OLLSCOIL NA GAILLIMHE
UNIVERSITY OF GALWAY

Provided by the author(s) and University of Galway in accordance with publisher policies. Please cite the published version when available.

Title	Speckle Reduction and Edge Detection in Ultrasound Imagery
Author(s)	Finn, Sean
Publication Date	2010-09-01
Item record	http://hdl.handle.net/10379/2726

Downloaded 2024-03-13T07:44:16Z

Some rights reserved. For more information, please see the item record link above.

Speckle Reduction and Edge Detection in Ultrasound Imagery

A thesis presented
by

Seán Finn

to
The College of Engineering and Informatics
in fulfillment of the requirements
for the degree of
Doctor of Philosophy
in the subject of

Electronic Engineering

National University of Ireland Galway
Galway, Ireland
September 2010

Research Director
Prof. W. G. Hurley

Supervisors
Dr. E. Jones, Dr. M. Glavin

Contents

Title Page i
Table of Contents ii
Declaration vi
Abstract vii
Acknowledgments viii
List of Figures ix
List of Tables xii
List of Abbreviations xiv

1 Introduction **1**
- 1.1 Motivation 1
- 1.2 Objectives and Contributions 2
  - 1.2.1 Objectives 2
  - 1.2.2 Contributions 3
  - 1.2.3 Publications 4
- 1.3 Chapter Summary 5

2 Background **7**
- 2.1 Ultrasound Image Acquisition Overview 8
  - 2.1.1 Ultrasound Fundamentals 8
  - 2.1.2 Scanning Systems 10
- 2.2 Characteristics of Speckle 12
- 2.3 Literature Review: Speckle Reduction 16
  - 2.3.1 Compounding Approaches 16
  - 2.3.2 Postacquisition Speckle Reduction 17
    - 2.3.2.1 Adaptive Filters 17
    - 2.3.2.2 Diffusion Filtering 19
    - 2.3.2.3 Multiscale Methods 21
  - 2.3.3 Speckle Reduction Reviews 26
- 2.4 In-Depth Description of Speckle Reduction Methods 31
  - 2.4.1 Anisotropic Diffusion Filters 31
    - 2.4.1.1 Perona and Malik Anisotropic Diffusion (PMAD) 31
    - 2.4.1.2 Speckle Reducing Anisotropic Diffusion (SRAD) 33
    - 2.4.1.3 Detail Preserving Anisotropic Diffusion (DPAD) 34

2.4.1.4 Coherence Enhancing Diffusion (CED) 34
2.4.1.5 Nonlinear Coherent Diffusion (NCD) 35
2.4.1.6 Oriented Speckle Reducing Anisotropic Diffusion (OSRAD) 36
2.4.1.7 Discretisation Schemes for Anisotropic Diffusion 37
2.4.2 Multiscale Wavelet Filters 40
2.4.2.1 Speckle Reduction Method of Zong et al. 40
2.4.2.2 The Generalized Likelihood Method (GLM) 41
2.4.2.3 Nonlinear Multiscale Wavelet Diffusion (NMWD) 42
2.4.3 Local Statistics SAR Filters 44
2.4.3.1 The Lee filter 44
2.4.3.2 The Kuan et al. filter 45
2.4.3.3 The Frost et al. Filter 45
2.4.3.4 The Enhanced Lee and Frost Filters (EnhLee, EnhFrost) . 45
2.4.4 Geometric Filter 46
2.5 Concluding Remarks 47

3 Objective and Subjective Quality Assessment in Echocardiography 49
3.1 Introduction 49
3.2 Speckle Reduction Methods 50
3.3 Quality Assessment using Objective Metrics 51
3.3.1 Pratt's Figure of Merit (FoM) 51
3.3.2 Mean Squared Error (MSE) 54
3.3.3 Edge Region MSE 54
3.4 Quality Assessment by Expert Analysis 54
3.5 Results 56
3.6 Conclusions 59

4 Evaluation of Speckle Filtering 63
4.1 Introduction 63
4.2 Ultrasound Simulation 64
4.3 Speckle Filter Evaluation 69
4.3.1 Simulated Ultrasound Images 69
4.3.2 Clinical Ultrasound Images 72
4.3.3 Image Quality Metrics 74
4.3.3.1 Structural Similarity (SSIM) 74
4.3.3.2 Contrast to Noise Ratio (CNR) 74
4.3.3.3 $\mathbf{SNR_A}$ 75
4.3.4 Computational Requirement Analysis 77
4.3.4.1 Anisotropic Diffusion Discretisation 77
4.3.4.2 Various Filter Implementation Considerations 77
4.4 Results 79
4.4.1 Simulated Images 79
4.4.2 Clinical Images 87
4.4.3 Computational Requirements 93
4.4.3.1 Diffusion Discretisation 93

4.4.3.2 Overall Filter Requirements 94
4.5 Conclusions 97

5 Edge Detection and Speckle Reduction using Gabor Zero Crossings 101
5.1 Introduction 101
5.2 Background 104
5.2.1 Gabor Filters 104
5.2.2 Stable Processes 111
5.2.3 Existing Edge Detection Methods 113
5.2.3.1 Coefficient of Variation Edge Detection 113
5.2.3.2 Ratio of Averages Edge Detection 114
5.2.3.3 Phase Congruency Edge Detection 116
5.2.3.4 Feature Asymmetry Edge Detection 117
5.3 Removal of Speckle Noise from Gabor Coefficients 118
5.3.1 Modelling Gabor Coefficients 118
5.3.1.1 SαS Modelling of Signal Gabor Coefficients 118
5.3.1.2 Combined SαS and Gaussian Model for Speckled Gabor Coefficients 125
5.3.2 Parameter Estimation Method 131
5.3.3 MAP Speckle Removal 135
5.4 Edge Detection Using Gabor Zero Crossings 139
5.4.1 Zero Crossing Orientational Spacing 145
5.4.1.1 Quantification 145
5.4.1.2 Edge Discrimination Properties 148
5.4.2 Zero Crossing Interscale Agreement 157
5.4.2.1 Quantification 157
5.4.2.2 Edge Discrimination Properties 161
5.4.3 Zero Crossing Edge Detection 164
5.5 Speckle Removal with ZC Edge Detector and Anisotropic Diffusion 167
5.6 Conclusions 170

6 Performance Evaluation 172
6.1 Introduction 172
6.2 Test Procedures 173
6.2.1 Stable Model Parameter Estimation 173
6.2.2 Edge Detection 174
6.2.2.1 Simulated Images 174
6.2.2.2 Phantom Images 179
6.2.2.3 Clinical Images 183
6.2.2.4 Edge Detector Performance Evaluation 186
6.2.3 Speckle Reduction 187
6.3 Results 190
6.3.1 Model Parameter Estimation 190
6.3.2 ZC Edge Detection Performance 197
6.3.2.1 Simulated Images 197

            6.3.2.2 Phantom Images 206
            6.3.2.3 Clinical Images 210
        6.3.3 Speckle Reduction using ZC Edges 213
        6.3.4 Simulated Images 213
        6.3.5 Clinical Images 217
        6.3.6 Computational Requirements 218
    6.4 Conclusions 222

7 Conclusions and Future Work **223**
    7.1 Summary of Thesis 223
    7.2 Contributions 224
    7.3 Suggestions for Future Work 226

Bibliography **227**

A Additional Data from Chapter 5 **244**
    A.1 Modelling Gabor Coefficients of Speckled Images 244
    A.2 Parameter Estimation: Computed t_{max} Values 255

B Supplementary Results from Chapter 6 **258**
    B.1 Parameter Estimation 258
    B.2 Edge Detection 268

C Publications **289**

I hereby declare that the work contained in this thesis has not been submitted by me pursuance of any other degree.

Signature: ....

Seán Finn

Abstract

This thesis investigates the related topics of speckle reduction and edge detection as applied to speckled ultrasonography. A strategy for the evaluation of speckle filters is presented. This includes their application to a large test set of both clinical and simulated images. Functional performance is measured by application of objective image quality metrics, and the relationship between these objective metrics and the subjective opinion of clinical experts is investigated. Finally a detailed analysis of the complexity of the filtering methods is performed to ascertain the computational requirements of the different filters.

Edge detection is investigated using Gabor filters to perform orientation-specific multiscale decompositions. The effect of speckle is reduced in the Gabor transforms by MAP estimation. This involves modelling the statistics of the speckle Gabor transform as a combined Gaussian and symmetric α-stable distribution. A technique is proposed for the estimation of the parameters of this model, and its accuracy is demonstrated. Edge detection is achieved by finding the zero crossings of the even-symmetric Gabor coefficients.

To suppress the zero crossings which do not correspond to images edges, two observed properties of the Gabor zero crossings are examined. These measurements are combined to form the final edge detector, which is shown to exceed the performance of other edge detectors for speckled imagery.

A close relationship exists between edge detection and speckle removal, especially for the diffusion class of filters. The proposed zero crossing edge detector is accordingly incorporated into a tensor-valued diffusion scheme. The performance of this speckle filter is evaluated using the framework proposed in this thesis, and is found to exhibit excellent functional performance.

Acknowledgments

I would like to thank a number of people for their help in the completion of the research presented in this thesis. Firstly, I am deeply grateful to my family for the extensive support which they provided throughout my education. My parents, brothers and sister have all helped to keep me focused on the work at hand.

My supervisors Dr. Edward Jones and Dr. Martin Glavin deserve a special thanks for their expertiese, patience and diligence. I am grateful both professionally and personally for their insight and advice.

I would like to thank the staff of Electrical and Electronic Engineering, in particular Professor Gearóid Ó'Laighin and Professor W. G. Hurley. A warm thanks is also extended to the technical staff of Aodh Dalton, Martin Burke and Myles Meehan. I would also like to express a particular gratitude to Mary Costello for her assistance in countless ways.

In the Cardiology department of University College Hospital, Galway, I would like to thank Paul Nolan and Dr. Keiran Daly for their assistance in conducting studies, and kindly providing both advice and data.

The assistance and camaraderie of my colleagues in Electrical and Electronic Engineering has been greatly appreciated, especially that of Ronan O'Malley, Dr. Martin O'Halloran, Dr. Ciaran Hughes, Diarmaid Ó'Cualain, Raquel Cruz Conceição, Dallan Byrne, Robert McFeely and Pat Hurney.

Finally a special word of thanks to my girlfriend Sinéad. Her love, support and understanding has brightened my days, and I dedicate this thesis to her.

List of Figures

2.1 Block Diagram of a B-mode ultrasound system. 10
2.2 Example echocardiographic image with speckle. 12
2.3 Geometric filter pixel selection. 46

3.1 Example speckle filtered long axis echocardiographic frame. 52
3.2 Example speckle filtered short axis echocardiographic frame. 53
3.3 Example of edge region selection for the Edge Region MSE metric. 55
3.4 Box-plots of each video assessment category, per expert. 60

4.1 Ultrasound system point spread functions. 65
4.2 Example simulated ultrasound image. 68
4.3 Simulated clinical ultrasound images. 71
4.4 Histograms of homogeneous regions in simulated and clinical image. 72
4.5 Clinical ultrasound images, used in speckle filter evaluation. 73
4.6 Regions selected for CNR metric calculation. 75
4.7 Speckle filter output for simulated image of Fig. 4.3(c). 85
4.8 Speckle filter output for simulated image of Fig. 4.3(d). 86
4.9 Speckle filter output for the clinical image of Fig. 4.5(a). 87
4.10 Speckle filter output for the clinical image of Fig. 4.5(b). 88
4.11 Speckle filter output for the clinical image of Fig. 4.5(c). 89
4.12 Speckle filter output for the clinical image of Fig. 4.5(d). 90

5.1 Edge detection scheme flowchart. 102
5.2 Example log Gabor spectra. 106
5.3 Example log Gabor spectra, of various bandwidths. 106
5.4 Example of a log Gabor filter in the time domain. 107
5.5 Ultrasound images used to demonstrate log Gabor transformation. 108
5.6 Gabor transform of the clinical image of Fig. 5.5(a). 109
5.7 Gabor transform of the simulated image of Fig. 5.5(b). 110
5.8 Example symmetric α stable probability density functions. 112
5.9 Local masks used in the RoA edge detection method. 114
5.10 Example of a Laplacian of Gaussian filter kernel. 115
5.11 Speckle free images to demonstrate SαS fitting. 119
5.12 Histograms of real-valued Gabor coefficients of Fig. 5.11(a). 120

5.13 Histograms of real-valued Gabor coefficients of Fig. 5.11(b). 121
5.14 Probability plots of the real-valued Gabor coefficients from the first scale of image 5.11(a). 125
5.15 Probability plots of the real-valued Gabor coefficients from the first scale of image 5.11(b). 126
5.16 Speckle corrupted images used to demonstrate model accuracy. 128
5.17 Observed and fitted densities of the real-valued Gabor coefficients from the first scale of Fig. 5.16(a). 129
5.18 Observed and fitted densities of the real-valued Gabor coefficients from the first scale of Fig. 5.16(b). 130
5.19 Comparison of the empirical and expected characteristic functions. 133
5.20 Example of Gabor transform MAP speckle removal. 136
5.21 Sample MAP input/output functions for various values of α. 137
5.22 Response of a LoG kernel to a step edge. 139
5.23 Spectra of 1D log Gabor filters with varying bandwidth. 140
5.24 Response of log Gabor filter to step edge. 142
5.25 Speckled images used in demonstration of Gabor zero crossing behaviour. . 142
5.26 Example Gabor zero crossings, from the image of Fig. 5.25(a). 143
5.27 Example Gabor zero crossings, from the image of Fig. 5.25(b). 144
5.28 Measuring zero crossing orientational spacing using rotation. 147
5.29 Orientational spacing of the Gabor zero crossings of Fig. 5.26. 148
5.30 Orientational spacing of the Gabor zero crossings of Fig. 5.27. 149
5.31 Example Short zero crossing removal. 150
5.32 Zero crossings determined to coincide with image edges of Fig. 5.25(a). . . 152
5.33 Zero crossings determined to coincide with image edges of Fig. 5.25(b). . . 153
5.34 Interscale agreement for Gabor zero crossings of Fig. 5.26. 159
5.35 Interscale agreement for Gabor zero crossings of Fig. 5.26. 160
5.36 Examples of the final Gabor zero crossing edge detector. 166

6.1 Examples of simulated ultrasound images of simple shapes. 175
6.2 Ideal edge maps for the simulated shape images of Fig. 6.1. 176
6.3 Examples of simulated clinical ultrasound images. 177
6.4 Ideal edge maps for the simulated clinical ultrasound images of Fig. 6.3. . 178
6.5 Phantom ultrasound images used in evaluating edge detector performance. 181
6.6 Ideal edge maps for the phantom ultrasound images of Fig. 6.5. 182
6.7 Clinical ultrasound images used in evaluating edge detector performance. . 184
6.8 Clinical ultrasound images used in evaluating edge detector performance. . 185
6.9 Edge detector output for the simulated edge images of Fig. 6.1(c). 198
6.10 Edge detector output for the simulated edge images of Fig. 6.3(a). 199
6.11 Edge detector output for the phantom ultrasound images of Figs. 6.5(b). . 207
6.12 Edge detector output for the phantom ultrasound images of Figs. 6.5a(f). 208
6.13 Edge detector output for the clinical ultrasound images of Fig. 6.7(c). . . . 211
6.14 Edge detector output for the clinical ultrasound images of Fig. 6.7(e). . . . 212
6.15 Filtered simulated images, including ZCAD. 214
6.16 Filtered clinical images, including ZCAD. 221

A.1 Histograms of Gabor coefficients of Fig. 5.11(a). 245
A.2 Histograms of Gabor coefficients of Fig. 5.11(b). 246
A.3 Probability plots of the real-valued Gabor coefficients of image 5.11(a). . . 247
A.4 Probability plots of the imaginary-valued Gabor coefficients of image 5.11(a). 248
A.5 Probability plots of the real-valued Gabor coefficients of image 5.11(b). . . 249
A.6 Probability plots of the imaginary-valued Gabor coefficients of image 5.11(b). 250
A.7 Observed and fitted densities of the real-valued Gabor coefficients of Fig. 5.16(a). 251
A.8 Observed and fitted densities of the imaginary-valued Gabor coefficients of Fig. 5.16(a). 252
A.9 Observed and fitted densities of the real-valued Gabor coefficients of Fig. 5.16(b). 253
A.10 Observed and fitted densities of the imaginary-valued Gabor coefficients of Fig. 5.16(b). 254

B.1 Edge detector output for the simulated edge images of Fig. 6.1(a). 269
B.2 Edge detector output for the simulated edge images of Fig. 6.1(b). 270
B.3 Edge detector output for the simulated edge images of Fig. 6.1(d). 271
B.4 Edge detector output for the simulated edge images of Fig. 6.3(b). 272
B.5 Edge detector output for the simulated edge images of Fig. 6.3(c). 273
B.6 Edge detector output for the simulated edge images of Fig. 6.3(d). 274
B.7 Edge detector output for the phantom ultrasound images of Figs. 6.5(a). . 275
B.8 Edge detector output for the phantom ultrasound images of Figs. 6.5(c). . 276
B.9 Edge detector output for the phantom ultrasound images of Figs. 6.5(d). . 277
B.10 Edge detector output for the phantom ultrasound images of Figs. 6.5(e). . 278
B.11 Edge detector output for the clinical ultrasound images of Fig. 6.7(a). . . . 279
B.12 Edge detector output for the clinical ultrasound images of Fig. 6.7(b). . . 280
B.13 Edge detector output for the clinical ultrasound images of Fig. 6.7(d). . . 281
B.14 Edge detector output for the clinical ultrasound images of Fig. 6.7(f). . . . 282
B.15 Edge detector output for the clinical ultrasound images of Fig. 6.8(a). . . . 283
B.16 Edge detector output for the clinical ultrasound images of Fig. 6.8(b). . . 284
B.17 Edge detector output for the clinical ultrasound images of Fig. 6.8(c). . . . 285
B.18 Edge detector output for the clinical ultrasound images of Fig. 6.8(d). . . 286
B.19 Edge detector output for the clinical ultrasound images of Fig. 6.8(e). . . . 287
B.20 Edge detector output for the clinical ultrasound images of Fig. 6.8(f). . . . 288

List of Tables

2.1 Despeckle Filter Summary. 32

3.1 Aggregate Distribution of Expert Scores, N=288. 58
3.2 Inter-expert Kruskal-Wallis Parameter Values. 59
3.3 Intra-expert association between scoring categories. 59
3.4 Intra-expert association between scoring categories and metrics. 59

4.1 Filter Parameters. 80
4.2 SNR_A for the simulated images of Figs. 4.3(c)-(d). 82
4.3 Image quality metrics for filtered simulated images of Figs. 4.3(c)-(d). . . . 83
4.4 Image quality metrics for speckle filtered clinical videos. 91
4.5 Discretisation Scheme Accuracy. 93
4.6 Speckle filter computational requirements. 95
4.7 Filter Computational Requirements, for $N = 512 \times 512$ pixels. 96
4.8 Summary of relative filter performance 98

5.1 Maximum likelihood α parameter estimates. 123
5.2 Maximum likelihood γ parameter estimates. 124
5.3 Correlation values for the probability plots of Gabor coefficients of Fig. 5.11. 127
5.4 r^2 goodness of fit measures for fitted densities for Fig. 5.16. 128
5.5 Variance reduction in simulated image regions due to MAP estimation. . . . 138
5.6 Gabor zero crossing orientational spacing discrimination for Fig. 5.25(a). . 154
5.7 Gabor zero crossing orientational spacing discrimination for Fig. 5.25(b). . 155
5.8 Gabor zero crossing interscale agreement discrimination for Fig. 5.25(a). . . 162
5.9 Gabor zero crossing interscale agreement discrimination for Fig. 5.25(b). . . 163

6.1 α parameter estimation errors, for various values of α, γ and σ. 193
6.2 γ parameter estimation errors, for various values of α, γ and σ. 194
6.3 σ parameter estimation errors, for various values of α, γ and σ. 195
6.4 α and γ estimation errors, reported in [172]. 196
6.5 List of abbreviations used in Table 6.4. 196
6.6 Edge detector FoM values for simulated images of Figs. 6.1(a)-(d). 202
6.7 Average FoM values for simulated images of Figs. 6.1(a)-(d). 203
6.8 Edge detector FoM values for simulated images of Figs. 6.3(a)-(d). 204

6.9 Average FoM values for simulated images of Figs. 6.3(a)-(d). 205
6.10 Edge detector FoM values for phantom images of Fig. 6.5. 209
6.11 SNR_A for the simulated images of Figs. 4.3(c)-(d). 215
6.12 Image quality metrics for simulated images of Figs. 4.3(c)-(d). 216
6.13 Image quality metrics for clinical videos. 217
6.14 Computational requirements for the ZCAD filter. 219
6.15 Filter Computational Requirements, for $N = 512 \times 512$ pixels. 220

A.1 Optimal t_{max} values for the proposed parameter estimation method. 255

B.1 α parameter estimation errors, for various values of α, γ and σ. 259
B.2 γ parameter estimation errors, for various values of α, γ and σ. 262
B.3 σ parameter estimation errors, for various values of α, γ and σ. 265

List of Abbreviations

CED	Coherence Enhancing Diffusion
CNR	Contrast to Noise Ratio
DPAD	Detail Preserving Anisotropic Diffusion
EnhFrost	Enhanced Frost *et al.* Filter
EnhLee	Enhanced Lee Filter
FoM	Pratt's Figure of Merit
Geo	Geometric Filter
GLM	Generalized Likelihood Method
LoG	Laplacian of a Gaussian
MAP	Maximum a Posteriori
MSE	Mean Squared Error
NCD	Nonlinear Coherent Diffusion
NMWD	Nonlinear Multiscale Wavelet Diffusion
OSRAD	Oriented Speckle Reducing Anisotropic Diffusion
PMAD	Perona and Malik Anistropic Diffusion
RoA	Ratio of Averages
SαS	Symmetric α Stable
SAR	Synthetic Aperture Radar
SNR_A	Amplitude Signal to Noise Ratio
SRAD	Speckle Reducing Anisotropic Diffusion
SSIM	Structural Similarity
ZC	Zero Crossing
ZCAD	Zero Crossing Anisotropic Diffusion

Chapter 1

Introduction

1.1 Motivation

The use of ultrasound imaging in medial applications has been investigated since the 1950s. Today ultrasonography is used in a wide range of clinical applications, including vascular imaging, obstetrics and gynaecology, and cardiology. The ultrasound imaging modality has a number of advantages which have led to its widespread usage in medicine. It operates without the use of ionising radiation, and so provides diagnostically important information without clinically significant biological effects on the patient [1]. The real-time nature of an ultrasound examination allows the analysis of fast moving structures such as the heart, while multi-planar imaging ability gives the physician flexibility in selection of an imaging plane.

Ultrasound imagery is affected by an artefact of granular appearance known as speckle. This results from the interference pattern experienced by acoustic waves which are randomly scattered in the imaged medium. A number of researchers have examined the negative consequences of the presence of speckle [2–9]. Both manual analysis and automated processing techniques are hindered by speckle effects, such as the introduction of false detail, the masking of borders of anatomical structure, and an overall reduction in contrast.

