Walsh transform, is also applied directly to the transformed data. **Tables 9.20** and **9.21** summarise the respective post-transform analysis methods and the texture features extracted for the Walsh and Slant transforms. These tables provide a reference between the features and the keys used in the graphs and tables presented throughout this section.

As with the standard implementations of the Walsh and Slant transform, the DC components in the image transforms must be zeroed to prevent them from dominating the subtler discriminating features of the texture. Both the Walsh and Slant transforms contain signed values, corresponding to a positive or negative contribution of the relevant basis function. In order to perform the four analysis methods described above, the magnitude of the transform data is taken for analysis.

Table 9.20 : Texture features derived from non-dir^n Walsh transformed data

Methods	Texture features	Key
Vertical run length	Short runs emphasis	NWVRL1
	Long runs emphasis	NWVRL2
	Grey level non-uniformity	NWVRL3
	Run length non-uniformity	NWVRL4
	Run percentage	NWVRL5
Horizontal run length	Short runs emphasis	NWHRL1
	Long runs emphasis	NWHRL2
	Grey level non-uniformity	NWHRL3
	Run length non-uniformity	NWHRL4
	Run percentage	NWHRL5
Cluster size	Small cluster emphasis	NWCS1
	Large cluster emphasis	NWCS2
	Grey level non-uniformity	NWCS3
	Cluster size non-uniformity	NWCS4
	Cluster percentage	NWCS5
Fourier transform	Energy	NWFT1
	Fractal dimension	NWFT2
Axial symmetry	Mean	NWAS1
	Standard deviation	NWAS2
	Kurtosis	NWAS3
	Skewness	NWAS4
Co-occurrence	Haralick features F1-F14	NWF1-NWF14

Table 9.21 : Texture features derived from non-dir^n Slant transformed data

Methods	Texture features	Key
Vertical run length	Short runs emphasis	NSVRL1
	Long runs emphasis	NSVRL2
	Grey level non-uniformity	NSVRL3
	Run length non-uniformity	NSVRL4
	Run percentage	NSVRL5
Horizontal run length	Short runs emphasis	NSHRL1
	Long runs emphasis	NSHRL2
	Grey level non-uniformity	NSHRL3
	Run length non-uniformity	NSHRL4
	Run percentage	NSHRL5
Cluster size	Small cluster emphasis	NSCS1
	Large cluster emphasis	NSCS2
	Grey level non-uniformity	NSCS3
	Cluster size non-uniformity	NSCS4
	Cluster percentage	NSCS5
Fourier transform	Energy	NSFT1
	Fractal dimension	NSFT2
Co-occurrence	Haralick features F1-F14	NSF1-NSF14

Both transforms have been applied with and without Fourier pre-processing, a process shown to render the process invariant under translation in the conventional pattern recognition application of these methods. Because the Walsh and Slant transforms are very sensitive to image structure, extreme care must be used to ensure that image texture regions are homogenous and contain no additional features eg. tissue boundaries or blood vessels.

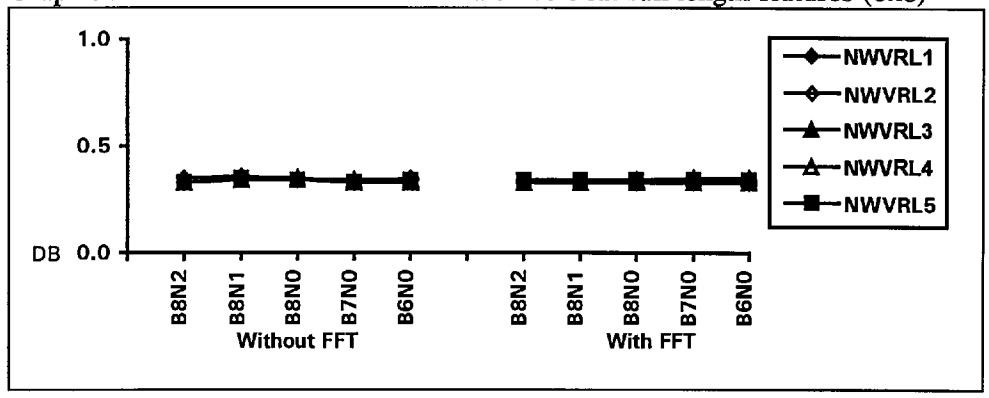
9.4.2 PERFORMANCE OVERVIEW

The aim of this overview is examine the performance trends of the texture features extracted from non-directional Walsh and Slant transformations under different pre-processing conditions. The best performing features identified as part of this study are examined in more detail in subsequent sections. The plots in **graphs 9.47** to **9.58** illustrate the discriminative behaviour of these texture features. For each point on the graphs, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (bits and normalisation scheme). The graph legend indicates the texture feature represented by each series of data points. The 'Thigh' MRI images used in this overview contain three distinct classes of tissue. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions (muscle, fat & tumour) on several images.

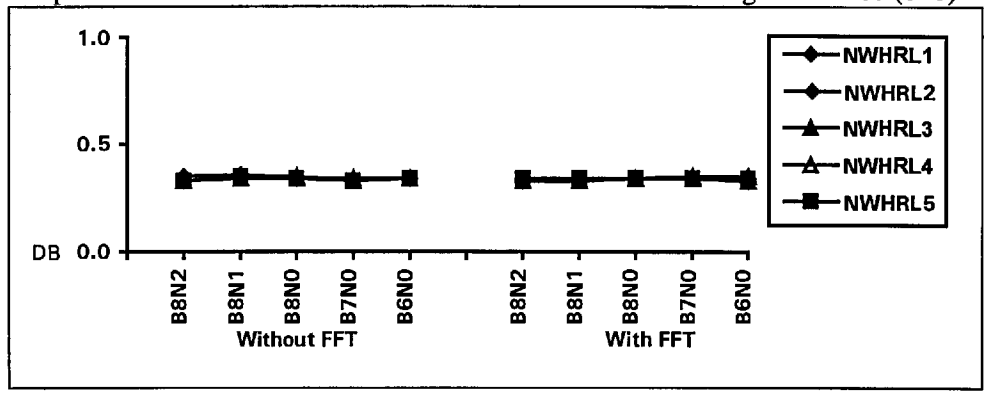
Graphs 9.47 to **9.52** show variation of the DB performance measures for the non-directional Walsh transform texture features grouped by method as summarised in **table 9.20**. **Graphs 9.53** to **9.58** show variation of the DB performance measures for the non-directional Slant transform texture features grouped by method as summarised in **table 9.21**. All twelve graphs show feature behaviour with number of grey level bits (6-8), normalisation scheme (0-2), and with the use of Fourier transform pre-processing. The data used in these graphs was acquired with 8x8 ROI size.

A large number of texture features are considered in this overview. The purpose of this study is to assess the discriminant ability of each feature and to consider whether it would provide useful texture information. Only those features that demonstrate a satisfactory DB performance measure ($DB > 0.5$) have been considered for closer examination. These features have been identified from **graphs 9.47** to **9.58**.

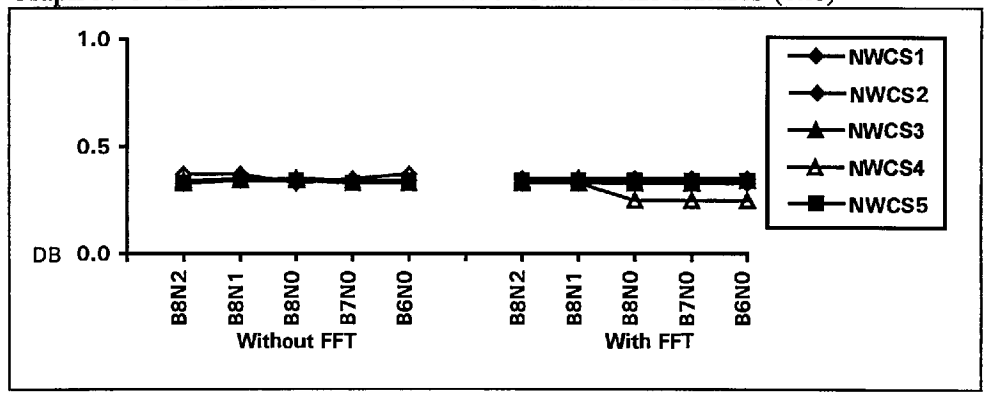
Graph 9.47 : Behaviour of non-dir^{n} Walsh vertical run length features (8x8)



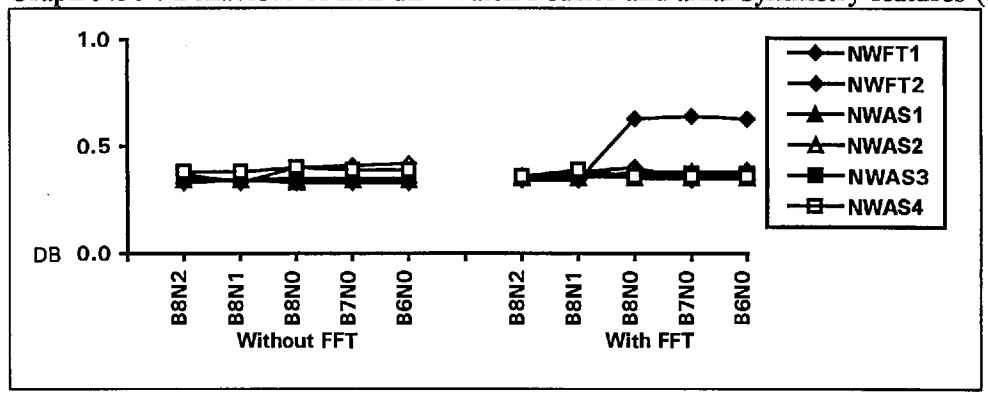
Graph 9.48 : Behaviour of non-dir^{n} Walsh horizontal run length features (8x8)



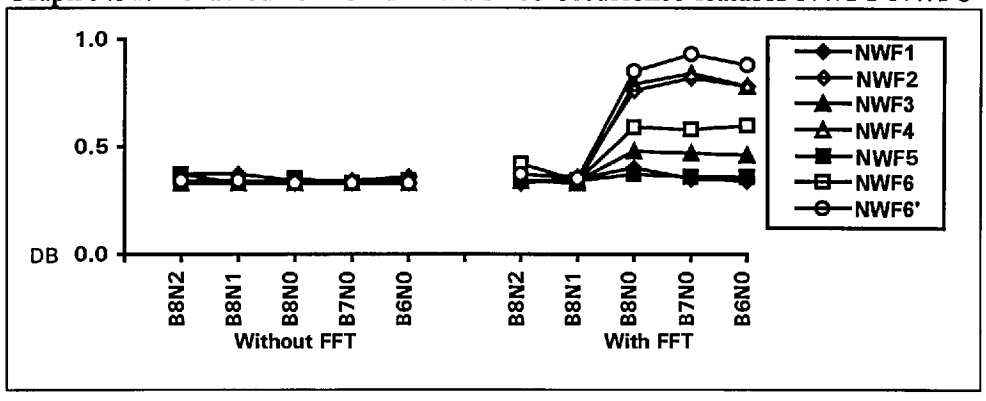
Graph 9.49 : Behaviour of non-dir^{n} Walsh cluster size features (8x8)



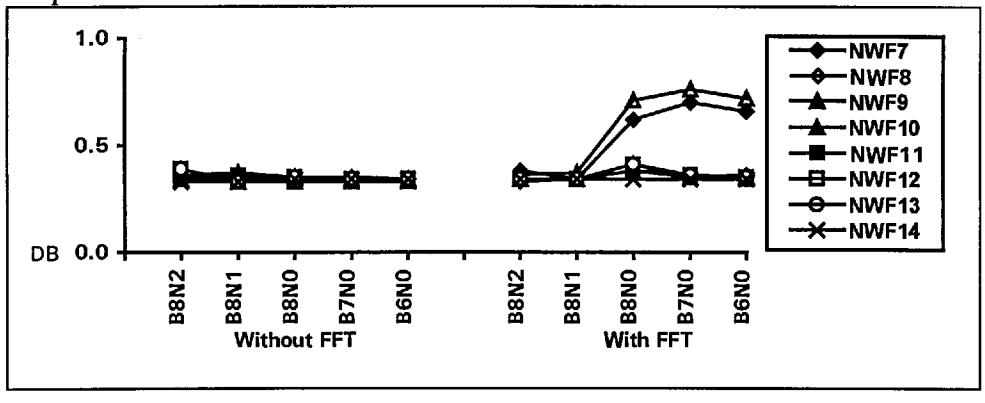
Graph 9.50 : Behaviour of non-dir^{n} Walsh Fourier and axial symmetry features (8x8)



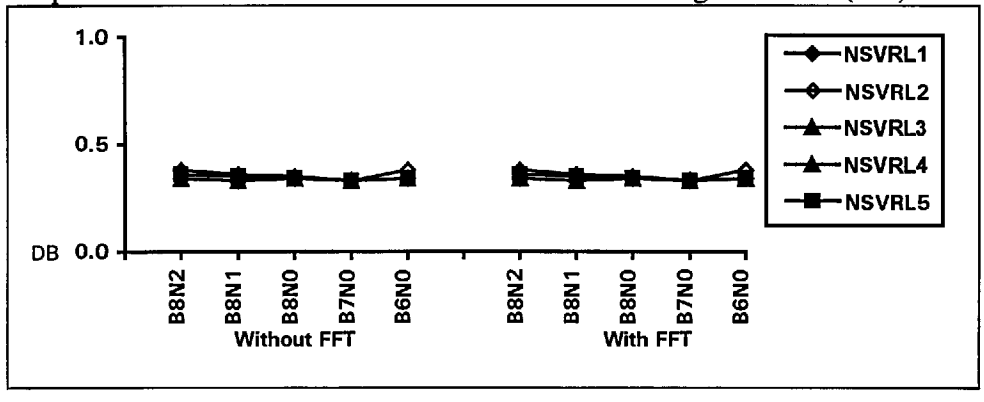
Graph 9.51: Behaviour of non-dir^n Walsh co-occurrence features NWF1-NWF6' (8x8)



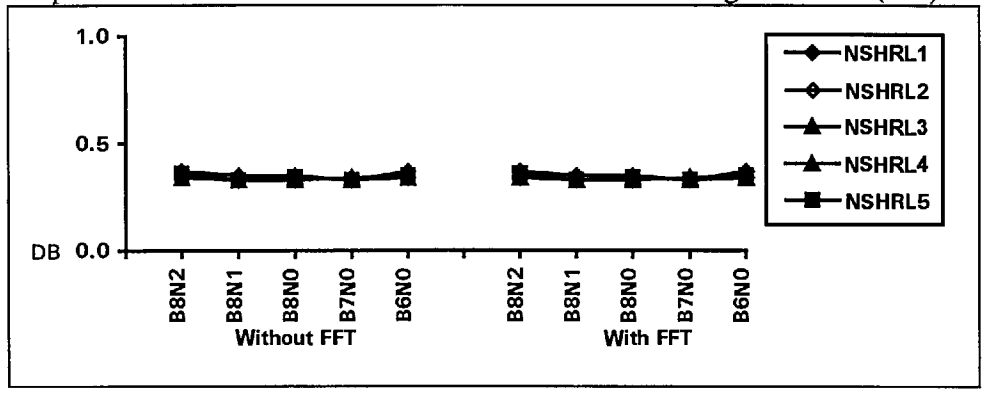
Graph 9.52: Behaviour of non-dir^n Walsh co-occurrence features NWF7-NWF14 (8x8)



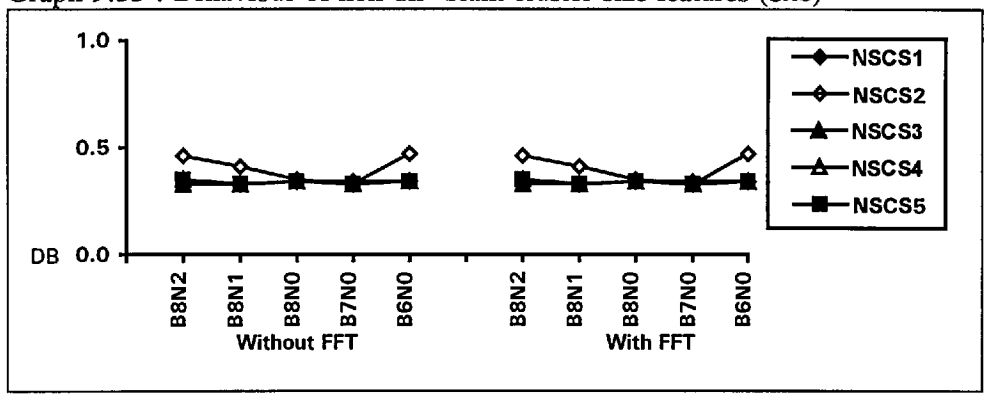
Graph 9.53 : Behaviour of non-dir^n Slant vertical run length features (8x8)



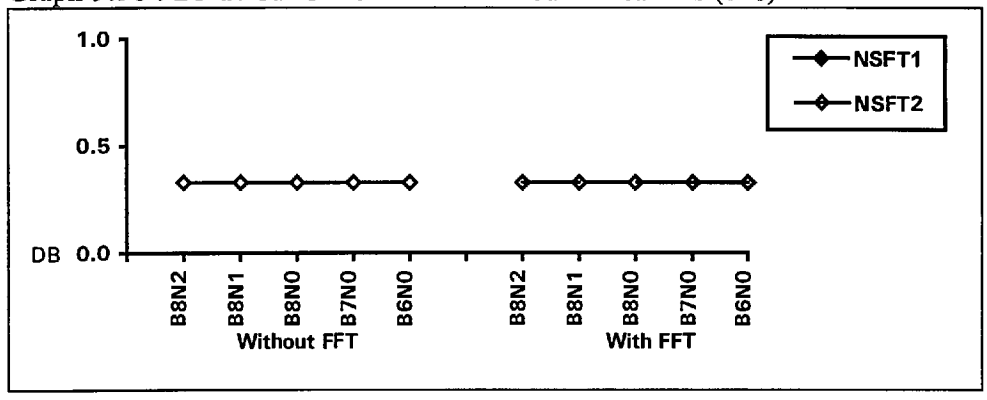
Graph 9.54 : Behaviour of non-dir^n Slant horizontal run length features (8x8)



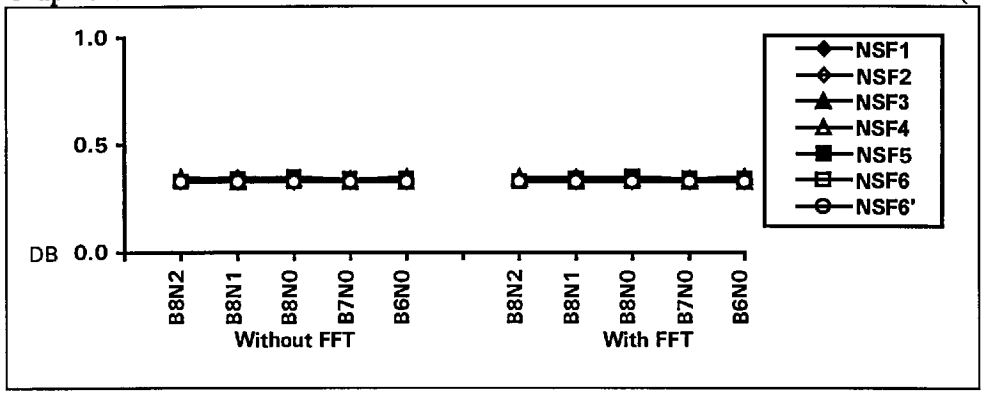
Graph 9.55 : Behaviour of non-dir^{n} Slant cluster size features (8x8)



Graph 9.56 : Behaviour of non-dir^{n} Slant Fourier features (8x8)



Graph 9.57 : Behaviour of non-dir^{n} Slant co-occurrence features NSF1-NSF6' (8x8)



Graph 9.58 : Behaviour of non-dir^{n} Slant co-occurrence features NSF7-NSF14 (8x8)

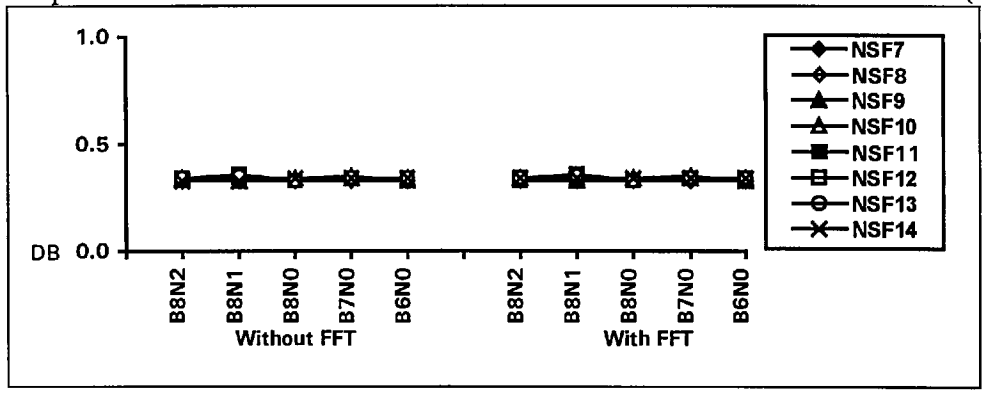
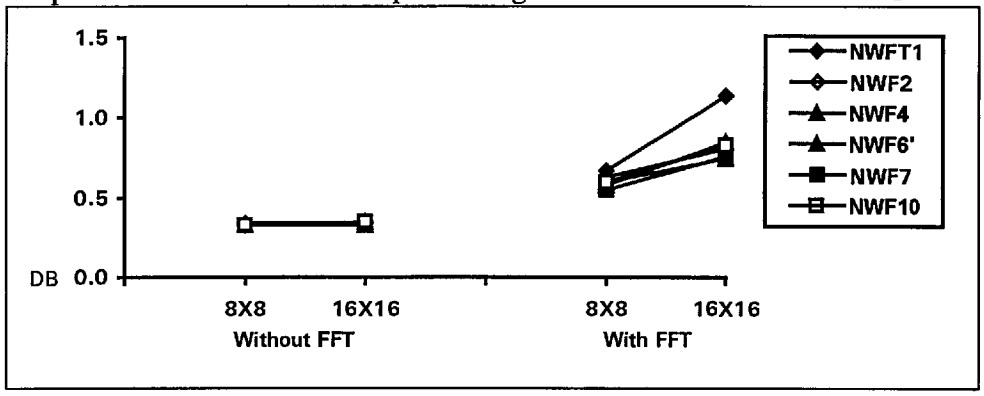


Table 9.22 provides a short-list of the texture features derived from the non-directional implementation of the Walsh transform that demonstrate a satisfactory level of performance. Optimum performances were achieved with 7-bit data and with Fourier pre-processing. None of the features derived from the non-directional form of the Slant transform are considered for further investigation because of their poor discriminative ability, as demonstrated in **graphs 9.53** to **9.58**. **Graph 9.59** summarises the performances of the selected Walsh features, while also demonstrating the influence of ROI size. Better discriminative performances were achieved for ROI of size 16x16 rather than 8x8.

Table 9.22 : Best performing non-directional Walsh features

Transform	Feature	Processing method	Texture feature
Walsh	NWFT1	Fourier	Energy
	NWF2	Co-occurrence	Haralick feature F2 (contrast)
	NWF4	Co-occurrence	Haralick feature F4 (variance)
	NWF6'	Co-occurrence	Haralick feature F6' (difference average)
	NWF7	Co-occurrence	Haralick feature F7 (sum variance)
	NWF10	Co-occurrence	Haralick feature F10 (difference variance)

Graph 9.59 : Behaviour of best performing non-directional Walsh features



9.4.3 CORRELATION STUDY

The behaviour of the five co-occurrence Walsh features summarised in **graph 9.59** are very similar, reflecting the redundancy between them. Redundant features can be eliminated from the feature set given in **table 9.22** on the basis of a study of the correlation between them. More objective tissue characterisation can be achieved with the smaller features set that results from this study.

The correlation between each pair of texture features is measured over 78 ROI evenly distributed among the three tissue classes on 'Thigh' image data. This is the same source of

data used for the discriminative performance evaluation in the previous section. The texture features are collected from ROI of size 8x8 and 16x16, with 7-bit data, Fourier pre-processing, and without data normalisation. The correlation coefficient is calculated in Microsoft EXCEL from the function libraries (see **chapter 7** for details). The coefficient varies between 0 (no correlation) and 1 (high correlation), and a negative sign indicates that the two features are inversely correlated.