A considerable volume of literature is dedicated to methods of speckle removal, which is summarised in Chapter 2 below. Approaches to this problem include locally adaptive filtering techniques (some of which were proposed for speckle removal in radar applications), anisotropic diffusion-based speckle suppression methods, and approaches based on multiscale analysis. Given the extensive coverage of this topic in the literature, which is still being expanded with improved methods from an active body of researchers, it is clear that this is still an active area of interest. As one of the main uses for ultrasound imaging is clinical diagnosis, the principal benefit of improved speckle reduction approaches are more accurate clinical analyses. This includes diagnosis from both manual interpretation by physicians, and also the use of automated methods such as segmentation and automated measurement.

This thesis addresses a number of issues surrounding the development of improved speckle reduction methods. Although many methods for speckle reduction have been published, as have a number of comparative evaluations, a single accepted evaluation strategy has not been established. A framework for evaluating the performance of speckle filtering is presented in this thesis, including an investigation of the relationship between objective image quality metrics and the opinion of clinical experts. The importance of accurate edge detection in speckle filtering methods is highlighted, and an accurate edge detection method is proposed for operation in speckled imagery. This edge detector is then used as the basis for a novel method of speckle removal.

1.2 Objectives and Contributions

1.2.1 Objectives

This section presents the main objectives and novel contributions of this thesis, based on the motivation of the previous section. The primary aim of this thesis is to develop an improved method of speckle removal. In pursuit of this aim, the main objectives of this thesis are:

- The development of a comprehensive strategy for evaluating the performance of speckle filtering. In particular objective methods that correlate with subjective clinical opinion are used;
- Development of an improved method of edge detection for speckled imagery;
- The construction of a speckle removal filter based on this new edge detector.

1.2.2 Contributions

A number of novel contributions were made in order to achieve the objectives summarised above.

- The relationship between subjective expert opinion and a set of objective quality metrics is explored statistically for the first time, for the assessment of quality in clinical ultrasound. A set of objective image quality metrics is found, that correlates well with various aspects of expert judgement of the diagnostic usefulness of speckle filtered echocardiography;
- A detailed evaluation is performed, comparing the suitability of a wide variety of speckle filtering approaches for clinical echocardiography. This study is more comprehensive than those seen to date in the literature, in terms of the wide range of recent filtering methods considered (some of which have not been included in comprehensive evaluations before), the large set of clinical test data, and the inclusion of realistically simulated ultrasound. A further unique aspect of this investigation is a detailed consideration of the computational requirements of each filtering method. The results of this evaluation rank speckle filtering methods by their suitability for inclusion in real-time echocardiographic systems. This is the first time such a detailed ranking has been carried out.
- The development of a new method for estimating the parameters of a statistical model of the coefficients of an orientation-specific multiscale image decomposition;
- The development of an accurate edge detection method for speckle-corrupted ultrasound, based on quantifying observed properties of the zero crossings of the same orientation specific multiscale image decomposition. This proposed method is shown to perform better than current edge detection methods when applied to echocardiography;
- Construction of a speckle suppression technique, based on the approach of anisotropic diffusion and the proposed edge detection method. This technique improves upon the existing anisotropic diffusion approach, in particular by not requiring iterative construction of the edge map.

1.2.3 Publications

The research presented in this thesis has resulted in the following publications (which are summarised in Appendix C):

Journal Publications:

- S. Finn, M. Glavin, and E. Jones, "Echocardiographic Speckle Reduction Comparison", *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 58, no. 1, pp. 82-2011, Jan. 2011; doi: 10.1109/TUFFC.2011.1776. Full text included.
- S. Finn, E. Jones, P, Nolan, K. Daly, and M. Glavin, "Ultrasound Edge Detection using Stable Modeling and Zero Crossings of Gabor Coefficients", *Med. Eng. Phys.*, 2011. Submitted;
- S. Finn, M. Glavin, and E. Jones, "Speckle Reducing Diffusion using a Static Edge Map", 2011. In preparation.

Conference Proceedings:

- S. Finn, E. Jones, and M. Glavin, "Objective and subjective evaluations of quality for speckle reduced echocardiography", in *Proc. IEEE Intl. Annu. Conf. Engineering Medicine and Biology Society*, Sep. 2009, pp. 503-506. doi: 10.1109/IEMBS.2009.5333572.

The structure of this thesis is presented in the following section.

1.3 Chapter Summary

This structure and content of the remainder of this thesis are presented in this section.

Chapter 2 presents background information relevant to the topics discussed in this thesis. Described here is the process of ultrasound image acquisition, including the physical basis of the imaging modality and the hardware and signal processing elements present in modern scanning systems. The nature and statistical properties of the speckle artefact observed in ultrasonography are detailed, as well as its relationship with the structure of the imaged medium. This chapter then presents a review of the literature which summarises the approaches which have been proposed for speckle filtering. The strategies used in the evaluation of these filters are also reported, as are a number of independent reviews comparing a set of speckle filters.

In Chapter 3, a study is presented in which two approaches to the evaluation of speckle filtering are compared. Filtering approaches are assessed subjectively by manual visual inspection by a number of clinical experts, and also by objective assessment by the computation of image quality metrics. The relationships between these two approaches are explored using statistical analysis, for clinical echocardiographic imagery.

Chapter 4 presents a comprehensive evaluation strategy for clinical speckle filtering. Speckle filters are applied to a large test set of clinical video frames, as well as a set of realistically simulated images. A set of objective image quality metrics are applied to the filtered clinical and simulated images. These metrics quantify different aspects of the filtering process, motivated by the study of Chapter 3 in which the relationships between image quality metrics and the evaluation of clinical experts are investigated. As well as functional performance, the filters are compared in terms of their computational requirements, which are important for implementational considerations. The evaluation strategy presented here is applied to a wide range of speckle reducing filters from the literature.

In Chapter 5, an accurate edge detector for operation in speckled imagery is presented. This operates using orientation specific multiscale decomposition achieved with Gabor filtering. The effect of the speckle component in the Gabor domain is first removed by MAP estimation. This method includes the use of an accurate model for the speckled Gabor coefficients. A method is proposed for the estimation of the parameters of this model. The edge detection method operates by quantifying observed properties of the zero crossings of Gabor coefficients. These properties are demonstrated using simulated and clinical images, and methods for their measurement are presented. The proposed Gabor zero crossing edge detector is then used in the removal of speckle, by inclusion in an anisotropic diffusion process.

Chapter 6 presents an evaluation of the performance of the edge detection and speckle reduction components of the technique presented in Chapter 5. Included here is a Monte Carlo simulation to determine the accuracy of estimation of the model parameters. The accuracy of the proposed edge detection method is quantified by comparison to a number of other edge detection methods from the literature, using simulated, phantom and clinical ultrasound images. The removal of speckle using the proposed method is verified by applying the evaluation framework proposed in Chapter 4. Results from each of these tests are presented and discussed in this chapter.

This thesis is concluded in Chapter 7. Here the contributions of the thesis are further examined, and suggestions for areas of future work are presented.

Chapter 2

Background

The objective of this thesis is to develop improved methods of speckle reduction and edge detection in speckled ultrasound imagery. A crucial step in the development of improved methods of speckle reduction is a comprehensive framework for the evaluation of their performance. This chapter presents background information on topics fundamental to the investigation of this area of image processing. An understanding of the physical process of ultrasound image formation and acquisition provides a crucial conceptual background for image processing applications in this area. These physical processes are also key to understanding the nature of the speckle artefact in ultrasonography. This chapter presents a detailed description of the process of ultrasound scanning, including an examination of the physical interactions upon which the modality is based, and also the hardware and signal processing components of ultrasound imaging systems. The statistical treatment of speckle in the literature is also described in this chapter. This includes details of the relationship between the characteristics of speckle and the properties of the reflecting medium. A review of the literature on speckle reduction methodologies is reported upon in this chapter. These methods cover a range of approaches to speckle filtering, including local statistics, anisotropic diffusion and multiscale methods. The evaluation strategies for assessing the performance of these filters are also reported, as are a number of independent reviews comparing a set of speckle filters. A detailed technical description of a number of these filters is also presented.

This Chapter is organised as follows: Section 2.1 describes the process of ultrasound image acquisition. The properties of speckle are described in Section 2.2. A review of the literature is presented in Section 2.3, focusing on methods for the removal of speckle.

Section 2.4 contains an in-depth description of a number of post processing methods which are evaluated further in this thesis. Section 2.5 then concludes this chapter.

2.1 Ultrasound Image Acquisition Overview

This section presents a review of the process of ultrasound image acquisition, including the fundamental principals of operation and the hardware and signal processing elements present in modern scanning systems.

2.1.1 Ultrasound Fundamentals

While the term *ultrasound* can technically be used to refer to all acoustics of frequency greater than the upper threshold of human audibility ($f > 20\,\mathrm{k\,Hz}$), clinical imaging is generally in the $1\,\mathrm{M\,Hz}$ to $20\,\mathrm{M\,Hz}$ range [10]. In some specific application, frequencies as high as $40\,\mathrm{M\,Hz}$ can be used [11]. Imaging is based on the transmission of acoustic pulses into the body, usually from a hand held transducer. The acoustic pulse interacts with the tissue medium, and echoes are reflected by interfaces between tissue of differing acoustic properties. These are detected by a receiver (usually the same transducer which launched the acoustic pulses). If the propagation velocity of sound waves in the imaged medium is known, the depth of interactions giving rise to the echoes can be determined. Characteristics of the returned signal (e.g. amplitude, phase) provide information on the interaction, and indicate the nature of the media involved. The amplitude of the reflected signal is used to produce ultrasound images, while the frequency shifts provide information on moving targets such as blood.

The propagation speed in tissue varies with tissue type, temperature and pressure. Assuming constant temperature and pressure in the body, only the tissue type is considered [11]. The mean speed of sound propagation in human soft tissue is generally taken as $1540\,\mathrm{m\,s^{-1}}$. Acoustic pulses transmitted into the body can experience:

Reflection: Also known a backscatter, reflection occurs when an acoustic pulse encounters an interface between tissues of differing acoustic impedances.

Refraction: When sound waves pass through an interface between media of different propagation speeds, a change in the direction of propagation occurs.

Absorption: Energy from the acoustic pulse is absorbed into the tissue, by conversion to thermal energy.

The acoustic impedance of a medium (Z) is the product of its density, ρ, and the speed of acoustic propagation in that medium, c:

$$Z = \rho c \tag{2.1}$$

The strength of reflection at an interface depends on the difference in acoustic impedance on each side of the boundary, as well as the size of the interface, its surface characteristics, and the angle of insonification. [1, 10]. The reflection coefficient at the interface is given in the same manner as the analogous case of electromagnetic propagation:

$$\Gamma = \frac{Z_2 - Z_1}{Z_2 + Z_1} \tag{2.2}$$

Depending on the nature of the interface, two types of reflection are observed:

- *Specular Reflection* occurs when the interface is large and smooth with respect to the ultrasound pulse wavelength, e.g. the diaphragm and a urine filled bladder [10]. Strong clear reflections are produced, in the same fashion as for a mirror. Detection of these echoes is highly dependent on the angle of insonification [10].
- *Scattering* results from interfaces much smaller than the wavelength of the ultrasound pulse. A volume of small scattering interfaces such as blood cells or a non smooth organ surface produce echoes which are scattered in all directions. This collection of scatterers are said to act as a diffuse reflector. The constructive and destructive interference of these scattered echoes results in a granular artefact known as *speckle*.

Refraction is governed by Snell's law:

$$\frac{\sin\theta_1}{\sin\theta_2} = \frac{c_1}{c_2} \tag{2.3}$$

where the quantities θ_1, θ_2, c_1, and c_2 are respectively the angles of incidence and refraction, and the corresponding speeds of propagation. Attenuation of the acoustic pulse is mainly due to absorption and scattering [11], and can be modelled by the equation [12]:

$$A(x,t) = A(x,0)e^{-\alpha f 2x} \tag{2.4}$$

where $A(x,t)$ is the amplitude at depth x and time t, f is the frequency and α is the attenuation coefficient. Generally attenuation is quite severe: in [11], a halving of intensity every 0.8 cm at 5 MHz operation is described.

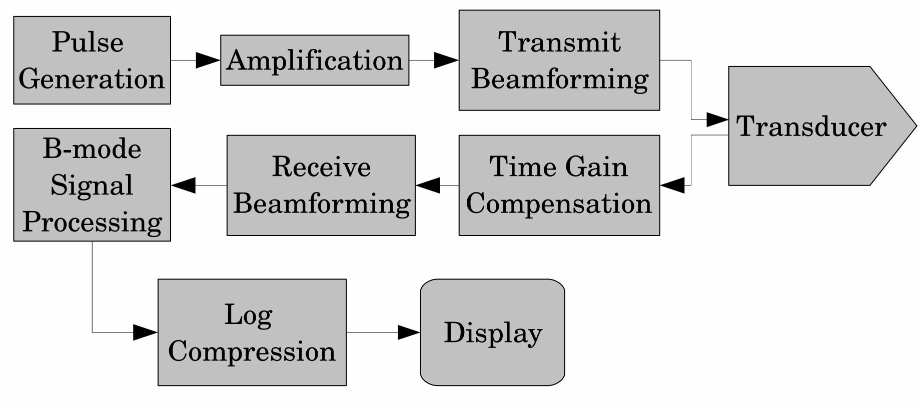

Figure 2.1: Block Diagram of a B-mode ultrasound system.

2.1.2 Scanning Systems

A single line of insonification (referred to an ultrasound beam) can be used to display the amplitude and position of a time-varying echo. This technique can be used to track the position of moving reflectors, and is known as M-mode ultrasound (M for motion). High frame rates are possible in this imaging mode, allowing the analysis of fast moving structures such as cardiac valve leaflets. By taking multiple lines of insonification (i.e. steering the ultrasound beam over and back across a target volume at sufficient speed) a two dimensional image can be constructed. This is referred to as B-mode ultrasonography (B for brightness), and is the focus of the research presented in this thesis. Other modes of ultrasound imaging are also in common usage. 3D ultrasound is essentially an extension of 2D mode, wherein either a 2D scanning array or software registration of images from multiple imaging planes is performed to construct the 3D volume. Modern systems are capable of 3D scanning in real-time, and have found usage in fetal and cardiac scanning. A different approach is presented by Doppler ultrasonography. This mode operates by using frequency information of the returned signals. As the name suggests, Doppler ultrasound exploits the Doppler frequency shift effect to determine information on the motion of moving reflectors.

A practical ultrasound system consists of a number of components, represented visually in Fig. 2.1

- *Pulsing System:* The pulsing system is a mechanism which provides the electrical signals necessary to drive sound wave production. The active elements in the transducer which generate the acoustic pulses require precisely timed driving signals. These signals determine the centre frequency of the acoustic pulses, the pulse repetition frequency (i.e. the interval between pulses) and the spectral and amplitude properties of the pulses. Different shapes of waveform can be used, from a simple square wave to waveforms of arbitrary complexity. The pulse repetition frequency is an important imaging parameter, as enough time must be allowed for pulses to return from the imaged depth before the next pulse is transmitted. The acoustic pulses are broadband in nature (i.e. very short duration) for anatomical imaging, to achieve maximal spatial resolution. For Doppler flow analysis, longer durations of pulses are used to provide better frequency resolution, up to continuous wave imaging.

- *Transducer:* In modern ultrasound systems, transducers contain active elements which are generally made from piezoelectric ceramic materials. These piezoelectric elements produce pressure (sound) waves when the polarity of an applied electric field changes. In the reverse situation, the elements are compressed as they are struck by returning sound waves. A measurable potential difference is then produced, which forms the basis of an ultrasound image. The simplest transducer design contains a single element, which is capable of producing accurate plots of a single scan line (line of insonification). Array transducers contain multiple active elements, which can be combined in a number of configurations. Pulses are generated in these transducers using groups of multiple active elements. By altering the timing and sequence of activation, the scan line can be steered directionally (allowing sweeping over the medium) and also focused at different depths.

- *Beamforming:* Ultrasound systems generally perform beamforming in both transmit and receive directions. The sequence of active element excitation in transmission allows the steering and focusing of the beam. Multiple transmission focal points can be achieved in many modern systems by the use of time multiplexing. The received signals from each transducer element are usually amplified to compensate for depth dependent attenuation, a process known as time gain compensation (TGC). The individual signals from the elements are then combined to form the composite received signal, generally using the delay and sum method [13].

- *Signal Processing:*

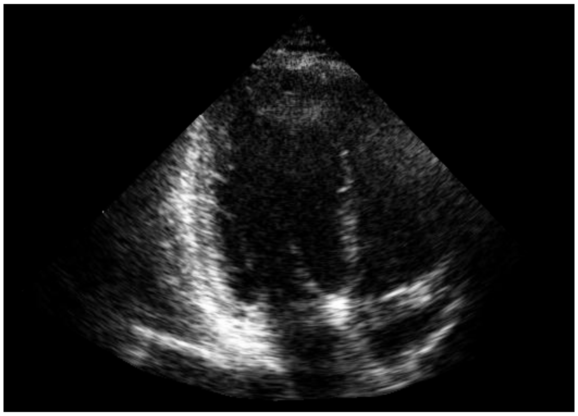

Figure 2.2: Example echocardiographic image with speckle.

After the received signals have been acquired, a number of processing steps are applied to facilitate image display. The steps involved differ between imaging mode (e.g. B-mode, Doppler, etc.) and the general processing steps applied for B-mode are [13]: Band-pass filtering is usually performed first to isolate the frequencies of interest, followed by a step referred to as *detection*. Detection refers to the process of extracting an amplitude image. This involves demodulation of the signals to baseband and envelope detection. The high frequency signals are known as radio frequency (RF) before demodulation and envelope detection. Finally, the dynamic range of the amplitude image is restricted to facilitate display. This is accomplished using log compression, generally by:

$$Y = \alpha log(X + 1) + \beta \tag{2.5}$$

where (X, Y, α, β) are respectively the amplitude image, log compressed image, and selectable contrast and brightness parameters [14]. The exact details of these signal processing techniques are generally held by manufacturers as proprietary information [11].

2.2 Characteristics of Speckle

This section details the properties and nature of the speckle pattern observed in coherent imagery, as well as its relationship with the underlying medium. As detailed in the description of ultrasound acquisition of Section 2.1 above, the interference experienced by reflections from diffuse scatterers result in a granular pattern known as speckle. Fig. 2.2 displays an example ultrasound image, containing a speckle component. Speckle is common to all imaging systems using coherent waves for illumination, including laser and radar imagery. In echocardiography, speckle noise is prominent in all cross-sectional views [15], and its effect is far more significant than additive noise sources such as sensor noise [2].

The basic description of ultrasound speckle in the literature is based on the characterisation of laser speckle by Goodman [16, 17]. This approach is extended to acoustic imagery by a number of authors [9, 18, 19]. Burckhardt [9] notes that despite its random appearance, speckle is essentially deterministic: scans under identical situations produce the same speckle pattern. This behaviour is in contrast to that of true stochastic processes such as electrical noise. There is no direct relationship between the imaged medium and the observed speckle pattern however. If the same object is imaged with different imaging parameters, the speckle pattern produced is quite different. Burckhardt thus concludes that the speckle pattern has only a tenuous relationship to the imaged medium, and is instead dependant on the parameters of the imaging system. The size of the speckle granules are of similar size to the resolution of the scanner, in both axial and lateral directions. Burckhardt justifies the treatment of pulsed acoustics as a coherent wave source in the situation where the each pulse contains a number of cycles of the carrier wave.

Each ultrasound pulse encloses a three dimensional volume which defines the smallest resolvable structure, which is known as the resolution cell. The nature of the speckle at an image location is determined by the number of diffuse scatterers which are present in the resolution cell at the relevant position in the imaged medium. If the number of scatterers is large and randomly positioned, the resulting pattern is known as fully formed (or fully developed) speckle. In this case the speckle pattern depends only on the imaging system. Wagner *et al.* [18] show that in the case of a mixture of diffuse and specular scattering, the speckle pattern is related to the underlying texture of the medium.

The randomly scattered echoes from the scatterers are summed within each resolution cell. These echoes are sinusoidal in nature. Expressed as phasors, the summation is described as a random walk of the real and imaginary components. In the case of fully developed speckle, and with uniformly distributed phase values between 0 and 2π, these components have a circular Gaussian distribution [18]:

$$p(a_r, a_i) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{a_r^2 + a_i^2}{2\sigma^2}\right) \tag{2.6}$$

where (a_r, a_i) are the real and imaginary components. The amplitude within each resolution cell, $A = \sqrt{a_r^2 + a_i^2}$, is then given by a Rayleigh distribution:

$$p(A) = \frac{A}{\psi} \exp\left(-\frac{A^2}{2\psi}\right), \quad A \geq 0 \tag{2.7}$$

The Rayleigh parameter ψ depends on the mean square scattering amplitude of the medium [16]. Burckhardt defined an SNR measure for the amplitude:

$$SNR_A = \frac{\bar{A}}{\sigma_A} \tag{2.8}$$

i.e. the ratio of mean to standard deviation. It can be shown that $SNR_A = 1.91$ for the Rayleigh distribution of fully developed speckle. Burckhardt explores the statistics of image compounded from multiple individual scans, and demonstrates a method of simulating a theoretically maximally speckle free image from a known structure. Wagner *et al.* explore the second order statistics of speckle, and extend this analysis to compounding images in [19]. Alternative distributions to Rayleigh have been proposed for situations not meeting the requirements of fully developed speckle, such as insufficient numbers of scatterers per resolution cell, or non random positioning. These include the Rician, Nakagami and K distributions [20–22].

The number of scatterers required for fully formed speckle varies in the literature. Ten or greater is a common figure [12, 23, 24], although it is stated in [25] that at least thirty scatterers should be present for the central limit theorem to hold. However, the work of Dantas *et al.* [7] extends that of Burckhardt, and shows that neither the number of scatterers or their random positioning are required for fully developed speckle governed by Rayleigh statistics. Based on the concept of effective scatterers, it is noted that a subset of scatterers can give a greater contribution the reflected signal strength. A set of equivalent scatterers are introduced, which provides the same fully formed speckle pattern as before. In contrast to the original set of scatterers, these equivalent scatterers are a sparse set and uniformly positioned.

A large volume of literature exists concerning methods of reducing or eliminating speckle (see Section 2.3 below). The common justifications for the removal of speckle in these works are the general reduction in image quality due to the presence of speckle.

However, the question of whether or not to remove speckle as noise in clinical imagery is an open one, and depends largely on the application. A number of specific negative effects of speckle, and benefits of its removal in clinical ultrasonography are noted:

- Speckle introduces spurious ‘false-fine’ structures, which give the appearance of resolution beyond that of the imaging system [7].
- The speckle pattern can mask small grey level differences of the imaged medium [9], which can obscure tissue boundaries [5, 7].
- Speckle reduces image contrast [5, 6].
- The presence of speckle can negatively impact human interpretation of ultrasonography [2, 4, 6], introducing a degree of subjectivity to this task [5]. The presence of speckle has been determined to be the cause of an eight-fold reduction in lesion detectability [8]. Reduction of echocardiographic speckle has been shown to positively affect subjective image quality and improve boundary definition [15].
- The effectiveness (speed and accuracy) of automated processing tasks is also reduced by speckle [2–4, 6], such as edge detection, segmentation and registration.
- While speckle can be viewed as deterministic [7], it does not contain information on the imaged structure in the fully developed speckle case.