Tables 9.23 to **9.24** present the correlation coefficients between the six Walsh texture features identified in **table 9.22** for ROI sizes 8x8 and 16x16, respectively. There is considerable variation in the correlation coefficients found between the features acquired with ROI of size 8x8, while they are all found to be highly correlated when acquired with ROI of size 16x16. The improved correlation between features observed as ROI size increases can be explained by the associated progression in discriminative performance. At ROI size 8x8 the textures are less well defined and are more heavily influenced by image noise, resulting in more variability in feature performance and less correlation between similar features. The result presented in **table 9.24** is therefore a more accurate description of the relationship between these features. The high correlation between all six features indicates high redundancy within the feature set. The feature set can therefore be reduced on the basis of relative discriminative performance.

Table 9.23 : Feature correlation (ROI 8x8)

	NWFT1	NWF2	NWF4	NWF6'	NWF7	NWF10
NWFT1	-	0.71	-0.92	0.91	0.66	0.49
NWF2	0.71	-	-0.88	0.84	0.75	0.95
NWF4	-0.92	-0.88	-	-0.97	-0.65	-0.69
NWF6'	0.91	0.84	-0.97	-	0.72	0.62
NWF7	0.66	0.75	-0.65	0.72	-	0.66
NWF10	0.49	0.95	-0.69	0.62	0.66	-

Table 9.24 : Feature correlation (ROI 16x16)

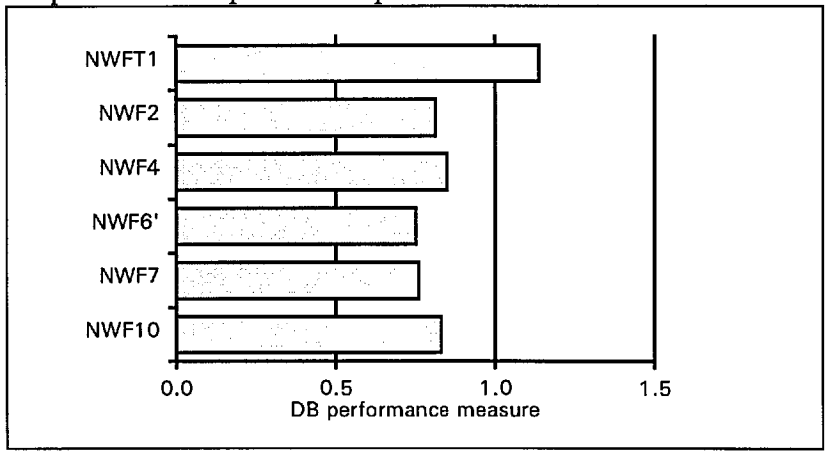
	NWFT1	NWF2	NWF4	NWF6'	NWF7	NWF10
NWFT1	-	0.94	-0.92	0.92	0.93	0.95
NWF2	0.94	-	-0.96	0.99	0.99	1.00
NWF4	-0.92	-0.96	-	-0.93	-0.92	-0.96
NWF6'	0.92	0.99	-0.93	-	0.99	0.98
NWF7	0.93	0.99	-0.92	0.99	-	0.99
NWF10	0.95	1.00	-0.96	0.98	0.99	-

9.4.4 OPTIMISED FEATURE SET

None of the texture features derived from the non-directional form of the Slant transform were capable of demonstrating adequate texture discrimination in the performance overview exercise reported in **section 9.4.2** of this chapter. None of the features derived from this method are considered in the optimised feature set for this reason.

Optimum discriminant performance was achieved for the Walsh transform features considered without data normalisation, with 7-bit data, and with Fourier pre-processing. Better performances were achieved from ROI of size 16x16 rather than 8x8. This result is consistent with the improved specificity of the transform matrix that is expected with larger ROI ie. more data points. A comparison of the DB performance of the best performing non-directional Walsh texture features acquired under optimum conditions are presented in **graph 9.60**. The best DB performance is clearly demonstrated by feature NWFT1 (Fourier transform energy), while the performance of the remaining features is comparable. The feature set can therefore be reduced to two features; NWFT1 and NWF4 (arbitrarily chosen). The performance of this feature set is examined more closely further in the tissue characterisation examples provided in the following section.

Graph 9.60 : Comparison of optimised non-directional Walsh features



9.4.5 EXAMPLES OF TISSUE CHARACTERISATION

This section provides examples of the application of the optimised non-directional texture features derived from the Walsh transforms for discrimination between tissue classes in a variety of MRI images. The two optimised texture features implemented in these examples

(NWFT1 and NWF4) are listed with their corresponding graph keys in **table 9.22** in the previous section. **Table 9.25** provides a list of the MRI image data sets investigated and the tissue classes they contain. The extent of tissue characterisation in the case of each data set is illustrated in **graphs 9.61** to **9.65**. A summary of the information presented in each graph is provided in **table 9.26**.

Table 9.25 : Description of image data

Data set	Tissue classes
Thigh	Muscle, Tumour, Fat
Abdomen	Fat, Liver, Tumour
Leg104	Muscle, Tumour
Leg192	Muscle, Tumour
Brains	Brain1, Brain2, Brain3

Table 9.26 : Examples of tissue segmentation with non-dir^n Walsh texture features

Graph	Data sample	X-axis feature	Y-axis feature	Conditions
9.61	Thigh	NWF4	NWFT1	B7N0, FFT, ROI = 16x16
9.62	Abdomen	NWF4	NWFT1	B7N0, FFT, ROI = 16x16
9.63	Leg104	NWF4	NWFT1	B7N0, FFT, ROI = 16x16
9.64	Leg192	NWF4	NWFT1	B7N0, FFT, ROI = 8x8
9.65	Brains	NWF4	NWFT1	B7N0, FFT, ROI = 16x16

Graph 9.61 shows the successful segmentation between the crucial muscle tissue and muscle tumour tissue classes in the 'Thigh' image data set. The relatively tight clustering of the three tissue groups demonstrates the proficiency of these texture features in this application.

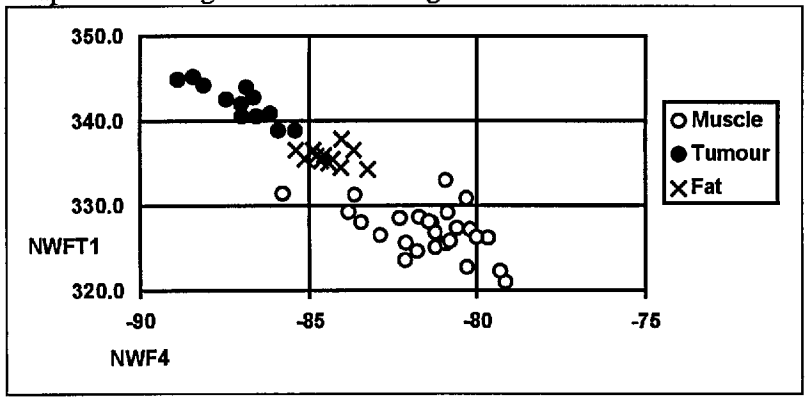
Less success was demonstrated by the two texture features in the segmentation of liver tissue and tumour tissue in the liver. **Graph 9.62** shows that, despite evidence of clustering of the data points corresponding to the liver tumour, there is no clearly defined spatial boundary between these points and the sprinkling of liver tissue data points. This may be due to the difficulty in finding homogenous ROI in the liver tissue. In addition to the main clusters of liver and tumour tissue, there appear to be two smaller clusters amongst the fat data points. These points may correspond to regions of the liver or liver tumour where there has been some fatty infiltration (striations). This phenomena could, of course, be due to partial volume effects.

Graphs 9.63 and **9.64** show two cases of the segmentation of muscle tumour from muscle tissue in the leg: 'Leg104' and 'Leg192' image data sets. In both examples there is evidence of

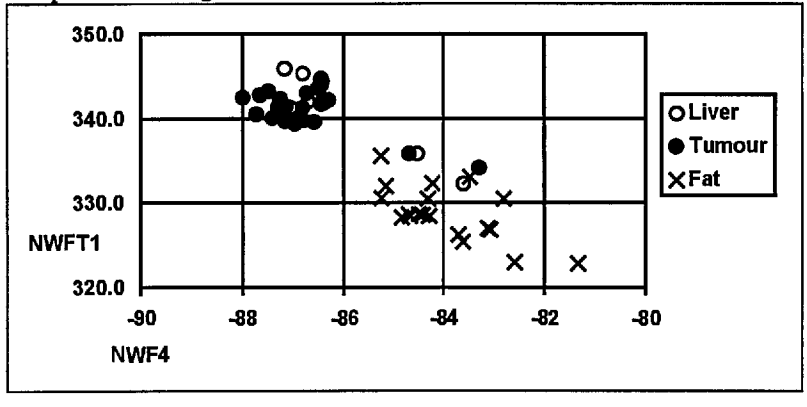
clustering of the data points originating from the same tissue regions, and a degree of segmentation between the two tissue classes.

Graph 9.65 shows the results for features extracted from grey matter in a series of mid-brain slices in three healthy volunteers. Feature stability can be assessed by examining the variability of data points both in each subject class and in all three classes combined. For good stability, one is looking for minimal variability firstly in each class then in all classes. According to these subjective criteria, neither non-directional Walsh features are unstable.

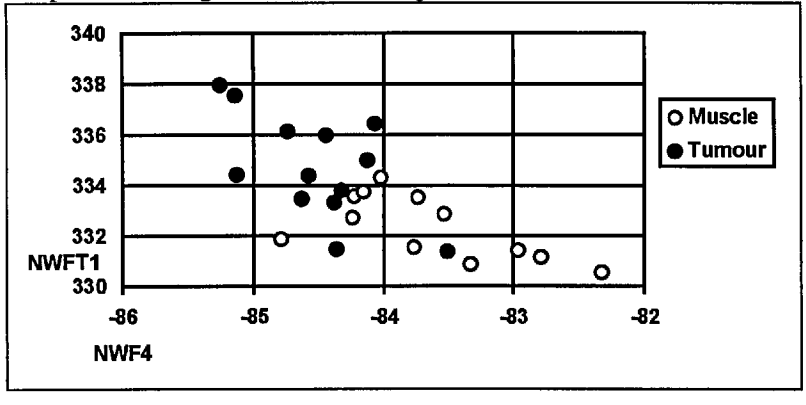
Graph 9.61 : Segmentation of 'Thigh' data - features NWF4 v NWFT1



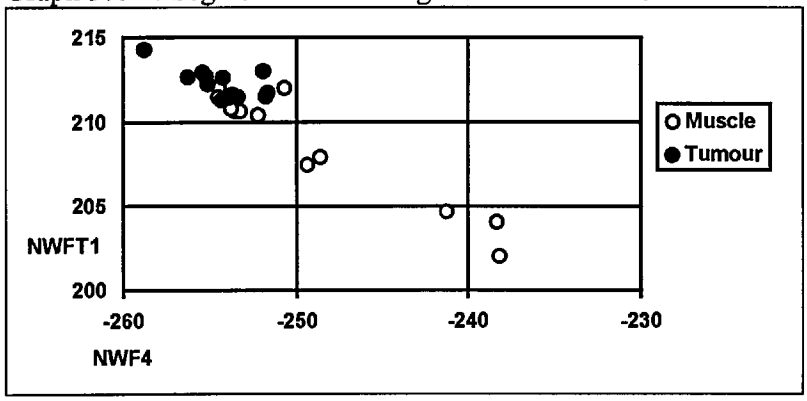
Graph 9.62 : Segmentation of 'Abdomen' data - features NWF4 v NWFT1



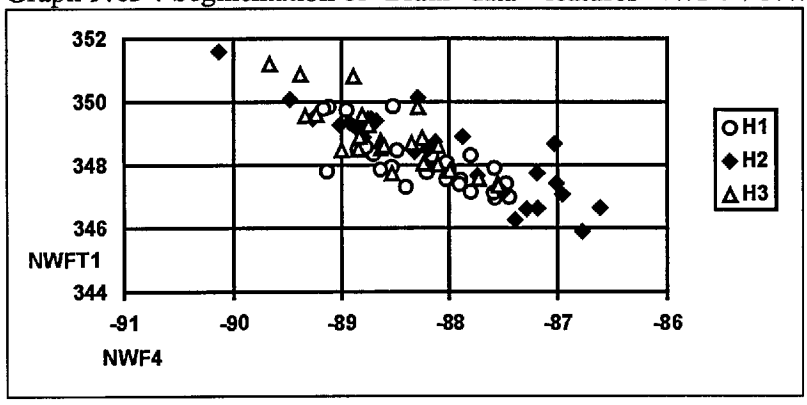
Graph 9.63 : Segmentation of 'Leg104' data - features NWF4 v NWFT1



Graph 9.64 : Segmentation of 'Leg192' data - features NWF4 v NWFT1



Graph 9.65 : Segmentation of 'Brain' data - features NWF4 v NWFT1



CHAPTER 10 - RESULTS: FRACTAL TEXTURE FEATURES

10.1 INTRODUCTION

As described in **chapter 6**, three methods of deriving the fractal dimension have been examined as part of this project; the power spectral density method, the blanket method and the box method. A summary of the texture features derived using these methods are given in **table 10.1** below. This table provides a reference between the texture features and the keys used in the graphs and tables presented throughout this chapter.

The results presented in this chapter are sub-divided into three main sections. The first section (**section 10.2**) provides an overview of the discriminative performance of each fractal texture feature and the optimum conditions under which it can be extracted. The second section (**section 10.3**) examines the correlation between the texture features, and the final section (**section 10.4**) gives examples of fractal tissue characterisation in a number of clinical applications.

Table 10.1 : Fractal texture features

Fractal method	Texture features	Key
Power spectral density (PSD)	Fractal dimension	FR1
Blanket	Fractal dimension	FR2
Box (or Tile)	Fractal dimension	FR3
	Average fractal dimension	FR4

The discriminative performance tests are carried out on several transverse MRI images of the thigh (see 'Thigh' image data in **chapter 7**). These images contain three distinct classes of tissue; (1) muscle, (2) fat and (3) tumour. The discriminative behaviour of each feature is quantified using the DB performance measure defined in **chapter 7** (based on the relationship between the statistics of tissue classes). Unless stated, the DB performance measure is calculated for the discrimination between all three tissue classes (1-3).

The regions of interest (ROI) in the image data are standardised before any texture information is extracted by fixing the number of grey level bits or by implementing some normalisation scheme. The pre-processing options are indicated by 'B' for the number of grey level bits (in the range of 4-8), and 'N' for the normalisation scheme; (1) none, (2)

histogram equalisation, and (3) statistical moments. A detailed description of normalisation schemes (2) and (3) are given in **chapter 7**.

The results of the performance overview attempt to demonstrate the effect of data pre-processing on the discriminative behaviour of each texture feature. Consideration is also given to the influence of region of interest (ROI) size. This section concludes with a description of the conditions under which each of the fractal texture descriptors can be implemented optimally.

The correlation study in **section 10.3** enables the relative behaviour of the fractal texture features to be investigated. Similar behaviour would be expected from highly correlated features. When two features are highly correlated, therefore, one feature can be discarded without any significant loss of information. Large numbers of texture features can be reduced to more manageable subsets by examining the correlation between each and every feature, and eliminating redundant features. This process is applied to the four fractal texture features in this section.

The tissue characterisation examples in the final section of this chapter (**section 10.4**) provide an illustration of the efficacy of the optimised fractal texture features with a number of sets of clinical images containing a variety of different tissue types.

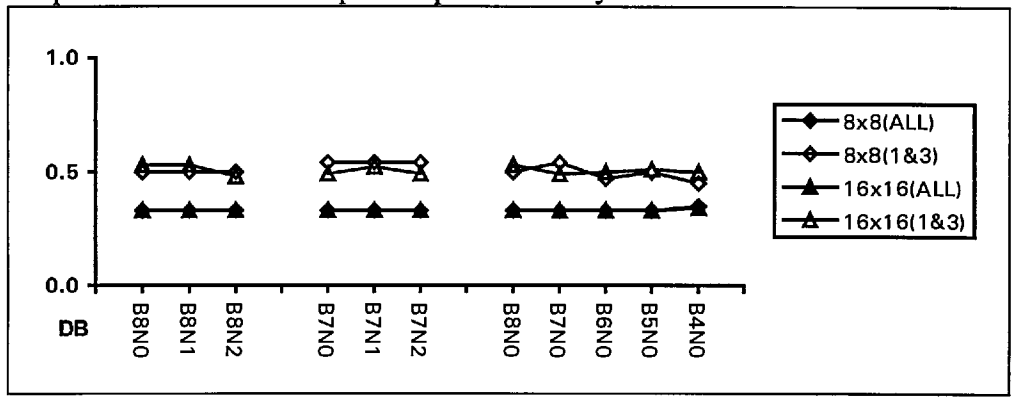
10.2 FRACTAL FEATURES PERFORMANCE EVALUATION

10.2.1 POWER SPECTRAL DENSITY METHOD PERFORMANCE OVERVIEW

The plots in **graph 10.1** indicate the discriminative performance trends of the power spectral density (PSD) fractal dimension parameter extracted from ROI under different pre-processing conditions. For each point on the graph, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (B for bits and N for normalisation scheme). The graph legend labels each series of data points plotted with the ROI size (8x8 or 16x16) and the tissue classes discriminated between (ALL classes or classes 1&3). The 'Thigh' MRI images used in this overview contain three distinct classes of tissue. Significance has been placed on the discriminative performance of this texture feature between tissue types 1 (muscle tissue) and 3 (tumour tissue within the muscle region) because this represents the most important segmentation in this example.

Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions on several images.

Graph 10.1 : Behaviour of power spectral density fractal feature



Graph 10.1 indicates that image data normalisation (N) has little or no effect on the discriminative performance of the PSD fractal texture feature. Little or no effect is also noted when the ROI size is increased from 8x8 to 16x16. A comparison of performance with the number of image grey levels (B) is equally in-conclusive. Although the discriminative performance of the PSD fractal dimension is better between tissues 1&3 than between all tissue classes, the overall performance of this descriptor under the conditions of this experiment is poor. This method is being penalised by the limiting size of homogenous tissue regions on the MRI images. A maximum ROI size of 16x16 provides only 8 data points on a 1D Fourier spectrum. Once the DC component and the highly DC influenced low frequency components have been eliminated, only half a dozen points remain from which an exponential fit must be made to extract the fractal dimension. Larger ROI resulting in larger data sets would enable texture information to be extracted from a segment of the power spectrum that neither contains the very low nor the very high frequency components that are respectively influenced by DC and random image noise.

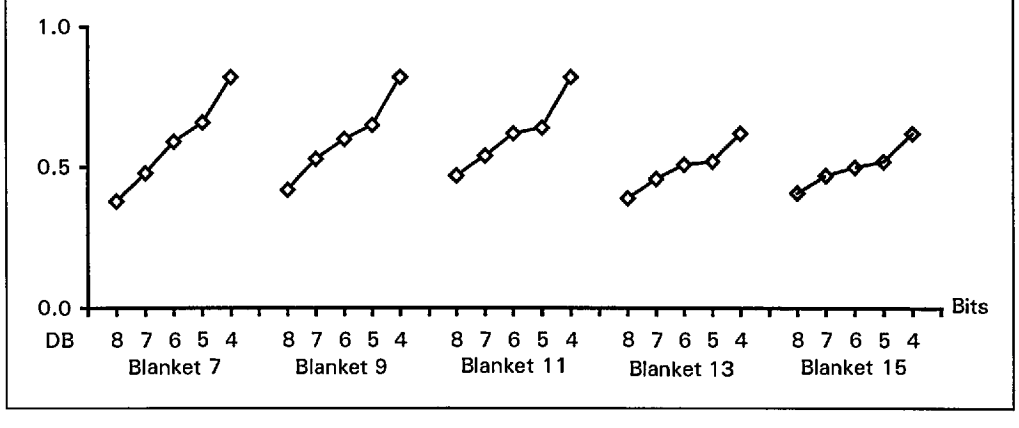
The conclusion that can be drawn from this study is that the potential of the PSD fractal dimension for texture discrimination cannot be realised until large ROI (at least 32x32) of homogenous tissue can be investigated. This kind of MRI data has not been available for this project. The poor discriminative performance of this parameter between MRI tissues is demonstrated in tissue characterisation examples provided in **section 10.4**. This section also includes an illustration of the potential discriminant ability of the PSD fractal parameter using large ROI in artificial texture samples.

10.2.2 BLANKET METHOD PERFORMANCE OVERVIEW

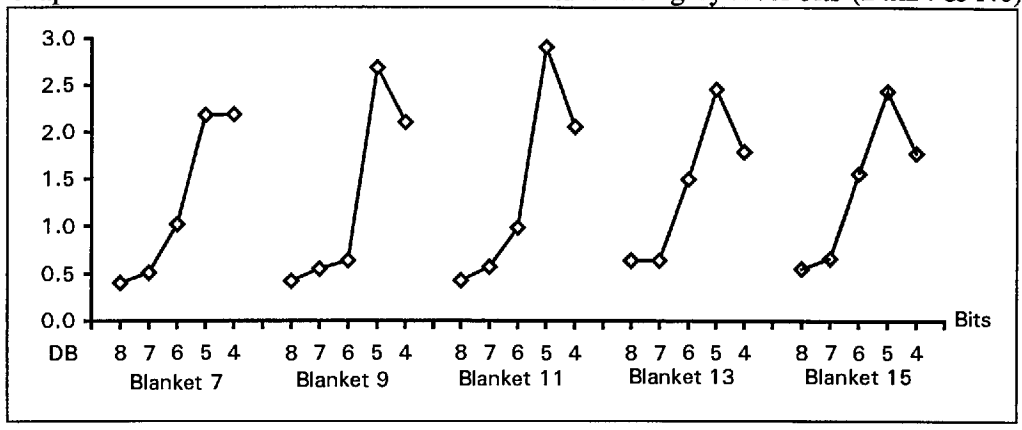
The plots in **graphs 10.2** to **10.5** indicate the discriminative performance trends of the blanket fractal dimension parameter extracted from ROI derived under different pre-processing conditions. For each point on the graph, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (either bits or normalisation scheme). The blanket method derives the fractal dimension from the changes in surface area associated with successive applications of a laplace averaging filter over the image surface. The number of iterations in this process is controlled by a parameter that shall be known as the blanket index. The blanket index corresponding to each point on the graph is also indicated on the x-axis. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions (muscle, fat & tumour) on several of images.

Graphs 10.2 and **10.3** show the variation of the DB performance measure for the blanket fractal dimension with the number of grey level bits (4-8) assigned to the image data by the pre-processing. This behaviour was examined for a series of different blanket indices. The data used in these graphs was acquired without normalisation (N0) and with both 12x12 and 24x24 ROI sizes. Optimum performance was achieved with a blanket index of 11 and with 4-bit data and 5-bit data respectively for the 12x12 and 24x24 ROI size. Better DB performance measures were attained for 24x24 than for 12x12 ROI size.

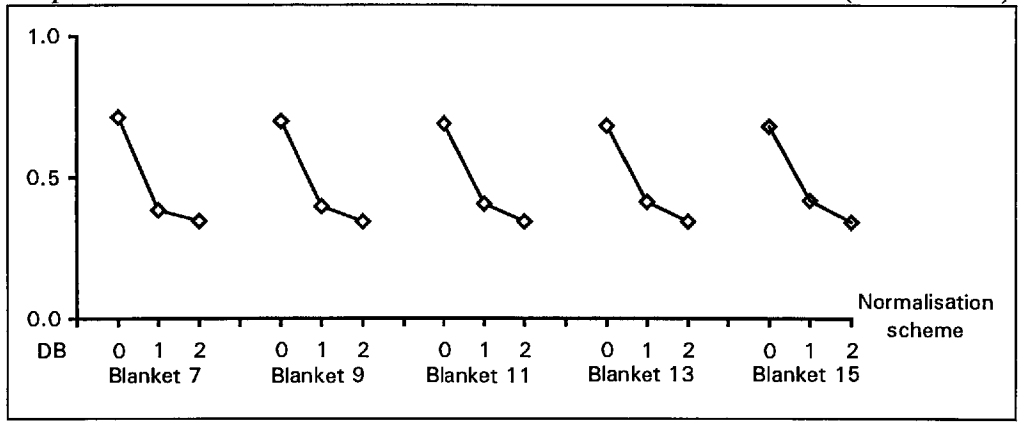
Graph 10.2 : Behaviour of blanket fractal feature with grey level bits (12x12 & N0)



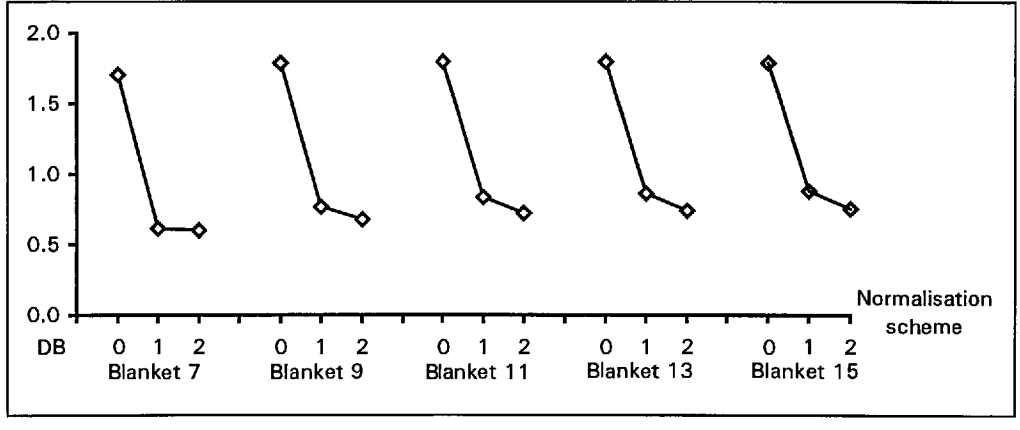
Graph 10.3 : Behaviour of blanket fractal feature with grey level bits (24x24 & N0)



Graph 10.4 : Behaviour of blanket fractal feature with normalisation (12x12 & B4)



Graph 10.5 : Behaviour of blanket fractal feature with normalisation (24x24 & B5)



Graphs 10.4 and **10.5** show the variation of the DB performance measure for the blanket fractal dimension with normalisation scheme (0-2). This behaviour was also examined for a series of different blanket indices. The data used in these graphs was acquired with 4-bit and 5-bit image data, respectively, and with both 12x12 and 24x24 ROI sizes. The effect of data normalisation generally proved to be detrimental to the DB performance measure. The

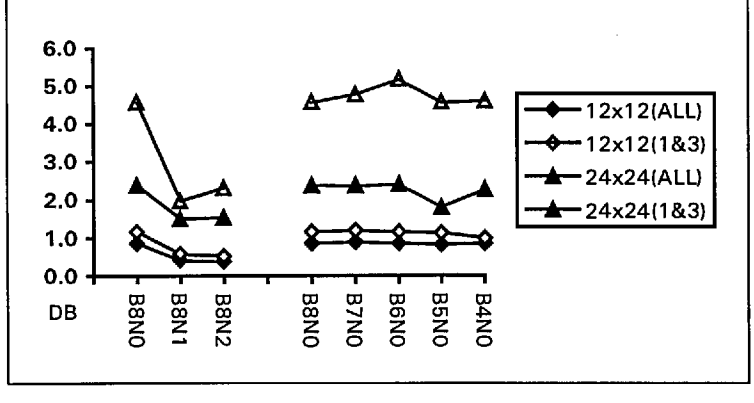
data acquired for these plots was optimised with respect to the number of grey levels used to represent the image data; 4-bit data for ROI of 12x12 and 5-bit data for ROI of 24x24.

10.2.3 BOX METHOD PERFORMANCE OVERVIEW

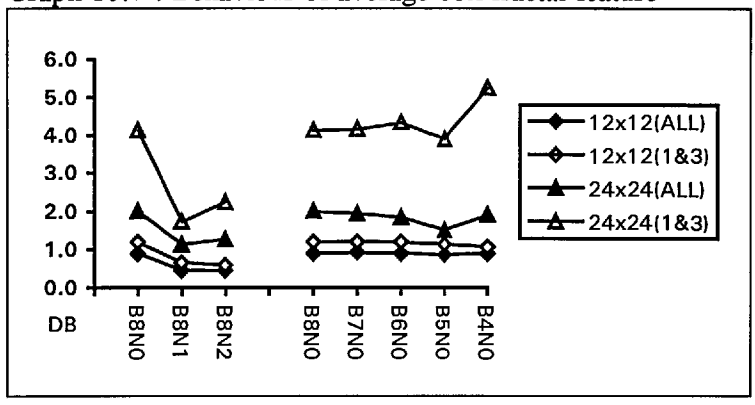
The plots in **graphs 10.6** and **10.7** indicate the performance trends of the box fractal and box average fractal dimension parameters extracted from ROI under different pre-processing conditions. For each point on the graph, the y-axis represents the DB performance measure and the x-axis the image ROI pre-processing implemented (B for bits and N for normalisation scheme). The graph legend labels each series of data points plotted with the ROI size (12x12 or 24x24) and the tissue classes discriminated between (ALL classes or classes 1&3). The 'Thigh' MRI images used in this overview contain three distinct classes of tissue. Significance has been placed on the discriminative performance of this texture feature between tissue types 1 (muscle tissue) and 3 (tumour tissue within the muscle region) because this represents the most important segmentation in this example. Each data point on the graph is calculated from 78 ROI (carefully placed by hand) evenly distributed among the three tissue regions on several images.

The DB performance measure for both box fractal dimension and average box fractal dimension was significantly degraded with the introduction of normalisation (N), dropping by as much as a factor of two. A comparison of performance with the number of image grey levels (B) indicates the best results are obtained with 6-bit data for the box fractal dimension, and with 4-bit data for the box average fractal dimension. An improvement in performance is also noted for both features when the ROI size is increased from 12x12 to 24x24.

Graph 10.6 : Behaviour of box fractal feature



Graph 10.7 : Behaviour of average box fractal feature



The overall discriminative performances of the two fractal dimension texture features between all three tissue classes and between tissues 1 & 3 correlate well, although the performances are markedly better between the two critical tissue classes. The performances of both texture features, as expected, appear to correlate well.

10.2.4 SUMMARY OF CONDITIONS FOR OPTIMAL PERFORMANCE

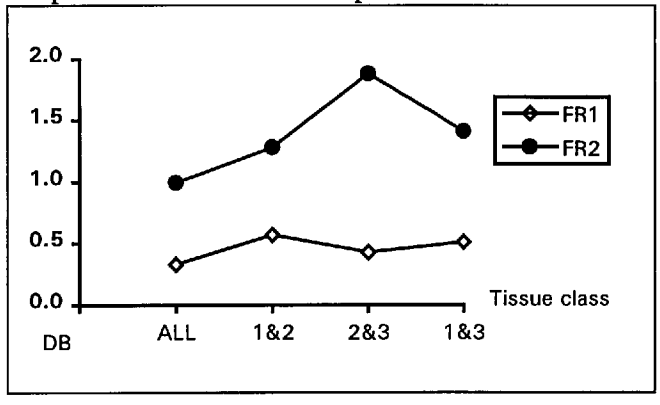
Table 10.2 provides a summary of the parameters found to enable optimal performance for the four fractal texture descriptors investigated. The number of grey level bits the image data is represented by before the features are extracted varies from method to method. In all cases image data normalisation was found to lower the overall performance of the features. As expected, better performance was achieved in each case with the larger of the two ROI sizes investigated. All the fractal dimension parameters considered in this study are derived in some way from the gradient of a plot of N data points, where N is related to the ROI size. Larger ROI provide more data points, which therefore enable better definition of the gradient. The limiting factor for ROI size in this study is the area of a particular tissue type available in the test images.

Table 10.2 : Optimal feature performance parameters

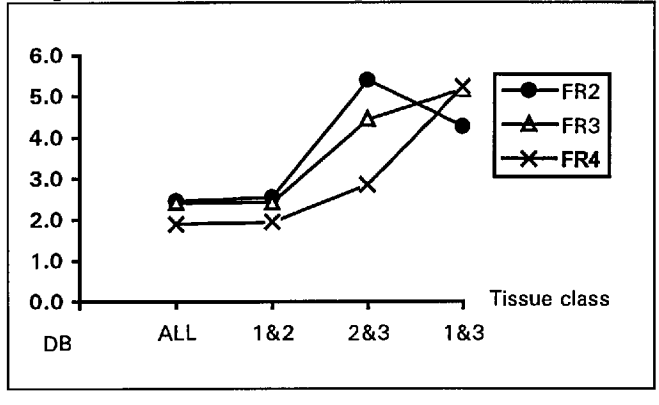
Fractal texture feature	Legend	Bits (B)	Norm. (N)	ROI size	Other
PSD fractal dimension	FR1	8	None (0)	>16x16	-
Blanket fractal dimension	FR2	5	None (0)	>24x24	Index=11
Box fractal dimension	FR3	6	None (0)	>24x24	-
Box average fractal dimension	FR4	4	None (0)	>24x24	-

Graph 10.8 compares the optimal performance of the PSD fractal dimension (FR1) and the blanket fractal dimension (FR2) as prescribed by **table 10.2**, but with ROI size 16x16. **Graph 10.9** compares the optimal performance of the blanket fractal dimension (FR2), the box fractal dimension (FR3) and the box average fractal dimension (FR4) as prescribed by **table 10.2**. The data for both graphs was collected as part of the feature performance evaluation exercise described in **sections 10.2.1** to **10.2.3** of this chapter. For each point on the graphs, the y-axis represents the DB performance measure and the x-axis the image tissue classes being discriminated between; (1) muscle, (2) fat, and (3) tumour.

Graph 10.8 : Behaviour comparison of features FR1&FR2 (ROI=16x16)



Graph 10.9 : Behaviour comparison of features FR2,FR3&FR4 (ROI=24x24)



Graph 10.8 shows the blanket fractal dimension to be several factors more effective in the discrimination between the three tissue groups than the PSD fractal dimension. More comparable performances are expected as the ROI size increases, although this cannot be demonstrated in the context of this study because of the limited available clinical MRI data because of the limited number of pixels covering any one tissue group. An example of the PSD fractal dimension applied to non-clinical texture samples using large ROI is given in **section 10.4** to demonstrate this parity.

The comparison made in **graph 10.9** shows equivalent performance between the blanket fractal dimension and the two box method fractal dimensions. Of the three texture features, the box average fractal dimension performs the least well.

10.3 CORRELATION BETWEEN ALL FRACTAL FEATURES

In order to establish the relationship between the four fractal dimension texture features, a study of the correlation between them has been carried out. High correlation between the features indicates that they are measuring the same properties. Under these circumstances, the redundant and less well performing features are eliminated from our feature set. All the texture features in an ideal set would be un-correlated and high performing.

The correlation between each pair of texture features is measured over 78 ROI evenly distributed among the three tissue classes on 'Thigh' image data. This is the same source of data as used for the performance evaluation in **section 10.2**. The texture features are calculated under the optimum conditions described in **table 10.2**. The correlation coefficient is calculated in Microsoft EXCEL from the function libraries (see **chapter 7** for details). The coefficient varies between 0 (no correlation) and 1 (high correlation), and a negative sign indicates that the two features are inversely correlated. **Tables 10.3** to **10.6** present the correlation coefficients between the four texture features for different ROI size.

Table 10.3 : Feature correlation (ROI 8x8)

8x8	FR1	FR2	FR3	FR4
FR1	-	-0.223	-	-
FR2	-0.223	-	-	-
FR3	-	-	-	-
FR4	-	-	-	-

Table 10.4 : Feature correlation (ROI 12x12)

12x12	FR1	FR2	FR3	FR4
FR1	-	-	-	-
FR2	-	-	0.893	0.883
FR3	-	0.893	-	0.995
FR4	-	0.883	0.995	-

Table 10.5 : Feature correlation (ROI 16x16)

16x16	FR1	FR2	FR3	FR4
FR1	-	0.002	-	-
FR2	0.002	-	-	-
FR3	-	-	-	-
FR4	-	-	-	-

Table 10.6 : Feature correlation (ROI 24x24)

24x24	FR1	FR2	FR3	FR4
FR1	-	-	-	-
FR2	-	-	0.846	0.844
FR3	-	0.846	-	0.992
FR4	-	0.844	0.992	-

Correlation studies cannot be made between the PSD fractal dimension method (FR1) and the box fractal dimension methods (FR3 and FR4) because of the ROI size constraints of each technique. The Fourier transform algorithm employed for the PSD fractal dimension

requires an ROI of size 2^n (n being an integer), while the box method requires the ROI size to be a number that can be divisible by a series of integers. There are no coincidences between the permitted ROI dimensions of these methods.

The coefficient of correlation between the PSD fractal dimension (FR1) and the blanket fractal dimension (FR2) is very low for both ROI sizes examined (12x12 and 24x24). The low correlation indicates that the behaviour of these features is very different. This result is contrary to the theory, and may in fact be due to the contrast in performance between the two features (**section 10.2.4**). The poor performance in the PSD fractal dimension (FR1) has been attributed to the clinically limiting ROI size. Both features have been applied to the characterisation of tissue in several example MRI images in **section 10.4** of this chapter. These examples include an artificial texture data set that has been used to demonstrate the influence of ROI size on the performance of the PSD fractal dimension.

Both fractal dimension features derived from the box method (FR3 and (FR4) are, predictably, highly correlated for both the ROI sizes examined (12x12 and 24x24). The box average fractal dimension feature is discarded because of its inferior performance (**section 10.2.3**). The correlation between the blanket fractal dimension (FR2) and the box fractal dimension (FR3) features is also significant for both ROI sizes. Both features, however, perform equally well and have been applied to the characterisation of tissue in several example clinical MRI images in **section 10.4**.

10.4 EXAMPLES OF FRACTAL TISSUE CHARACTERISATION

This section provides examples of the application of the optimised fractal texture features for discrimination between tissue classes in a variety of MRI images. The three optimised texture features (as described in the previous section) are listed in **table 10.7** with their corresponding graph keys. **Table 10.8** provides a list of the MRI image data sets investigated and the tissue classes they contain. The degree to which the texture features characterise the tissue classes in the case of each data set is illustrated in **graphs 10.10** to **10.30**. A summary of the conditions under which the data for each of these graphs was acquired is provided in **table 10.9**.

Table 10.7 : Fractal texture features - graph key

Fractal texture feature	Key
PSD fractal dimension	FR1
Blanket fractal dimension	FR2
Box fractal dimension	FR3

Table 10.8 : Description of image data

Data sample	Tissue classes
Thigh	Muscle, Tumour, Fat
Abdomen	Fat, Liver, Tumour
Leg104	Muscle, Tumour
Leg192	Muscle, Tumour
Brains	Brain1, Brain2, Brain3
Brodatz texture samples	Cloud, Water, Wood, Granite

Table 10.9 : Examples of tissue segmentation with fractal texture features

Graph	Data sample	ROI size	X-axis feature	Y-axis feature
10.10	Thigh	8x8	FR1 - B8N0	FR2 - B5N0
10.11		16x16	FR1 - B8N0	FR2 - B5N0
10.12		12x12	FR2 - B5N0	FR3 - B6N0
10.13		24x24	FR2 - B5N0	FR3 - B6N0
10.14	Abdomen	8x8	FR1 - B8N0	FR2 - B5N0
10.15		16x16	FR1 - B8N0	FR2 - B5N0
10.16		12x12	FR2 - B5N0	FR3 - B6N0
10.17	Leg104	8x8	FR1 - B8N0	FR2 - B5N0
10.18		16x16	FR1 - B8N0	FR2 - B5N0
10.19		12x12	FR2 - B5N0	FR3 - B6N0
10.20		24x24	FR2 - B5N0	FR3 - B6N0
10.21	Leg192	8x8	FR1 - B8N0	FR2 - B5N0
10.22		16x16	FR1 - B8N0	FR2 - B5N0
10.23		12x12	FR2 - B5N0	FR3 - B6N0
10.24		24x24	FR2 - B5N0	FR3 - B6N0
10.25	Brains	8x8	FR1 - B8N0	FR2 - B5N0
10.26		16x16	FR1 - B8N0	FR2 - B5N0
10.27		12x12	FR2 - B5N0	FR3 - B6N0
10.28		24x24	FR2 - B5N0	FR3 - B6N0
10.29	Brodatz texture samples	32x32	FR1 - B8N0	FR2 - B5N0
10.30		64x64	FR1 - B8N0	FR2 - B5N0

Graphs 10.10 to **10.13** show good examples of successful segmentation between the crucial muscle tissue and muscle tumour tissue classes in the 'Thigh' image data set. It was not, however, possible to clearly differentiate between muscle and fat tissue in any of these

examples. The best discriminative performance, indicated by the degree of clustering of same class data points, is shown by features FR2 and FR3 in **graph 10.13.**

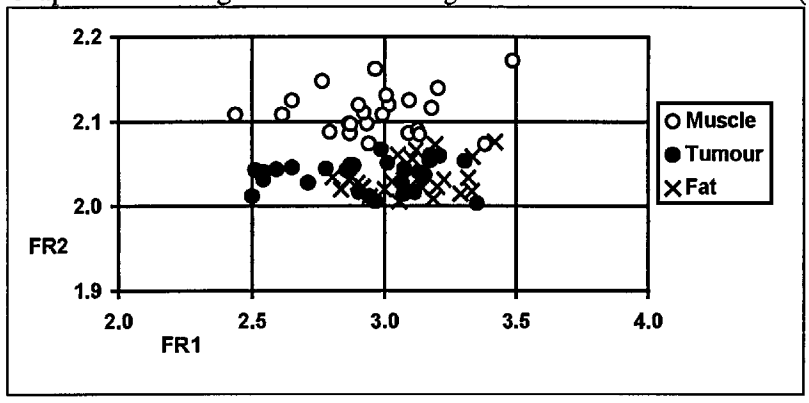
Graphs 10.14 to **10.16** show poor discrimination between the crucial liver tissue and liver tumour tissue classes from the 'Abdomen' data set. Despite this unsatisfactory result, there were discernible boundaries in feature space between fat tissue and the two types of liver tissue in all three examples. Of all the combinations of texture features and ROI sizes applied to this data set, the best results were demonstrated in **graph 10.16** with fractal features FR2 and FR3 calculated from ROI of size 12x12.

Graphs 10.17 to **10.20** show the first of two cases of segmentation of muscle tumour from muscle tissue in the leg (data set 'Leg104'). For both combinations of texture features, discriminant performance is clearly improved with larger ROI size. This was demonstrated in **graphs 10.18** and **10.20**, where the two tissue classes were completely segmented. Where a smaller ROI was used (**graphs 10.17** and **10.19**) the data points from the two tissue classes overlap.

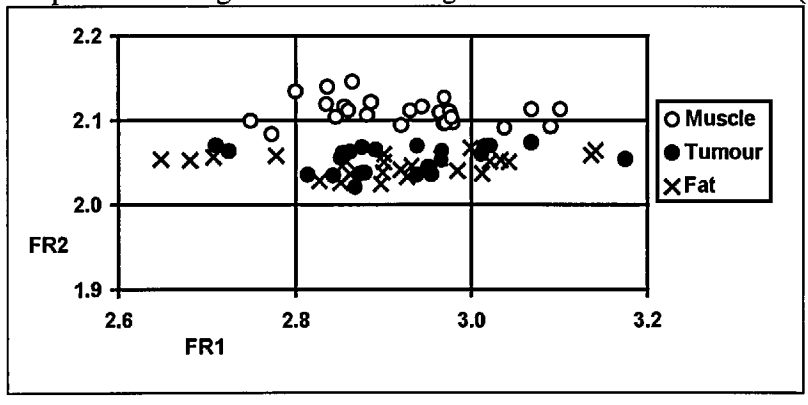
Graphs 10.21 to **10.24** show the second case of segmentation of muscle tumour from muscle tissue in the leg (data set 'Leg192'). In all four examples the two tissue classes are clearly separated in feature space. These graphs confirm the relationship between ROI size and performance, and also the superior performance of features FR2 and FR3 over feature FR1.

Graphs 10.25 to **10.28** show the results for texture features extracted from grey matter in a series of mid-brain slices in three healthy volunteers; 'Brain' image data. Feature stability can be assessed by examining the variability of features both in each subject class and in all three classes combined. For good stability, one is looking for minimal variability firstly in each class and then in all classes. The evidence provided by **graphs 10.25** and **10.26** indicates that feature FR1 is a very unstable parameter. By contrast, features FR2 and FR3 demonstrate good stability in all four graphs.

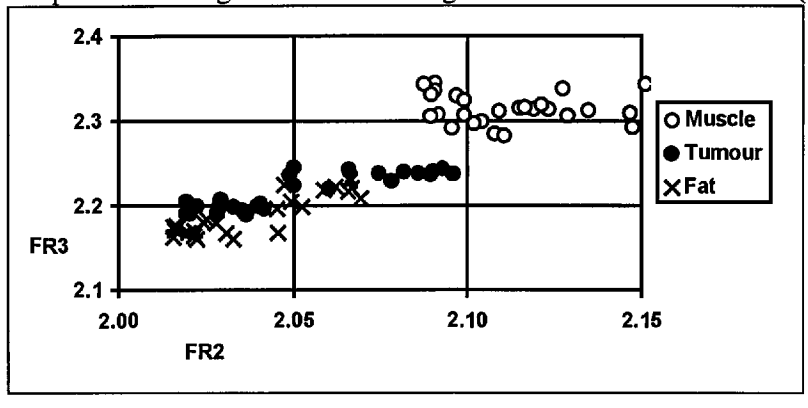
Graph 10.10 : Segmentation of 'Thigh' data - features FR1vFR2 (ROI 8x8)



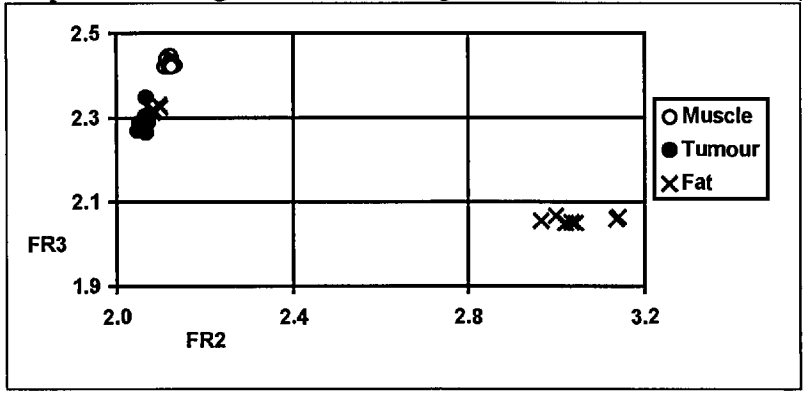
Graph 10.11 : Segmentation of 'Thigh' data - features FR1vFR2 (ROI 16x16)



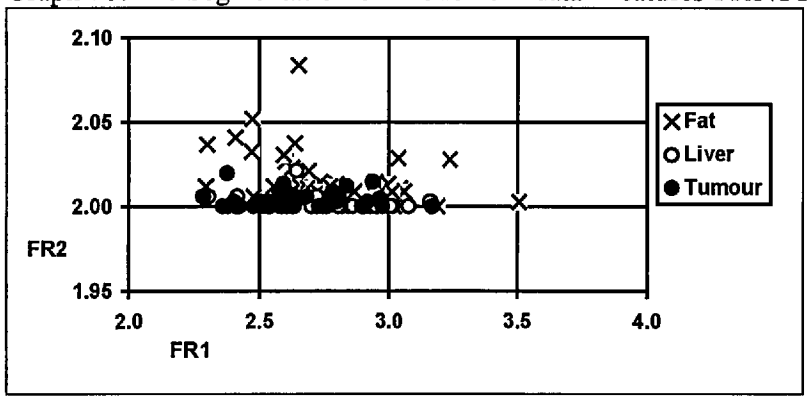
Graph 10.12 : Segmentation of 'Thigh' data - features FR2vFR3 (ROI 12x12)