However there are situations where speckle preservation is desired. In clinical situations, it is noted that clinicians may prefer the original speckled images in some cases [6]. Dantas and Costa [5] noted that when speckle is removed, the loss of false fine detail can lead to a perceived reduction in image sharpness, even if the boundaries of anatomical structures are not blurred. The speckle pattern is seen as having diagnostic utility in specific conditions, such as diffuse liver diseases [26] in abdominal imaging and hypertrophic cardiomyopathy in echocardiography [15]. Some automated processing tasks such as feature tracking [27] and tissue characterisation, some recent examples of which can be found in [28–30].

So while some particular applications are best served by preserving speckle, others benefit from its removal. A pragmatic approach was taken by Zhang *et al.* [6], who promote the idea of the speckle reduced image as a complementary addition to the original image, rather than a replacement.

2.3 Literature Review: Speckle Reduction

This section presents a review of the most widely cited and recent approaches to the removal of speckle in the literature. Approaches to speckle reduction can be broadly grouped into compounding and postacquisition methods. An important related issue is the question of how to evaluate the performance of these speckle reduction techniques. The methods of evaluation used in the literature are also reviewed here.

This section describes the broad state of literature in the area. A number of the speckle reduction methods described here are evaluated in Chapter 4 of this thesis. These techniques are described in an in-depth fashion in Section 2.4 below, including all of the technical information required for the implementation of these methods.

2.3.1 Compounding Approaches

Compounding techniques combine two or more images of the same imaged area. The measurements from image structure will be partially correlated, i.e. they all contain features from the imaged structure. The speckle pattern in each will be different in each however, so compounding these images (e.g. by averaging) results in an image with enhanced structure, but with a reduced speckle pattern. A number of different scanning methods can be used to produce the images to be compounded. Frequency compounding uses images with separate frequency ranges within the transducer bandwidth [31–34]. A common technique is spilt spectrum processing (SSP) [35–37], in which the wideband RF signal is split into a number of subbands using bandpass filters. Envelope detection of these RF subbands yields amplitude data, which is combined to produce an image with enhanced structure and reduced speckle component. A recent work in this area is that of Dantas and Costa [5]. These authors proposed a related method which is applied to the entire 2D RF image, as opposed to some methods of split spectral processing (SSP), which are applied to each 1D RF scan line individually. The RF image was decomposed into a number of orientation specific subbands by use of a bank of modified log Gabor filters. Each subband RF image was used to generate an amplitude image, and the final speckle reduced image was produced by compounding these amplitude images. The proposed method was tested by application to simulated images. Simulation involved modelling the echogenicity of the imaged medium as a distribution of complex point scatterers, with differing densities in separate regions. The method of generating a maximally noise free reference proposed by Burckhardt [9] was adapted to this method of simulation. The output of the proposed speckle reduction method was compared to the noisy input and the speckle free image using visual inspection, and the SNR_A metric of Burckhardt.

Other compounding approaches include spatial compounding, which combine multiple images from different scan directions [38–40]. Burckhardt [9] showed that in order for these scans to be independent of each other (and so having uncorrelated speckle patterns), the transducer must be translated by approximately half its element width. Temporal compounding (frame averaging) operates as the name suggests by combining scans performed over time. This approach suffers from a dependence on motion: in a still medium the speckle pattern will not change. Conversely, fast moving organs such as the heart may appear smeared with this method.

2.3.2 Postacquisition Speckle Reduction

Postacquisition methods operate on the image after it has been envelope detected. These methods are a practical alternative to compounding approaches, as they do not require a specific mode of scanning, or access to the RF data (which in most commercial systems is inaccessible to the user). In addition, a simulation study comparing postacquisition filters to spatial compounding reported better image improvement for filtering, in terms of speckle reduction and image quality [41]. The number of postacquisition speckle reduction methods in the literature is large, and the selection presented here are grouped according to their general approach.

2.3.2.1 Adaptive Filters

Adaptive filters attempt to adjust the level of filtering at each image location. The Lee filter [42, 43], Kuan *et al.* filter [44] and Frost filter [45] were proposed for the task of speckle removal in synthetic aperture radar (SAR). These assume a multiplicative model for speckle noise. Enhanced versions of the Lee and Frost filter were proposed by Lopes *at al.* [46]. These are improved by the classification of image pixels into a number of classes, and filtering accordingly. The method of Bamber and Daft [8] extended this adaptive approach to ultrasound images by varying the degree of smoothing according to a local estimate of the level of speckle.

Median filtering was extended to the case of speckle removal by Loupas *et al.* [47].

This approach replaced pixels with the weighted median of a window of pixels around the pixel of interest, and is known as the adaptive weighted median filter (AWMF). Dynamically variable window sizes are also employed. Region growing techniques such as [48–51] have been applied to ultrasonography. Pixels are grouped in these methods, according to similarity of intensity and connectivity. Spatial filtering is performed to extend these regions, and the challenge is the selection of appropriate similarity criteria.

A recently proposed method by Tay *et al.* [52, 53] used an iterative technique of speckle reduction by the removal of outliers. The locale around each pixel is examined, and local extrema are replaced with a local average. This process is repeated until no further outliers are found. Thus the filter reduces the local variance around each pixel, and is referred to by the authors as the "Squeeze Box" filter. The proposed filter was compared with the SRAD filter of [3], using a simulated phantom generated with the Field II software package [54]. This phantom contained regions of differing contrast relative to a background. To quantify performance, a contrast metric, Q, was proposed. This measures the degree to which the variance within homogeneous regions is reduced, and also the separation between these regions. These regions of differing contrast, etc., need to be known beforehand. The filters were also compared on a single clinical image by visual inspection.

The speckle filter of Coupé *et al.* [55] operates in an adaptive non-local means fashion. A non-local means approach estimates the true value of each pixel as the weighted sum of the windowed averages of within a search volume centred at the pixel of interest. The method proposed here extended this approach to speckled ultrasound by incorporation of a multiplicative noise model by Bayesian estimation. The proposed method was compared against the Squeeze Box filter of [52, 53], the SRAD method of [3], and the original non-local means filter. Two ultrasound simulation approaches were used. Firstly, simple artificial speckle in the form of multiplicative Gaussian noise was added to noise free reference image. The filtered images were compared with the reference using the SNR measure of Sakrison [56]. The second simulated was identical to that used in the method of Tay *et al.* above, including comparison using the Q contrast metric. The proposed method was also compared with the SRAD and Squeeze Box filters using a small number of clinical images, both 2D and 3D, by visual inspection.

Massay *et al.* [15] proposed a method of speckle reduction using local statistics. This is similar to approaches suggested for the reduction of speckle in SAR images, adapted so that the level of smoothing is determined by an estimate of he level of speckle at each location. The proposed method was applied to a phantom image, and the output visually inspected. The effect of speckle reduction on clinical echocardiographic images was investigated by measuring the mean and standard deviation in regions of interest, and quantifying the change in these values due to filtering. A total of 24 pathological and 12 symptomatic images were used. The clinical implications of the change seen in the images due to filtering were discussed. Manual measurements were made of the valve opening in short axis images and of the cavity area in long axis images. It was noted that the cavity wall becomes clearer and easier to delineate, and that the measured area is generally larger after filtering. The valve boundaries were noted to be clearer as well. A note was made that the speckle pattern of the myocardium can be useful in the diagnosis of hypertrophic cardiomyopathy, and so speckle reduction should be used cautiously in this case.

2.3.2.2 Diffusion Filtering

Anisotropic diffusion is an iterative method of smoothing an image, based on the concept of heat diffusion. Perona and Malik [57] introduced the first anisotropic diffusion method for additive noise. A diffusion function is calculated at each iteration, with the aim of inhibiting smoothing across image edges and permitting it in homogeneous areas. This was later enhanced by a number of authors, including Catté *et al.* [58] who introduced a regularising step. This improves stability in the presence of noise, and is practically implemented by computing the diffusion function on a smoothed version of the input image. Further details on anisotropic diffusion (equations and methods of discretising the partial differential equations) can be found in Section 2.4.1 below.

The diffusion function used to control image smoothing in the diffusion method of Perona and Malik is based on the image gradient, and is unsuited to multiplicative speckle situations. Yu and Acton [3] extended the anisotropic diffusion approach to speckled applications. The gradient based edge discriminator in the Perona and Malik formulation is replaced with a discriminator better suited to speckle. For this purpose, the coefficient of variation discriminator used in synthetic aperture radar is adapted to the diffusion process. A practical method of explicit discretisation was detailed in this work. The proposed method was compared the enhanced Lee and enhanced Frost filters [46] from the field of SAR imaging. A homomorphic version of Perona and Malik anisotropic diffusion was also proposed. This involves a log operation prior to filtering, and a subsequent exponential operation. The filters were applied to a set of simulated carotid ultrasound images. The method of simulation proposed models the point spread function (PSF) of the ultrasound system as space invariant (constant throughout the image). The imaged medium's echogenicity is modelled as the product of a smooth reference image with a Gaussian field to provide diffuse scattering. Calculation of the simulated amplitude image was performed by convolving the PSF with the echogenicity model, and then taking the magnitude of the Hilbert transform of the result. The filters were applied to images simulated in this manner, and the results were evaluated by measuring the mean and variance in homogeneous regions. Edge preservation was measured using Pratt's figure of merit (FoM) [59]. The filters were also applied to a clinical carotid ultrasound and a SAR image, and the results visually inspected.

The work of Aja-Fernandez and Alberola-Lopez [60] proposed a number of improvements to the SRAD technique, known as detail preserving anisotropic diffusion (DPAD). A superior edge discriminator from the field of radar imagery was incorporated, and the use of larger windows to calculate local statistics was introduced. A simple yet accurate method for estimating the noise coefficient of variation was also presented. The method purposed here was compared against the SRAD method, and also the regularised version of Perona and Malik's anisotropic diffusion. Simulated images were generated by corrupting a noise free reference image with a simple approximation of speckle, as a multiplicative uniformly distributed field. The filtered images were compared by visual inspection, and by the structural similarity for images (SSIM) measure of Wang *et al.* [61]. The proposed method was also compared with the SRAD filter by application to a single clinical image, assessed by visual inspection.

The nonlinear coherent diffusion (NCD) method of Abd-Elmoniem *et al.* [4] is a tensor valued anisotropic diffusion scheme for the removal of speckle. Tensor valued schemes, proposed by Weickert [62], allow the strength of diffusion to vary directionally at each location. Similar to the approach of Weickert, this method makes use of the structure tensor to describe image gradient. The method of determining diffusion strength is adapted to speckle removal. The eigenvalues of the diffusion tensor define the directional strength of diffusion at each image location, and these are chosen to reflect an estimate of the strength of speckle at that location. Thus the diffusion is termed speckle adaptive. The proposed method was evaluated by comparison with the adaptive weighted median filter of Loupas *et al.* [47], and two wavelet shrinkage methods: the soft threshold of Donoho [63], and the thresholding and enhancement method of Zong et al. [2]. The methods were applied to a commercial contrast phantom, which had a number of regions of differing contrast. A speckle free reference was manually constructed by selecting the mean intensity value in each of the regions, and assigning this value to the entire region. The method by which the borders of these regions is determined was not reported. Evaluation was performed by visual inspection, with the aid of profile line through the images. The MSE metric was also used to compare each speckle filtered phantom image with the speckle free reference. Finally, the speckle filters were applied to a clinical echocardiographic image, and also a kidney and a liver ultrasound. Evaluation was again performed by visual inspection with profile lines.

The orientated SRAD (OSRAD) proposed by Krissian *et al.* [24] provides the capability of varying diffusion with direction to speckle adaptive diffusion filtering. This method adapts the SRAD approach into a tensor diffusion framework, similar to the NCD filter above and the work of Weickert [62]. Thus the strength of filtering can vary in contour and curvature directions. The improved edge discriminator mentioned above is also employed, and a numerical scheme for solution of the diffusion equation is presented. The proposed method was compared to median filtering, the Lee [42] and Kuan [44] SAR filters, Perona and Malik's anisotropic diffusion both with and without a homomorphic transform, the regularised diffusion of Catté *et al.*, the filters of Rudin *et al.* [64,65], the DPAD and SRAD filters, and previous work by the Krissian [66]. A simulated 2D image was generated by taking a speckle free reference and applying artificial speckle, approximated as multiplicative Gaussian noise. A 3D simulated image was generated using the same technique. The evaluation strategy of SRAD was followed here, i.e. the filters were compared to the speckle free reference images by measuring the mean and standard deviations in homogeneous regions, and also by the use of the FoM metric. A metric was proposed which combines the mean and standard deviation values in homogeneous regions into a single value for each filter. A clinical 3D liver image was processed with the proposed method, and the results visually inspected.

2.3.2.3 Multiscale Methods

The analysis of signals and images using multiscale decompositions is common place in signal processing. A number of such methods have been applied to the reduction of speckle. Two distinct methods of decomposition are used: the wavelet and pyramid transforms. Both of these techniques capture spatial and frequency information of an image. The wavelet transform achieve this through subband coding, using a quadrature mirror pair of filters to find the high and low pass images (known as the approximation and detail images, respectively). By this approach a series of approximations at increasingly coarser scales can be generated. The multiresolution pyramid transform [67] achieves a similar series of approximations, without the need for a quadrature pair of filters: the detail image is found at each scale as the difference between the original and low pass approximation images. In both cases the filtering of noise is achieved by modifying the coefficients after decomposition. Reconstruction is the reverse process, synthesising the denoised image from the adjusted coefficients. Further information on the multiresolution pyramid transform can be found in [68].

Wavelet techniques can be grouped into those which operate by thresholding, Bayesian estimation, or correlation between coefficients. Many techniques use the soft thresholding approach of Donoho [63] Some of the earliest adaptions of the thresholding approach to the case of speckle noise were proposed by Guo *et al.* [69] and Moulin [70]. These methods treat the multiplicative nature of the speckle noise by a homomorphic transformation: a logarithmic operation is applied before filtering to allow the noise to be treated as additive, and a complementary exponential operation is applied after filtering. Hao *et al.* [71] proposed a speckle filtering method, which operates by thresholding wavelet coefficients. The image is split into two images: the first being the output of the AWMF filter, and the second the difference between the AWMF filter output and the original image. The AMWF filter is considered by these authors to produce blurred output, so the difference image contains the removed high pass information. Most of the speckle will be in this image, as will high frequency information removed by the AWMF method. Both of these images are decomposed using a wavelet transform, which are thresholded by the method of Donoho. The thresholded wavelet images are reconstructed, and summed to produce the final image. The proposed method was compared with the AWMF filter and a number of wavelet thresholding approaches. These were applied to simulated speckled images, generated by applying multiplicative signal dependent white Gaussian noise to a reference image. The MSE and Sattar's β metric were used to compare the methods. The methods were also applied to scans of a pigs heart, and are evaluated by visual inspection with the aid of profile lines.

The wavelet domain method of Zong *et al.* [2] performs thresholding to remove noise. Both soft and hard thresholding are applied, and a method for selecting threshold values was presented. A nonlinear processing step is also applied to the wavelet coefficient, with the aim of enhancing image features. The method proposed here was compared with median filtering, an extension to median filtering [72, 73], homomorphic Weiner filtering and wavelet soft thresholding. The methods considered were applied to clinical echocardiograms, and the results evaluated by visual inspection. A clinical case study was also undertaken, in which the accuracy and consistency of manually defined borders by experts was investigated. Two experts participated in the study, and a sequence of sixty short axis echocardiograms were used. The difference between the manually defined borders for each image was evaluated by calculating the mean difference in border difference, and the mean difference in the areas defined by the borders. A statistically significant improvement in the consistency of the defined borders for was seen with speckle filtered images. The correlation between experts of the areas enclosed by the defined borders were plotted, showing a closer agreement between experts for the filtered images.

In addition to thresholding, noise removal in the wavelet domain is also performed using Bayesian denoising. These approaches model the distributions of the wavelet coefficients, and use this a priori information to infer the noise free coefficients. Gupta *et al.* [74] modelled the wavelet coefficients of the underlying speckle free image using a generalised Laplacian distribution. This method simultaneously removed the speckle component and performs image compression using a quantisation function. The quantisation function was adjusted to account for the estimated level of speckle, the estimated parameters of the generalised Laplacian signal distribution, and the desired rate of compression.

Achim *et al.* [75] proposed a method for statistically removing ultrasound speckle in the wavelet domain. The wavelet coefficients of the log transformed speckled image are modelled as the convolution of a symmetric α stable (SαS) distribution for the speckle free coefficients, and a zero mean Gaussian distribution for the speckle noise. The variance of the Gaussian noise is estimated using the median absolute deviation (MAD) of wavelet coefficients. The parameters of the SαS distribution are then estimated by least squares fitting of the observed density function spectrum to the empirical characteristic function. After estimation of the parameters of the statistical model, a shrinkage function is found by numerical calculation of the maximum a posteriori (MAP) estimation curve. The wavelet coefficients are modified by the shrinkage function, and the inverse wavelet transform (followed by an exponential operation to reverse the initial log transform) produces the denoised image. The proposed method was compared to median filtering, the homomorphic Weiner filter [76], and wavelet shrinkage by the standard soft and hard thresholding methods. Evaluation was performed using a hybrid of clinical images with added artificial speckle. Clinical test images of the liver, kidney, pancreas and gall bladder were acquired, and those with a low speckle content are selected for testing. The remaining speckle in these images was removed using the homomorphic Weiner filter. Artificial speckle was then added, approximated as multiplicative noise in the form of the magnitude of a complex Gaussian random field, which was low pass filtered to introduce spatial correlation. Speckle was removed with the methods considered, and the resulting images were compared to the 'speckle free' reference using the MSE and signal to MSE metrics. The β metric of Sattar *et al.* [77] was used here, which measures edge region correlation between the original and filtered images.

The technique of Rabani *et al.* [78] models the distribution of the noise free wavelet coefficients using a local mixture of either Gaussian or Laplacian distributions. The speckle noise is assumed to be either Gaussian or Rayleigh in nature. For all combinations of local mixture distributions and speckle noise distributions, both the MAP and minimum mean squared error (MMSE) estimators are derived analytically. A recently proposed method by Fu *et al.* [79] accomplishes bivariate shrinkage of the wavelet coefficients. This approach models the joint density of the coefficients in each scale with their parents in the next coarser scale.

The third type of wavelet noise removal is the correlation based method proposed by Pižurica *et al.* [80]. This generalised likelihood method for medical imaging removes noise in the wavelet domain by utilising the correlation of useful wavelet coefficients across scales. Thus the technique does not rely on a model for the image noise, but rather locates the signals of interest based on their interscale persistence. This initial classification step is followed by empirical estimates of the signal and noise probability density functions (PDFs). These are used to define a shrinkage map which suppress those wavelet coefficients resulting from noise. The inverse wavelet transform then produces the denoised image. The proposed method was tested on simulated ultrasound, in which speckle was approximated as multiplicative, and generated as the low pass filtered magnitude of a complex Gaussian field. The reference images used were a simple test shape image, and a clinical image which had been treated to reduce speckle with the proposed method. The method presented in this work was compared to the homomorphic Weiner filter [76], and also to the noise free reference with an SNR measure.

The nonlinear multiscale wavelet diffusion (NMWD) method of Yue *et al.* [81] combines the approach of multiscale wavelet methods with anisotropic diffusion. This method aims to utilise the discrimination between signal and noise offered by multiscale methods, and also the iterative noise removal available through anisotropic diffusion. This iterative procedure performs wavelet decomposition in each step. A shrinkage function is applied the coefficients to suppress speckle, which is found using diffusion. The parameters of the diffusion process are calculated based on image statistics. The simulation approach of Sattar *et al.* was employed here, i.e. the application of multiplicative lowpass filtered magnitude of a complex Gaussian random field to a noise free reference. The proposed method was compared to the SRAD filter, and also the generalised likelihood method of Pižurica *et al.*. The filtered images were quantified relative to the noise free reference by use of the FoM metric and the ρ measure of Sattar *et al.*. The filters were also applied to clinical images of the heart and liver. Evaluation was performed by visual inspection in this case, with the aid of profile lines taken throughout the images.

As well as wavelet based methods, multiresolution pyramid methods of speckle reduction have been proposed. The approach of Aiazzi *et al.* [82] extends the approach of the Kuan filter to process each layer in the multiscale pyramid decomposition. Sattar *et al.* [77] presented a method which both reduces speckle and enhances image edges. Multiscale decomposition is performed using a pyramid transform. An edge detector is applied to the lowpass image, the output of which determines the coefficients from the high pass image which are included in reconstruction. The success of this method is therefore dependant on accurate operation of the edge detector, and a number of different techniques are tested. The proposed method was evaluated by comparison with the approach of Nagao and Matsuyama [83] by application to simulated images. These were constructed by corrupting a noise free reference with multiplicative speckle, approximated as the lowpass filtered magnitude of a complex Gaussian random field. The filters were compared to the noise free reference using the SNR, computed in a manually defined region of interest. The correlation between the filtered and reference images within the region of interest was also measured. A second correlation measure was proposed, which measures similarity between high pass filtered version of the filtered and original image. This measure is used elsewhere in the literature to quantify edge preservation, and is referred to as the β edge correlation metric here. The proposed method was also evaluated by application to a contrast phantom image and a clinical echocardiographic image, and assessed by visual analysis.

The approach of Zhang *et al.* [6] combines the approaches of multiscale analysis and anisotropic diffusion, and so is similar in approach to the NMWD method above. In this technique, multiscale decomposition is performed using Laplacian pyramid decomposition. Band pass images are then smoothed using Perona and Malik's diffusion, using the extension of Catté *et al.*. The diffusion gradient threshold of the diffusion function is estimated using a robust median absolute deviation (MAD) method. The method proposed here was compared to a version of the Weiner filter, in which filtering is performed on the individual layers of the pyramid decomposition. Also compared are Perona and Malik's diffusion as mentioned, and the SRAD filter. The filtering methods were compared by application to a simulated image, generated using the method proposed for evaluation of the SRAD method. Filtered output was compared to the noise free reference using the MSE and SSIM metrics, the β metric of Sattar *et al.*, and the contrast to noise ratio (CNR). An image from a contrast phantom was also used, and paired t-tests were used to determine if there is a statistically significant difference between the proposed method and the other approaches considered. Finally, the filters were applied to a clinical liver ultrasound and a carotid ultrasound. Evaluation for the clinical images was performed by visual analysis with the aid of profile lines taken from the images.

2.3.3 Speckle Reduction Reviews

As described above, many of the works proposing methods of speckle filtering include a comparison of the proposed method with some other from the literature. A select number of independent reviews have also been published. These vary in the precise nature and depth of the investigation performed, and some of these are summarised here.

The work of Thakur and Anand [84] compared the suitability of a number of different wavelets for the reduction of speckle in ultrasound imagery. Both the discrete wavelet transform (DWT) and the wavelet packet transform were considered. The latter is a method of decomposition similar to the DWT; the difference is the both the high and low pass components are decomposed at each level in the packet transform, while only the low pass component is decomposed in the DWT. Results of filtering using a wide range of wavelets were reported for a liver ultrasound image. These were quantified relative to the speckled input image by measuring the MSE, the normalised MSE, the PSNR (peak signal to noise ratio) and a correlation measure. It was concluded that the biorthogonal wavelet 'bior3.9' is optimal for the image considered. The wavelet filtering methods were compared to the AWMF filter by visual inspection.