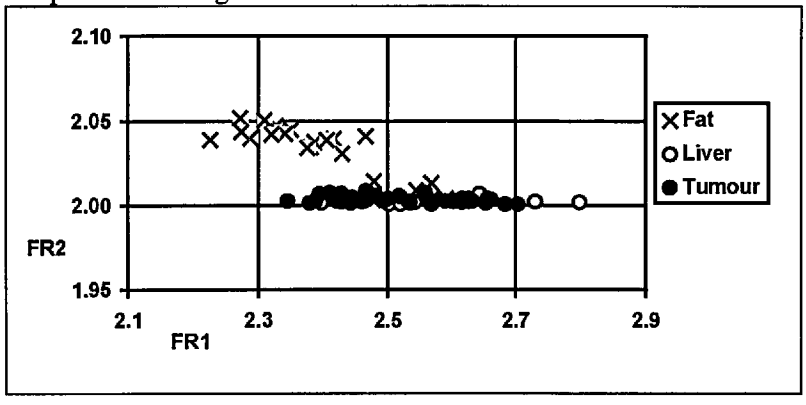
Graph 10.13 : Segmentation of 'Thigh' data - features FR2vFR3 (ROI 24x24)



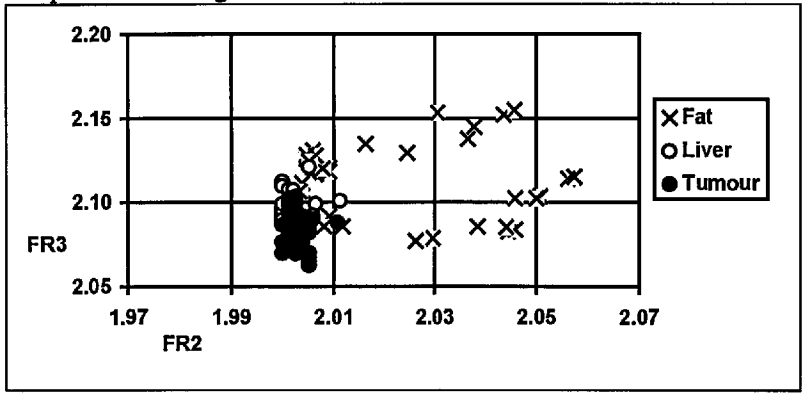
Graph 10.14 : Segmentation of 'Abdomen' data - features FR1vFR2 (ROI 8x8)



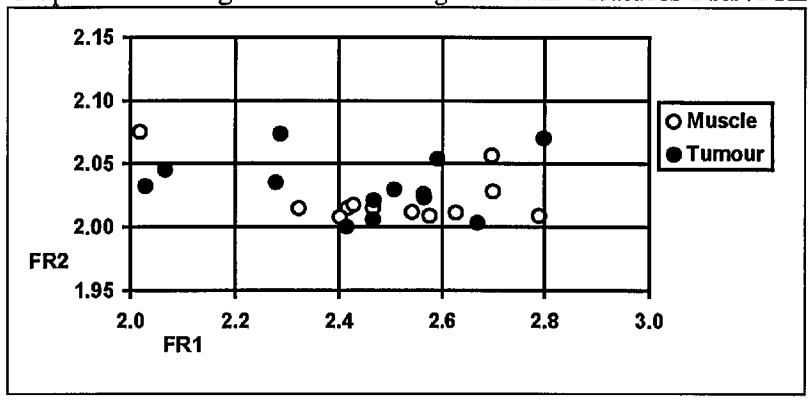
Graph 10.15 : Segmentation of 'Abdomen' data - features FR1vFR2 (ROI 16x16)



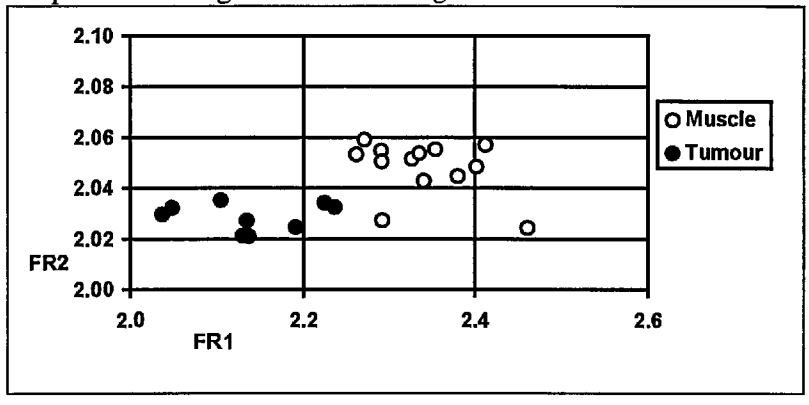
Graph 10.16 : Segmentation of 'Abdomen' data - features FR2vFR3 (ROI 12x12)



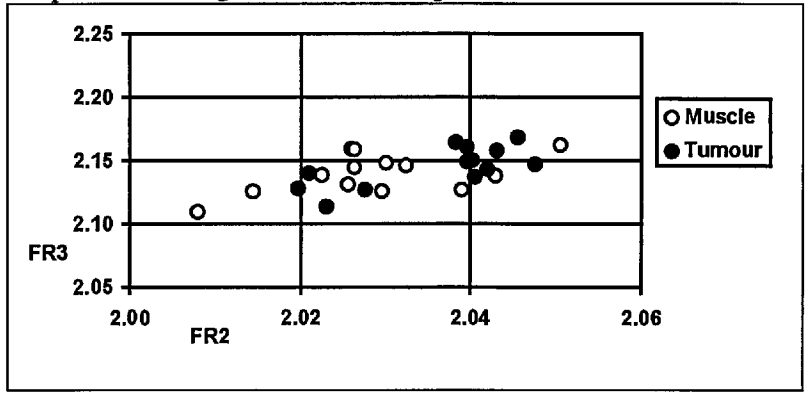
Graph 10.17 : Segmentation of 'Leg104' data - features FR1vFR2 (ROI 8x8)



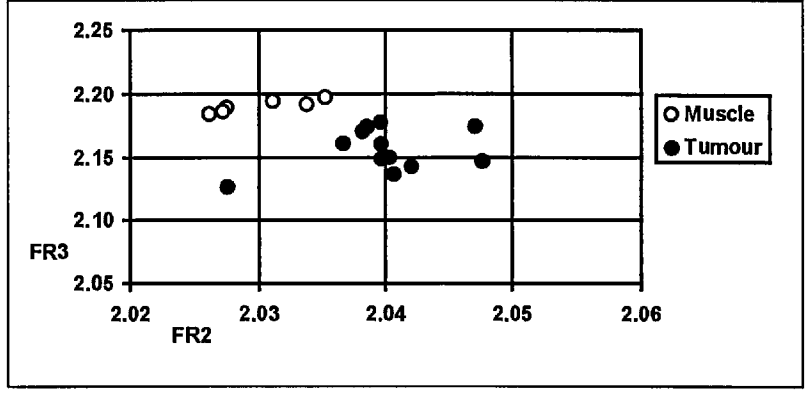
Graph 10.18 : Segmentation of 'Leg104' data - features FR1vFR2 (ROI 16x16)



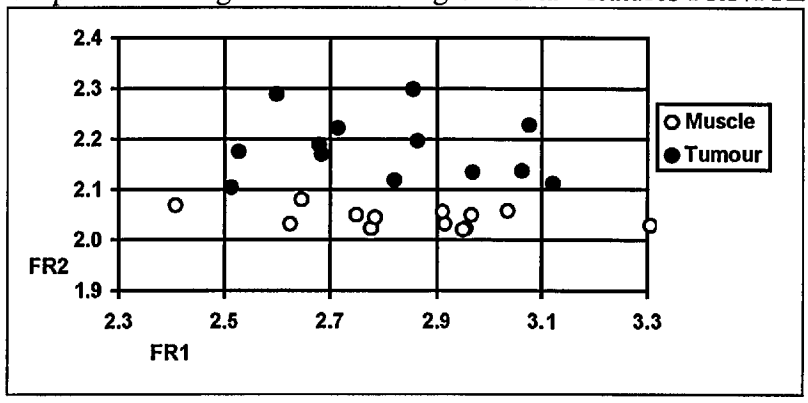
Graph 10.19 : Segmentation of 'Leg104' data - features FR2vFR3 (ROI 12x12)



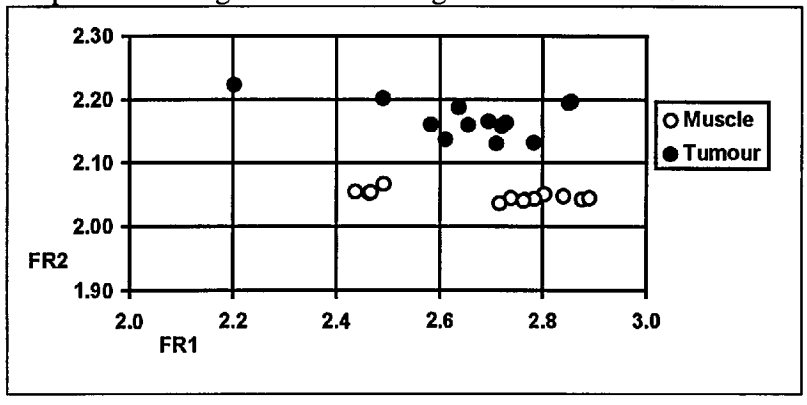
Graph 10.20 : Segmentation of 'Leg104' data - features FR2vFR3 (ROI 24x24)



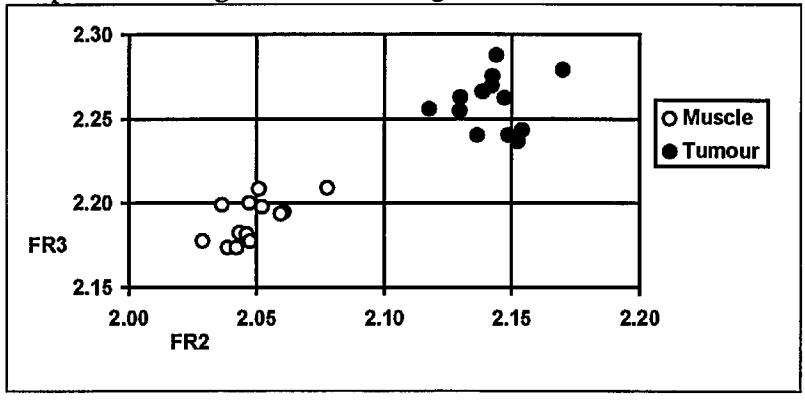
Graph 10.21 : Segmentation of 'Leg192' data - features FR1vFR2 (ROI 8x8)



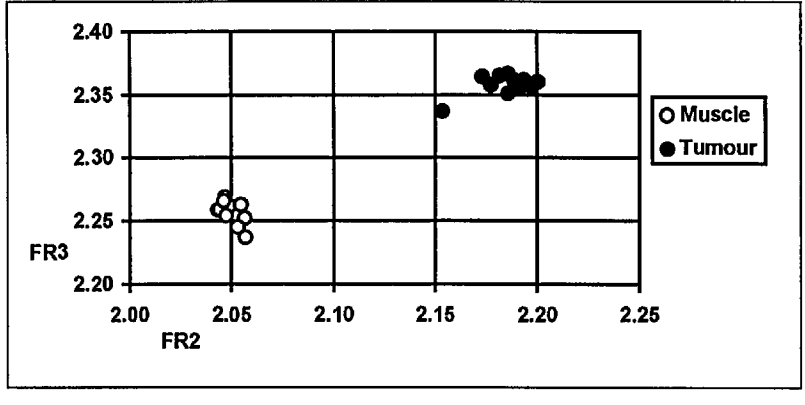
Graph 10.22 : Segmentation of 'Leg192' data - features FR1vFR2 (ROI 16x16)



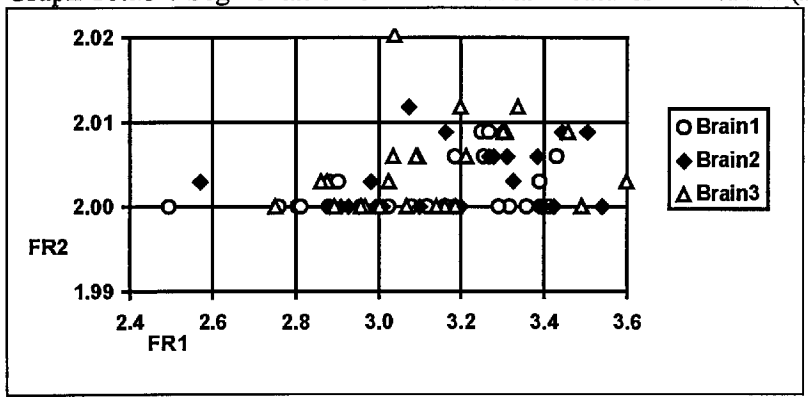
Graph 10.23 : Segmentation of 'Leg192' data - features FR2vFR3 (ROI 12x12)



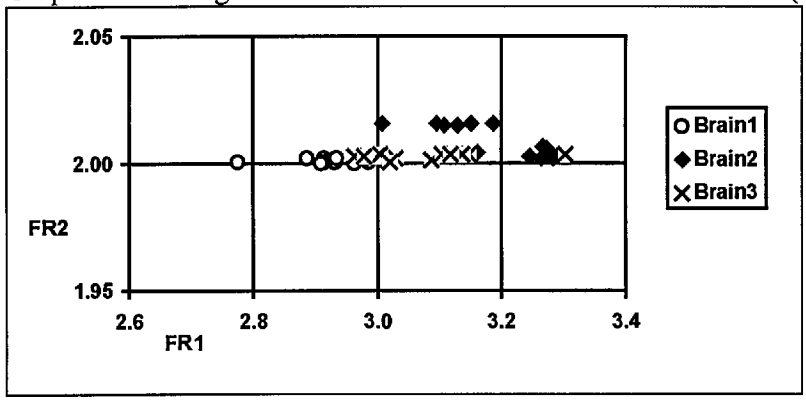
Graph 10.24 : Segmentation of 'Leg192' data - features FR2vFR3 (ROI 24x24)



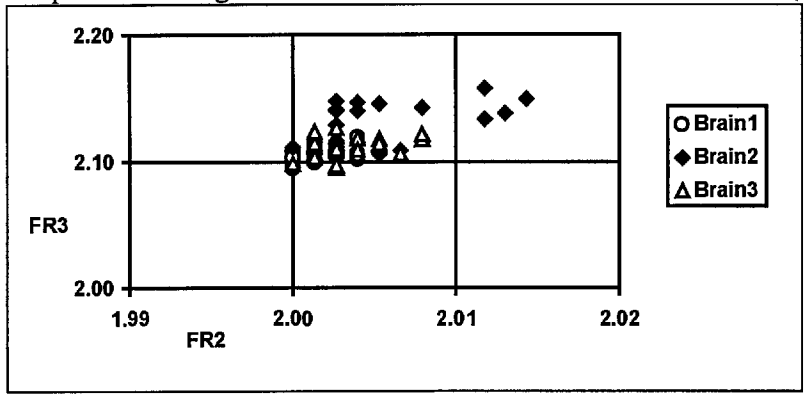
Graph 10.25 : Segmentation of 'Brain' data - features FR1vFR2 (ROI 8x8)



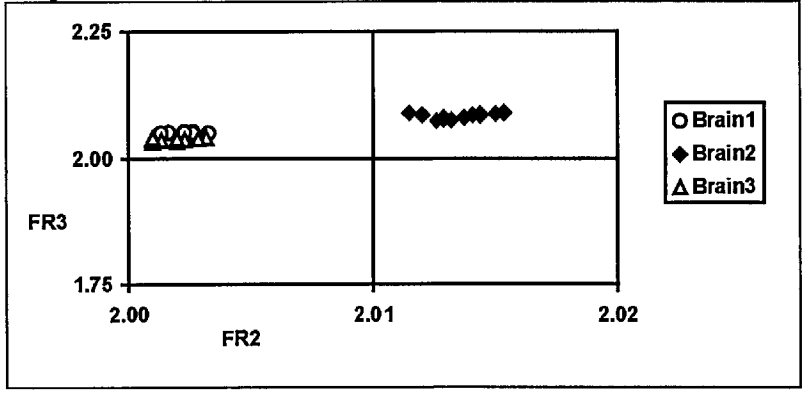
Graph 10.26 : Segmentation of 'Brain' data - features FR1vFR2 (ROI 16x16)



Graph 10.27 : Segmentation of 'Brain' data - features FR2vFR3 (ROI 12x12)



Graph 10.28 : Segmentation of 'Brain' data - features FR2vFR3 (ROI 24x24)



In theory, the performance of the PSD fractal dimension (FR1) is related to the ROI size. The clinical MRI images examined as part of this project limit the ROI size, and so far the potential of this parameter has not been demonstrated. **Graphs 10.29** to **10.30** show the results of a study of the discriminative performance of texture features FR1 and FR2 with large ROI applied to 'Brodatz' data (artificial texture samples). The performance of feature FR1 was shown to be at least as good as that of FR2. **Tables 10.10** and **10.11** present the DB performance measures for both features, along with the correlation coefficient between them. Excellent texture segmentation was achieved for both ROI sizes 32x32 and 64x64.

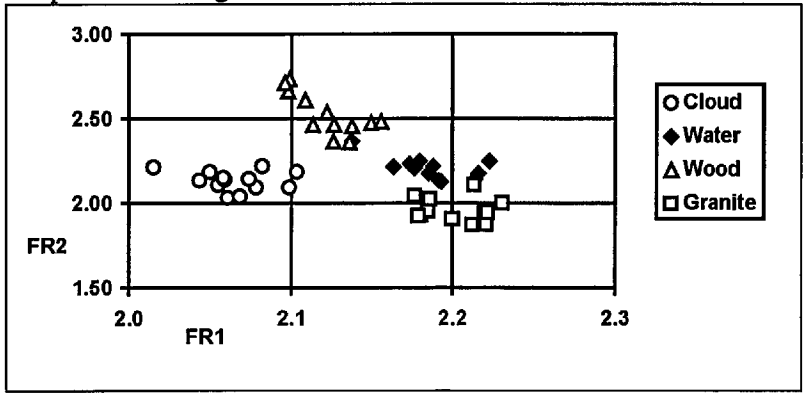
Table 10.10 : Comparison between PSD and blanket fractal features (ROI 32x32)

Parameter	Conditions	DBall
FR1	B8N0	0.70
FR2	B5N0 - blanket 11	0.68
Correlation coefficient = -0.34		

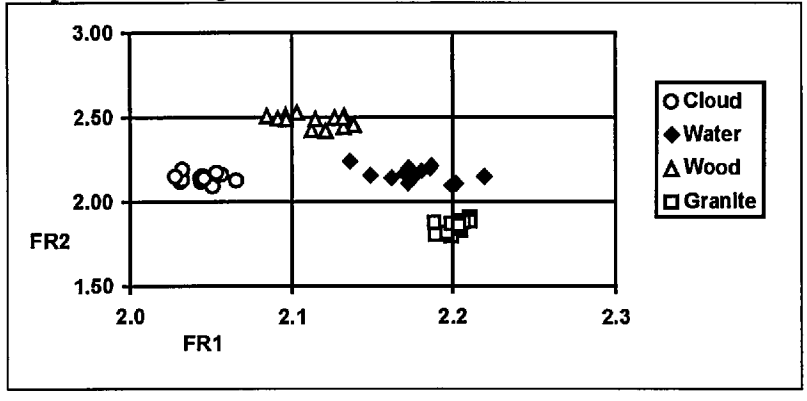
Table 10.11 : Comparison between PSD and blanket fractal features (ROI 64x64)

Parameter	Conditions	DBall
FR1	B8N0	1.63
FR2	B5N0 - blanket 11	1.09
Correlation coefficient = -0.46		

Graph 10.29 : Segmentation of 'Brodatz' data - features FR1vFR2 (ROI 32x32)



Graph 10.30 : Segmentation of 'Brodatz' data - features FR1vFR2 (ROI 64x64)



CHAPTER 11 - DISCUSSION OF RESULTS

11.1 INTRODUCTION

The first section of this final chapter provides a summary of the results obtained from the discriminative performance studies reported in **chapters 8** to **10** on the features derived from statistical, transform, and fractal texture analysis methods. The best performing texture features from all three analysis methods are identified, and the optimum image parameters for each of these features are summarised in the form of a table.

The second section of this chapter examines three methods of measuring the discriminative ability of texture features between distinct tissue regions. The relative performance of the best nineteen optimised texture features are examined using all three measurement methods. The purpose of this study is to validate the use of the DB performance measure against established statistical quantities, and to provide a global comparison of the best performing texture features. The relationship between these texture features is further established following a study of the correlation between them. The discussion section is a review of what has been learnt from these experiments, and includes some interpretation of the results and suggests possible sources of error.

The penultimate section of this chapter presents the overall conclusions of this project. This section comments on the degree of success this project has achieved with the goals set out in the introduction. From the results obtained with this project, it contemplates whether textural analytic methods have the potential to derive robust pathological indicators for MRI images, and to what degree this has been demonstrated. The relationships between the three approaches to texture analysis are also discussed, and how best they can be used for tissue characterisation.

The final section suggests future work to further clarify the results obtained for this project. For example, a number of additional experiments have been suggested to explore the influences of imaging factors such as signal-to-noise ration and resolution on texture feature performance.

11.2 SUMMARY OF TEXTURE FEATURE EVALUATION

11.2.1 STATISTICAL METHODS

The results of the performance evaluation of texture features from four approaches to extracting statistical texture information are presented in **chapter 8**. The performance of these features was examined for different ROI sizes (3x3, 5x5, 8x8, 12x12 &16x16) and with the image data scaled to a range of grey level bits (4-8). The influence of two methods of normalisation on feature performance was also examined: (1) histogram equalisation and (2) fixed first order image statistics ie. mean and standard deviation.

Five texture parameters calculated from first order image statistics were examined: mean, standard deviation, skewness, kurtosis, and coarseness. The overall discriminative performance of all features was shown to improve with ROI size, with the best performance achieved with ROI of 16x16. The effect of varying the number of data bits did not influence discriminative performance apart from for the skewness feature, which varied erratically. It is proposed that these features be extracted from 8-bit data. The second normalisation method (fixed first order image statistics) when applied to image data was found to enhance the performance of the skewness parameter. In the other four parameters, both data normalisation methods were found to be detrimental to performance. A high correlation coefficient was found between the mean, coarseness and standard deviation parameters. The former two features were eliminated on the basis of the superior discriminative performance of the latter. The kurtosis was also eliminated because of its comparatively poor performance with the skewness, with which it was found to be highly correlated.

Three implementations of run length analysis were examined as part of this project: vertical run length, horizontal run length and cluster size (2D run length). In each case, five texture features were extracted: short runs emphasis, long run emphasis, grey level non-uniformity, run length non-uniformity and run percentage. The overall discriminative performance of all features was shown to improve with ROI size, with the best performance achieved with ROI of 16x16. The effect of varying the number of data bits was also shown to influence discriminative performance. Performance was seen to improve as the data bits were reduced from 8 to 4, with optimum performance generally being

achieved with 5-bit data. Data normalisation of the image data was not found to improve the performance of the subsequently calculated texture features. In all three implementations of run length the discriminative performance of the grey level non-uniformity proved to be significantly better than the other four texture features. All three grey level non-uniformity parameters were found to be highly correlated. Of the three, the horizontal run length parameter is arbitrarily chosen as the best performing run length texture feature.

Three texture parameters calculated from the grey tone difference were examined: contrast, second angular moment, and entropy. The overall discriminative performance of all features was shown to improve with ROI size, with the best performance achieved with ROI of 16x16. The performance of the features was also shown to be influenced by the number of data bits, with optimum performance generally being achieved with 7-bit data. The effect of data normalisation was shown to be detrimental to the performance of all three features. As well as demonstrating comparable discriminative performance, all three texture features were found to be highly correlated. Two out of three features are redundant and therefore should be discarded. The contrast parameter is arbitrarily retained.

The grey tone difference method proposed as part of this project is derived from the aggregate of implementations in four individual directions. This approach attempts to eliminate directional sensitivity from the grey tone difference method, but it is not clear if vital texture information is being lost. In a comparison of the performance of this implementation against those in the individual directions, better overall performance was shown with the aggregate method, indicating that that no texture information is being lost.

Fifteen texture parameters calculated from the co-occurrence matrix have been examined: the standard fourteen features F1-F14 and an additional feature suggested as part of this project F6'. More descriptive names for these texture features are provided in **table 8.1** in **chapter 8**. The discriminative performance of all features generally improved with ROI size, with the best performance achieved with ROI of 16x16. It was very difficult to establish any trends of performance with data bits. Optimum performance was generally achieved with 7-bit data, despite minimal performance fluctuations in the range of 4-8 bits. The effect of data normalisation was shown to be detrimental to discriminative performance for all fifteen texture features. In order to reduce the feature set, seven features were eliminated on the basis of their relative discriminative performance: F3, F4, F5, F6, F12,

F13, and F14. As a results of a correlation study, the remaining eight features were divided into four groups of highly correlated features. All but the best performing feature in each group were discarded as being redundant. The four texture features that remained after this process of elimination are F1 (homogeneity), F11 (difference entropy), F8 (sum entropy), and F9 (entropy).

The co-occurrence method implemented as part of this project is derived from the aggregate of implementations with the co-occurring image pixels in four different directions. This approach attempts to eliminate directional sensitivity, but it is not clear if vital texture information is being lost. In a comparison of the performance of this multi-directional implementation against those in the individual directions, better overall performance was shown with the aggregate method, indicating that that no texture information is being lost. Another factor that affects the identification of the co-occurring pixels is the distance between them, know as the displacement vector distance. In texture analysis this is normally considered to be one pixel, because texture is a property of the relationship between a pixel and its neighbours. A comparison of the performance of the texture features calculated with a displacement vector distance of one and of two pixels was carried out. There was no discernible variation in feature performance.

11.2.2 TRANSFORM METHODS

The results of the performance evaluation of texture features extracted from the power spectrum and two implementations of the Walsh and Slant transforms are presented in **chapter 9**. The performance of these features was examined for different ROI sizes and with the image data scaled to a range of grey level bits. The influence of two methods of normalisation on feature performance was also examined: (1) histogram equalisation and (2) fixed first order image statistics ie. mean and standard deviation.

Six texture parameters calculated from the power spectrum were examined: energy, cut-off frequency, mean, standard deviation, skewness, and kurtosis. The cut-off frequency feature was found not to be influenced at all by image texture. Better (by a factor of two) discriminative performance was demonstrated by the energy, mean, and standard deviation texture features for ROI size 16x16 than for size 8x8. The performance of the skewness and the kurtosis were comparable for both ROI sizes. The effect of varying the number of data bits (4-8) did not influence discriminative performance for any of the features. It is

proposed that these features be extracted from 8-bit data. The effect of data normalisation was shown to be detrimental to the performance of all features. As the result of a correlation study, the power spectrum features were separated into two groups of highly correlated features. The features in these groups also demonstrated comparable performance. The energy and skewness texture features are arbitrarily chosen as being representative of the two groups of features, while the rest are discarded.

Thirty six texture features extracted from the Walsh transform and thirty two texture features extracted from the Slant transform were examined as part of this study. The resulting matrix from each transform was treated as raw image data, and texture features were extracted using run length analysis, co-occurrence analysis and spectrum analysis. Texture features were also derived from the symmetry properties of the Walsh transform. It was suggested in **chapter 5** that variations in the transform caused by texture translation could be eliminated using Fourier pre-processing. The effect on feature performance of applying the Fourier transform to the image data prior to the transformation was also investigated.

Of the large number of texture features examined from the Walsh and Slant transforms, only a small sub-set demonstrated the ability to discriminate clearly between the texture classes in the images examined. The five Walsh features identified were grey level non-uniformity (vertical & horizontal run length), difference average (co-occurrence), difference entropy (co-occurrence) and energy (Walsh spectrum analysis). The seven Slant features identified were grey level non-uniformity (vertical & horizontal run length), sum average (co-occurrence), sum entropy (co-occurrence), entropy (co-occurrence), difference entropy (co-occurrence), and energy (Slant spectrum analysis).

The discriminative performance of these features has been examined over a range of image data bits (6-8). In the absence of Fourier pre-processing, performance was found to be optimum with 6-bit image data. Where Fourier pre-processing was used, the effect of varying the image data bits did not influence performance. It was therefore proposed that under this condition, 8-bit data should be used. The performance of the Walsh and Slant features were also examined with two ROI sizes: 8x8 and 16x16. For all features, a significant improvement in performance was noted for those features extracted from the larger ROI (16x16). It was not possible to improve discriminative performance using either of the two data normalisation methods implemented. The optimum texture features

collected with and without Fourier pre-processing have been grouped separately. The overall performance of the features collected with Fourier pre-processing was shown to be significantly better. This group is made up of two Walsh features: difference average (co-occurrence) and energy (Walsh spectrum analysis), and two Slant features: entropy (co-occurrence) and difference entropy (co-occurrence).

The sixty eight Walsh and Slant transform texture features have also been extracted from the non-directional implementations of the same transforms that have been developed as part of this project. A similar evaluation of discriminative performance was also carried out on these features. Only six non-directional Walsh texture features demonstrated the ability to discriminate clearly between the texture classes in the images examined. No non-directional Slant features met this performance criteria. The non-directional Walsh features identified were contrast (co-occurrence), variance (co-occurrence), difference average (co-occurrence), sum variance (co-occurrence), difference variance (co-occurrence) and energy (Walsh spectrum analysis). It was only possible to obtain good discriminative performance with these features when Fourier pre-processing was used. The effect of varying the image data bits (6-8) was subtle, with the best performance achieved with 7-bit data. For all six features, a significant improvement in performance was noted for those features extracted from the larger ROI ie. 16x16 rather than 8x8. It was not possible to improve discriminative performance using either of the two data normalisation methods implemented. As all six features were highly correlated, it was possible to reduce the feature set to two parameters on the basis of performance ie. variance (co-occurrence) and energy (Walsh spectrum analysis).

11.2.3 FRACTAL METHODS

The results of the performance evaluation of the fractal dimension texture features derived from three different methods are presented in **chapter 10**. The performance of these features was examined for different ROI sizes and with the image data scaled to a range of grey level bits (4-8). The influence of two methods of normalisation on feature performance was also examined: (1) histogram equalisation and (2) fixed first order image statistics ie. mean and standard deviation.

The power spectral density (PSD) fractal dimension was calculated from the characteristic decay of the power spectrum of a particular image region. For the two ROI sizes

investigated (8x8 and 16x16), the performance of this texture feature was neither influenced by grey level bits nor the use of normalisation. The overall performance of this feature under these conditions was poor. It was suggested that better performance would only be achieved with much larger ROI (at least 32x32), providing more points on the power spectrum. It was not possible to find homogenous regions of a single tissue class of this size in the MRI data available for this project. A favourable illustration of this potential was demonstrated with ROI sizes 32x32 and 64x64 on artificial texture samples in **section 10.4** in **chapter 4**.

The fractal dimension can be defined by the difference in image information with successively decreasing spatial resolution. The blanket fractal dimension method uses a surface averaging technique to achieve this decrease in resolution. The number of successive averages carried out on the image is indicated by the blanket index. A variation in the performance of the blanket texture feature was seen over a range of blanket index values (in the range 7-15). Optimum performance was achieved with a blanket index of 11. The performance of the blanket texture feature was shown to improve with ROI size and vary inversely with grey level bits over the range (4-8). The use of data normalisation was shown to be detrimental to the performance of this feature.

The box method derives the fractal dimension from the difference in image surface area between successive versions of the same image at different resolutions. The performance of the box fractal dimension was examined for ROI of sizes 12x12 and 24x24 and over a range of grey level bits (4-8). Optimum performance was attained with the larger ROI and with 6-bit data, although there was little variation over the range of grey levels experimented with. The use of data normalisation was shown to be detrimental to the performance of this feature.

11.2.4 OPTIMISED TEXTURE FEATURES

The best performing texture features derived from the three approaches to texture examined as part of this project are presented in **table 11.1**. This table provides a summary of the optimum conditions under which the eight statistical features (1-8), eight transform features (9-16) and three fractal features (17-19) identified in this section are extracted. Each feature is catalogued by both the texture method used and its full name, with details of the image data grey levels and normalisation scheme required for optimum performance.

Table 11.1 : Optimised texture features

No.	Texture method	Texture feature	Grey levels	Normalisation	Test matrix	Comments
1	First order statistics	Standard deviation	8 bit	None	16x16	
2	First order statistics	Skewness	8 bit	Method 2	16x16	
3	Run length - horizontal	Grey level non-uniformity	5 bit	None	16x16	
4	Grey tone difference	Contrast	7 bit	None	16x16	Non-dir^n
5	Co-occurrence	F1 - homogeneity	7 bit	None	16x16	Non-dir^n, d = 1
6	Co-occurrence	F11 - difference entropy	7 bit	None	16x16	Non-dir^n, d = 1
7	Co-occurrence	F8 - sum entropy	7 bit	None	16x16	Non-dir^n, d = 1
8	Co-occurrence	F9 - entropy	7 bit	None	16x16	Non-dir^n, d = 1
9	Power spectrum	Energy	8 bit	None	16x16	
10	Power spectrum	Skewness	8 bit	None	16x16	
11	Walsh transform	Co-occurrence (F6' - difference average)	8 bit	None	16x16	Fourier pre-process
12	Walsh transform	Energy (Walsh spectrum)	8 bit	None	16x16	Fourier pre-process
13	Slant transform	Co-occurrence (F9 - entropy)	8 bit	None	16x16	Fourier pre-process
14	Slant transform	Co-occurrence (F11 - difference entropy)	8 bit	None	16x16	Fourier pre-process
15	Non-dir^n Walsh transform	Co-occurrence (variance)	7 bit	None	16x16	Fourier pre-process
16	Non-dir^n Walsh transform	Energy (Walsh spectrum)	7 bit	None	16x16	Fourier pre-process
17	Fractal	Power spectral density fractal dimension	8 bit	None	16x16	
18	Fractal	Blanket fractal dimension	5 bit	None	12x12	Blanket index = 11
19	Fractal	Box fractal dimension	6 bit	None	12x12	

For all nineteen features, performance has been shown to improve with ROI size. The ROI size entry in **table 11.1** is not optimal, but the largest possible due to the limited size of same class tissue regions in the MRI images examined. Further information about how the texture method is implemented is provided under the comments heading.

11.3 COMPARISON BETWEEN OPTIMISED TEXTURE FEATURES

11.3.1 PERFORMANCE COMPARISON

The three methods of measuring the discriminative performance of texture features that have been considered for this project are the Student's t-test, the Fisher linear classifier and the DB performance measure. All three methods are discussed in **chapter 7**. The DB performance measure has been used exclusively in all the texture analysis studies carried out as part of this project because of its simplicity. This section provides a comparison of the performance of the nineteen optimised texture features identified in **table 11.1** using all three measurement methods. In addition to providing a global comparison of the performance all the texture features considered, this study enables the behaviour of the DB performance measure to be checked against the Student's t-test and Fisher linear classifier measures. All nineteen texture features were acquired from 78 ROI evenly distributed among the three distinct tissue regions (muscle, fat & tumour) on several slices of the 'Thigh' MRI image data (see **chapter 7**).

Graphs 11.1 to **11.4** show the DB performance measure calculated for all nineteen features, firstly between all three tissue classes, and then between each pair of tissue classes. The higher the value, the better the discriminative performance. The parameters with the same shading are derived using the same texture method eg. parameters 5-8 are derived using the co-occurrence method. Similar performance trends were seen between the texture features for each pair of tissue classes.

Graphs 11.5 to **11.7** show the Fisher linear classifier calculated for all nineteen features between each pair of tissue classes. The higher the value, the more sensitive the texture feature. Similar performance trends were seen across the nineteen texture features for the discrimination between muscle and fat and between muscle and tumour (**graphs 11.5** and

11.7). The performance behaviour of the texture features in the discrimination between fat and tumour (**graph 11.6**) was, however, very different.

Graphs 11.8 to **11.10** show the student's t-test calculated for all nineteen features between each pair of tissue classes. A logarithmic plot of the t-test value enables comparison of the between the performance of these features. The more negative the value on the graph the better the performance of the feature. Similar performance trends were seen across the nineteen texture features for the discrimination between muscle and fat and between muscle and tumour (**graphs 11.8** and **11.10**), and again different performance behaviour was noted in the discrimination between fat and tumour (**graph 11.9**).

To justify the use of the simpler DB performance measure instead of the Student's t-test or Fisher linear classifier method, a correlation study of the behaviour of these three methods has been carried out. The results of this study are presented in **table 11.2**.

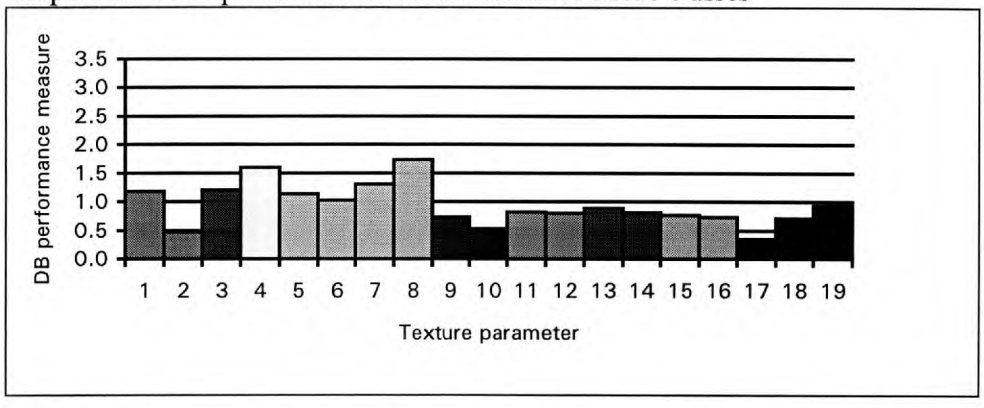
Table 11.2 : Correlation between performance measures

Measure 1	Measure 2	Correlation coefficient
Fisher classifier (classes 1&2)	T-test (classes 1&2)	-0.88
Fisher classifier (classes 2&3)	T-test (classes 2&3)	-0.98
Fisher classifier (classes 1&3)	T-test (classes 1&3)	-0.87
DB measure (classes 1&2)	Fisher classifier (classes 1&2)	0.97
DB measure (classes 2&3)	Fisher classifier (classes 2&3)	-0.11
DB measure (classes 1&3)	Fisher classifier (classes 1&3)	0.94
DB measure (classes 1&2)	T-test (classes 1&2)	-0.91
DB measure (classes 2&3)	T-test (classes 2&3)	-0.07
DB measure (classes 1&3)	T-test (classes 1&3)	-0.91

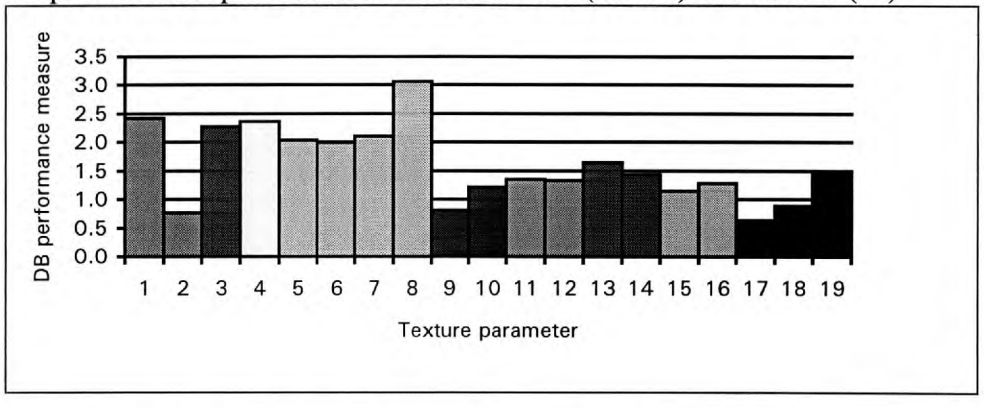
The correlation coefficient between two measures is calculated from the performance values obtained for all nineteen texture features. The Fisher linear classifier and the t-test performance measures were shown to be highly correlated for all three combinations of tissue classes. The DB performance measure was shown to be highly correlated with both the Fisher linear classifier and the t-test measure in the discrimination between muscle and fat (classes 1&2) and between muscle and tumour (classes 1&3), but un-correlated with both in the discrimination between fat and tumour (classes 2&3). The results of this correlation study indicate that apart from the disparity in the discrimination between tissue

classes 2&3, the DB parameter is as good a measure of discriminative performance as either the Student's t-test or the Fisher linear classifier. Although the disparity cannot be explained from the data in this limited study, it must be noted that the DB parameter was the only measure of the three that demonstrated consistent behaviour across all nineteen texture features for all combinations of tissue classes (**graphs 11.2** to **11.10**).

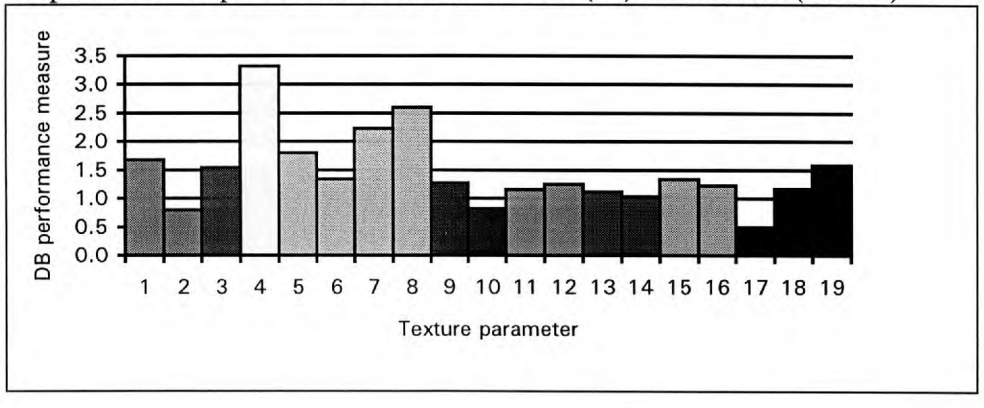
Graph 11.1 : DB performance between all three tissue classes



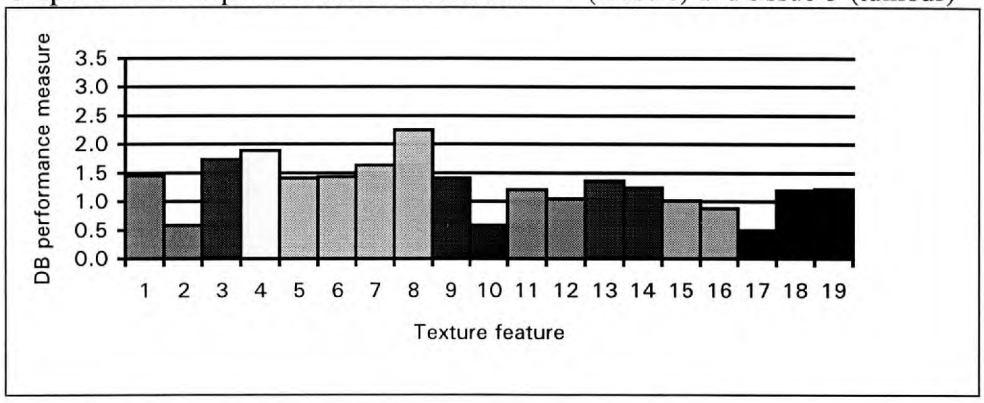
Graph 11.2 : DB performance between tissue 1 (muscle) and tissue 2 (fat)



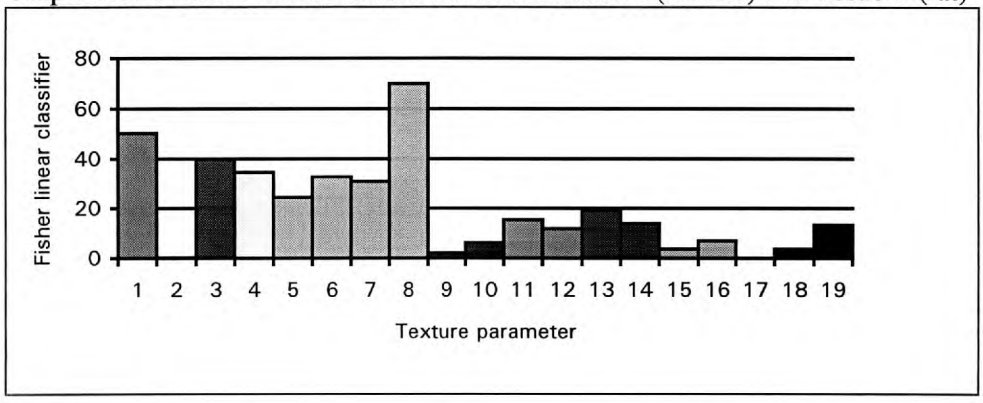
Graph 11.3 : DB performance between tissue 2 (fat) and tissue 3 (tumour)



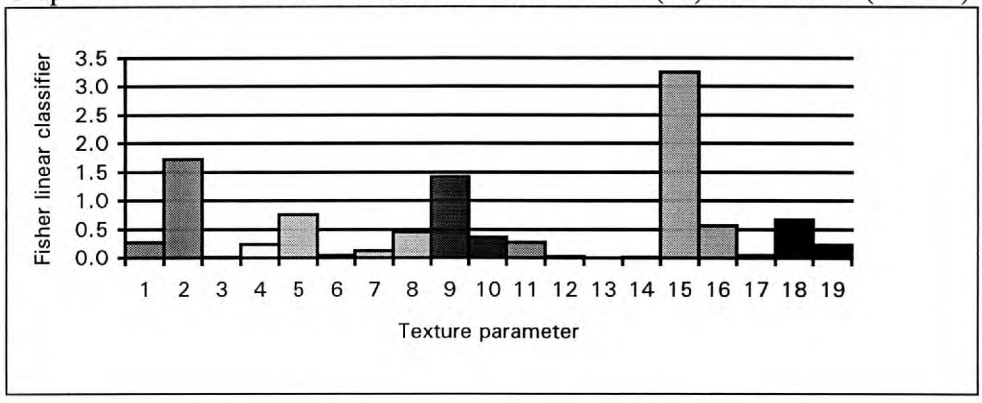
Graph 11.4 : DB performance between tissue 1 (muscle) and tissue 3 (tumour)



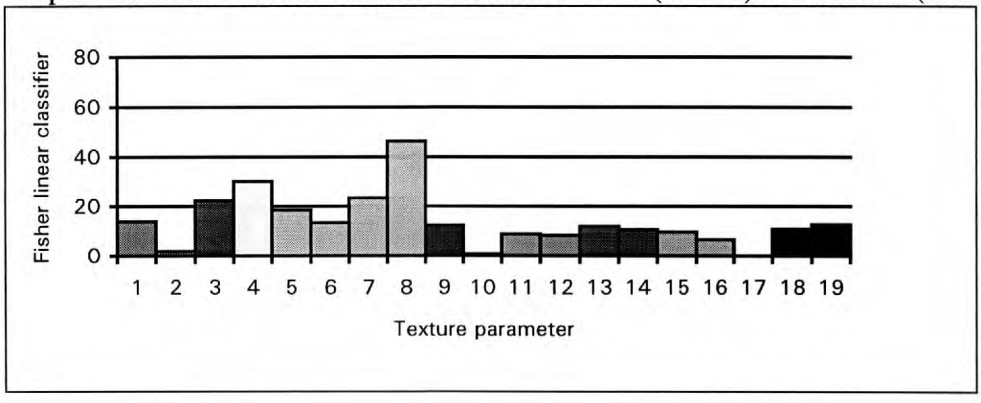
Graph 11.5 : Fisher linear classifier between tissue 1 (muscle) and tissue 2 (fat)



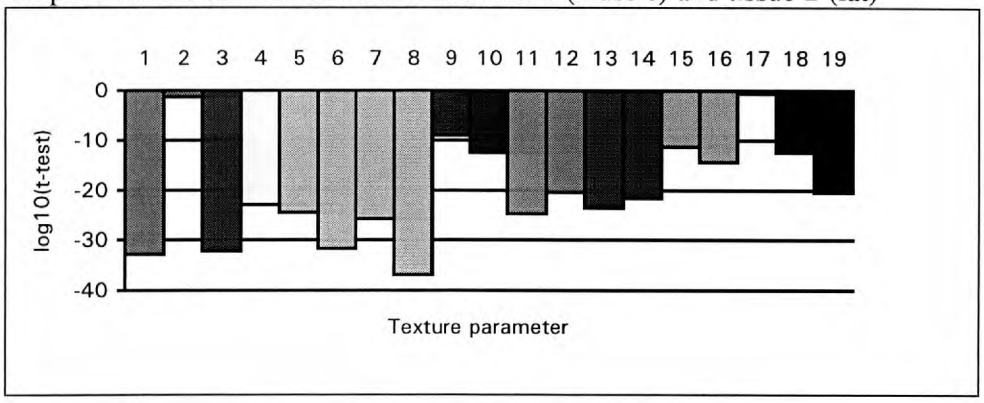
Graph 11.6 : Fisher linear classifier between tissue 2 (fat) and tissue 3 (tumour)