The work of Adam *et al.* [41] looked at the effect of a combination of spatial compounding and postacquisition filtering. The transducer design reported in [85] was used, which can acquire three independent images simultaneously without loss of frame rate due to the arrangement of the receiving elements. Different methods of compounding these images were evaluated in this work, in combination with a number of postacquisition filters. The compounding of three images (I_1, I_2, I_3) is denoted by $I_c = f(I_1, I_2, I_3)$, where $f(.)$ is the compounding function. The strategies considered were:

$$\begin{aligned} f(I_1, I_2, I_3) &= I_2 \\ f(I_1, I_2, I_3) &= \sqrt[3]{I_1 \cdot I_2 \cdot I_3} \\ f(I_1, I_2, I_3) &= \sqrt{\overline{(I_1 \cdot I_2)} \cdot I_3} \\ f(I_1, I_2, I_3) &= \sqrt{\overline{I_1 \cdot I_2^2 \cdot I_3}} \\ f(I_1, I_2, I_3) &= \overline{I_1 \cdot I_2 \cdot I_3} \end{aligned} \tag{2.9}$$

where I_2 is the image received by the main centre element array, and I_1 and I_3 are those from elements placed to each side. The operator $\overline{x}$ denotes the average value. The compounded image in each case were processed using the SRAD filter and the non-linear Gaussian filter of [86]. Each combination of compounding and filtering method considered was applied to simulated kidney images generated using the Field II software package [54]. A number of regions of interest (ROIs) were measured, and within each three metrics were evaluated: the contrast to noise ratio (CNR), the lesion signal to noise ratio (LSNR) [87], and an SNR measure. The statistics of these metrics over the set of ROIs were examined to produce mean values and confidence intervals. Without the filtering step, it was shown that the compounding methods all provide improvement in visibility and smoothing over the single image case, although some speckle was still visible. Without spatial compounding, it was shown that both filtering methods improve the images to a greater degree than the compounding stage alone. The SRAD filter was the better of the two in this case. The simulation here was verified by use of an experimental setup based around a phantom, with a suitable transducer arrangement. The same compounding strategies and filtering methods were applied, and the combination of compounding by averaging with the nonlinear Gaussian filter was judged to give the best performance.

A recent comparison by Mateo and Fernández-Caballero [88] evaluated:

- The median filter
- The AWMF filter
- Frequency domain filtering: high frequency speckle noise is reduced using the 'ideal' lowpass filter (i.e. that with the sharpest cut-off) and a Butterworth filter.
- A wavelet domain lowpass filtering approach, in which the detail images are discarded completely before reconstruction.
- The above filters, with the exception of the median filter, were also applied after a homomorphic transform. This attempts to convert the multiplicative speckle noise into an additive form by application of a logarithmic operation before filtering. A complementary exponential operation is applied after filtering is complete.

These filters were applied to a simulated ultrasound image, generated using same method as in the review of Adam *et al.* above. The MSE, SNR, PSNR and the β edge correlation metric were used to compare each filtered output to the noisy input image. The effect of varying the parameters was also investigated in this work. Varying results were seen for the filters. As one might expect, the simple wavelet method was not capable of removing the speckle with blurring the image. Frequency domain filtering using the ideal lowpass filter led to Gibbs phenomena. Lowpass filtering with the Butterworth filter was deemed by the authors to be the best of the methods considered, in both the linear and homomorphic cases.

A comprehensive comparison of a large number of speckle reduction filters for application to clinical carotid ultrasonography was presented by Loizou *et al.* [89]. The methods compared include the Lee and Weiner filters, the median filter of [90], the homomorphic filter of [83], the homogeneity filter of [91], the geometric filter [92]. The homomorphic filter of [93, 94] which operates in the Fourier domain was also compared, as are the NCD and Perona and Malik diffusion methods. A wavelet shrinkage method was also compared, which uses the soft thresholding of Donoho [63], with thresholds calculated as in the approach of Zong *et al.* [2]. Evaluation was performed both by automated analysis, and also using classification by experts. In both cases, the focus was on the diagnosis of atherosclerosis (thickened artery walls due to plaque deposit). A set of clinical images from patients deemed to be at risk of this condition was used in filter evaluation, and these were divided into a symptomatic set (from patients who have displayed symptoms of this condition, such as stroke incidents), and an symptomatic set. For the test by automated analysis, 440 images were used, divided equally between the two sets. As image texture is deemed to be clinically important in the diagnosis of this condition, the automated analysis proceeded by calculating a large number (56) of texture features for each filtered image. The level of separability between the symptomatic and symptomatic sets was then analysed using a number of approaches. Firstly, a distance metric was defined, measuring the discrimination between the two image sets for each feature. Secondly, statistical analysis using the Wilcoxon rank sum test was employed. This determined if a significant difference in each metric exists between the image sets. Finally, a kNN (k nearest neighbour) classifier was used to classify images as symptomatic or symptomatic based on the texture scores. A set of image quality metrics were also applied, comparing the filtered image to the unfiltered version in each case:

- Distance metrics: MSE, RMSE, a set of Minkowski distances, and the Geometric Average Error [95].
- The SNR and peak to SNR measures of Sakrison [56].
- The universal quality index (UQI) of [96], and the SSIM.

The expert test was performed with 100 images, split evenly between symptomatic and symptomatic groups. The filtered and unfiltered images were shown to two experts at random, and each expert rates the quality of the image on a scale of one to five. The results of this study are interesting, and relevant to the expert assessment performed in Chapter 3 below. Only one method (the Geometric filter) was seen to improve the image quality as perceived by both experts across the entire dataset. The Lee filter was deemed to have the next best performance. A difference in the evaluations between the experts was noted, although this was not investigated statistically. The authors accounted for this by reference to the differing clinical specialities of the experts.

This section has presented a comprehensive review of the most recent and commonly used methods of reducing speckle from the literature. A number of compounding methods are available which although generally simple in operation require either specialist hardware, or access to the RF signals internal to ultrasound scanning system. Postacquisition methods operate on the final image produced by the scanning system. Numerous techniques of this nature have been proposed, and these can be broadly grouped into adaptive filtering, diffusion methods, and multiscale methods. The methods of evaluation used to compare these filters varies throughout the literature, but some common elements are:

- The use of simulated images, which allow the quantification of filtering effects relative to a noise free reference. Approaches vary in their level of realism, and these range from simplistic models which corrupt a reference image with multiplicative noise of various types, to methods which model physical properties of ultrasound scanners and represent the imaged medium by complex point scatterers.
- Real world ultrasound images are also commonly used, and can be taken from an ultrasound phantom or a clinical image of anatomical structure. Filtering methods are typically compared by visual inspection of the resulting images.
- The quantification of the changes made by filtering is commonly found by calculating objective image quality metrics. A range of metrics can be applied to capture different aspects of the filtering process:
    1. Distance metrics such as the MSE, SNR and correlation (and derivatives thereof) quantify the mean level of distortion through filtering. These can be applied in the image as a whole, or restricted to selected regions of interest.
    2. Measures which attempt to quantify the desirable properties of the speckle removal process, namely the reduction of variance in homogeneous regions the preservation of mean values within these regions. Examples include are the mean and variance measures used in the SRAD method above, and the related Q metric of Tay *et al.*. Homogeneous regions must be selected beforehand.
    3. Metrics which quantify the preservation of edges by the filtering process; namely the FoM and β metrics.
    4. The contrast to noise ratio (CNR) which measures the contrast in a region of interest relative to that in a background area.
    5. Metrics which attempt to measure preservation of images using a model of criteria important to the human visual system, namely the universal quality index (UQI) and structural similarity index measure (SSIM) of Wang **et al.** [61, 96].

- Many of these metrics are comparative in nature. The question of what to compare filtered simulated images is easily addressed as speckle free reference images are typically available. For real-world images, no such reference is available however. Metrics can be applied so as to measure aspects of filtering relative to the noisy input (e.g. a reduction in local variance and the preservation of image edges by filtering are commonly used as criteria for successful filtering). In the case of phantom images, some authors manually generate a noise free reference (e.g. the NCD method above), although it relation to the underlying structure is unclear.

2.4 In-Depth Description of Speckle Reduction Methods

This section presents a detailed description of a number of post processing speckle reduction methods found in the literature. The speckle reduction techniques detailed here use the most common approaches of anisotropic diffusion, wavelet denoising and local statistics. A nonlinear method, known as the Geometric filter [92, 97], is also described here as it deemed to have good performance in the review of Loizou *et al.* [89]. These methods are evaluated in depth in Chapter 4 of this thesis. Table 2.1 displays a summary of the filters in question, including references for the papers in which they are proposed, and the abbreviations used to refer to them.

As post-processing filters, these methods operate on the observed ultrasound image. As described in Section 2.2 above, the observed image is best described as a random variable, composed of the sum of the reflections from the imaged medium. However, many of the filters investigated here use the simple multiplicative model of [76]. The effect of additive noise is considered negligible, leading to:

$$I(x, y) \;=\; R(x, y)n(x, y) \tag{2.10}$$

where (x, y) are the spatial coordinates, $I(x, y)$ is the observed image, $R(x, y)$ represents the signal or reflectivity, and $n(x, y)$ the speckle noise.

2.4.1 Anisotropic Diffusion Filters

2.4.1.1 Perona and Malik Anisotropic Diffusion (PMAD)

Anisotropic diffusion, developed by Perona and Malik [57], is a method of selectively smoothing an image while preserving edges. Diffusion takes place according to the

Table 2.1: Despeckle Filter Summary. AD = Anisotropic Diffusion, W = Wavelet.

Method	**Type**	**Refrences**	**Abbreviation**
Perona and Malik Diffusion	AD	[57]	PMAD
Speckle Reducing Anisotropic Diffusion	AD	[3]	SRAD
Detail Preserving Anisotropic Diffusion	AD	[60]	DPAD
Coherence Enhancing Diffusion	AD	[98]	CED
Nonlinear Coherent Diffusion	AD	[4]	NCD
Oriented Speckle Reducing Anisotropic Diffusion	AD	[24]	OSRAD
Zong *et al.* Filter	W	[2]	Zong
Generalized Likelihood Method	W	[80]	GLM
Nonlinear Multiscale Wavelet Diffusion	W	[81]	NMWD
Lee Filter	SAR	[42]	Lee
Frost *et al.* Filter	SAR	[45]	Frost
Kuan *et al.* Filter	SAR	[44]	Kuan
Enhanced Lee Filter	SAR	[46]	EnhLee
Enhanced Frost *et al.* Filter	SAR	[46]	EnhFrost
Geometric Filter	-	[97]	Geo

following Partial Differential Equation (PDE):

$$\begin{aligned} \frac{\partial I(x,y;t)}{\partial t} &= \nabla \cdot \{c(|\nabla I_\sigma(x,y;t)|).\nabla I(x,y;t)\} \\ I(x,y;0) &= I_0(x,y) \end{aligned} \tag{2.11}$$

where $I(x,y;t)$ is the image under diffusion, t is an artificial time dimension representing the progress of diffusion, I_0 is the observed image, ∇ and $\nabla \cdot ()$ are the gradient and divergence operators, and $|\cdot|$ represents magnitude. The diffusion function $c(\cdot)$ controls the level of diffusion at each image position. Smoothing is inhibited across image edges by choosing a monotonically decreasing function of gradient magnitude for $c(|\nabla I(x,y;t)|)$, such as:

$$c(x) = e^{-(x/k)^2} \tag{2.12}$$

where k is an edge threshold, set to 90% of the absolute gradient histogram integral [57].

This method has been extended, most notably by Catté *et al.* [58], who regularised the calculation of the diffusion function $c(\cdot)$. This was accomplished by computing the diffusion function using the gradient of a smoothed version of the evolving image, i.e. replacing $c(|\nabla I(x,y;t)|)$ with $c(|\nabla I_\sigma(x,y;t)|)$ in (2.11). Here I_σ is a smoothed version of I, obtained by convolution with a Gaussian of variance σ^2. This makes (2.11) mathematically well-posed, having a unique solution. The value of σ is chosen based on the diffusion time at iteration t, i.e. $\sigma = T$ (where T is the current level of diffusion), so this modification does not require the selection of an additional parameter. While this anisotropic diffusion is capable of intra-region smoothing with edge preservation for images corrupted with additive noise, it has been noted that its effect on images corrupted with multiplicative noise such as speckle is less than satisfactory [3].

2.4.1.2 Speckle Reducing Anisotropic Diffusion (SRAD)

Yu and Acton [3] developed a diffusion approach better suited to multiplicative speckle noise removal, performed according to the PDE:

$$\begin{aligned} \frac{\partial I(x,y,t)}{\partial t} &= \nabla.[c(q).\nabla I(x,y,t)] \\ I(x,y,0) &= I_0(x,y) \end{aligned} \tag{2.13}$$

In contrast to (2.11), the diffusion function $c(.)$ is not a function of the gradient magnitude, but rather of the Instantaneous Coefficient of Variation (ICOV) q. The ICOV is based on the variation coefficient used in SAR filtering methods as a signal/edge discriminator (see Section 2.4.1), and is defined as:

$$q(x,y,t) = \sqrt{\frac{(0.5)\left(|\frac{\nabla I}{I}|\right)^2 - (0.25)^2\left(\frac{\nabla^2 I}{I}\right)^2}{[1+(0.25)\left(\frac{\nabla^2 I}{I}\right)]^2}} \tag{2.14}$$

where ∇^2 is the Laplacian operator. The diffusion function $c(.)$ used here is:

$$c\left[q(x,y,t), q_0(t)\right] = \left(1 + \frac{q^2(x,y,t) - q_0^2(t)}{q^2(x,y,t)(1+q_0^2(t))}\right)^{-1} \tag{2.15}$$

where q_0 is the 'speckle scale function', a diffusion threshold controlling the level of smoothing, equivalent to the noise variation coefficient C_n of the SAR filters. For the purposes of this study the median estimator of [60] is used to calculate q_0, as detailed in the next section.

2.4.1.3 Detail Preserving Anisotropic Diffusion (DPAD)

In [60], a filter was proposed which improves upon the operation of the SRAD filter. Noting that (2.15) is derived from the Lee filter, this is replaced with the following, which is derived from the Kuan filter:

$$c\left[q(x,y,t), q_0(t)\right] = \frac{1 + 1/q^2(x,y,t)}{1 + 1/q_0^2(t)} \tag{2.16}$$

The second alteration made to the SRAD approach concerns the calculation of the ICOV. It was shown that (2.14) is equivalent to the ratio of local standard deviation and mean estimators, similar to C_I in Section 2.4.1. These local estimators were calculated in (2.14) using the four nearest neighbours of each pixel. The approach of [60] is to use a larger neighbourhood for more accurate estimates.

$$q(x,y,t) = \sqrt{\frac{\sigma_I^2(x,y,t)}{\bar{I}(x,y,t)^2}} = \sqrt{\frac{\frac{1}{|\eta_{x,y}|-1} \sum_{p \in \eta_{x,y}} \left(I_p - \bar{I}(x,y,t)\right)^2}{\bar{I}(x,y,t)^2}} \tag{2.17}$$

where $\eta_{x,y}$ is a square $Z \times Z$ neighbourhood, and $\bar{I}(x,y,t) = (1/|\eta_{x,y}|) \sum_{p \in \eta_{x,y}} I_p$. This is shown to be more accurate than the formulation of (2.14). The third contribution of [60] related to the estimation of $q_0(t)$, calculated as $median\{q(x,y;t)\}$. This approach requires less computation than the previously proposed method of [99], and produced similar results.

2.4.1.4 Coherence Enhancing Diffusion (CED)

Weickert's CED method [98] aims to enhance the smooth curves within an image, such as those often present in medical images. Weickert introduced a tensor-valued diffusion function, allowing the level of smoothing to vary with direction. This diffusion tensor is based on the structure tensor, a widely used method for describing flow-like structures.

The gradient of an image $I(x,y)$ is a two-element vector, $\nabla I = (I_x, I_y)$, where I_x, I_y are the gradients in the x and y directions. The phase of this vector is the direction of maximum gradient at that point, and the magnitude is the strength of that gradient. The structure tensor T presents this information as:

$$T \quad = \nabla I \otimes \nabla I^T = \begin{pmatrix} I_x^2 & I_{xy} \\ I_{xy} & I_y^2 \end{pmatrix} \tag{2.18}$$

To make the gradient description robust to small noise fluctuations, local averaging of the observed image is performed as $I_\sigma = K_\sigma * I$, where $*$ represents the convolution operator, and K_σ is a Gaussian kernel of variance σ^2. The gradient ∇I_σ represents only information from image details larger than $O(\sigma)$. Thus σ is referred to as the *noise scale*. The structure tensor is formed using a second level of Gaussian smoothing:

$$T_\rho = K_\rho * (\nabla I_\sigma \nabla I_\sigma^T) = \begin{pmatrix} T_{11} & T_{12} \\ T_{21} & T_{22} \end{pmatrix} \tag{2.19}$$

The values of T_ρ then represent image information from a neighbourhood defined by ρ, the *integration scale*. While the structure tensor is simply another representation of the image gradient ∇I, and contains no more information than ∇I, it has the advantage of allowing local averaging as in (2.19) without cancellation.

The eigenvectors of T_ρ are denoted by $(\vec{\omega}_1, \vec{\omega}_2)$, and the corresponding eigenvalues by (μ_1, μ_2). If the eigenvalues are ordered so that $\mu_1 \geq \mu_2$, then $(\vec{\omega}_1, \vec{\omega}_2)$ give the directions of maximum and minimum local variation, respectively. These are the directions normal and tangent to the local image gradient, referred to as the gradient and contour directions. The corresponding eigenvalues give the strength of the gradient in these directions, and also provide information on the local coherence or anisotropy. A measure of local coherence is defined as $\kappa = (\mu_1 - \mu_2)^2$ [98]. The CED diffusion process is described by the PDE:

$$\frac{\delta I(x,y,t)}{\delta t} = \nabla.(D \nabla I(x,y,t)) \tag{2.20}$$

where D is the diffusion matrix, constructed with the same eigenvectors as T_ρ ($\vec{e}_1 = \vec{\omega}_1, \vec{e}_2 = \vec{\omega}_2$) and eigenvalues (λ_1, λ_2) given by:

$$\begin{aligned} \lambda_1 &= c_1 \\ \lambda_2 &= \begin{cases} c_1, & \text{if } \mu_1 = \mu_2 \\ c_1 + (1 - c_1) \exp\left(\frac{-c_2}{\kappa}\right), & \text{otherwise} \end{cases} \end{aligned} \tag{2.21}$$

where c_1 and c_2 are parameters constrained by $0 < c_1 \ll 1$ and $c_2 > 0$.

2.4.1.5 Nonlinear Coherent Diffusion (NCD)

The NCD method of Abd-Elmoniem *et al.* [4] attempts to simultaneously remove speckle noise and enhance coherent structures in ultrasound images. The filtering model aims to discriminate between different levels of speckle, based on the similarity to fully developed speckle. Image regions closely resembling fully developed speckle are mean filtered, while those dissimilar remain unaltered.

Similar to the CED method, this approach uses a tensor-valued diffusion function, calculated as a component-wise convolution of a Gaussian kernel with the structure tensor:

$$J_\rho = K_\rho * (\nabla I \ \nabla I^T) = \begin{pmatrix} K_\rho * I_x^2 & K_\rho * I_{xy} \\ K_\rho * I_{xy} & K_\rho * I_y^2 \end{pmatrix} \tag{2.22}$$

Here, the initial stage of smoothing performed in (2.19) is not used, so the structure matrix represents gradient information from image details of even the smallest size. As in the CED method, ρ is integration scale, the window size over which the orientation information is averaged.

Eigen decomposition of J_ρ yields eigenvectors $(\vec{\omega}_1, \vec{\omega}_2)$ (giving the gradient and contour directions), and corresponding eigenvalues (μ_1, μ_2) (the strength of the variations in these directions). The diffusion process is described by the PDE (2.20), using a diffusion tensor D constructed so as to have the same eigenvectors as J_ρ, but with eigenvalues λ_1, λ_2 defined as:

$$\begin{aligned} \lambda_1 &= \begin{cases} \alpha\left(1 - \frac{\kappa}{s^2}\right), & \text{if } \kappa \ \leq \ s^2 \\ 0, & \text{otherwise} \end{cases} \\ \lambda_2 &= \alpha \end{aligned} \tag{2.23}$$

where α is a parameter determining the level of smoothing in regions of fully developed speckle, and s^2 is a heuristically chosen 'stopping level'.

2.4.1.6 Oriented Speckle Reducing Anisotropic Diffusion (OSRAD)

Krissian *et al.* [24] extended the SRAD method to a matrix diffusion scheme. This allows speckle adaptive diffusion to vary in strength in the contour and curvature directions. The improvements of the DPAD method are also used in this method, such as the use of a larger window to estimate $q(x, y; t)$, and the median estimation of $q_0(t)$. The OSRAD diffusion function $c(q)$ is based on the Kuan *et al.* filter, as in (2.16).

Krissian *et al.* showed that the local directional variance is related to the local geometry of the image. The extension of the SRAD method to a matrix scheme is performed by finding the local directions of gradient and curvature. This can be performed by the use of the structure tensor, as in [98] and [4], or by using the Hessian matrix as in [66]. In this implementation the structure tensor approach is employed, to facilitate direct comparison with the CED and NCD methods. In the 2-D case, $(\vec{\omega}_1, \vec{\omega}_2)$ are the eigenvectors of T_ρ from (2.19), and so represent the directions of local gradient and curvature. These vectors are used as the basis of the diffusion matrix, D. The eigenvalues of D, which determine the strength of diffusion in the gradient and curvature directions, are given as:

$$\lambda_1 \quad = c_{srad} \quad , \quad \lambda_2 \quad = c_{tang} \tag{2.24}$$

Where c_{srad} is the SRAD diffusion ($c(q)$), and c_{tang} is a constant. Diffusion is then performed as per (2.20).

2.4.1.7 Discretisation Schemes for Anisotropic Diffusion

Anisotropic diffusion methods can be grouped according to the use of either scalar valued diffusion functions (PMAD, SRAD, DPAD), or tensor valued functions (CED, NCD, OSRAD). In both cases, the method of discrete approximation must ensure accurate and stable approximation of the continuous diffusion PDEs. Scalar diffusion is defined in the continuous case by:

$$\begin{aligned} \frac{\partial I(x,y,t)}{\partial t} &= \nabla.[c\nabla I(x,y,t)] \\ I(x,y,0) &= I_0(x,y) \end{aligned} \tag{2.25}$$

where $c(.)$ can be a function of image gradient (as in PMAD) or the ICOV in SRAD and similar methods. All diffusion methods considered here are implemented with central differences used to approximate the gradient. Equation (2.25) can be approximated in the simplest case using a Euler-forward scheme:

$$\frac{I(i,j;t+\tau) - I(i,j;t)}{\tau} = \sum_{n\in\eta_{i,j}} \frac{1}{2}\{c(n;t) + c(i,j,t)\}\{I(n;t) - I(i,j;t)\} \tag{2.26}$$

where (i,j) are the discrete spatial coordinates, τ is the time step, and $\eta_{i,j}$ are the four-neighbours of the pixel at (i,j). This is known as an explicit scheme, as the next value $I(i,j;t+\tau)$ can be found directly from $I(i,j;t)$ without solving a system of equations. It is convenient to write this in matrix-vector notation in the 2-D case by rearranging the $N \times M$ matrix I as a vector of length NM, denoted v. Then (2.26) is equivalent to [100, 101]:

$$v(k; t + \tau) = [J + \tau\,(A_x\{v(k;t)\} + A_y\{v(k;t)\})]v(k;t) \tag{2.27}$$

where k is the spatial coordinate in this vector, A_x and A_y are sparse tridiagonal matrices of size $NM \times NM$, and J is the identity matrix of the same size. A_x and A_y correspond to derivatives along the x and y axes. The elements of A_x, denoted $a_{n,m}$ are given by:

$$a_{n,m} = \begin{cases} \frac{1}{2}\{c(n) + c(m)\}, & \forall n \in \eta_m \\ -\sum\limits_{p \in \eta_m} \frac{1}{2}\{c(p) + c(m)\}, & \text{if } n = m \\ 0, & \text{otherwise} \end{cases} \tag{2.28}$$

where η_n are the nearest two-neighbours of pixel n in the x direction. The elements of A_y are generated in the same manner, excepting the use of neighbours in the y direction. Weickert *et al.* [100] demonstrated that the stability of this scheme is dependant on a number of conditions. A particular condition which must be met is that:

$$\tau \leq \frac{1}{\max\,(A_x + A_y)} \tag{2.29}$$

This places an upper limit on the timestep which may be used. For diffusion functions with a maximum value of unity, the values of τ permissible for stability are $\tau \leq \frac{1}{4}$. As noted in [24], the SRAD diffusion function derived from Kuan's filter (2.16) (used in the DPAD and OSRAD methods here) is unbounded at 0^+. This places a further restriction on the choice of τ for these filtering techniques.

The practical alternative to explicit discretisation schemes is the following semi-implicit scheme, which splits the 2D problem into its 1D components using 'Additive Operator Splitting' (AOS): defined as [100]:

$$v(k; t + \tau) = \frac{1}{2} \sum_{l \in (x,y)} [J - 2\tau A_l\{v(k;t)\}]^{-1} v(k;t) \tag{2.30}$$

The two linear systems in this scheme can be solved efficiently using the Thomas algorithm, a variant of Gaussian elimination for tridiagonal systems. The semi implicit approach is unconditionally stable, even for arbitrarily large values of τ, permitting diffusion to take place in fewer iterations.

An alternate method which does not require rearranging the image matrix into a vector is proposed in [24]. Using the Jacobi approach, the following discretisation is obtained:

$$I(i,j;t) = \frac{I(i,j;t) + \tau \sum_{p \in \eta_{i,j}} \frac{1}{2}[c(i,j;t) + c(n;t)]I(p;t)}{1 + \tau \sum_{p \in \eta_{i,j}} \frac{1}{2}[c(i,j;t) + c(n;t)]} \tag{2.31}$$

This approach has the advantage of parallelisation potential, which iterative solution schemes such as the Thomas algorithm do not.

Diffusion methods which employ a tensor-valued diffusion function, such as CED, NCD and OSRAD, are governed by the PDE (2.20). This can rewritten for two dimensional images as [98]:

$$\frac{\delta I(x,y,t)}{\delta t} = \nabla.(D\nabla I(x,y,t)) = \sum_{n,m=1}^{2} \frac{\partial}{\partial x_n}\{d_{nm}\frac{\partial}{\partial x_m} I(x,y,t)\} \tag{2.32}$$

where x_1, x_2 represent the x and y dimensions, and the elements of D at each image location are denoted as:

$$D = \begin{pmatrix} d_{11} & d_{12} \\ d_{12} & d_{22} \end{pmatrix} \tag{2.33}$$

The explicit discretisation of (2.32) can then be written in matrix vector notation as:

$$v(k;t+\tau) = [J + \tau \sum_{n,m}^{2} A_{nm}]v(k;t) \tag{2.34}$$

where A_{nm} is the approximation to the operator $\frac{\partial}{\partial x_n} d_{nm} \frac{\partial}{\partial x_m}$. The approximations A_{nm} are generated using (2.28) by replacing the scalar diffusivities c with the appropriate d_{nm} values.

Weickert also derived a semi-implicit scheme tensor-valued diffusion schemes, which is given in the two dimensional case by the equations:

$$\begin{aligned} z(k,t) &= [J + \tau\{A_{12} + A_{21}\}]v(k,t) \\ w_1(k,t+\tau) &= [J - 2\tau A_{11}]^{-1} z(k,t) \\ w_2(k,t+\tau) &= [J - 2\tau A_{22}]^{-1} z(k,t) \\ v(k,t+\tau) &= \frac{1}{2}\{w_1(k,t+\tau) + w_2(k,t+\tau)\} \end{aligned} \tag{2.35}$$

The matrix inversions in (2.35) can again be performed using the efficient Thomas algorithm for tridiagonal systems.

2.4.2 Multiscale Wavelet Filters

2.4.2.1 Speckle Reduction Method of Zong et al.

The speckle reduction method of Zong *et al.* [2] models speckle noise according to (2.10). To separate the multiplicative noise component from the noise-free image I, a log transform of (2.10) is taken. Using the notation $I^l = log(I)$, (2.10) becomes:

$$I^l(i,j) = R^l(i,j) + n^l(i,j) \tag{2.36}$$

after transformation to the discrete domain by uniform sampling. Here (i,j) are discrete spatial coordinates. $n^l(i,j)$ is approximated as additive white noise and wavelet based denoising is then applied.

The discrete Dyadic Wavelet Transform (DWT) formulated by Mallat and Zhong [102] is used in this method. This DWT has favourable edge detection properties, and its non-decimated nature allows calculation of useful cross-scale products. The K-level DWT of $I(i,j)$ is represented as: $W[I(i,j)] = \{(W_k^d[I(i,j)])_{1\leq k\leq J}^{d=1,2}, S_K[I(i,j)]\}$ where $W_k^d\{I(i,j)\}$ is the set of wavelet coefficients at scale 2^k (level k) and spatial orientation d ($d = 1$ for horizontal and $d = 2$ for vertical). The approximation coefficients at the coarsest scale K are denoted by $S_K\{I(i,j)\}$. Four decomposition levels are used. Taking the DWT of (2.36) gives $W\{I^l(i,j)\} = W\{R^l(i,j)\} + W\{n^l(i,j)\}$. Assuming the speckle noise lies in the higher spatial frequencies, the method of Zong *et al.* uses a combination of hard and soft thresholding. The soft thresholding method of Donoho [63] is applied to the finer scales (levels one and two). Coefficient dependent thresholds are defined for each scale and orientation as:

$$t_k^d = \begin{cases} T_{max} - \alpha(k-1)\sigma_k^d, & \text{if } T_{min} - \alpha(k-1) > T_{min} \\ T_{min}\sigma_k^d, & \text{otherwise} \end{cases} \tag{2.37}$$

where σ_k^d is the standard deviation of $W\{I_k^d(i,j)\}$, and T_{min} and T_{max} are the minimum and maximum values for σ_k^d. The parameter α is heuristically chosen.

A nonlinear function is applied to the scales three and four. This incorporates hard thresholding and nonlinear contrast enhancement, and is defined by:

$$E(v) = \begin{cases} 0, & \text{if } \quad |v| < T_1 \\ sign(v)T_2 + \bar{u}, & \text{if } \quad T_2 \leq |v| \leq T_3 \\ 1, & \text{otherwise} \end{cases} \tag{2.38}$$

where $v \in [-1, 1]$ represents each wavelet coefficient, and the three thresholds are related as $0 \leq T_1 \leq T_2 \leq T_3 \leq 1$. The value of $\bar{u}$ is determined according to $\bar{u} = a(T_3 - T_2)\{f\left[c(u-b)\right] - f\left[-c(u+b)\right]\}$. The sigmoid function is represented by f, and

$$u = \frac{sign(v)(|v| - T_2)}{T_3 - T_2}, a = \frac{1}{f\left[c(1-b)\right] - f\left[-c(1+b)\right]} \tag{2.39}$$

Thus the operator $E(v)$ is dependant on five parameters: b, c, T_1, T_2, T_3.

2.4.2.2 The Generalized Likelihood Method (GLM)

The GLM of Pižurica *et al.* [80] uses a K-level wavelet transform, calculated using the quadratic spline wavelet. The wavelet coefficients at decomposition scale 2^k are represented as ω_k^d. Three orientations are calculated corresponding to the horizontal, vertical, and diagonal directions ($d \in LH, HL, HH$). The lowpass image at the coarsest level K is represented by u_K. The wavelet transform used is non-decimated, and as with the DWT of [102] this allows calculation of cross-scale products for edge detection.

Central to this method is a preliminary classification step which estimates which wavelet coefficients represent signal (and which noise), based on the assumption of [103] that useful wavelet coefficients persist well across the scales of decomposition. A binary mask $\hat{x}_k^d \in \{0, 1\}$ is calculated, in which $\hat{x}_k^d = 1$ for pixels considered to represent signal:

$$\hat{x}_k^d = \begin{cases} 0, & \text{for } |\omega_k^d||\hat{y}_{k+1}^d| \; < \; (\alpha\hat{\sigma}_k^d)^2 \\ 1, & \text{for } |\omega_k^d||\hat{y}_{k+1}^d| \; \geq \; (\alpha\hat{\sigma}_k^d)^2 \end{cases} \tag{2.40}$$

where $\hat{y}_{k+1}^d$ is the estimate for the corresponding denoised coefficient at the next-coarser scale. To initialize the method, $\hat{y}_K^d = \omega_K^d$ for all d. The thresholds are composed of α, a heuristically chosen parameter, and $\hat{\sigma}_k^D$, an estimation of the noise level in ω_k^d.

Once this classification has been performed, conditional densities ξ_k and μ_k are estimated as:

$$\xi_k^d = \frac{\hat{p}_{m_k^d|x_k^d}(m_k^d|1)}{\hat{p}_{x_k^d|x_k^d}(m_k^d|0)}, \quad \mu_k^d = \hat{r}\frac{\hat{p}_{e_k^d|x_k^d}(e_k^d|1)}{\hat{p}_{e_k^d|x_k^d}(e_k^d|0)} \tag{2.41}$$

Here, m_k^d denotes the wavelet coefficient magnitudes, (i.e. $m_k^d = |\omega_k^d|$), and e_k^d is the *local spatial activity indicator*, calculated as a windowed average of $(\omega_k^d)(\omega_{k+1}^d)$. Further details regarding calculation these quantities can be found in [80]. The estimated ratio of prior probabilities $\hat{r}$ represents the ratio of pixels in ω_k^d deemed to represent signal in the original image, and is calculated as:

$$\hat{r} = \frac{\sum_{k=1}^{N} \widehat{x_k^d}}{N - \sum_{k=1}^{N} \widehat{x_k^d}} \tag{2.42}$$

The noise free coefficients are then estimated as:

$$\hat{y}_k^d = \frac{\xi_k^d \mu_k^d}{1 + \xi_k^d \mu_k^d} \omega_k^d \tag{2.43}$$

2.4.2.3 Nonlinear Multiscale Wavelet Diffusion (NMWD)

The NMWD method [81] aims to combine wavelet analysis with anisotropic diffusion. Wavelet based signal/noise discrimination is employed to overcome the shortcomings of the image gradient (as used in the PMAD, CED and NCD methods) for discrimination in speckled images. The gradient cannot always precisely separate the image and noise in ultrasound images, as variations due to speckle noise may be larger than those corresponding to underlying image [81]. The image is decomposed using the DWT of Mallat and Zhong, as in the Zong *et al.* method. The modulus of the wavelet coefficients at each scale is defined as:

$$M_k\{I(i,j)\} = \sqrt{\sum_{d=1}^{2} |W_k^d\{I(i,j)\}|^2_{k=1,2,..,K}} \tag{2.44}$$

The normalised wavelet modulus was found to be large in edge-related regions and small for noise and texture, and so is used for signal/noise discrimination in this method. For amplitude images the wavelet modulus is normalised by:

$$\tilde{M}_k I = \frac{M_k\{I(i,j)\}}{\mu_Z}, \quad k = 1, 2, .., K \tag{2.45}$$

Here μ_Z is a local mean calculated using a widow size $Z \times Z$. The histogram of $\tilde{M}_k I$ is modelled as a Rayleigh mixture distribution, composed of the sum of edge and noise pixels as:

$$\tilde{M}_k I \simeq n_k p_{n,k}(x, \lambda_n) + (1 - n_k) p_{e,k}(x, \lambda_e) \tag{2.46}$$

where $p_{n,k}$ is a Rayleigh distribution with parameter λ_n, representing the noise pixels. Similarly, $p_{e,k}$ are the edge pixels, Rayleigh distributed with parameter λ_e. The proportion of the mixture distribution resulting from noise-related values is given as n_k. The parameters of the normalised modulus distribution ($\sigma_{n,k}$, $\sigma_{e,k}$ and n_k) are estimated using the Expectation-Maximisation (EM) method [104]. These parameters are used to determine the set of homogeneous pixels in each scale U_k, as:

$$U_k = \begin{cases} 1, & \text{if } \left[(1 - U_{k+1})\tilde{M}_{k+1} I\right] \tilde{M}_k I < K^2 T_k T_{k+1} \\ 0, & \text{otherwise} \end{cases} \tag{2.47}$$

where K is a heuristically chosen parameter. T_k is a threshold value for each scale, determined by the parameters of the Rayleigh mixture model as:

$$T_k = \sqrt{\frac{2(log\frac{\sigma^2_{n,k}}{\sigma^2_{e,k}}) + \frac{n_k}{1-n_k}}{\left|\frac{1}{\sigma^2_{e,k}} - \frac{1}{\sigma^2_{n,k}}\right|}} \tag{2.48}$$

To initialise the calculation of (2.47), it is assumed that the coarsest scale K contains only edge pixels, i.e. $U_J = 0$. A diffusion function is calculated for each scale as:

$$g_k(\tilde{M}_k I) \quad = \quad \begin{cases} 1, & \text{if } \eta_k \leq 0 \\ 1 - exp\left[\dfrac{-3.315}{\left(\tilde{M}_k I/\lambda_k\right)^4}\right], & \text{otherwise} \end{cases} \tag{2.49}$$

where λ_k is the diffusion threshold calculated as:

$$\lambda_k = \frac{Mean(U_k \tilde{M}_k)}{\sqrt{2}k'} \tag{2.50}$$

The value of k' is empirically chosen as $k' = 0$ for $k = 1$ and $k' = k$ for $k \geq 2$. The removal of noise from the wavelet coefficients uses the regularisation function $p_k\left(\tilde{M}_k I\right) = 1 - g_k\left(\tilde{M}_k I\right)$. The regularisation process is performed as:

$$\tilde{W}^d_k\{I(i,j)\} = p_k\left(\tilde{M}_k I\right) W^d_k\{I(i,j)\} \tag{2.51}$$

The steps involved in this filtering method are performed in the following sequence:

1. Wavelet decomposition using the DWT of [102];
2. Calculation of normalised moduli $\tilde{M}_k\{I(i,j)\}$;
3. Estimation of Rayleigh mixture parameters using EM method. These parameters are then used to calculate thresholds T_k (2.48), which in turn are used in calculation of homogeneous regions U_k (2.47);

4. Computation of diffusion thresholds λ_k, and diffusion functions $g_k(\tilde{M}_k I)$;
5. Regularisation of wavelet coefficients and reconstruction using inverse DWT.

This is an iterative method, and the steps above are repeated to achieve the desired level of filtering. Step 3 is only performed during the first iteration.

2.4.3 Local Statistics SAR Filters

Synthetic Aperture Radar (SAR) suffers speckle degradation similar to ultrasound imaging, and so methods proposed for SAR speckle reduction are also applicable to medical ultrasonics. Five such methods are considered here, based on a multiplicative model of speckle noise as in (2.10) [46, 105]. These methods assume statistical independence of $n(x,y)$ and $R(x,y)$ [105, 106]. These filters make the assumption that the ratio of noise standard deviation to mean is constant throughout the image. This simplification is only strictly true in some situations. The spatially-varying mean and variance of the observed image are denoted $\overline{I(x,y)}$ and $\sigma_I^2(x,y)$. The coefficients of variation of the image and noise are given as:

$$C_I(x,y) \;=\; \frac{\sigma_I(x,y)}{\overline{I(x,y)}}, \quad C_n \;=\; \frac{\sigma_n}{\bar{n}} \tag{2.52}$$

$C_I(x,y)$ is known to be an effective descriptor of textural information and image homogeneity [107]. The filters in this section operate as a test based on this descriptor. In this evaluation we use the approach of [60], and estimate C_n^2 as the median of $C_I^2(x,y)$ over the image.

2.4.3.1 The Lee filter

The multiplicative Lee filter [42] approximates (2.10) with a linear model to obtain the signal estimate $\hat{R}$. The formulation of [46] is presented here:

$$\hat{R}(x,y) \;=\; I(x,y)W(x,y) \;+\; \overline{I(x,y)}\{1 - W(x,y)\} \tag{2.53}$$

where $W(x,y)$ is a weighting function given by:

$$W(x,y) \;=\; 1 - C_n^2/C_I^2(x,y) \tag{2.54}$$

2.4.3.2 The Kuan et al. filter

The filter proposed by Kuan *et al.* [44] is derived by transforming (2.10) into a signal-dependant additive noise formulation instead of the linear approximation used in the Lee filter. The same general form as the Lee filter (2.53) is used, but with a weighting function given by:

$$W(x,y) \;=\; \frac{1 - C_n^2/C_I^2(x,y)}{1 + C_n^2} \tag{2.55}$$

2.4.3.3 The Frost et al. Filter

The filter proposed by Frost *et al.* [45] estimates the noise-free image by convolving the observed image with a spatially-varying kernel as $\hat{R}(x,y) = I(x,y) * m(x,y)$.The kernel $m(x,y)$, centred at the pixel at location (x_0, y_0), is given as:

$$m(x,y) = K_1 exp\left(-KC_I^2(x_0,y_0)|x,y|\right) \tag{2.56}$$

where the parameter K controls the dampening rate, $|x,y|$ represents the distance of each pixel within the window to (x_0, y_0), and K_1 is a normalizing constant. This method does not require an estimate for C_n. Parameter K must be chosen such that mean filtering is performed for homogeneous regions and filtering is inhibited at edges.

2.4.3.4 The Enhanced Lee and Frost Filters (EnhLee, EnhFrost)

Lopes *et al.* [46] expands the Lee and Frost *et al.* approaches by adapting to local heterogeneity. Based on the C_I, the image is split into three regions of different filtering, using thresholds C_n and C_2, as:

$$\hat{R}(x,y) = \begin{cases} \overline{I(x,y)}, & \text{for } C_I(x,y) \leq C_n \\ \overline{I(x,y)}W(x,y) + I(x,y)(1 - W(x,y)), & \text{for } C_n < C_I(x,y) < C_2 \\ I(x,y), & \text{for } C_I(x,y) \geq C_2 \end{cases} \tag{2.57}$$

Threshold C_2 is chosen to reflect how much of the image is to remain unfiltered. The weighting function $W(x,y)$ is given by:

$$W(x,y) \;=\; exp\left(\frac{-K\left(C_I(x,y) - C_n\right)}{C_2 - C_I(x,y)}\right) \tag{2.58}$$

As with the Frost *et al.* filter, K is an empirically determined damping constant. The enhanced version of the filter of Frost *et al.* is given by the impulse response:

$$m(x,y) \; = \; K_1 exp\left(-K\phi\{C_I^2(x_0,y_0)\}|x,y|\right) \tag{2.59}$$

where $\phi\{C_I^2(x_0,y_0)\}$ is a function given by:

$$\phi\{C_I^2(x_0,y_0)\} = \begin{cases} 0, & \text{for } C_I(x_0,y_0) \leq C_n \\ \dfrac{C_I(x_0,y_0)-C_n}{C_2-C_I(x_0,y_0)}, & \text{for } C_n \leq C_I(x_0,y_0) \leq C_2 \\ \infty, & \text{for } C_I(x_0,y_0) > C_2 \end{cases} \tag{2.60}$$

2.4.4 Geometric Filter

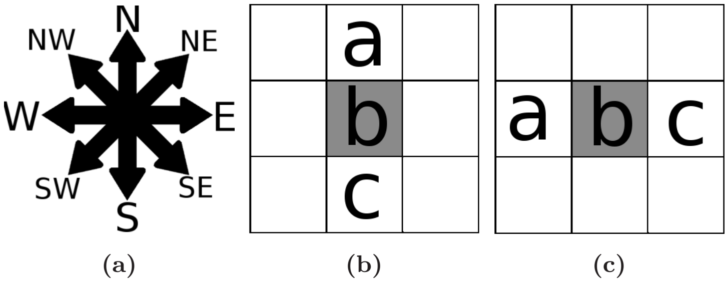

(a) (b) (c)

Figure 2.3: Geometric filter: (a) pixel directions, (b) selected pixels for the North-South direction, (c) selected pixels for the West-East direction.

The geometric filter of Crimmins [97], and improved by Busse *et al.* [92], is derived from geometric concepts. An image is viewed as a surface in which features appear as wide high regions, and speckle noise as narrow walls and valleys to be flattened. Each pixel is considered with two of its eight-neighbours, in the horizontal, vertical, and diagonal directions (the N-S, W-E, NW-SE and NE-SW directions of Fig. 2.3(a)), according to:

1. The N-S direction is selected, as in Fig. 2.3(a).

2. Pixel b is adjusted by the following rules [89]:

$$\begin{array}{llll}
\text{if} & a \geq b+2 & \text{then} & b = b+1 \\
\text{if} & a > b \text{ and } b \leq c & \text{then} & b = b+1 \\
\text{if} & c > b \text{ and } b \leq a & \text{then} & b = b+1 \\
\text{if} & c \geq b+2 & \text{then} & b = b+1 \\
\text{if} & a \geq b-2 & \text{then} & b = b-1 \\
\text{if} & a < b \text{ and } b \geq c & \text{then} & b = b-1 \\
\text{if} & c < b \text{ and } b \geq a & \text{then} & b = b-1 \\
\text{if} & c \geq b-2 & \text{then} & b = b-1
\end{array}$$

3. Repeat steps 1 and 2 for W-E (Fig. 2.3(c)), WN-SE and NE-SW directions.

4. Steps 1 to 3 are performed for all pixels, and the process is repeated for a given number of iterations to achieve the desired level of filtering.

2.5 Concluding Remarks

This chapter has examined a number of topics of interest to the development of image processing techniques applied to speckled ultrasound imagery. Ultrasound image acquisition has been described in detail, including both the physical basis for the imaging methodology, and also the hardware and signal processing necessary to perform scanning. Speckle itself is also examined. The relationship between speckle and the properties of the scattering medium are presented, and the statistics of speckle are described.

A review of the methods contributed to the literature for the filtering of speckle is presented, covering the range of approaches commonly used. The question of how best to evaluate the performance of speckle filtering methods is addressed in this thesis, and so the evaluation strategies used in these prior works are also described. A detailed technical description of a number of speckle reduction methods from the literature was presented.

A number of reviews which compare these various groups of speckle filters were reported in this chapter. Common approaches to evaluating performance of speckle filters include their application to simulated and real world images, and quantification of filtering effects through the use of objective image quality metrics. A small number of methods also evaluate the effects of speckle filtering through the use of expert evaluation. The relationship between objective image quality metrics and the subjective clinical opinion of experts has not been the focus of research in the literature. The following chapter presents a study into this relationship between these objective and subjective assessments of clinical image quality.