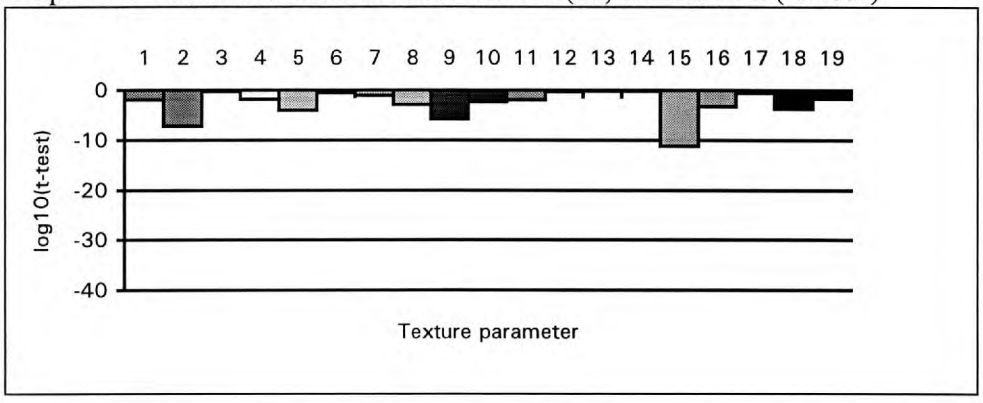
Graph 11.7 : Fisher linear classifier between tissue 1 (muscle) and tissue 3 (tumour)



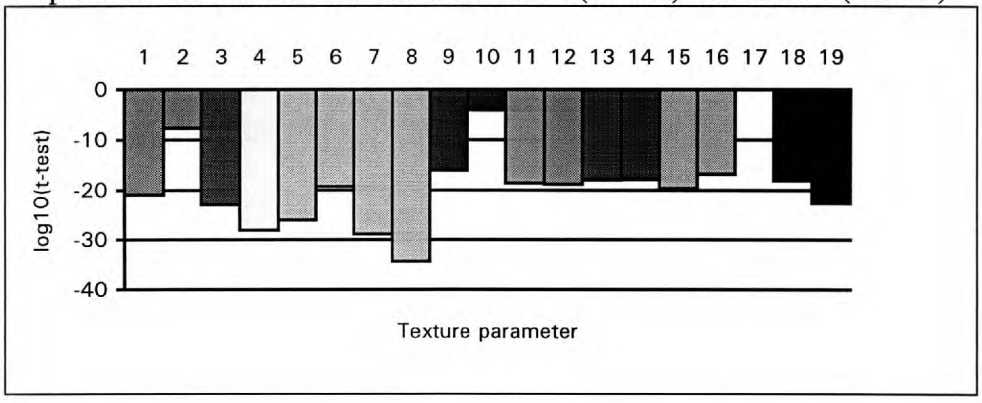
Graph 11.8 : Student's t-test between tissue 1 (muscle) and tissue 2 (fat)



Graph 11.9 : Student's t-test between tissue 2 (fat) and tissue 3 (tumour)



Graph 11.10 : Student's t-test between tissue 1 (muscle) and tissue 3 (tumour)



Graph 11.1 shows the DB performance measure for each of the nineteen optimised texture features identified in table 11.1 between three classes of tissue: muscle, fat and tumour. Assuming that the DB measure is a legitimate meter of discriminative performance, it is possible to comment on the relative abilities of the texture features. With the exception of the first order statistics parameter skewness, the statistical texture features perform much better than any of the transform or fractal texture features. On aggregate, the two best performing statistical texture features are probably the co-occurrence method entropy

feature and the grey tone difference method contrast feature. This result is also reflected by the Fisher linear classifier and student's t-test parameters.

11.3.2 CORRELATION

The nineteen best performing statistical, transform and fractal texture parameters identified as part of this project, and the conditions under which they perform optimally are summarised in **table 11.1**. The results of the performance evaluation carried out with the DB performance measure suggests that the statistical texture features demonstrate the best discriminative performance. The implication of this result is that the value of transform and fractal texture methods must depend on whether they can provide addition textural information to the statistical methods. The worth of these methods can therefore be assessed by examining the correlation coefficients between all nineteen texture features, as presented in **table 11.3**. A low correlation coefficient between features would indicate that they are measuring different properties of the texture.

With the exception of the first order statistics skewness feature, the worst performer, all eight statistical texture features are highly correlated ie. with correlation coefficients greater than 0.9. This set of texture features can therefore be reduced on the basis of performance alone without any texture information being lost. The most useful non-statistical texture features appear to be the non-directional Walsh transform co-occurrence variance parameter and the power spectrum features, which have correlation coefficients with the statistical texture features in the range of 0.60 to 0.84. Unfortunately neither of these features perform as well as any of the statistical features. The four Walsh and Slant transform features, the non-directional Walsh transform energy feature, and the blanket and box fractal dimension features are all highly correlated with the statistical texture features, and are therefore redundant. The low correlation of the power spectral density fractal dimension with all other parameters is accounted for by its comparatively poor performance.

Table 11.3 : Correlation between optimised texture features

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
1		0.34	-0.97	0.94	-0.92	0.97	0.95	0.94	-0.76	-0.82	0.91	-0.95	0.95	0.94	0.79	-0.90	-0.07	0.78	0.86
2	0.34		-0.43	0.50	-0.54	0.41	0.52	0.49	-0.52	0.01	0.40	-0.37	0.36	0.36	0.55	-0.44	0.04	0.44	0.43
3	-0.97	-0.43		-0.96	0.97	-0.97	-0.98	-0.96	0.82	0.73	-0.92	0.93	-0.96	-0.95	-0.82	0.89	0.04	-0.83	-0.89
4	0.94	0.50	-0.96		-0.97	0.91	0.98	1.00	-0.83	-0.64	0.86	-0.94	0.90	0.89	0.80	-0.84	0.14	0.85	0.93
5	-0.92	-0.54	0.97	-0.97		-0.91	-0.99	-0.97	0.84	0.64	-0.85	0.92	-0.90	-0.90	-0.84	0.86	-0.05	-0.86	-0.92
6	0.97	0.41	-0.97	0.91	-0.91		0.94	0.91	-0.77	-0.76	0.95	-0.91	0.97	0.95	0.80	-0.91	-0.17	0.77	0.82
7	0.95	0.52	-0.98	0.98	-0.99	0.94		0.98	-0.82	-0.68	0.89	-0.94	0.92	0.91	0.82	-0.86	0.05	0.85	0.92
8	0.94	0.49	-0.96	1.00	-0.97	0.91	0.98		-0.83	-0.63	0.86	-0.93	0.90	0.89	0.82	-0.85	0.14	0.88	0.94
9	-0.76	-0.52	0.82	-0.83	0.84	-0.77	-0.82	-0.83		0.49	-0.76	0.81	-0.81	-0.81	-0.76	0.75	-0.10	-0.87	-0.82
10	-0.82	0.01	0.73	-0.64	0.64	-0.76	-0.68	-0.63	0.49		-0.68	0.78	-0.72	-0.72	-0.59	0.75	0.32	-0.47	-0.56
11	0.91	0.40	-0.92	0.86	-0.85	0.95	0.89	0.86	-0.76	-0.68		-0.81	0.94	0.92	0.71	-0.81	-0.08	0.76	0.80
12	-0.95	-0.37	0.93	-0.94	0.92	-0.91	-0.94	-0.93	0.81	0.78	-0.81		-0.89	-0.89	-0.80	0.88	0.02	-0.80	-0.86
13	0.95	0.36	-0.96	0.90	-0.90	0.97	0.92	0.90	-0.81	-0.72	0.94	-0.89		0.99	0.77	-0.88	-0.11	0.79	0.84
14	0.94	0.36	-0.95	0.89	-0.90	0.95	0.91	0.89	-0.81	-0.72	0.92	-0.89	0.99		0.79	-0.88	-0.10	0.80	0.84
15	0.79	0.55	-0.82	0.80	-0.84	0.80	0.82	0.82	-0.76	-0.59	0.71	-0.80	0.77	0.79		-0.93	-0.21	0.77	0.76
16	-0.90	-0.44	0.89	-0.84	0.86	-0.91	-0.86	-0.85	0.75	0.75	-0.81	0.88	-0.88	-0.88	-0.93		0.31	-0.72	-0.76
17	-0.07	0.04	0.04	0.14	-0.05	-0.17	0.05	0.14	-0.10	0.32	-0.08	0.02	-0.11	-0.10	-0.21	0.31		0.16	0.19
18	0.78	0.44	-0.83	0.85	-0.86	0.77	0.85	0.88	-0.87	-0.47	0.76	-0.80	0.79	0.80	0.77	-0.72	0.16		0.91
19	0.86	0.43	-0.89	0.93	-0.92	0.82	0.92	0.94	-0.82	-0.56	0.80	-0.86	0.84	0.84	0.76	-0.76	0.19	0.91	

11.4 DISCUSSION

All the texture features examined performed better with larger ROI. The largest size that could be used consistently with the MRI test image data was 16x16, ie. 256 data points or 8 data points on a power spectrum. In some cases it was also possible to obtain ROI of size 24x24. Without the appropriate image data, one can only presume that the texture features would perform better with ROI larger than this.

The influence of normalising the image data prior to feature extraction was generally shown to be detrimental to feature performance. This can probably be accounted for by the lack of second order statistical differences between the tissue classes in the 'Thigh' data being compared. Most of the information enabling texture discrimination was in the first order statistics, which was conspicuous by the differences in image intensity of the tissue groups in the image data. After the first order image statistics have been normalised, the differences in second order statistics between tissue classes may be limited further by the presence of image noise.

More value may be found from normalisation methods when there is greater difference between the second order statistics of the tissue classes, or when the ROI are from different sources eg. they have been acquired with different receiver gain settings. Normalisation is an important application for the separation of pathologies that may have similar first order statistics ie. appearance, but different in 2nd order statistics. It is important to limit image noise, because these differences may not be seen if the signal-to-noise ratio is too low.

Apart from the skewness feature, whose erratic behaviour cannot easily be explained, none of the first order statistical features are influenced by the number of data bits the image data is scaled to. This is because varying the number of grey levels serves only to scale the frequency histogram. Run length feature performance improves as the data bits are reduced, peaking at 5-bits. The method is probably more stable under these circumstances, because the small variations due to the additive image noise are more likely to be smoothed out with less image grey levels. The performances of the texture features derived from the grey tone difference method and the co-occurrence matrix are not significantly influenced by the number of image grey levels.

None of the texture features derived from the Fourier transform are influenced by the number of grey levels. Bit rate differences in image data correspond with differences in high frequency components in the Fourier transform. These high frequency components do not make large contributions to the image texture. This applies to the power spectrum texture features, the power spectral density fractal dimension and all forms of the FFT pre-processed Walsh and Slant transforms. For the remaining features derived from the Walsh and Slant transforms without pre-processing, and the blanket and box fractal dimensions, better performance was generally achieved with fewer grey levels.

It has been difficult to establish a clear relationship between the performance of the texture features and the image grey levels. This again must be due to the high levels of noise in the test data, and depends much on the noise sensitivity of the texture feature. It is probably the case that the more sensitive the feature is to noise, the more dramatic the improvement in performance it will experience with reduction in grey level.

It is very likely that MRI image texture may vary in orientation due to the natural shape of tissue structures in the body. We are particularly interested in texture methods that are insensitive to texture orientation for this reason. Experiments with both standard and non-directional forms of the grey tone difference and co-occurrence methods have been carried out to explore whether texture information is lost at the expense of eliminating orientation sensitivity. In both cases performance was not compromised for non-directionality. The non-directional implementation of the Walsh transform also performed comparably with its standard counterpart.

Although these results are encouraging, it is not clear how much directional texture information is contained in the MRI test images. More appropriate test data are probably required to validate the results of this study. It is not inconceivable that the orientation of tissue texture may change with the onset of disease. Although our primary concern is to identify texture features that are not sensitive to texture orientation, there will still be a need for orientation sensitive texture methods for such applications.

Early pattern recognition studies have shown that the Walsh and Slant transforms are not positionally invariant. In the context of texture analysis, this means that the transforms will vary to some degree depending on where the ROI is placed in the texture sample. This effect will become more significant as the size of the texture primitives increases. This

dependency can be eliminated by Fourier processing the image data before applying the image transform. A measurable improvement in the discriminative performance of Walsh and Slant texture features was noted when this pre-processing method was applied.

All the MRI test images were acquired in the period 1990/91. By today's standards these images were collected using old technology, and to a certain degree this is reflected in the quality of the images. The images appear noisy, and in the case of the abdomen images, there are respiratory motion artefacts. A certain level of image quality is required for the extraction of texture information ie. images with good signal-to-noise ratio which are free of artefacts. The role of MRI Quality Assurance (QA) is therefore important in a site that intends to carry out texture analysis. At the time of publication of this thesis, image quality is generally much better. Gating, motion compensation, and breath-hold techniques make it possible to obtain good quality abdomen images. Hi-resolution imaging techniques also enable larger ROI to be used as input to the texture analysis algorithms.

The following section comments on the overall performance from the segmentation examples provided with each set of optimised texture features in the results chapters: statistical features (**chapter 8, section 8.7**); power spectrum features (**chapter 9, section 9.2.5**); Walsh and Slant features (**chapter 9, section 9.3.5**); Non-directional Walsh features (**chapter 9, section 9.4.5**); fractal features (**chapter 10, section 10.4**).

In the 'Thigh' image data, all optimised features are capable of clearly differentiating between the crucial muscle tissue and muscle tumour tissue groups with ROI of at least 16x16. A clearly more difficult task was to differentiate between the muscle tumour and fat tissue groups, which have the same average image intensity ie. first order statistics. This would firstly indicate that there are differences in second order statistics between these two groups, and secondly that the texture features can measure these second order differences. The large degree of overlap of the two clusters of data points in feature space indicates that the statistical, power spectrum and fractal texture features all fail to segment these two tissue groups. The Walsh and Slant texture features performed slightly better; there was only partial overlap of the clusters of data points. In the case of the non-directional Walsh texture features, these tissue groups are almost completely separated in feature space. To demonstrate the improved performance with ROI size, it was possible to clearly segment all three tissue classes with the fractal features extracted from a ROI of size 24x24.

In the 'Abdomen' image data, the best discrimination between the crucial liver tissue and liver tumour tissue was achieved with the power spectrum texture features. There was good clustering of both liver tissues and fat tissue, and all three tissues were segmented in feature space. The only other set of texture features that were able to separate the two liver tissue groups were the Walsh features. These features also separated the fat tissue data points into two sub-groups in feature space; one sub-group formed a well defined cluster that was completely segmented from the liver tissue groups, and the other was more loosely clustered and overlapped with the two liver tissue classes. It is likely that the latter of these sub-groups is due to respiratory motion artefacts.

It was not possible for any of the statistical, Slant, non-directional Walsh, or fractal texture features to separate liver tumour tissue from liver tissue in feature space, despite evidence of grouping of same class data points. The Slant, non-directional Walsh, and fractal features were all able to separate the fat tissue from the liver tissues. The statistical texture features were not able to clearly separate any of these three tissue classes, despite some grouping of class data points. In the case of the statistical and non-directional Walsh texture features, an additional cluster containing several points from both types of liver tissue was found amongst the fat data points. This cluster probably corresponds to regions taken in the liver with fat infiltration (striation). Like the Walsh transform features, the statistical, Slant, and fractal texture features also separated the fat tissue data points into two sub-groups in feature space; one authentic fat tissue cluster, and one containing tissue samples probably influenced by respiratory motion artefacts.

In the 'Leg104' image data, only the fractal dimension texture features could completely separate the tumour and muscle tissue groups in feature space. Although it was possible to segment the two tissue types using statistical texture features, there was clearly more than one muscle tissue cluster in feature space. In the case of the power spectrum, Walsh, Slant, and Non-directional Walsh texture features, there were multiple clusters for both tissue classes, making it difficult to clearly segment the two tissue types. This phenomena is caused by a combination of the inherent variability of the image texture in such a large field of view, and the poor image quality.

The results obtained with the 'Leg192' image data were much better than for the previous data set. It was possible to clearly segment muscle tissue and tumour tissue in feature space with the statistical, power spectrum and fractal texture features. It was not possible to

clearly separate the tissue classes for the Walsh, Slant and non-directional Walsh texture features, because there was some overlap between their data points. For all optimised texture features, there was evidence of multiple clusters in the muscle tissue data points.

Feature stability has been established by examining the variability of the texture features extracted in mid-brain slices of three healthy volunteers. In all the slices examined from the 'Brain' image data, the ROI were consistently positioned in the grey matter, avoiding the any of the obvious structures. The most stable texture features were derived from the power spectrum and the Walsh transform. The least stable features were some of the statistical texture features. The stability of a texture feature, however, can only be considered as a benefit if it is matched by good discriminative performance.

The main source of error with textural tissue characterisation is the heterogeneity of the texture sample. The effect of texture homogeneity has been illustrated in the 'Abdomen' and 'Leg104' segmentation examples above. In the 'Abdomen' image data, the texture parameters for several liver tissue samples are shifted towards those of fat, probably as a results of fat deposits (striation) in the liver. Respiratory motion artefacts have also been shown to cause the texture parameter values for fat samples to be shifted. In the 'Leg104' image data, both poor image quality and partial volume effects contribute to sub-clustering of the texture feature values. These errors can be minimised with an improvement in image quality and with more careful positioning of ROI in the texture sample.

The best nineteen texture features derived from the three approaches to texture analysis have been identified with the optimum conditions for their extraction in **section 11.2.4**. The comparative DB performance measures calculated for each of these features using 'Thigh' image data indicate that statistical texture features perform better than any of the transform or fractal methods (**graphs 11.1** to **11.4**). We have also shown from the segmentation examples in this discussion, however, that the relative performance of the optimised features is dependent on the image data being examined. The process of identifying the best performing texture features is therefore a non trivial process with the MRI test data used for this project. The disparity in performance may be explained by differences in second order statistical information content between the image sets, perhaps due to different levels of image noise.

11.5 PROJECT CONCLUSIONS

This project considers the design and implementation of texture analysis methods to assist clinicians in their interpretation of images from MRI. Three approaches to texture analysis have been identified from a large review of clinical and non-clinical applications. Sets of statistical, transform and fractal texture feature algorithms have been successfully implemented in the *MAIVIS* image visualisation environment. *MAIVIS* has been designed to allow all three sets of texture features to be easily extracted from the clinical MRI test data. Feature values are easily exported to the PC from this platform for discriminative performance studies.

The discriminative performance of each set of texture features was studied with the influence of ROI size, normalisation, bit rate, texture orientation, and texture translation. From this study it was possible to identify the optimum conditions under which each feature could be extracted. The feature set was reduced to a selection of the best performing features from each method (nineteen in all) using correlation to eliminate redundancy.

The project specifically explores the use of texture analysis to derive robust pathological indicators for MRI images. The nineteen optimised texture features were examined more closely to assess whether they provide enough diagnostic information to be clinically useful. Varying degrees of tissue segmentation were achieved between overtly different tissue classes in the MRI test images, demonstrating at least some clinical potential. It was clear from these studies that texture feature performance was very sensitive to image noise, and that the quality of the test data was an important factor in their clinical value.

It was proposed by (HARALICK *et al*, 1973) that all the texture information in an image could be extracted using the co-occurrence matrix. The relative performance of the nineteen optimised texture features was seen to depend on both the image data set and the tissue groups being separated. This variation in parameter performance is probably due to image noise influencing the second order statistical information content of the texture samples. It was therefore not possible to further reduce the feature set, nor was it possible to examine the claim of HARALICK *et al*, with the current set of MRI test data.

The results of this project work clearly demonstrate the clinical potential of texture analytic tools for tissue characterisation in MRI. With the advancement of MRI technology since the test

images were acquired in 1990/91, coupled with a well implemented site QA programme, it is now possible to obtain very high quality clinical images. With such data, it may be possible to improve on the excellent results obtained with this project, and fully realise this clinical potential of texture analysis.

11.6 FUTURE WORK

MORE COMPREHENSIVE MRI TEST DATA

- High quality images - good signal-to-noise ratio and minimal artefacts
- Hi-resolution images - matrix of 512x512 (or larger) with small FOV
- Identify tissue samples that contain second order statistical information
- Apply to clinical applications - both overt and subtle differences between tissue groups

POSSIBLE IMPROVEMENTS TO EXPERIMENTS

- Extend feature segmentation from 2D to a multi-dimensional approach
- Refine optimisation criteria for feature extraction - utilise high quality images
- Comparative study of nineteen identified features - utilise high quality images
- Apply optimised texture features to specific clinical studies
- Examine performance of transform methods with large ROI - utilise hi-resolution images
- Examine performance of fractal methods with large ROI - utilise hi-resolution images

FURTHER STUDY OF THE EFFECTS OF ORIENTATION AND TRANSLATION

- Deliberately rotate texture samples
- Examine the effect of non-directional texture methods
- Deliberately offset texture samples
- Examine the effect of positionally invariant texture methods

FURTHER STUDY OF PARAMETER DEPENDENCY

- Examine the effect of SNR degradation
- Examine the effect of resolution: field of view and slice width (voxel size)
- Examine the effect of acquisition sequence
- Examine the effect of issues relating to scanner QA
- Examine the effect of different scanners - different models and field strengths

POST FEATURE EXTRACTION NORMALISATION

- Develop a method to normalise texture parameters between repeated trials
- Measure signal-to-noise ratio and texture features from a test object
- Modify imaging bandwidth to compensate for reference SNR
- Compare reference texture features with those from tissue samples
- Normalise tissue texture values

TEXTURE FEATURE DATABASE

- Attempt to uniquely identify certain tissues from normalised texture features
- Make 'black box' texture algorthms freely available to clinical research projects
- Make a survey of the results being published using these methods
- Make a large database of texture feature values
- Assess the clinical efficacy of texture analysis from this data

BIBLIOGRAPHY

ADAMS, A. H., PALING, M. R., ABBITT, P. L. and MERICKEL, M. B. (1982): 'Image processing as applied to magnetic resonance image of human breast tissue', *Proc. IEEE Ninth Annual Conference of Engineering in Medicine and Biology Society.*

ALEXANDRIDIS, N. A. (1971a): 'Relations among sequency, axis symmetry, and period of walsh functions', *IEEE Trans. Inf. Theory*, **17**, pp. 495-497.

ALEXANDRIDIS, N. A. and KLINGER, A. (1971b): 'Walsh orthogonal functions in geometrical feature extraction', *Proc. Symp. Walsh Function Applications*, Wash. D.C., pp. 18-25.

ALLINSON, A. A. and LAWSON, M. (1990): 'Accurate textural characterisation using fractal techniques', *IEE Colloquim on the Application of Fractal Techniques in Image Processing,* London, pp. 6.1-6.6.

ALVAREZ, R. E., MARSHALL, W. H. and LEWIS, R. (1981): 'Tissue characterization using energy-selective computed tomography', *SPIE. Proc.*, **273**, pp. 301-307.

ANDREWS, H. C. and PRATT, W. K. (1969): 'Transform image coding', *Proc. Symp. Computer Processing in Communications*, pp. 63-84.

ANDREWS, H. C. and PRATT, W. K. (1970): 'Digital image transform processing', *Proc. Symp. Walsh Function Applications*, Wash. D.C., pp. 183-194.

ANDREWS, H. C. (1971): 'Walsh functions in image processing, feature selection and pattern recognition', *Proc. Symp. Walsh Function Applications*, Wash. D.C., pp. 26-32.

BALABAN, R. S., CHESNICK, S., HEDGES, K., SAMAHA, F. and HEINEMAN, F. W. (1991): 'Magnetization transfer contrast in MR imaging of the heart', *Radiology*, **180**, pp. 671-675.

BARNSLEY, M. F. (1988): 'The science of fractal images' (Springer-Verlag, New York).

BARTLETT, M. L. (1991): 'Comparison of methods for measuring fractal dimension', *Australas. Phys. Eng. Sci. Med.*, **14**, pp. 146-152.

BOYCE, D. W. M., BEAN, J. P., FREE, S. L. and STRAUGHAN, K. (1991): 'Development of an integrated diagnosis workstation for the display and analysis of multi-modal medical images', World Congress on Medical Physics and Biomedical Engineering, Kyoto.

BOTTOMLEY, P.A., FOSTER, R.E., ARGERSINGER, R.E. and PFEIFER, L. (1984): 'A review of normal tissue hydrogen NMR relaxation times and relaxation mechanisms from 1 to 100MHz: Dependence on tissue type, NMR frequency, temperature, species, excision, and age', *Med. Phys.*, **4**, pp. 425-448.

BRODATZ, P. (1966): 'Textures', Dover Publications, New York.

CADY, E. B., GARDENER, J. E. and EDWARDS, R. H. T. (1983): 'Ultrasonic tissue characterisation of skeletal muscle', *Eur. J. Clin. Invest.*, **13**, pp. 469-473.

CALDWELL, C. B., STAPLETON, S. J., HOLDSWORTH, D. W., JONG, R. A., WEISER, W. J., COOKE, G. and YAFFE, M. J. (1990): 'Characterisation of mammographic parenchymal pattern by fractal dimension', *Phys. Med. Biol.*, **35**, pp. 235-247.

CHANDRASEKARAN, K., AYLWARD, P. E., FLEAGLE, S. R., BURNS, T. L., SEWARD, J. B., TAJIK, A. J., COLLINS, S. M. and SKORTON, D. J. (1989): 'Feasibility of identifying amyloid and hypertrophic cardiomyopathy with the use of computerized quantitative texture analysis of clinical echocardiographic data', *J. Am. Coll. Cardiol.*, **13**, pp. 832-840.

CHIEN, Y. P. and FU, K. S. (1974): 'Recognition of x-ray picture patterns', *IEEE Trans. Syst. Man. Cybern.*, **4**, pp. 145-156.

CHING-MING, W., YUNG-CHANG, C. and KAI-SHENG, H. (1992): 'Texture features for classification of ultrasonic liver images', *IEEE Trans. Med. Imaging*, **11**, pp. 141-152.

CLARKE, R. J. (1985): 'Transform coding of images' (Academic Press, London)

CLOOSTERMANS, M. J., MOL, H., VERHOEF, W. A., THIJSSEN, J. M. and KUBAT, K. (1986): 'In vitro estimation of acoustic parameters of the liver and correlations with histology', *Ultrasound Med. Biol.*, **12**, pp. 39-51.

COCHRAN, W. T., COOLEY, J. W., FAVIN, D. L., HELMS, H. D., KAENEL, R. A., LANG, W. W., MALING, G. C., NELSON, D. E., RADER, C. M. and WELCH, P. D. (1967): 'What is the fast Fourier transform?', *Proc. IEEE*, **55**, pp. 1664-1674.

COLEMAN, A. J., TONGE, K. A. and RANKIN, S. C. (1982): 'The power spectral density as a texture measure in computed tomographic scans of the liver', *Br. J. Radiol.*, **55**, pp. 601-603.

CONNERS, R. W. and HARLOW, C. A. (1980): 'A theoretical comparison of texture algorithms', *IEEE Trans. Pattern. Anal. Machine Intel.*, **2**, pp. 202-222.

COOLEY, J. W. and TUKEY, J. W. (1965): 'An algorithm for the machine calculation of complex Fourier series', *Math. of Comput.*, **19**, pp. 297-301.

COOLEY, J. W., LEWIS, P. A. W. and WELCH, P. D. (1967): 'Historical notes on the fast Fourier transform', *Proc. IEEE*, **55**, pp. 1675-1677.

CRAWFORD, D. C., MORRIS, D. T., FENTON, D. W. and PRYCE, W. I. (1985): 'Possible applications of digital analysis of ultrasonic images of the placenta', *Ultrasound Med. Biol.*, **11**, pp. 79-84.

CRAWLEY, E. O. (1990): 'In vivo tissue characterization using quantitative computer tomography: a review', *J. Med. Eng. Technol.*, **14**, pp. 233-242.

DE CERTAINES, J. D., HENRIKSEN, O., SPISNI, A., CORTSEN, M. and RING, P. B. (1993): 'In vivo measurments of proton relaxation times in human brain, liver and skeletal

muscle: a multicenter MRI study', *Magn. Reson. Imaging*, **11**, pp. 841-850.

DE WILDE, J. P. (1994): 'Parameterisation of MR system performance: Towards optimised measures of image quality', PhD thesis, University of London.

DODD, N. (1987): 'Multispectral texture synthesis using fractal concepts', *IEEE Trans. Pattern. Anal. Machine Intel.*, **9**, pp. 703-707.

DUERINCKX, A. J. and MACOVSKI, A. (1979): 'Information and artifact in second order statistics from computed tomography (CT) images', *SPIE. Proc.*, **173**, pp. 250-257.

ENG, J., CECKLER, T. L. and BALABAN, R. S. (1991): 'Quantitative 1H magnetisation transfer imaging *in vivo*', *Magn. Reson. Med.*, **17**, pp. 304-314.

FEDER, J. (1988): 'Fractals' (Plenum Press, New York).

FINETTE, S., BLEIER, A. R., SWINDELL, W. and HABER, K. (1983): 'Breast tissue classification using diagnostic ultrasound and pattern recognition techniques: II. Experimental results', *Ultrason. Imaging*, **5**, pp. 71-86.

FORTIN, C., KUMARESAN, R., OHLEY, W. and HOEFER, S. (1992): 'Fractal dimension in the analysis of medical images', *IEEE Eng. Med. Biol.*, June, pp. 65-71.

FREE, S. (1994): 'The enhancement, extraction and parameterisation of diagnostic features im mammograms', PhD thesis, University of London.

GALLOWAY, M. (1974): 'Texture analysis using gray level run lengths', *Comput. Graph. Image Process.*, **4**, pp. 172-199.

GERSONDE, K., TOLXDORFF, T. and FELSBERG, L. (1985): 'Identification and characterization of tissues by T2-selective whole-body proton NMR imaging', *Magn. Reson. Med.*, **2**, pp. 390-401.

GLAZER, G. M., WOOLSEY, E. J., BORRELLO, J., FRANCIS, I. R., AISEN, A. M., BOOKSTEIN, F., AMENDOLA, M. A., GROSS, M. D., BREE, R. L. and MARTEL, W. (1986): 'Adrenal tissue characterization using MR imaging', *Radiology*, **158**, pp. 73-79.

GOLDBERGER, A. L. (1992): 'Fractal mechanisms in the electrophysiology of the heart', *IEEE Eng. Med. Biol.*, June, pp. 47-52.

GONZALEZ, R. C. and WINTZ, P. (1987): 'Digital Image Processing', Second edition (Addison-Wesley, Massachusetts).

GOTLIEB, C. C. and KREYSZIG, H. E. (1990): 'Texture descriptors based on cooccurrence matrices', *Comput. Vision Graphics Image Processing*, **51**, pp. 70-86.

HARALICK, R. M., SHANMUGAM, K. and DINSTEIN, I. (1973): 'Textural features for image classification', *IEEE Trans. Syst. Man. Cybern.*, **3**, pp. 610-621.

HARALICK, R. M. (1979): 'Statistical and structural approaches to texture', *Proc. IEEE*, **67**, pp. 786-804.

HARMUTH, H. F. (1968): 'A generalized concept of frequency and some applications', *IEEE Trans. Inf. Theory*, **14**, pp. 375-382.

HO, B. K., RATIB, O. and HORII, S. C. (1991): 'PACS workstation design', *Comput. Med. Imaging Graph.*, **15**, pp. 147-55.

HORII, S. C., LEVINE, B. A., GOGER, G., MUN, S. K., FIELDING, R., GARRA, B. S., LO, S. C., KRASNER, B. and BENSON, H. (1992): 'A comparison of case retrieval times: film versus picture archiving and communication systems', *J. Digit. Imaging*, **5**, pp. 138-143.

HU, B. S., CONOLLY, S. M., WRIGHT, G. A., NISHIMURA, D. G. and MACOVSKI, A. (1992): 'Pulsed saturation transfer contrast', *Magn. Reson. Med.*, **26**, pp. 231-240.

INSANA, M. F., WAGNER, R. F., GARRA, B. S., BROWN, D. G. and SHAWKER, T. H. (1986a): 'Analysis of ultrasound texture via generalised Rician statistics', *Optical Eng.*, **25**, pp. 743-748.

INSANA, M. F., WAGNER, R. F., GARRA, B. S., MOMENAN, R. and SHAWKER, T. H. (1986b): 'Pattern recognition methods for optimising multivariate tissue signatures in diagnostic ultrasound', *Ultrason. Imaging*, **8**, pp. 165-180.

JASPAN, T., WORTHINGTON, B. S. and HOLAND, I. M. (1988): 'A comparative study of magnetic resonance imaging and computed tomography-assisted myelography in spinal dysraphism', *Br. J. Radiol.*, **61**, pp. 445-453.

JENKINS, J. P. R., BARGANZA, J. M., HICKEY, D. S., ISHERWOOD, I. and MACHIN, M. (1987): 'Quantitative tissue characterisation in pancreatic disease using magnetic resonance imaging', *Br. J. Radiol.*, **60**, pp. 333-341.

JOHNS, P. C. and YAFFE, M. J. (1987): 'X-ray characterisation of normal and neoplastic breast tissue', *Phys. Med. Biol.*, **32**, pp. 675-695.

JONES, O. (1989): 'Introduction to the X-windows system' (Prentice Hall, New Jersey).

JOYNT, L., BOYLE, D., RAKOWSKI, H., POPP, R. and BEAVER, W. (1979): 'Identification of tissue parameters by digital processing of real-time ultrasonic clinical cardiac data', *In* LINZER, M. (Ed.): 'Ultrasonic Tissue Characterisation II', National Bureau of Standards (U.S. Government Printing Office, Washington, D.C.), **525**, pp. 267-273.

JOYNT, L., MARTIN, R., MACOVSKI, M. D. and MACOVSKI, A. (1980): 'Techniques for in vivo tissue charaterization', *Acoust. Imaging*, **8**, pp. 527-538.

JULESZ, B. (1975): 'Experiments in the visual perception of Texture', *Scient. Am.*, **232**, pp. 34-43.

JULESZ, B. (1981): 'Textons, the elements of texture perception, and their interactions', *Nature*, **290**, pp. 91-97.

JUST, M., HIGER, H. P., SCHWARZ, M., BOHL, J., FRIES, G., PFANNENSTIEL, P. and THELEN, M. (1988a): 'Tissue characterization of benign brain tumours: use of NMR-

tissue parameters', *Magn. Reson. Imaging*, **6**, pp. 463-472.

JUST, M. and THELEN, M. (1988b): 'Tissue characterization with T1, T2, and proton density values: results in 160 patients with brain tumors', *Radiology*, **169**, pp. 779-785.

KAIZER, H. (1955): 'A quantification of textures on aerial photographs', Boston Univ. Res. Lab., Boston Mass., Tech. Note 121.

KATSURAGAWA, S., DOI, K., MACMAHON, H., NAKAMORI, N., SASAKI, Y. and FENNESSY, J. J. (1990): 'Quantitative computer aided analysis of lung texture in chest radiographs', *Radiographics*, **10**, pp. 257-269.

KELLER, J. M., CROWNOVER, R. M. and CHEN, R. Y. (1987): 'Characteristics of natural scenes related to the fractal dimension', *IEEE Trans. Pattern. Anal. Machine Intel.*, **9**, pp. 621-627.

KERNIGHAN, B. W. and RICHIE, D. M. (1988): 'The C programming language', Second edition (Prentice Hall, New Jersey).

KITNEY, R. I. and DEWEY, C. F. (1993): 'New strategies for medical imaging and networking' *In* Proceedings of Computers in Cardiology 1993 (IEEE Computer Society Press, California), pp. 823-826.

KJAER, L., THOMSEN, C., GJERRIS, F., MOSDAL, B. and HENRIKSEN, O. (1991): 'Tissue characterization of intracranial tumors by MR imaging: In vivo evaluation of T1- and T2-relaxation behaviour at 1.5 T', *Acta Radiol.*, **32**, pp. 498-504.

KOUTCHER, J. A., GOLDSMITH, M. and DAMADIAN, R. (1978): 'NMR in cancer. X. A malignancy index to discriminate normal and cancerous tissue', *Cancer*, **41**, pp. 174-182.

KRATZIK, C., SCHUSTER, E., HAINZ, A., KUBER, W. and LUNGLMAYR, G. (1988): 'Texture analysis - a new method for discriminating prostatic carcinoma from prostatic hypertrophy', *Urol. Res.*, **16**, pp. 395-397.

KRUGER, R. P., THOMPSON, W. B. and TURNER, A. F. (1974): 'Compter diagnosis of pneumoconiosis', *IEEE Trans. Syst. Man. Cybern.*, **4**, pp. 40-49.

KUBE, P. and PENTLAND, A. (1988): 'On the imaging of fractal surfaces', *IEEE Trans. Pattern. Anal. Machine Intel.*, **10**, pp. 704-707.

KUC, R. (1980): 'Clinical application of an ultrasound attenuation coefficient estimation technique for liver pathology characterization', *IEEE Trans. Biomed. Eng.*, **27**, pp. 312-319.

LARSEN, H. G. (1976): 'An algorithm to compute the sequency ordered Walsh transform', *IEEE Trans. Acoust., Speech, Signal Processing*, **24**, pp. 334-335.

LARSEN, H. G. and LAI, D. C. (1980): 'Walsh spectral estimates with applications to the classification of EEG signals', *IEEE Trans. Biomed. Eng.*, **27**, pp. 485-492.

LERSKI, R. A., BARNETT, E., MORLEY, P., MILLS, P. R., WATKINSON, G. and MACSWEEN, R. N. (1979): 'Computer analysis of ultrasonic signals in diffuse liver disease',

Ultrasound Med. Biol., **5**, pp. 341-350.

LERSKI, R. A., SMITH, M. J., MORLEY, P., BARNETT, E., MILLS, P. R., WATKINSON, G. and MACSWEEN, R. N. (1981a): 'Discriminant analysis of ultrasonic texture data in diffuse alcoholic liver disease. 1. Fatty liver and cirrhosis', *Ultrason. Imaging*, **3**, pp. 164-172.

LERSKI, R. A., MORLEY, P., MILLS, P. R. and WATKINSON, G. (1981b): 'Texture analysis of ultrasonic signals - a comparison of the use of radiofrequency and demodulated A-scan data', *Ultrason. Imaging*, **3**, pp. 369-377.

LERSKI, R. A. (1982): 'Ultrasonic tissue charcterisation', *Diagn. Imaging*, **51**, pp. 238-248.

LERSKI, R. A. and STRAUGHAN, K. (1987): 'Identification and characterisation of biological tissue by NMR', *Eurospin Quarterly*, **12**, pp. 77-89.

LERSKI, R. A., MCROBBIE, D. W., STRAUGHAN, K., WALKER, P. M., DE CERTAINES, J. D. and BERNARD, A. M. (1988): 'Multi-center trial with protocols and prototype test objects for the assessment of MRI equipment', *Magn. Reson. Imaging*, **6**, pp. 201-214.

LERSKI, R. A. and DE CERTAINES, J. D. (1993a): 'Peformance assessment and quality control in MRI by Europsin test objects and protocols', *Magn. Reson. Imaging*, **11**, pp. 817-833.

LERSKI, R. A., STRAUGHAN, K., SCHAD, L. R., BOYCE, D., BLUML, S. and ZUNA, I. (1993b): 'MR image texture analysis - an approach to tissue characterisation', *Magn. Reson. Imaging*, **11**, pp. 873-887.

LEVITTT, R. G., SAGEL, S. S., STANLEY, R. J. and JOST, R. G. (1977): 'Accuracy of computed tomography and biliary tract', *Radiology*, **124**, pp. 123-128.

LIGIER, Y., RATIB, O., FUNK, M., PERRIER, R., GIRARD, C. and LOGEAN, M. (1992): 'Portable image manipulation software: what is the extra development cost?', *J. Digit. Imaging*, **5**, pp. 176-184.

LINNETT, L. M. and CLARKE, S. J. (1990): 'The use of fractal modelling in texture analysis', *IEE Colloquim on the Application of Fractal Techniques in Image Processing*, London, pp. 7.1-7.5.

LIPINSKI, J. K. and CREMIN, B. J. (1986): 'Ultrasound and computed tomography of the infant brain: a clinical comparison', *Clin. Radiol.*, **37**, pp. 365-369.

LIPKIN, B. S. and ROSENFELD, A. (1970): 'Picture processing and psychopictorics' (Academic Press, New York).

LIU, S. H. (1992): 'Formation and anomalous properties of fractals', *IEEE Eng. Med. Biol.*, June, pp. 28-39.

MAGNIN, I. E., CLUZEAU, F., ODET, C. L. and BREMOND, A. (1986): 'Mammographic texture analysis: an evaluation of risk for developing breast cancer', *Optical*

Eng., **25**, pp. 780-783.

MANDELBROT, B. B. (1977): 'Fractals : form, chance and dimension' (W.H.Freeman, San Francisco).

MANDELBROT, B. B. (1982): 'The fractal geometry of nature' (W.H.Freeman, San Francisco).

MATHUR-DE VRE, R., BINET, J., BOVEE, W. M. M. J. and FOSTER, M. A. (1988): 'Multi-center trial with an in vitro NMR protocol', *Magn. Reson. Imaging*, **6**, pp. 185-194.

McCRONE, J. (1993): 'Computer chaos at medicine's cutting edge', *New Scientist*, **1892**, pp. 25-29.

McCULLOUGH, E. C. and MORIN, R. L. (1983): 'CT-Number variability in thoracic geometry', *Am. J. Roentgenol.*, **141**, pp. 135-140.

MEDIONI, G. G. and YASUMOTO, Y. (1984): 'A note on using the fractal dimension for segmentation', *Computer Vision: Representation and control, Proceedings of 2nd IEEE workshop*, **257**, pp. 25-30.

MELTZER, B., SEARLE, N. H. and BROWN, R. (1967): 'Numerical specification of biological form', *Nature*, **216**, pp. 32-36.

MELTZER, B. (1968a): 'Speculations on the use of orthonormal functions in the study of morphogenesis', *Nature*, **217**, pp. 196.

MELTZER, B. and SEARLE, N. H. (1968b): 'Impotence principle in description of morphology', *Nature*, **217**, pp. 1289-1290.

MITCHELL, O. R., MYERS, C. R. and BUYNE, W. (1977): 'A max-min measure for image texture analysis', *IEEE Trans. Comput.*, **26**, pp. 408-414.

MORRIS, D. T. (1988): 'An evaluation of the use of texture measurements for the tissue characterisation of ultrasonic images of *in vivo* human placenta', *Ultrasound Med. Biol.*, **14**, pp. 387-395.

NICHOLAS, D., NASSIRI, D. K., GARBUTT, P. and HILL, C. R. (1986): 'Tissue characterisation from ultrasonic B-scan data', *Ultrasound Med. Biol.*, **12**, pp. 135-143.

NYE, A. and O'REILLY, T. (1988): 'The definitive guides to the X window system. Volume four: X toolkit intrinsics programming manual' (O'Reilly & Associates, Inc., Sebastopol CA).

PELEG, S., NAOR, J., HARTLEY, R. and AVNIR, D. (1984): 'Multiple resolution texture analysis and classification', *IEEE Trans. Pattern. Anal. Machine Intel.*, **6**, pp. 518-523.

PENTLAND, A. (1984): 'Fractal based descriptions of natural scenes', *IEEE Trans. Pattern. Anal. Machine Intel.*, **6**, pp. 661-674.

PODO, F. (1988): 'Tissue characterization by MRI: a multidisciplinary and multi-centre challenge today [editorial]', *Magn. Reson. Imaging*, **6**, pp. 185-194.

PODO, F., ORR, J. S., BOVEE, W. M. M. J., DE CERTAINES, J. D. and LEIBFRITZ, D. (1993): 'Tissue characterization by magnetic resonance spectroscopy and imaging: Results of a concerted research project of the European Economic Community. Introduction, objectives, and activities', *Magn. Reson. Imaging*, **11**, pp. 809-815.

POWELL, G. F., DOI, K. and KATSURAGAWA, S. (1988): 'Localization of inter-rib spaces for lung texture analysis and computer aided diagnosis in digital chest images', *Med. Phys.*, **15**, pp. 581-587.

PRATT, W. K., CHEN, W. H. and WELCH, L. R. (1974): 'Slant transform image coding' *IEEE Trans. Commun.*, **22**, pp. 1075-1093.

PRATT, W. K. (1978): 'Digital image processing' (Wiley-Interscience, New York), pp. 232-227.

PRENDERGAST, D. J., HICKEY, D. S., JENKINS, J. P. and ISHERWOOD, I. (1987): 'Increased potential for tissue discrimination in quantitative magnetic resonance imaging [letter]', *Br. J. Radiol.*, **60**, pp. 1142-1143.

PRESS, W. H., FLANNERY, B. P., TEUKOLSKY, S. A. and VETTERLING, W. T. (1991): 'Numerical recipes in C, the art of scientific computing' (Cambridge University Press, Cambridge).

PRETEUX, F., LAVAL-JEANTET, A. M., ROGER, B. and LAVAL-JEANTET, M. H. (1985): 'New prospects in CT image processing via mathematical morphology', *Eur. J. Radiol.*, **5**, pp. 313-317.

PRICE, R. R., JONES, T. B., GODDARD, J. and JAMES, A. E. JR. (1980): 'Basic concepts of ultrasonic tissue characterisation', *Radiol. Clin. North. Am.*, **18**, pp. 21-30.

PULLAN, B. R., FAWCITT, R. A. and ISHERWOOD, I. (1978): 'Tissue characterisation by an analysis of the distribution of attenuation values in computed tomography scans: A preliminary report', *J. Comput. Assist. Tomogr.*, **2**, pp. 49-54.

RAETH, U., SCHLAPS, D., LIMBERG, B., ZUNA, I., LORENZ, A., VAN KAICK, G., LORENZ, W. J. and KOMMERELL, B. (1985): 'Diagnostic accuracy of computerized B-scan texture analysis and conventional ultrasonography in diffuse parenchymal and malignant Disease', *J Clin. Ultrasound*, **13**, pp. 87-99.

RAITHEL, H. J., FUCHS, T., DREXLER, J. and ALEXANDER, J. (1988): 'Remarks on the evaluation of computerised tomograms in fibrotic lung diseases with the aid of histograms - a pilot study', *electromedica*, **56**, pp. 87-96.

RITCHINGS, R. T., PULLAN, B. R., LUCAS, S. B., FAWCITT, R. A., BEST, J. J. K., ISHERWOOD, I. and MORRIS, A. I. (1979): 'An analysis of the spatial distribution of attenuation values in computed tomographic scans of the liver and spleen', *J. Comput. Assist. Tomogr.*, **3**, pp. 36-39.

SCHAD, L. R., BLUML, S. and ZUNA, I. (1993): 'MR Tissue characterization of intercranial tumors by means of texture analysis', *Magn. Reson. Imaging*, **11**, pp. 889-896.

SCHALKOFF, R. J. (1989): 'Digital image processing and computer vision' (Wiley-Interscience, New York).

SCHEPERS, H. E., VAN BEEK, J. H. G. M. and BASSINGTHWAIGHTE, J. B. (1992): 'Four methods to estimate the fractal dimension from self-affine signals', *IEEE Eng. Med. Biol.*, June, pp. 57-64.

SCHOLZ, T. D., CECKLER, T. L. and BALABAN, R. S. (1993): 'Magnetization transfer characterization of hypertensive cardiomyopathy: significance of tissue water content', *Mag. Reson. Med.*, **29**, pp. 352-357.

SCHUSTER, E., KNOFLACH, P. and GRABNER, G. (1988): 'Local texture analysis: an approach to differentiating liver tissue objectively', *J. Clin. Ultrasound*, **16**, pp. 453-461.