Chapter 3

Objective and Subjective Quality Assessment in Echocardiography

3.1 Introduction

This chapter presents a study on the evaluation of the performance of speckle reduction filtering in clinical images. As detailed in Chapter 2 above, methods of speckle filtering are commonly evaluated by application to clinical images. These images can be assessed subjectively by manual visual inspection, and objectively by comparing the filtered and input images using image quality metrics. The study presented in this chapter explores the relationship between these two approaches, for the case of clinical echocardiography.

In the study presented here, subjective visual assessment is performed by a group of six clinical experts, all of whom are experienced cardiac physicians or technicians. The majority of the evaluation strategies described in the literature review of the previous chapter report results of visual inspection by the authors themselves, rather than the opinion of clinical experts. Exceptions to this trend can be found in [2, 89]. The basic qualities which are generally held to constitute favourable speckle filtering (homogeneous variance reduction, mean intensity preservation and edge preservation) can be readily determined visually, and do not require clinical expertise or training. However, the assessment of clinical images for diagnostic purposes does require such training and expertise. The study presented in this chapter explores the clinical opinion of the quality of speckle filtered echocardiographic images, as judged by experts in the field.

A large set of speckle filtered echocardiographic videos are produced by application of a number of speckle reduction filters. The quality of these filtered videos is evaluated using a number of objective image quality metrics. Videos were evaluated to quantify each expert's opinion on the change in speckle level, change in the clarity of diagnostically important details, and change in overall quality due to the filtering. Overall trends in these areas are examined, along with an investigation of inter-expert difference in video evaluation. The relationship between each evaluation category is also determined for each expert. These two data sets are then compared statistically to investigate relationships between the objective and subjective quality evaluations. The comparison of, and relationships between, subjective expert opinion with objective quality metrics has not been addressed in the literature to date.

The organisation of this chapter is as follows: Section 3.2 discusses the speckle reduction methods used in this study. The objective image quality metrics used to quantify the effects of filtering are presented in Section 3.3. Section 3.4 details the criteria used in expert assessment of the speckle reduced videos. The statistical methods used to analyse the expert assessment data, and to compare these with the objective image metrics, are also presented in this section. Section 3.5 presents the results of this study, while Section 3.6 concludes the chapter.

3.2 Speckle Reduction Methods

Four speckle reduction methods are chosen to filter the test set of echocardiographic videos. A thorough description of a large set of speckle reduction filters, including those used here, is presented in Section 2.4.

The choice of filters to perform speckle reduction for this study was made after some experimentation. The time consuming nature of subjective testing placed a practical limit on the number of filtering methods which could be considered. The filters chosen were:

- The SRAD filter proposed by Yu and Acton [3], described in Section 2.4.1.2 above;
- The NCD filter of Abd-Elmoniem *et al.* [4], which is described in Section 2.4.1.5;
- The GLM filter proposed by Pižurica *et al.* [80]. This is described in detail in Section 2.4.2.2 above;

- The NMWD method of Yue *et al.* [81], the details of which are provided in Section 2.4.2.3.

Example echocardiographic frames from a short and long axis video are displayed in Figs. 3.1-3.2. The original images and the results of filtering by these speckle reduction methods considered are displayed. Differing levels of speckle removal can be seen, as well as different image characteristics. The speckle reduction filters chosen are based on multiscale wavelet and anisotropic diffusion approaches, two of the most widely used methodologies in the recent literature.

The speckle reduction filters were applied to clinical echocardiographic videos, to produce a set of forty eight filtered videos. These videos were taken from the transthoracic short and long axis views, and were supplied by the Cardiology Department at University College Hospital Galway, Ireland. The videos are completely anonymised, with no subject information provided, and ethical approval was obtained for their usage. Scanning is performed with a General Electric Vivid 7 Series scanner (GE Healthcare, Piscataway, NJ, USA). The videos are subject to log compression, however no image optimisation or speckle reduction methods (such as compounding from multiple angles) are performed.

3.3 Quality Assessment using Objective Metrics

A set of image quality metrics are calculated to objectively quantify the performance of the speckle filters. Metrics are calculated on a frame-by-frame basis, and the metrics for each video are found as the average over all frames.

3.3.1 Pratt's Figure of Merit (FoM)

The FoM [59] measures edge pixel displacement between each filtered image I_{filt} and the original image I_{orig}:

$$FoM(I_{filt}, I_{orig}) = \frac{1}{max(N_{filt}, N_{orig})} \sum_{i=1}^{\hat{N}} \frac{1}{1 + d_i^2 \alpha} \tag{3.1}$$

where N_{filt} and N_{orig} are the number edge pixels in edge maps of I_{filt} and I_{orig}. Parameter α is set to a constant $\frac{1}{9}$ [3], and d_i is the Euclidean distance between the ith detected edge pixel and the nearest ideal edge pixel. Similar to assessments included in [3, 24], the edge maps are found using the Canny edge detector [108], with a standard deviation of 0.1 and a threshold of 0.5. The FoM metric thus measures the quality of the contours in the filtered image relative to the original image, i.e. how well the edges are preserved through out the

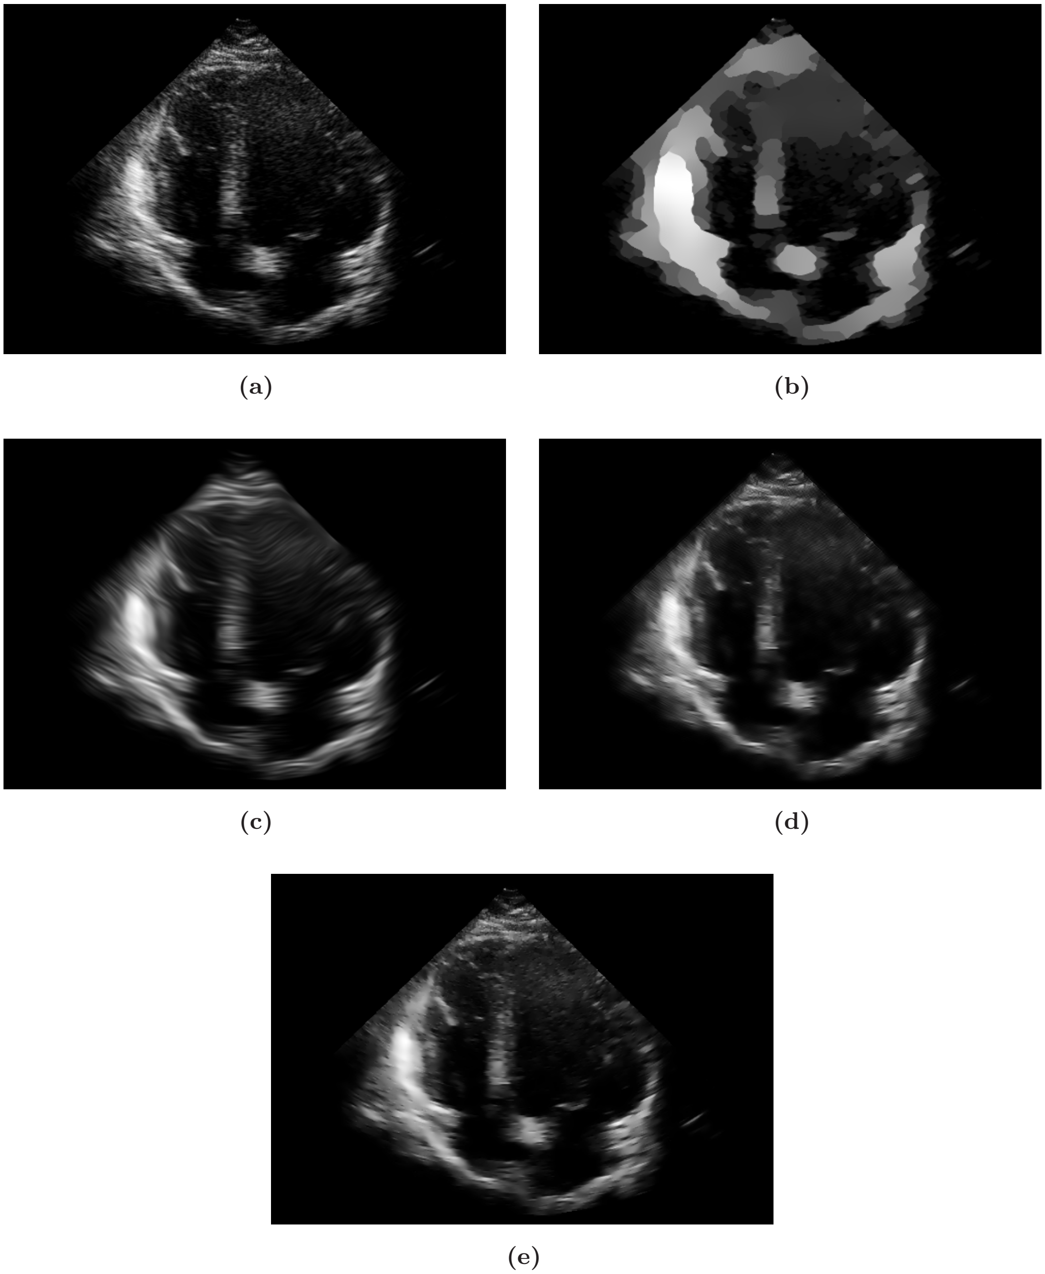

Figure 3.1: Example frame from an echocardiographic video, from the long axis view. (a) Unfiltered, (b) SRAD, (c) NCD, (d) GLM, (e) NMWD.

filtering process. The FoM has a range between 0 and 1, with unity representing perfect edge preservation.

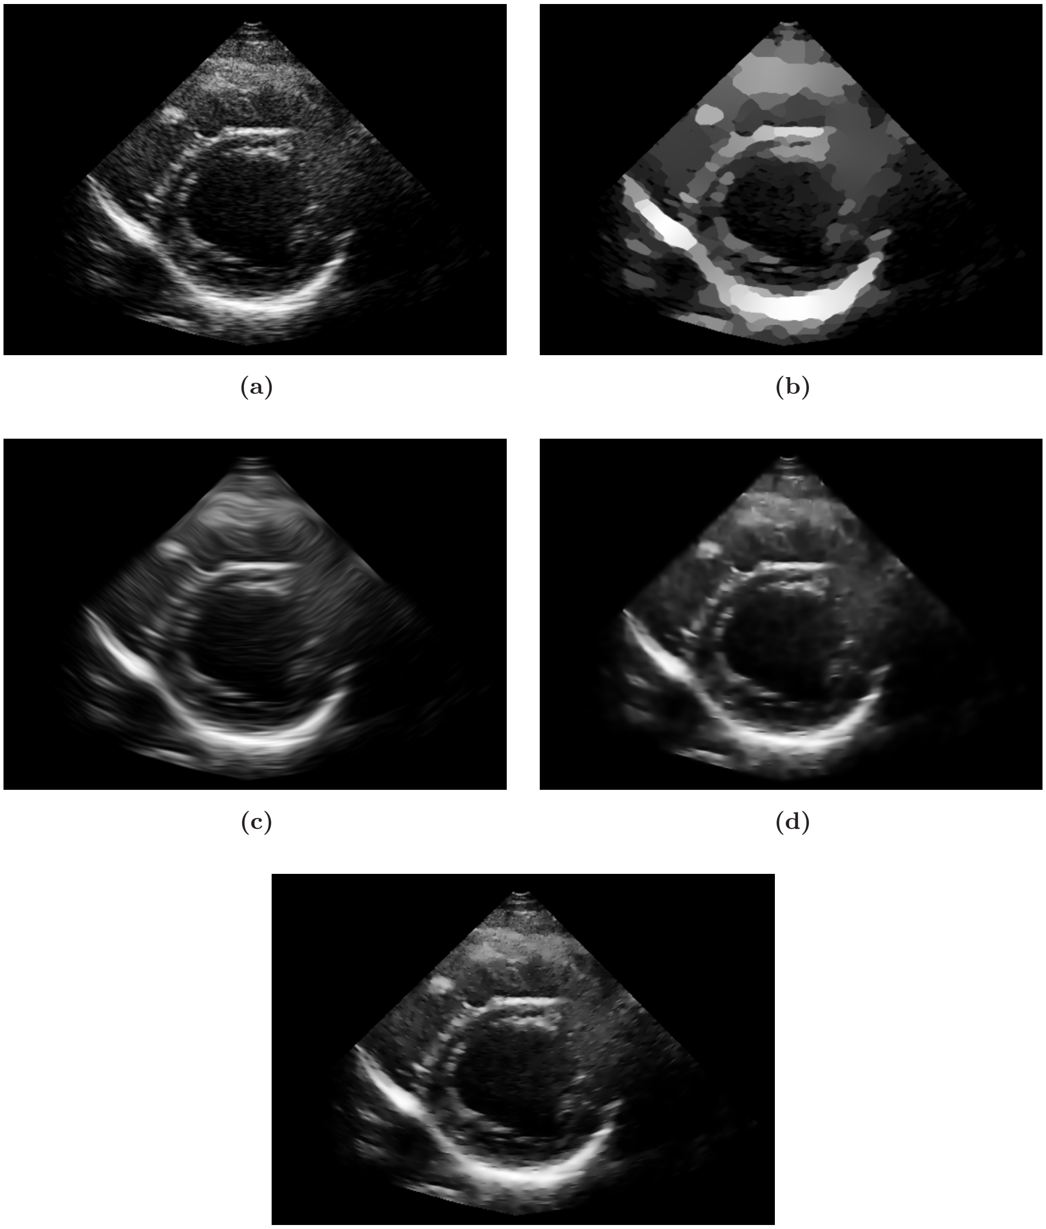

(a) **(b)** **(c)** **(d)** **(e)**

Figure 3.2: Example frame from an echocardiographic video, from the short axis view. (a) Unfiltered, (b) SRAD, (c) NCD, (d) GLM, (e) NMWD.

3.3.2 Mean Squared Error (MSE)

The MSE measures the average absolute difference between two the original and filtered images:

$$MSE(I_{filt}, I_{orig}) = \frac{1}{XY}\sum_{i=1}^{Y}\sum_{j=1}^{X}(I_{filt}(i,j) - I_{orig}(i,j))^2 \tag{3.2}$$

where both images are of size $X \times Y$. Higher values of this metric correspond to greater levels of image distortion due to filtering.

3.3.3 Edge Region MSE

The edge region MSE measures the average difference in edge regions considering only pixels close to image edges. It is calculated using (3.2) above, excepting that only pixels in the vicinity of image edges are considered. A number of edge detection approaches for speckled images are available, including [109–113]. In this study the method of [109] is employed. Once a binary edge map is found, a morphological dilation operation is performed using a kernel of size three. This produces a binary mask which selects those pixels close to image edges. Fig. 3.3 displays an example of this edge region selection method.

3.4 Quality Assessment by Expert Analysis

Six cardiac technicians, experienced in the analysis and interpretation of echocardiographic videos, assessed the set of videos. The three subjective criteria for evaluation were based on the opinion of a senior clinical expert as to what constitute important clinical factors:

Speckle Level This quantifies the expert's assessment of the level of speckle in each video.

Detail Clarity Quantifies the subjective resolvability of diagnostically important details.

Overall Quality This quantifies the overall quality of the video, including any other clinical considerations not covered by the other criteria.

Each expert was shown the filtered and original unfiltered videos in random order, and asked to assign a score in each category on a scale of one to ten.

Statistical analysis of these scores is performed in order to determine:

1. If there are any significant differences between the experts in each of the scoring categories.
2. If there are any statistically significant relationships between the three scoring categories for each expert.
3. If there are significant relationships between the subjective expert scores and any of the image quality metrics.

The statistical methods used to examine the data are determined by the nature of the data being investigated. Specifically, the expert scores are an ordinal categorical data set, which restricts the applicable methods. To investigate if there are inter-expert differences in each of the scoring categories, the non-parametric Kruskal-Wallis test is employed [114]. For each of the three assessment categories, the scores of the six experts represent independent groups of sample data. The Kruskal-Wallis test determines if these sample groups can be said to belong to populations with the same median. If this is the case, there is said to be no difference between the experts in that scoring category, i.e. the

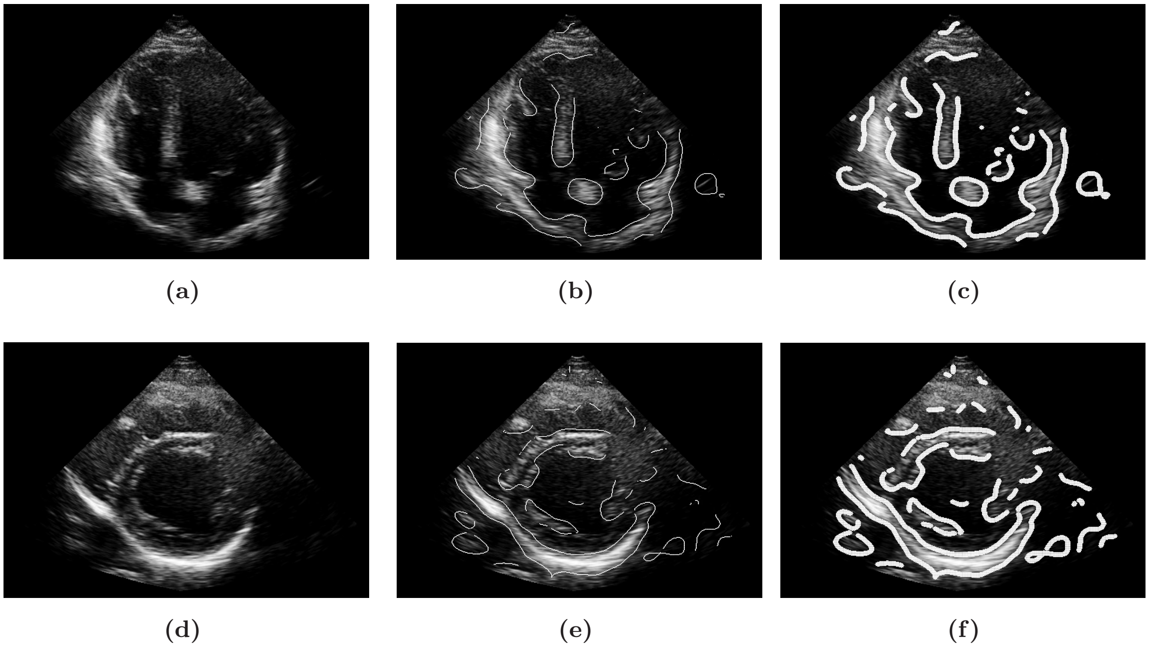

Figure 3.3: Example of edge region selection for the Edge Region MSE metric. (a),(d) Unfiltered, (b),(e) Detected edges, (c),(f) Resulting edge regions.

null hypotheses are that there are no differences between experts in the scoring categories. This technique is analogous to the commonly used one-way analysis of variance (ANOVA), which tests if groups of sample quantitative data belong to populations with the same mean. ANOVA assumes that the population is normally distributed however, making it unsuitable for ordinal data.

In the second and third analyses above, the goal is to determine if there are any relationships between variables: between each of the scoring categories, and between the scoring categories and the quality metrics. These are both examined individually for each expert. To determine if there is a relationship between two quantitative variables, the Pearson correlation coefficient is commonly used [115]. Due to the ordinal nature of the data considered here however, Spearman's rank correlation coefficient (ρ) is used instead [116]. This is a non-parametric measure of statistical dependence. Thus when examining the relationship between the scores for each expert, the null hypothesis is that there is no relationship between those scores. Similarly, when examining the relationship between the expert scores and the image quality metrics, the null hypothesis is that there is no relationship exists. These tests are performed using the SPSS software package, and the level of significance is chosen as 1% ($p = 0.01$) throughout.

3.5 Results

The filtered echocardiographic videos described in Section 3.2 are evaluated by the group of experts, as detailed in Section 3.4, and also by calculation of the image quality metrics as presented in Section 3.3 above.

An initial analysis of the global trends in the expert scores is performed. Table 3.1 displays a breakdown of each expert score into three categories, corresponding to an increase, decrease, or no change. These are measured for each expert relative to the same expert's score for the unfiltered version of that video, and scores are aggregated over all experts to determine general trends. Almost two thirds of Speckle Level scores indicate a reduction in perceived Speckle Level due to filtering. In 36% of cases, the Overall Quality before and after filtering was deemed unchanged, and in a very small number of cases Overall Quality was judged to be improved by speckle reduction. A similar distribution is observed for the perceived Detail Clarity, with less than 10% exhibiting an increase.

To investigate differences between the experts in each of the scoring categories, a graphical summary of the data is presented in the box-plots of Fig. 3.4. In these diagrams, the box is bounded by the lower quartile from above and the upper quartile from below (lowest 25% and 75% of the data, respectively). The whiskers extending from these boxes reach to the extrema of the data. A degree of inter-expert variability in all three scores is observed. To test the null hypotheses that there are no differences in expert scores in each of the categories, the Kruskal Wallis test is applied. The alternative hypothesis in each case is the logical negation of the null hypothesis, i.e. that there are significant inter-expert differences. Results of the Kruskal-Wallis test are presented in Table 3.2. It can be seen that the test statistic H is larger for the Detail Clarity and speckle level scoring categories, than for the Overall Quality category. This agrees with what can be observed in the distributions of Fig. 3.4, i.e. the Overall Quality scores are seen to have a greater degree of uniformity between experts. The significance values p in this table result in the acceptance of the null hypothesis for the Overall Quality scoring category, i.e. there is no statistically significant difference between the experts scoring of overall quality at a 1% level of significance. The significance of the Kruskal-Wallis H statistic for the Detail Clarity and Speckle Level scoring categories is seen to be in the critical region (i.e. less than $p = 0.01$) and so in these cases the null hypotheses are rejected, and the alternative hypotheses accepted. Thus significant differences are found between the experts in scoring the videos in the Detail Clarity and Speckle Level categories, at a 1% significance level.

The investigation of relationships between the scoring categories for each of the experts using Spearman's method produces values displayed in Table 3.3. The test statistic ρ and corresponding significance values p quantify the strength of relationship between the scoring categories. The value of ρ is bounded by $1 \leq \rho \leq -1$, where $\rho = 1$ and $\rho = -1$ correspond respectively to strong positive and negative relationships. For all of the experts, the relationship between Overall Quality and Detail Clarity is strongly positive, ranging from $0.72 \geq \rho \geq 0.95$. The significance values here are in all cases lower than the 1% significance of the test, placing the test statistic in the critical region in all cases. Thus the individual null hypotheses for the each of the experts are rejected for the relationship between Overall Quality and Detail Clarity, i.e. a statistically significant relationship is found between the Overall Quality and Detail Clarity scores for all experts. The ρ values calculated for the relationship between Overall Quality and Speckle Level show positive relationships in all cases. These are in general weaker than the relationships between the Overall Quality and Detail Clarity as quantified by this statistic. The significance values indicate place the test statistic ρ in the critical region for four of the six experts, therefore the null hypothesis is rejected in these cases. Thus in the case of four experts a statistically significant relationship between the Overall Quality and Detail Clarity scoring categories are observed at the 1% significance level, while for the other two experts no significant relationship is observed at this level. Finally, ρ values quantifying the strength of the relationship between the Detail Clarity and Speckle Level again indicate a positive relationship, and in general the strengths are compatible to those reported for the relationship between Overall Quality and Speckle Level. The significance levels of these test statistics are lower than the 1% level for all but one of the experts. The null hypothesis is therefore rejected for five of the six experts, i.e. there are significant relationships observed between Detail Clarity and Speckle Level scores at the 1% significance level for five of the six experts.

The second type of relationship examined are those between the expert scores and the objective image quality metrics. Again, Spearman's method is performed, and the resulting test statistics and significance values are displayed in Table 3.3. The relationships investigated are those between Overall Quality and FoM, Detail Clarity and Edge Region MSE, and Speckle Level and Overall MSE. A strong positive relationship is observed between the FoM metric and the Overall Quality scoring category. Strong negative relationships are observed between the Detail Clarity category and Edge Region MSE metric, and between the Speckle Level category and the Overall MSE metric. In all cases, the significance values are less than the 1% level of the test, putting the test statistic in the critical region. The null hypotheses are therefore rejected in all of these cases, i.e. for all experts, a statistically significant relationship is observed at a 1% level of significance between the Overall Quality and FoM, the Detail Clarity and edge region MSE, and between the Speckle Level and MSE.

Table 3.1: Aggregate Distribution of Expert Scores, N=288.

	Increase	**No Change**	**Decrease**
Overall Quality	16 (5.6%)	104 (36.1%)	168 (58.3%)
Detail Clarity	22 (7.6%)	113 (39.3%)	153 (53.1%)
Speckle Level	15 (5.2%)	86 (29.9%)	187 (64.9%)

Table 3.2: Inter-expert Kruskal-Wallis Parameter Values (H), with associated significance (p).

	Overall Quality	**Detail Clarity**	**Speckle Level**
Kruskal-Wallis	$H = 17.76$	$H = 37.19$	$H = 51.53$
H values	NS	$p = 547.8\times10^{-9}$	$p = 673.9\times10^{-12}$

NS=Not Significant

Table 3.3: Intra-expert association between scoring categories, using Spearmans correlation (ρ), with significance.

	Expert 1	**Expert 2**	**Expert 3**	**Expert 4**	**Expert 5**	**Expert 6**
Overall Quality/	ρ=0.947	ρ=0.888	ρ=0.715	ρ=0.936	ρ=0.783	ρ=0.825
Detail Clarity	$p=3.97\times10^{-15}$	$p=365.2\times10^{-19}$	$p=113.2\times10^{-10}$	$p=195.9\times10^{-24}$	$p=474.0\times10^{-13}$	$p=487.6\times10^{-20}$
Overall Quality/	ρ=0.625	ρ=0.355	ρ=0.609	ρ=0.217	ρ=0.571	ρ=0.759
Speckle Level	$p=156.1\times10^{-8}$	NS	$p=411.8\times10^{-8}$	NS	$p=225.5\times10^{-7}$	$p=627.4\times10^{-16}$
Detail Clarity/	ρ=0.644	ρ=0.454	ρ=0.534	ρ=0.233	ρ=0.493	ρ=0.777
Speckle Level	$p=389.6\times10^{-9}$	$p=117.7\times10^{-5}$	$p=935.7\times10^{-7}$	NS	$p=374.9\times10^{-6}$	$p=514.8\times10^{-16}$

NS=Not Significant

Table 3.4: Intra-expert association between scoring categories and metrics, using Spearmans correlation (ρ), with significance.

	Expert 1	**Expert 2**	**Expert 3**	**Expert 4**	**Expert 5**	**Expert 6**
Overall Quality/	ρ=0.74	ρ=0.59	ρ=0.53	ρ=0.72	ρ=0.55	ρ=0.82
FOM	$p=176.5\times10^{-11}$	$p=123.0\times10^{-7}$	$p=124.1\times10^{-6}$	$p=607.0\times10^{-11}$	$p=452.7\times10^{-7}$	$p=992.7\times10^{-15}$
Detail Clarity/	ρ=-0.67	ρ=-0.62	ρ=-0.49	ρ=-0.76	ρ=-0.49	ρ=-0.83
Edge Region MSE	$p=164.0\times10^{-9}$	$p=240.0\times10^{-8}$	$p=402.0\times10^{-6}$	$p=346.8\times10^{-12}$	$p=332.9\times10^{-6}$	$p=205.6\times10^{-15}$
Speckle Level/	ρ=-0.69	ρ=-0.47	ρ=-0.85	ρ=-0.47	ρ=-0.64	ρ=-0.87
MSE	$p=531.1\times10^{-10}$	$p=726.0\times10^{-6}$	$p=306.8\times10^{-16}$	$p=710.1\times10^{-6}$	$p=992.1\times10^{-9}$	$p=214.3\times10^{-17}$

3.6 Conclusions

This chapter has presented a study comparing objective image quality metrics and subjective expert opinion on the quality of speckle reduced echocardiography. Clinical echocardiographic videos were filtered using four speckle filtering techniques from the literature. A set of forty eight speckle filtered videos were produced in this manner, with varying levels of speckle reduction. Three commonly used image quality metrics were applied to these frames, and average values were taken for each filtered image. Evaluation of the video quality was also performed by six experts in the analysis of echocardiography. Experts were asked to rate each video in three categories, namely Overall Quality, Detail Clarity, and Speckle Level. These criteria were suggested by one of the experts (a senior cardiac technician) in a prior consultation.

The results of the subjective evaluation are presented in Section 3.5. The effect of

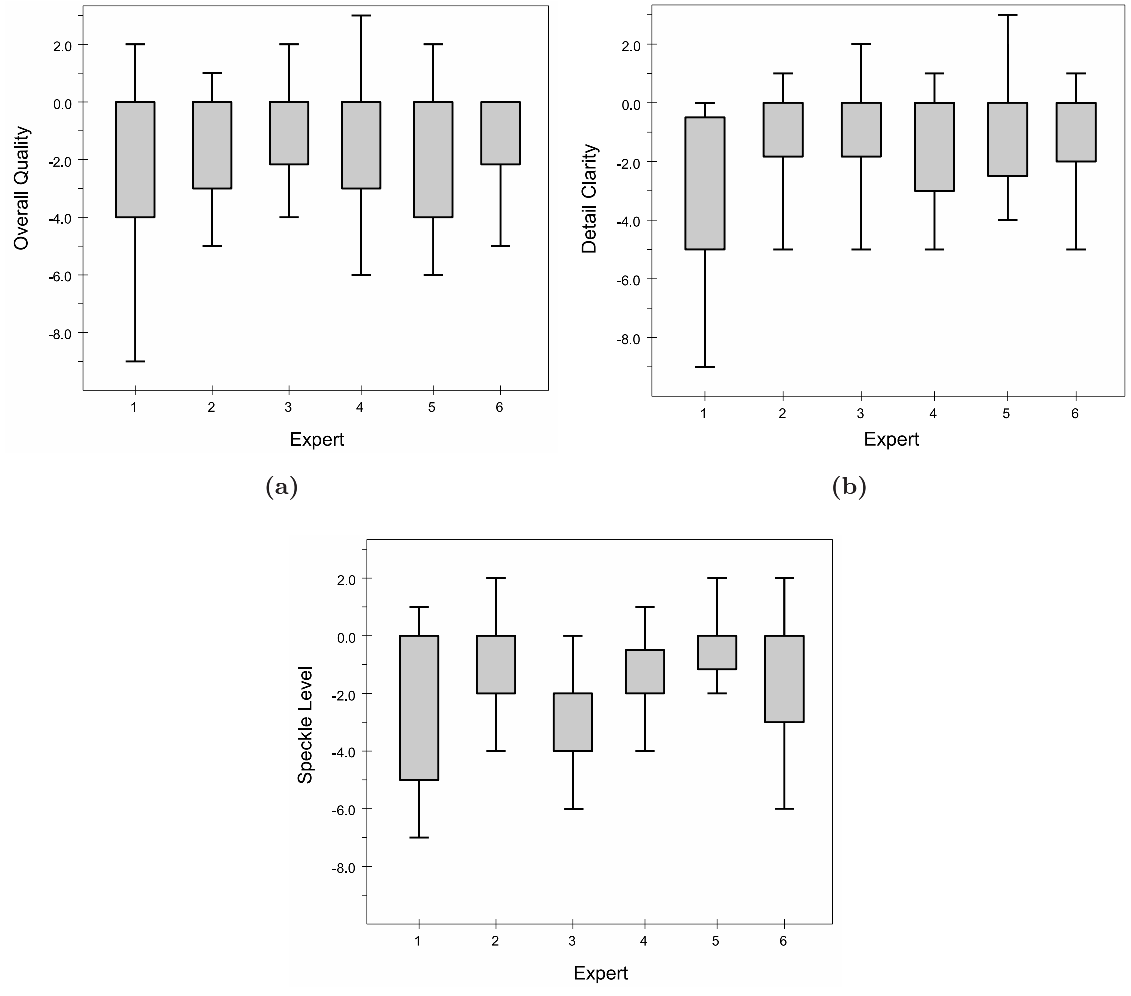

(a) **(b)**

(c)

Figure 3.4: Box-plots of each video assessment category, per expert. (a) Overall Quality, (b) Detail Clarity, (c) Speckle Level.

speckle reduction filtering on both the Overall Quality and Detail Clarity scores was negative in over half of cases. Based on these aggregate scores, the expert participants do not judge a reduction in speckle to have a positive effect on the Overall Quality of echocardiographic video. This is in general agreement with the results of Loizou *et al.* [89], in which expert assessment of speckle filtered carotid ultrasound was not rated as having an improvement in perceived quality for most of the considered cases. The study presented in this chapter was limited to the quantification of expert opinion, and its analysis by statistical methods to determine trends in these scores. The particular subjective reasons for the perceived reduction in quality are beyond the scope of this thesis, but the results of Dantas and Costa [5] appear to be relevant to discussion on this issue. In particular, these researchers noted that while filtering of speckle does not necessarily lead to a loss of clarity in the details of the image, it does remove the 'false-fine' structures introduced by the presence of speckle. These structures give the appearance of spurious fine detail, of a scale beyond the resolving power if the imaging system. Despite the fact that these details do not represent the underlying structure, their removal can lead to a perceived reduction in sharpness. It is noted however that the assessment of the experts is not universally negative, nor does this result preclude the further analysis of the video assessment metrics.

The expert scores were investigated to determine if there are statistically significant differences between the experts in each of the scoring categories. At a significance level of 1%, significant differences were observed between the experts in the Detail Clarity and Speckle Level categories. This suggests a variability in how experts subjectively asses the visibility of diagnostically relevant details, and also the strength of the speckle component in echocardiographic video. These differences highlight the subjective nature of appraisal by visual inspection, even by trained and experienced experts. However no significant difference was observed between the experts in their rating of Overall Quality, which indicates a similarity in the assessment of general quality for clinical applications.

The expert scores were investigated further to determine if there are any significant relationships between the scores themselves for each of the experts. The notable result of this analysis was the strong association between the experts perception of Overall Quality and Detail Clarity. This effect was noted for all of the expert participants, and was statistically significant at the 1% level. The relationship between these evaluation criteria indicates that subjective Detail Clarity is a strong indicator of diagnostic quality in echocardiography, and perhaps the primary consideration in expert evaluation of echocardiographic usefulness.

In addition, statistically-significant positive associations between Speckle Level and both Overall Quality and Detail Clarity was observed for most of the experts. This suggests perceptions of Overall Quality and Detail Clarity are higher for videos with less speckle reduction. These associations were not significant for all experts however, and were weaker than the association between Overall Quality and Detail Clarity.

The relationships between the expert scores and the image quality metrics were investigated. The correlation between Pratts FoM metric and the Overall Quality scoring category was statistically significant for all experts. Thus in general, filtered videos with higher FoM values tend to be rated higher in Overall Quality. This metric is thus suitable for evaluating the general quality of echocardiography. A significant negative correlation between edge region MSE and Detail Clarity was observed in all experts ratings. This means that higher Edge MSE implies lower Detail Clarity scores. This metric is therefore a good indicator of expert opinion on Detail Clarity. Finally, a significant negative correlation was observed between the overall MSE metric and the and Speckle Level as perceived by the expert participants. The MSE as applied here quantifies the difference between the filtered and unfiltered video, i.e. it measures the level of filtering performed. The relationship between this metric and the Speckle Level indicates that expert perception of the degree of speckle noise decreases with greater levels of filtering. In general, the agreement between expert scores and objective quality metrics verifies the use of these objective metrics.

The following chapter builds upon the results of the study presented in this chapter, to produce a comprehensive framework for the evaluation of performance in speckle filtering.

Chapter 4

Evaluation of Speckle Filtering

4.1 Introduction

This chapter presents a framework for evaluating the performance of a wide range of speckle reduction filters, when applied to speckle ultrasonography. As described in Chapter 2 above, a number of negative issues arise from the presence of speckle in ultrasonography. Chief among these are the introduction of spurious detail in the form of structures known as 'false-fine', whose resolution is beyond that of the imaging system [7]. These speckle structures can obscure details of the imaged medium (e.g. tissue boundaries) through the masking of small gray-level differences [9]. The presence of speckle has been determined to be the cause of an eight-fold reduction in lesion detectability in [8]. The accuracy of both human interpretation and automated diagnostic functions have also been found to be reduced by speckle [4]. In echocardiography, speckle noise is prominent in all cross-sectional views [15], and its effect is far more significant than additive noise sources such as sensor noise [2]. Reduction of echocardiographic speckle has been shown to positively affect subjective image quality and improve boundary definition [15].

Although a large number of speckle reduction methods are found in the literature, very few studies exist which compare a wide variety of these methods on a common set of test data. The main contribution of this chapter is the presentation of an evaluation framework to evaluate speckle reduction filters, from the perspective of both functional and computational performance for application to real-time echocardiography. A comprehensive evaluation of a wide range of existing speckle removal techniques is performed, taking into account performance when applied to both clinical and simulated ultrasound images, and also the computational requirements of each method. The fifteen despeckling filters considered in the study of this chapter are those detailed in Section 2.4, and are summarised in Table 2.1 above. These include anisotropic diffusion, wavelet denoising and local statistics.

The filters are compared experimentally by application to a large number of clinical echocardiographic images, and also on simulated images generated using a realistic model for image acquisition. The filtered output in both cases is examined using qualitative visual inspection and quantitative quality metrics, some of which are shown to be indicators of clinical perception of image quality in Chapter 3 above. A detailed examination of computational requirements is also undertaken. A previous review compares speckle reduction techniques for the application of carotid artery imaging [89]. The study presented here differs in the consideration of the echocardiographic application, the inclusion of a wide range of recent methods (particularly anisotropic diffusion based approaches) and also in the focus on evaluation based on computational performance.

This chapter is organised as follows: Section 4.2 presents the method used to produce simulated ultrasound images for filter evaluation. This method provides a realistic model of the image acquisition process, and permits the use of a noise free reference image for comparison. The tests comprising the method of speckle filter evaluation are presented in Section 4.3. This section includes examples of the simulated and clinical images used, and discusses the differences in the use of these images. Also described in this section are a set of image quality metrics used to quantify speckle reduction and other filtering characteristics. The details of the computational requirements analysis are presented in this section, including a speed-accuracy analysis of discretisation methods for anisotropic diffusion filters. The results of these filter evaluation tests are presented in Section 4.4, including the application of the filters to the simulated and clinical image sets, the results of image quality metrics calculation, and the computational requirement analysis. Finally, Section 4.5 concludes the chapter.

4.2 Ultrasound Simulation

This section presents a method of simulating realistically speckled ultrasound images, based on simulating acoustic interaction with scattering elements. Various approaches have been used in the literature to simulate ultrasonic images, the simplest of which follow (2.10) directly. A reference image $R(x, y)$ is multiplied with an arbitrary noise component

$n(x, y)$. For example, n could have a Gaussian distribution, with a variance controlling the degree of degradation. Sattar *et al.* [77] expanded this approach by introducing spatial correlation to $n(x, y)$. A complex Gaussian random field is lowpass filtered, and the magnitude is taken as noise component $n(x, y)$ in (2.10). Achim *et al.* [75] specified that lowpass filtering with a kernel of size three is sufficient to introduce the desired short term correlation to the speckle.

Yu and Acton [3] used a simulation method which models in greater detail the image acquisition process. An ultrasound pulse is modelled by a linear space invariant point spread function (PSF). The medium is modelled as the product of a smooth cross section and a Gaussian white noise field, which accounts for subresolution variances. The RF data is produced by convolving the PSF with the medium representation. A complex analytic version of the simulated RF data is obtained with the Hilbert transform. This approach was also employed elsewhere [6].

Ng *et al.* [25] explored the treatment of ultrasound imaging as a linear process. The backscattered RF signal is modelled as the linear filtering of an echogenicity map with a transfer function, similar to [117]. Such a linear model is only valid for the far-field of an unfocused transducer, or within the focal zone of a focused transducer [7]. In addition, Ng *et al.* concentrated on a rectilinear scanning geometry, for which the PSF of the imaging system varies only in the axial (depth) direction. In [118], Ng *et al.* used an axially-varying PSF for simulation, calculating the PSF at a number of axial depths. The PSF is treated as being piecewise constant, similar to [119].

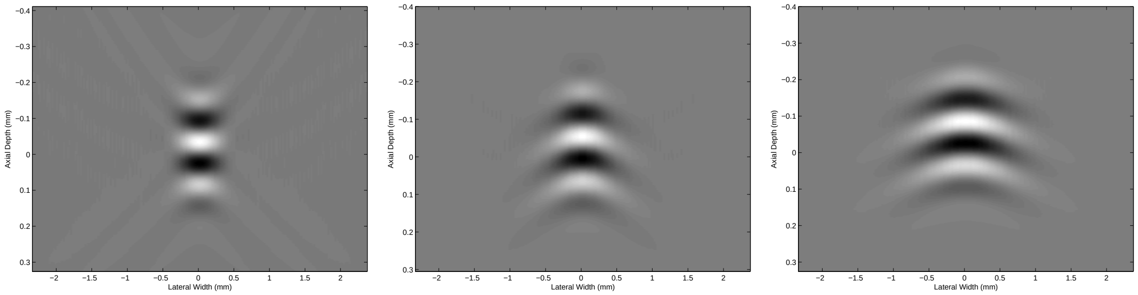

(a) **(b)** **(c)**

Figure 4.1: Ultrasound system point spread functions used in simulation, shown at various depths. (a) 6 mm, (b) 13 mm, (c) 20 mm.

In this thesis, an axially varying PSF is also employed. Linear interpolation is used to determine the PSF at depths between those calculated, similar to the work of [118]. PSFs for the simulated imaging system are generated using the Field II simulation software [54] and demodulated to baseband. The PSFs are calculated at fifty equally spaced depths covering the axial axis of the medium. A rectangular transducer is used, with a centre frequency of 6.5 MHz. Both the transmit and receive subapertures consist of 64 elements of width 0.3 mm and height 6 mm, with a lateral gap of 0.1 mm between elements. Hamming apodization is employed, and the speed of soundwave propagation is set at a constant 1540 m/s. To generate each PSF a single point scatterer is placed at the desired depth, and both transmit and receive foci are set to this depth. Fig. 4.1 displays example PSF functions computed using the technique described here. These show the PSFs at various axial depths (distances from the transducer face). It can be seen that the PSFs become wider in extent for greater depths, a result noted in [118]. This corresponds to a widening of the imaging system resolution cell, i.e. a decrease in axial resolution for objects farther from the transducer.

While the PSF models the imaging system, the imaged medium is modelled as an echogenicity map $h(x, y)$ composed of complex point scatterers, similar to [5]. A template is used to define regions of differing scatterer density within an image. Enough scatterers are then placed with each region to achieve the desired density level. Positioning of the scatterers within each region is randomly performed. The phase values of the scatterers follows a uniform distribution, varying from $0 \rightarrow 2\pi$ rad, while scatterer magnitude follows a Gaussian distribution with unity mean and $\sigma = 0.1$. Fig. 4.2(a) shows an example of an echomap with two regions of differing scatterer densities, with density values of 40% and 10%.

Echogenicity maps are filtered to generate a simulated RF image as:

$$\begin{aligned} RF(x, y) &= h(x, y) * p(x, y) \\ I(x, y) &= |RF(x, y)| \end{aligned} \tag{4.1}$$

where p is the analytic form of the PSF function at the correct depth, and $*$ denotes convolution. $I(x, y)$ is then the speckled amplitude image. The simulated amplitude image generated using this technique for the echomap of Fig. 4.2(a) is shown in Fig. 4.2(b), and a granular speckle pattern is observed.

Burckhardt's Maximum Writing (MW) technique [9] is used to generate a maximally speckle free version of each of the twenty simulated images. The above convolution is performed multiple times, randomly varying the phase values of the scatterers while leaving their location and magnitude unaltered. This produces a series of images with the same structure but differing speckle patterns, the maximum of which is the speckle free image:

$$I_{MW} = \max\{|p(x,y) * (|h(x,y)| \exp[\phi_n])\,|\} \tag{4.2}$$

where ϕ_n is the phase of the nth set of scatterers generated. In practice, this technique is applied until the contribution of the nth amplitude image is smaller than 0.1% of average image brightness, similar to the approach of [5]. When the MW method is applied to the example echomap of Fig. 4.2(a), the resulting maximally speckle free image is displayed in Fig. 4.2(c). Note that the image still contains some texture, i.e. some regions of differing intensity are present due to the non uniform random positioning of the scatterers.

This section has presented a technique used to generate simulated speckled ultrasound images. These images, and their speckle free MW reference images are using in the speckle filter evaluation framework presented in this chapter. The next section outlines in detail the tests performed in the filter evaluation.