SIEGELMAN, S. S., ZERHOUNI, E. A., LEO, F. P., KHOURI, N. F., STITIK, F. P. (1980): 'CT of the solitary pulmonary nodule', *Am. J. Roentgenol.*, **135**, pp. 1-30.

SMITH, F. W., BAYLISS, A. P., HUSSEY, J. K., ROBERTSON, E. M., WEIR, J. and CROSHER, G. A. (1989): 'Low-field (0.08T) magnetic resonance imaging of the pancreas: comparison with computed tomography and ultrasound', *Br. J. Radiol.*, **62**, pp. 796-802.

SOMMER, F. G., JOYNT, L. F., CARROLL, B. A. and MACOVSKI, A. (1981): 'Ultrasonic characterization of abdominal tissues via digital analysis of backscattered waveforms', *Radiology*, **141**, pp. 811-817.

SOMMER, F. G., JOYNT, L. F., HAYES, D. L. and MACOVSKI, A. (1982): 'Stochastic frequency-domain tissue characterisation: application to human spleens *in vivo*', *Ultrasonics*, **20**, pp. 82-86.

STEPHENS, D. H., SHEEDY, P. F., HATTERY, R. R. and MACCARTY, R. L. (1977): 'Computed tomography of the liver', *Am. J. Roentgenol.*, **128**, pp. 579-590.

SUTTON, R. N. and HALL, E. L. (1972): 'Texture measures for automatic classification of pulmonary disease', *IEEE Trans. Comput.*, **21**, pp. 667-676.

THERRIEN, C. W. (1987): 'Decision estimation and classification' (John Wiley & Sons, New York).

TAYLOR, D. G. and BUSHELL, M. C. (1986): 'NMR characterisation of healthy rat tissue *in vivo*', *Clin. Phys. Physiol. Meas.*, **7**, pp. 57-62.

UNSER, M. (1986): 'Sum and difference histograms for texture classification', *IEEE Trans. Pattern. Anal. Machine Intel.*, **8**, pp. 118-125.

WAGNER, R. F., INSANA, M. F. and BROWN, D. G. (1985): 'Progress in signal and texture discrimination in medical imaging', *SPIE. Proc.*, **535**, pp. 57-64.

WAGNER, R. F., INSANA, M. F. and BROWN, D. G. (1986): 'Unified approach to the detection and classification of speckle texture in diagnostic ultrasound', *Optical Eng.*, **25**, pp. 738-742.

WALACH, E., LIU, C. N., WAAG, R. C. and PARKER, K. J. (1986): 'Quantitative tissue characterisation based on pulsed echo ultrasound scans', *IEEE Trans. Biomed. Eng.*, **33**, pp. 637-643.

WATMOUGH, D. J., MORAN, C. and WATMOUGH, J. A. (1988): 'Son et limiere: a new combined optical and Doppler ultrasound approach to the detection of breast cancer', *J. Biomed. Eng.*, **10**, pp. 119-123.

WELLS, P. N. T. (1990): 'Choices in medicine: illustrations from imaging', *J. Med. Eng. Technol.*, **14**, pp. 225-232.

WESZKA, J. S., DYER, C. R. and ROSENFELD, A. (1976): 'A comparative study of texture measures for terrain classification', *IEEE Trans. Syst. Man. Cybern.*, **6**, pp. 269-285.

WIMMER, M. S., BLAINE, G. J., JOST, R. G., MOORE, S. M. and STUDT, J. W. (1991): 'A distributed approach to integrated inquiry and display for radiology', *J. Med. Syst.*, **15**, pp. 299-309.

WOLFF, S. D. and BALABAN, R. S. (1989): 'Magnetisation transfer contrast (MTC) and tissue water proton relaxation *in vivo*', *Mag. Reson. Med.*, **10**, pp. 134-144.

WOLFMAN, N. T., BOEHNE, J. M., CHOPLIN, R. H. and BECHTOLD, R. E. (1992): 'Evaluation of PACS in ultrasonography', *J. Ultrasound Med.*, **11**, pp. 217-223.

YOUNG, P. (1984): 'Recursive estimation and time-series analysis' (Springer-Verlag, Berlin), pp. 9-54.

APPENDIX 1 - AN EXAMPLE OF MENU FUNCTION

A1.1 INTRODUCTION

The menu structure developed for use with the *MAIVIS* system is based on the concept of a linked list of data structures, more commonly seen in the Pascal programming language. This appendix provides a working example of the menu structure to help describe its function.

Figure A1.1 shows the simple menu tree that has been used to illustrate this example. Each node of the tree has a screen window associated with it. The Parent of the window associated with the node at the top of the tree (*MENU*) is the *Main Menu and Selection Window* (*MMSW*) screen window, **figure 3.4, chapter 3**. When *MENU* is offered as an option, its window behaves like an icon. When this icon is selected, the window expands allowing the four windows corresponding to the nodes on the next level of the tree to be mapped onto it as icons. This process continues until the node at the end of a branch of the tree is reached and a program function is called. A more complicated menu tree would have a greater number of levels and a greater number of nodes on each level. The distribution of tree branches would also be less regular.

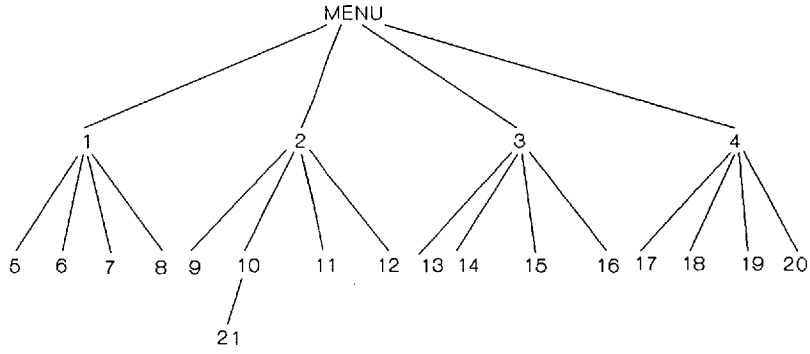


Figure A1.1 : A simple menu tree

In the software implementation of this idea, each node of the tree node is represented by a data structure called *WinMenu*[*i*], where *i* is the node index. Details of this data structure are given in **section 3.5.3**. In this example the node indices have been allocated along each level of the tree, as in **figure A1.1**.

A1.2 MOVING DOWN THE MENU

The top node of the tree is called *MENU*. The data structure for this node (index i=0) is defined in **figure A1.2**. The *No_children* element indicates this node has four children. The node indices of the children are stored in the elements *Children[j]*. Because the *MENU* node sits at the top of the menu tree, the *Parent_No* element has a node index of '-1' which corresponds to the *MMSW* screen window.

```
WinMenu[0].No_children          = 4
WinMenu[0].Parent_No            = -1
WinMenu[0].Children[0]          = 1
WinMenu[0].Children[1]          = 2
WinMenu[0].Children[2]          = 3
WinMenu[0].Children[3]          = 4
WinMenu[0].Screen_Menu_no       = 1
WinMenu[0].MenuName             = "Menu"
```

Figure A1.2 : *MENU* node data structure

When *MAIVIS* is started up, the screen windows are created and mapped onto the display, and the menu is activated. This starts with the window corresponding to the top node of the menu tree being mapped onto the *MMSW* screen window as an icon. The top node for this example being the node *MENU*. Menu options can then be selected by clicking the mouse pointer over the appropriate icon. The following section describes the workings of the menu structure as the user moves down the tree to select a program function at a termination node. All the termination nodes of a tree correspond to program functions (eg. indices i=5-9 & i=11-21 in **figure A1.1**).

If the *MENU* option icon is selected, its window is made the same size as its parent to obscure the parent window from view. This process is illustrated in **figure A1.3**. On the next level of the example menu tree there are four nodes. Four windows corresponding to these nodes are created with the *MENU* option window as their parent. These new menu options occupy the data structures *WinMenu*[i], where i = 1, 2, 3, 4. The new windows are mapped onto the re-sized *MENU* option window as option icons. The icons are labelled *Option1*, *Option2*, *Option3*, and *Option4*, from the *MenuName* element of their data structures, as shown in **figure A1.4**.

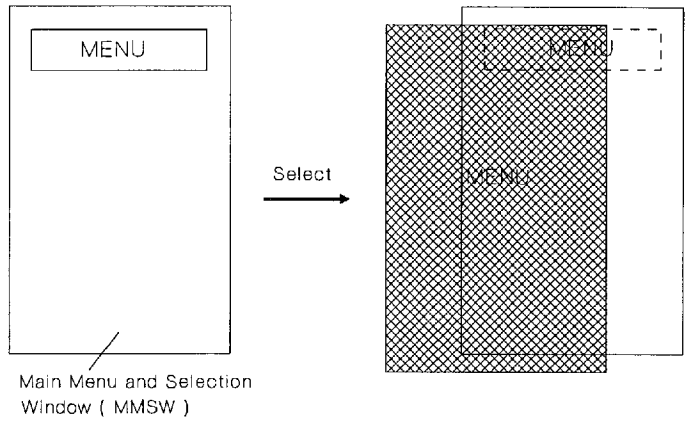


Figure A1.3 : *MENU* option selected and re-sized

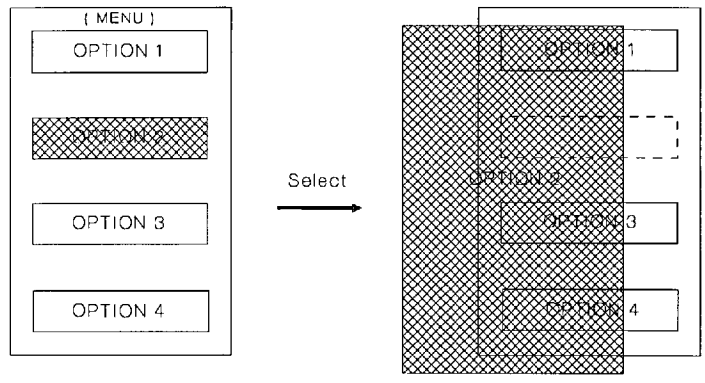


Figure A1.4 : *Option2* selected and re-sized

Figure A1.5 shows the data structure of the second node on the next level down from the top node *MENU*, node index i=2. This data structure gives the node indices of its parent (i=0), and its four children (i=9,10,11,12). The title of the menu option, *Option2*, is also given.

```
WinMenu[2].No_children          = 4
WinMenu[2].Parent_No            = 0
WinMenu[2].Children[0]          = 9
WinMenu[2].Children[1]          = 10
WinMenu[2].Children[2]          = 11
WinMenu[2].Children[3]          = 12
WinMenu[2].Screen_Menu_no       = 2
WinMenu[2].MenuName             = "Option2"
```

Figure A1.5 : *Option2* node data structure

If *Option2* is selected from the four available menu options, then its corresponding icon window is enlarged to obscure its parent window as shown in **figure A1.4**. Following the same procedure as before, windows are created for the nodes on the next level down the menu tree as its children. These windows are then mapped as the option icons *Option9*, *Option10*, *Option11* and *Option12*, as in **figure A1.7**.

If *Option10* (node index i=10) is selected from the four available menu options, then its corresponding icon window is enlarged to obscure its parent window as shown in **figure A1.7**. The menu tree example in **figure A1.1** and the data structure in **figure A1.6** both show that the node *Option10* only has one child, the node with index i=21. A single window is created and mapped as an option icon for this node with the title *Option21*, as in **figure A1.8.** Because this node is at the end of a branch of the menu tree, it is called a termination node. This indicates that if this option is selected then a program function will be called.

```
WinMenu[10].No_children         = 1
WinMenu[10].Parent_No           = 2
WinMenu[10].Children[0]         = 21
WinMenu[10].Screen_Menu_no      = 2
WinMenu[10].MenuName            = "Option10"
```

Figure A1.6 : *Option10* node data structure

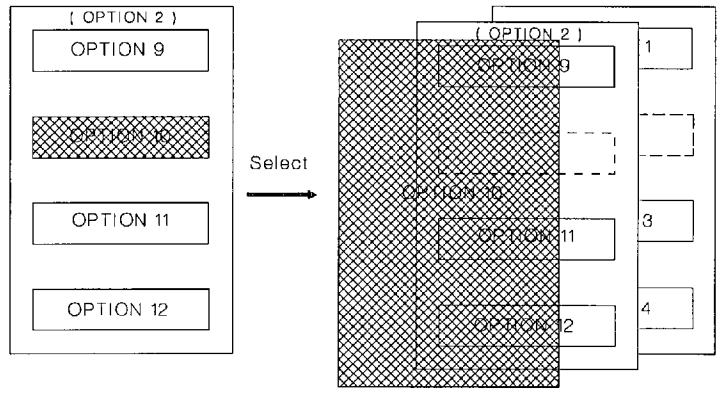


Figure A1.7 : *Option10* selected and re-sized

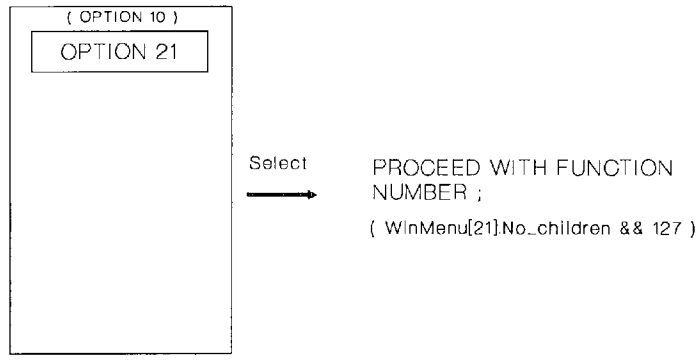


Figure A1.8 : *Option21* selected and function called

The data structure for the termination node *Option21* is given in **figure A1.9**. The 8th bit of the structure element *No_children* is set to indicate the node represents a function call. The identity of this function is held in the 7 least significant bits of this element. The C-language implementation of the function selection process is shown in **figure A1.10**. It can be seen that if *Option21* was selected, the program would call *function_1*.

```
WinMenu[21].No_children        =  128 + 1
WinMenu[21].Parent_No          =  10
WinMenu[21].Screen_Menu_no     =  1
WinMenu[21].MenuName           =  "Option21"
```

Figure A1.9 : *Option21* node data structure

```
int  choice = ( WinMenu[21].No_children && 127 );

        switch (choice)
        {
                case(1) :  function_1();  break;
                case(2) :  function_2();  break;
                .
                case(n) :  function_n();  break;
        }
```

Figure A1.10 : C-language for function selection

A1.3 MOVING UP THE MENU

Although not shown on **figure A1.1**, there are special nodes at each level of the tree called *EXIT* nodes. The only information contained in the data structures of these nodes are the node indices of their parents. These nodes are identifiable by a zero *No_Children* element.

The consequence of choosing and *EXIT* option is that its icon window and those of its siblings are destroyed. The parent window is then re-sized and re-positioned as an option icon, and its siblings become visible and active again. The titles of the options, lost when the windows were obscured, are then re-displayed for the move up one level of the menu to be complete.

APPENDIX 2 - RUN LENGTH TEXTURE MEASURES

The following five parameters are texture descriptors based upon the run length joint probability density function $P(i,j)$ for runs of length j and grey level i. The *Short Runs Emphasis*, *Long Runs Emphasis*, *Grey Level Non-Uniformity*, *Run Length Non-Uniformity*, and *Run Percentage* texture measures are defined in **equations A2.1** to **A2.5** respectively, where N_g is the number of grey levels, N_r is the max run length, and N is the image size.

These five parameters have also been used with the cluster size joint probability density function $P(i,j)$ for clusters of size j and grey level i.

$$SRE = \frac{\sum_{i=1}^{N_g}\sum_{j=1}^{N_r}\frac{p(i,j)}{j^2}}{\sum_{i=1}^{N_g}\sum_{j=1}^{N_r}p(i,j)} \tag{A2.1}$$

$$LRE = \frac{\sum_{i=1}^{N_g}\sum_{j=1}^{N_r}j^2 p(i,j)}{\sum_{i=1}^{N_g}\sum_{j=1}^{N_r}p(i,j)} \tag{A2.2}$$

$$GLN = \frac{\sum_{i=1}^{N_g}\left(\sum_{j=1}^{N_r}p(i,j)\right)^2}{\sum_{i=1}^{N_g}\sum_{j=1}^{N_r}p(i,j)} \tag{A2.3}$$

$$RLN = \frac{\sum_{j=1}^{N_r}\left(\sum_{i=1}^{N_g}p(i,j)\right)^2}{\sum_{i=1}^{N_g}\sum_{j=1}^{N_r}p(i,j)} \tag{A2.4}$$

$$RPC = \frac{\sum_{i=1}^{N_g}\sum_{j=1}^{N_r}p(i,j)}{N*N} \tag{A2.5}$$

APPENDIX 3 - CLUSTER IDENTIFICATION FLOWCHART

Figure A3.1 shows the flowchart for the cluster identification process.

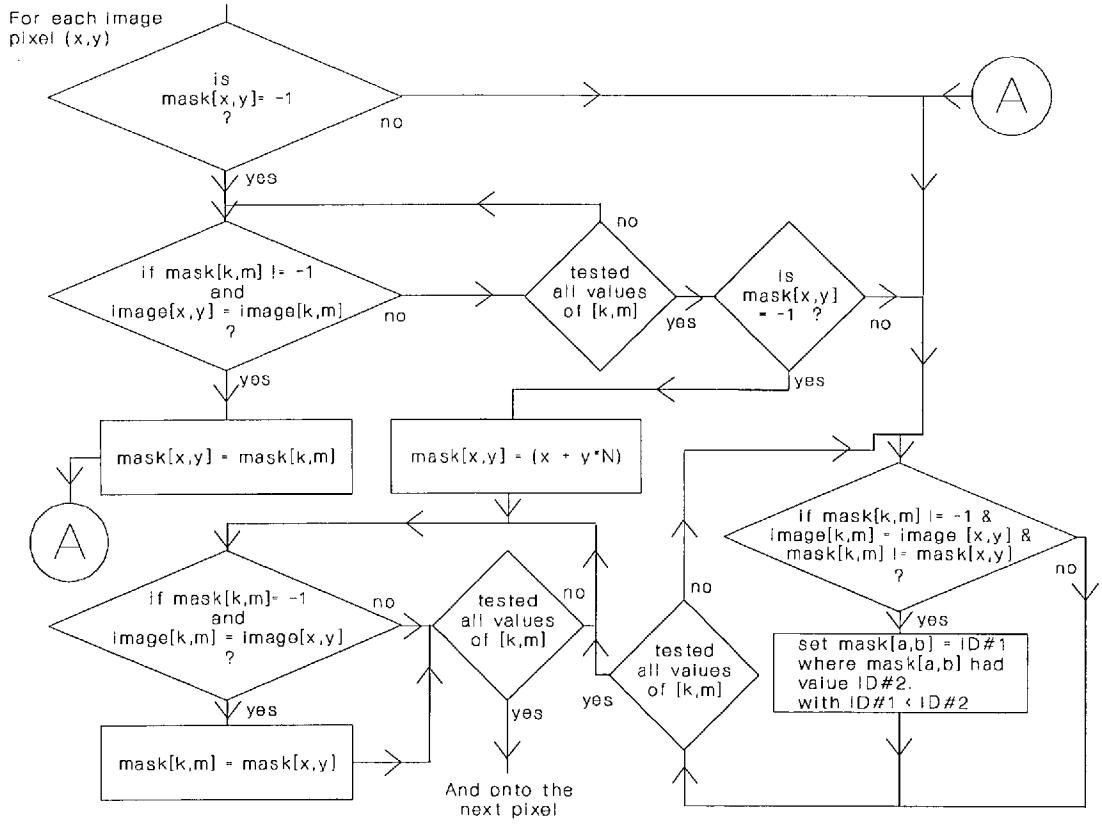


Figure A3.1 : Cluster identification

APPENDIX 4 - CO-OCCURRENCE TEXTURE FEATURES

A4.1 INTRODUCTION

The set of fourteen textures features *F1-F14* given in this appendix can be extracted from the joint probability density function of the grey levels of co-occurrent image pixels $p(i,j)$. The spatial dependency between each pair of pixels is defined by the vector (d,θ). The number of distinct grey levels is given by N_g.

Equations A4.1 to **A4.4** show a set of first order probability density functions derived from the co-occurence probability density function. These functions are used in the definition of several of the texture features.

$$p_x(i) = \sum_{j=1}^{N_g} p(i,j) \qquad \text{(A4.1)}$$

$$p_y(j) = \sum_{i=1}^{N_g} p(i,j) \qquad \text{(A4.2)}$$

$$p_{(x+y)}(k) = \sum_{i=1}^{N_g} \sum_{\substack{j=1 \\ i+j=k}}^{N_g} p(i,j) \qquad \text{(A4.3)}$$

$k = 2,3, \ldots ,2N_g$

$$p_{(x-y)}(k) = \sum_{i=1}^{N_g} \sum_{\substack{j=1 \\ i-j=k}}^{N_g} p(i,j) \qquad \text{(A4.4)}$$

$k = 0,1, \ldots ,Ng\text{-}1$

A4.2 TEXTURAL FEATURES

F1 Angular Second Moment:

$$F1 = \sum_i \sum_j \{p(i,j)\}^2 \qquad \text{(A4.5)}$$

F2 Contrast:

$$F2 = \sum_{n=0}^{N_g-1} n^2 \left\{ \sum_{i=1}^{N_g} \sum_{\substack{j=1 \\ |i-j|=n}}^{N_g} p(i,j) \right\} \tag{A4.6}$$

F3 Correlation:

$$F3 = \frac{\sum_i \sum_j (i-\mu_x)(j-\mu_y) p(i,j)}{\sigma_x \sigma_y} \tag{A4.7}$$

where μ_x, μ_y, σ_x, and σ_y are the means and standard deviations of the pdfs $p_x(i)$ and $p_y(j)$.

F4 Sum of Squares - Variance:

$$F4 = \sum_i \sum_j (i-\mu)^2 p(i,j) \tag{A4.8}$$

F5 Inverse Difference Moment:

$$F5 = \sum_i \sum_j \frac{1}{1+(i-j)^2} p(i,j) \tag{A4.9}$$

F6 Sum Average:

$$F6 = \sum_{i=2}^{2N_g} i p_{x+y}(i) \tag{A4.10}$$

F6' Difference Average:

$$F6' = \sum_{i=2}^{N_g-1} i p_{x-y}(i) \tag{A4.11}$$

F7 Sum Variance:

$$F7 = \sum_{i=2}^{2N_g} (i - F6)^2 p_{x+y}(i) \tag{A4.12}$$

F8 Sum Entropy:

$$F8 = -\sum_{i=2}^{2N_g} p_{x+y}(i) \log\{ p_{x+y}(i) \} \tag{A4.13}$$

F9 Entropy:

$$F9 = -\sum_i \sum_j p(i,j) \log\{ p(i,j) \} \tag{A4.14}$$

F10 Difference Variance:

$$F10 = variance of\ \ p_{x-y} \tag{A4.15}$$

F11 Difference Entropy:

$$F11 = -\sum_{i=0}^{N_g-1} p_{x-y}(i) \log\{ p_{x-y}(i) \} \tag{A4.16}$$

F12 Information Measure of Correlation:

$$F12 = \frac{HXY - HXY1}{\max\{HX, HY\}} \tag{A4.17}$$

F13 Information Measure of Correlation:

$$F13 = (1 - \exp[-2.0\ (HXY2 - HXY)])^{\frac{1}{2}} \tag{A4.18}$$

where *HX* and *HY* are the entropies of the pdfs $p_x(i)$ and $p_y(j)$. The other entropy measures *HXY*, *HXY1*, and *HXY2* are defined in **equations A4.19** to **A4.21**.

$$HXY = -\sum_i \sum_j p(i,j) \log\{ p(i,j) \} \tag{A4.19}$$

$$HXY1 = -\sum_i \sum_j p(i,j) \log\{ p_x(i) p_y(j) \} \tag{A4.20}$$

$$HXY2 = -\sum_i \sum_j p_x(i) p_y(j) \log\{p_x(i) p_y(j)\} \tag{A4.21}$$

F14 Maximal Correlation Coefficient:

$$F14 = (Second largest\ eigenvalue of Q)^{1/2} \tag{A4.22}$$

where

$$Q(i,j) = \sum_k \frac{p(i,k) p(j,k)}{p_x(i) p_y(k)} \tag{A4.23}$$