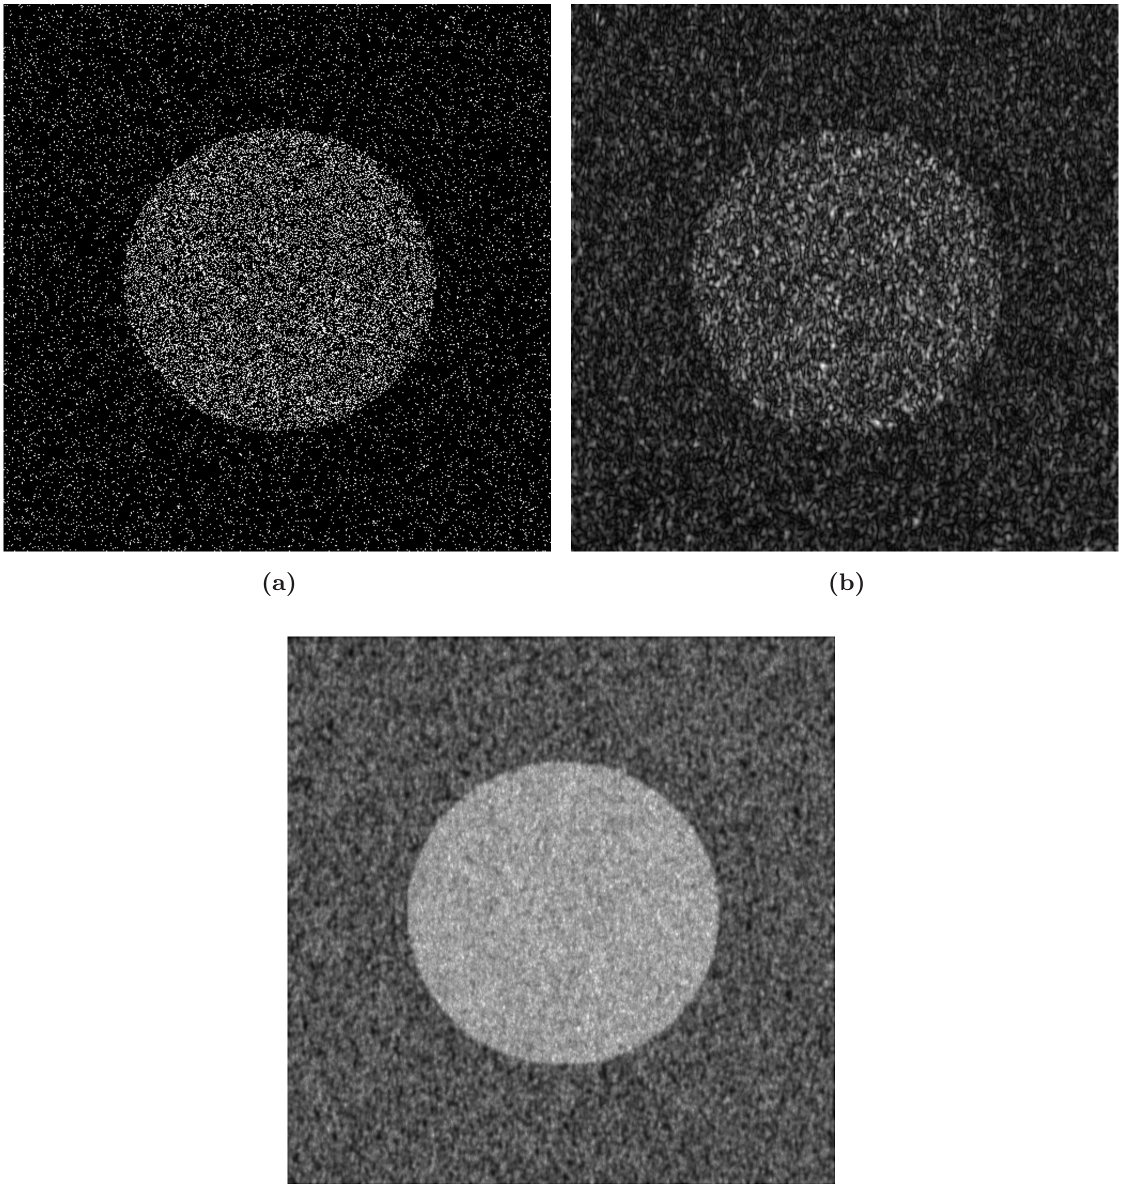

(a)

(b)

(c)

Figure 4.2: Example simulated ultrasound image. (a) Echomap, (b) Speckled amplitude image, (c) Speckle free MW image.

4.3 Speckle Filter Evaluation

This section details the tests used to evaluate filter performance in the evaluation framework proposed in this chapter.

- The filters are applied to a range of ultrasound images, both simulated and clinical.
- A set of image quality metrics are presented, for quantifying speckle filter performance when applied to the simulated and clinical image sets. The simulated and clinical tests quantify the effects of filtering in different ways: the filtered clinical images are evaluated relative to the original speckled input, while the filtered simulated images are compared to a noise free reference.
- A detailed analysis of the computational requirements of each filter is presented.

4.3.1 Simulated Ultrasound Images

Simulated ultrasound images are generated using the technique described in Section 4.2 above. Two simulation templates are used to generate simulated images for speckle filter evaluation. Ten images are produced using each template, each of which has the same underlying structure. However as the scatterer position, phase and magnitude are randomly determined, these ten images have different speckle content. The echogenicity templates define simplified representations of the shapes seen in clinical echocardiography, specifically in the transthoracic parasternal short and long axis views. These contain three regions of differing scatterer densities, with values of 25%, 5% and 1%.

Fig. 4.3 displays example images from the simulated image templates. Included here are examples of the echomaps used (Figs. 4.3(a)-(b)) which show the differing concentration of scatterers in different image regions. Brighter regions correspond to increasing scatterer densities. The resulting speckled images, generated according to (4.1), are displayed in Figs. 4.3(c)-(d). The maximally speckle free MW simulated images are then shown in Figs. 4.3(e)-(f). While these have theoretically the lowest possible level of speckle, some texture in the form of lighter and darker regions can be seen. This effect is caused by the non-uniform distribution of scatterers, due to the random nature of their positioning.

Ten images are produced for each template, so a total of twenty simulated images are used in speckle filter evaluation. The simulated images are generated with size 1024 ×

1024 pixels. The corresponding physical size of the simulated imaged media is defined by the sampling of the simulated ultrasound systems PSF. Here this leads to simulated images representing approximately 32×32 mm.

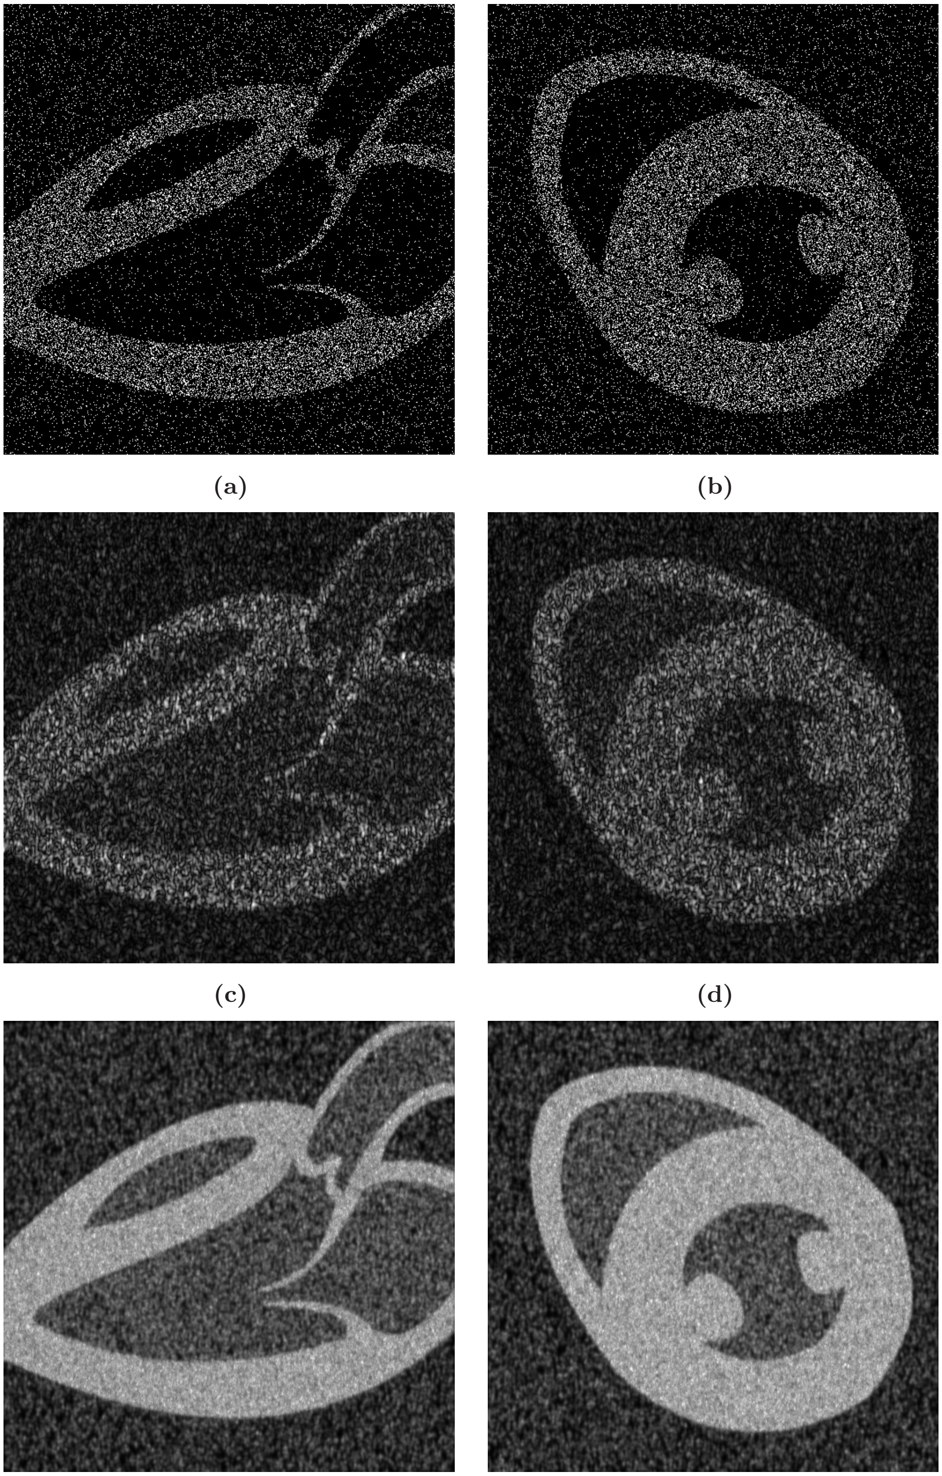

(a) (b)

(c) (d)

(e) (f)

Figure 4.3: Simulated clinical ultrasound images. (a)-(b) Echomaps, (c)-(d) Amplitude images, (e)-(f) Speckle free MW images.

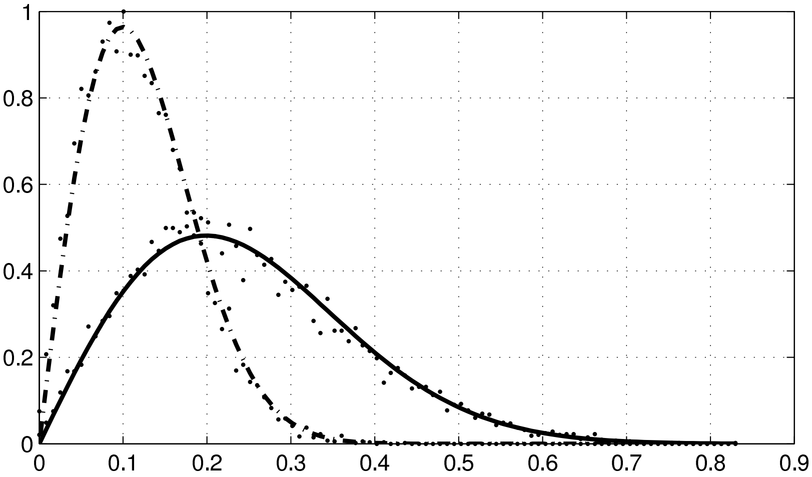

Figure 4.4: Empirical histograms of homogeneous regions in amplitude images, showing fitted Rayleigh distributions. Solid line represents simulated image, dashed line clinical image.

4.3.2 Clinical Ultrasound Images

A set of clinical images are used to evaluate the performance of the speckle reduction filters. A series of frames are taken from clinical echocardiographic videos, which are part of a pre-existing clinical database, supplied by the Cardiology Department at University College Hospital Galway, Ireland. They were scanned using a General Electric Vivid 7 Series scanner (GE Healthcare, Piscataway, NJ, USA). These images are subject to log compression, however no image optimization or speckle reduction (such as compounding from multiple angles) is performed. A total of 500 frames are used to evaluate filtering performance, taken from 100 videos from 40 patients.

Amplitude representations of these images are generated using inverse log compression. The model of log compression used is $Y = \alpha log(X+1)+\beta$ [14], where (X, Y, α, β) are respectively the amplitude image, log compressed image, contrast parameter and brightness parameter. Parameters α and β are estimated for clinical frames using the method of [14], and amplitude images can then be found as $X = \exp\{(Y-\beta)/\alpha\} - 1$. Example frames are displayed in Fig. 4.5, which have been log compressed for display. The average size of video frames is 419×511 pixels.

To demonstrate that the simulated and clinical images both exhibit Rayleigh statistics in regions of fully developed speckle, the empirical histograms of such regions from an example of each are displayed in Fig. 4.4. The regions selected are of size 100×100 pixels, and the clinical data is converted to amplitude form by inverse log compression as described above. Rayleigh distributions are fitted to these histograms, with maximum likelihood estimates of the Rayleigh parameters found as 0.129 ± 0.0017 and 0.188 ± 0.0024 for the simulated and clinical images, respectively, at a 1% confidence level. Using 128 point histograms, the goodness of fit is evaluated by calculating the r^2 value as 0.993 and 0.991, respectively, indicating a close fit to Rayleigh distributions. This shows the similar statistics of speckle in the simulated and clinical images.

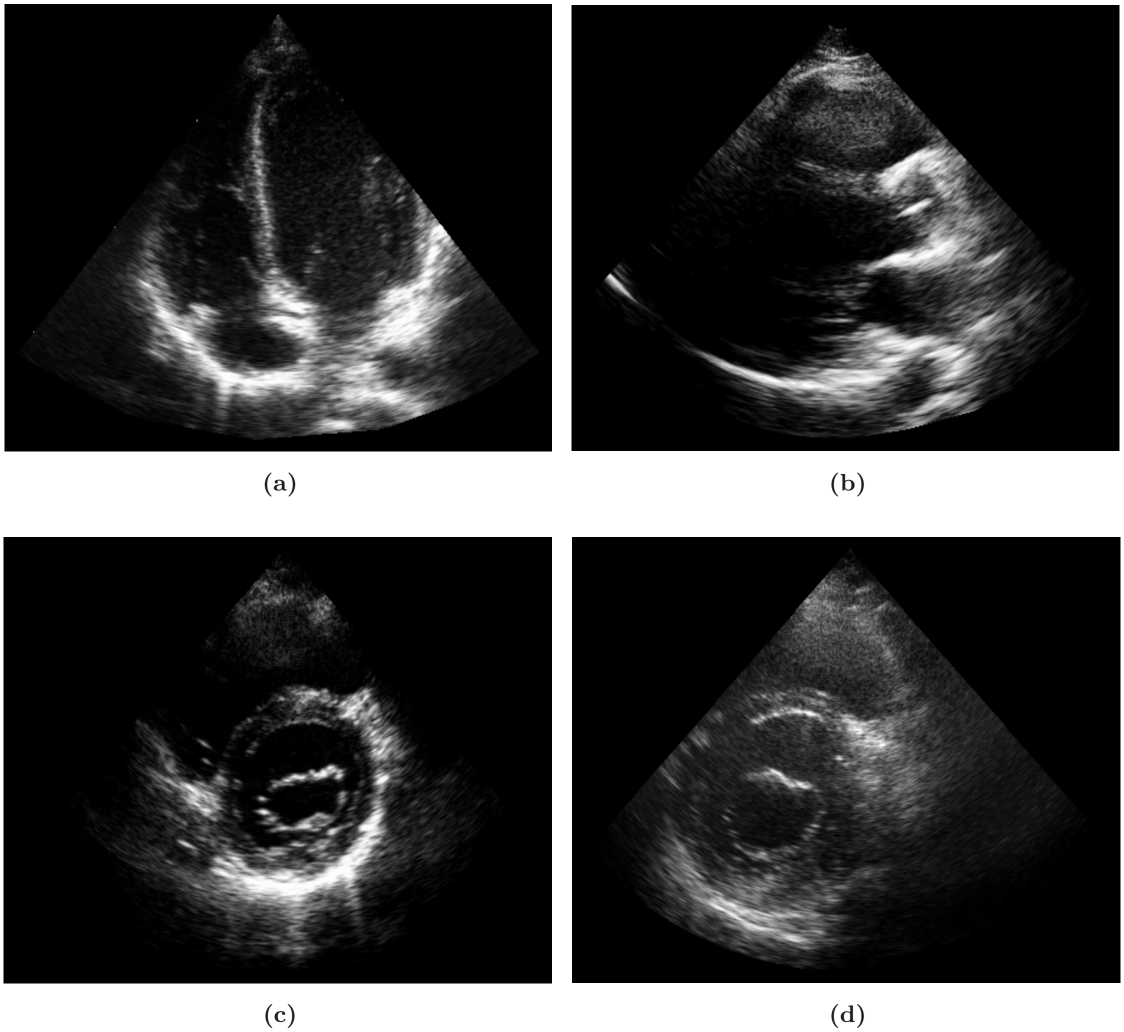
(a) **(b)** **(c)** **(d)**

Figure 4.5: Clinical ultrasound images, used in speckle filter evaluation. (a) Image 1, (b) Image 2, (c) Image 3, (d) Image 4.

4.3.3 Image Quality Metrics

Five image quality metrics are applied to both the simulated and clinical echocardiographic images. Of these, three (FoM, SSIM, and edge MSE) are comparative in nature, considering two images. For the simulated images, each filtered image is compared to the noise free MW reference. No noise-free reference exists for the clinical images, so the metrics are applied using the unfiltered equivalent as the reference. Thus in the simulated case filtering is quantified relative to the ideal result, while in the clinical case it is quantified relative to the noisy input. The CNR metric operates on a single image. To measure the change due to filtering the difference between the filtered and reference value is used. This is normalized relative to the value in the reference image. Finally, the SNR_A metric is used to measure the decrease in speckle due to filtering. The validity of these metrics is shown in a previous study [120], where a statistically significant relationship between the FoM metric and expert assessment of the overall quality of clinical echocardiographic videos is demonstrated. A significant relationship also exists between the edge region MSE and the clarity of important details as perceived by experts. The FoM and the edge MSE are detailed in Chapter 3 (page 3.3) above.

4.3.3.1 Structural Similarity (SSIM)

The SSIM measure [61] to asses the preservation of structural information in the filtering process:

$$SSIM = \frac{1}{M} \sum \frac{(2\mu_1\mu_2 + C_1)(2\sigma_{12} + C_2)}{(\mu_1^2 + \mu_2^2 + C_1)(\sigma_1^2 + \sigma_2^2 + C_2)} \tag{4.3}$$

where μ_1, μ_2 and σ_1, σ_2 are the means and standard deviations of the images being compared, and σ_{12} is the covariance between them. These quantities are calculated using local statistics within a total of M windows, the average of which is taken in (4.3). Constants $C_1, C_2 \ll 1$ ensure stability [61], and M is chosen as 32. The SSIM has values in the $0 \rightarrow 1$ range, with unity representing structurally identical images.

4.3.3.2 Contrast to Noise Ratio (CNR)

Similar to [6], the Contrast to Noise Ratio (CNR) quantifies the level of contrast between a region of interest and the background, and is calculated as:

$$CNR = \frac{|\mu_1 - \mu_2|}{\sqrt{\sigma_1^2 + \sigma_2^2}} \tag{4.4}$$

where μ_1 and σ_1^2 are the mean and variance of a region of interest, and μ_2 and σ_2^2 are the mean and variance of a similar sized region in the image background. For each image the largest homogeneous background region available is selected. A similar sized region of interest is also chosen, containing structure of interest. Fig. 4.6 display examples of these regions for the simulated and clinical images used in filter evaluation. Solid and dashed lines are used to respectively indicate the background and region of interest. The relative change in CNR is then calculated as the CNR difference between the filtered and reference images, normalized relative to the reference value.

4.3.3.3 $\mathbf{SNR_A}$

Burckhardt's [9] SNR_A quantifies the level of speckle as the ratio of mean to standard deviation of the amplitude values. This has a theoretical value of 1.91 in regions of fully developed speckle. Similar to [5], SNR_A is calculated for each of the three regions

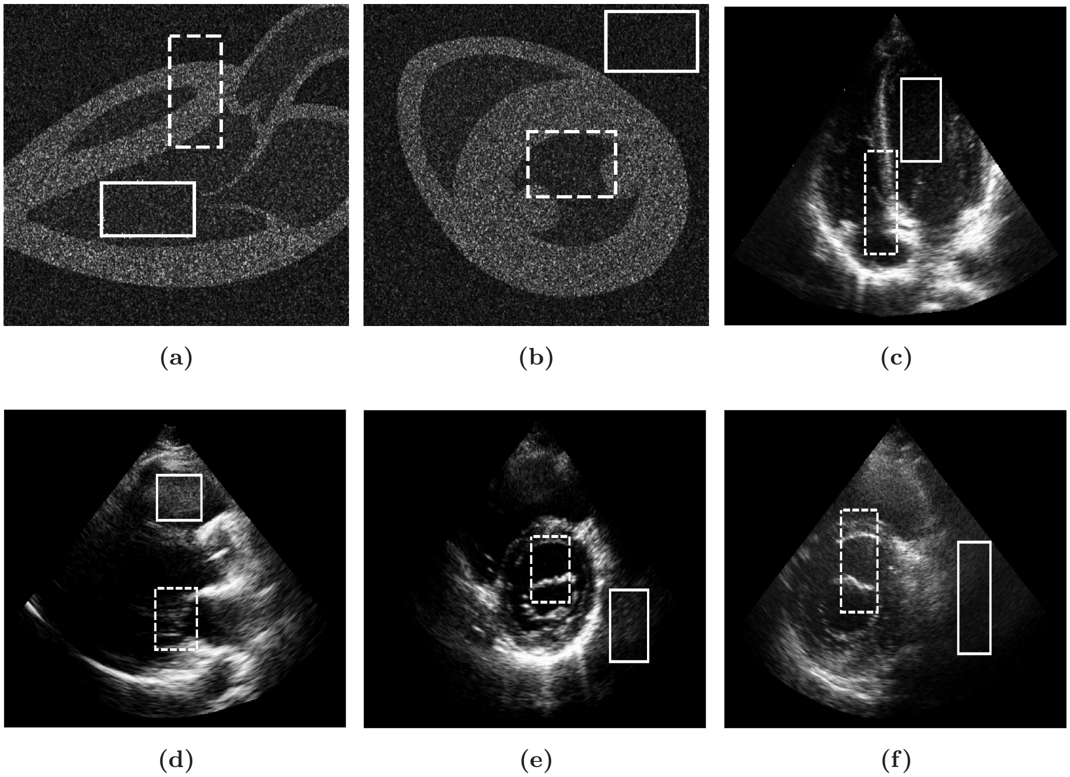
(a) (b) (c)

(d) (e) (f)

Figure 4.6: The regions selected for CNR metric calculation. (a)-(b): simulated images of Figs. 4.3(c)-(d), (c)-(f): clinical images of Fig. 4.5.

of the simulated images. This allows quantification of the level of speckle before and after each type filtering, in regions of differing scatterer density. In the clinical case SNR_A is calculated using the entire image.

4.3.4 Computational Requirement Analysis

The computational requirements for each of the filtering methods considered in Section 2.4 are determined by calculation of the number of multiplications, additions, and look-up table operations. A number of considerations are detailed here.

4.3.4.1 Anisotropic Diffusion Discretisation

The choice of discretisation method for anisotropic diffusion methods has a major impact on computational requirements. A given level of diffusion is represented by the diffusion 'time' $t = \tau Q$, where Q is the number of iterations required. A larger τ value thus permits fewer iterations, and lower computing requirements. As noted in Section 2.4.1.7, explicit methods are only accurate for small values of timestep τ. Semi-implicit methods are stable for arbitrary large τ, however this comes at the expense of accuracy. In [100], the authors compare the explicit method of discretisation of (2.26) and (2.27) to the AOS scheme of (2.30), using a diffusion function from [101].

Here a similar analysis is performed, and is extended to compare the explicit (2.27), AOS (2.30), and Jacobi (2.31) schemes. To determine if the accuracy is affected by the choice of diffusion function, we implement two diffusion methods: Perona and Malik's method, described in Section 2.4.1.1, and the DPAD method described in Section 2.4.1.3. Accuracy is measured for both filters, using a reference calculated using the explicit method and a very small timestep, which gives maximum accuracy. For both methods, diffusion is performed using the three discretisation methods discussed above for $t = 5s$, and the reference image is calculated with $\tau = 0.025$ and $Q = 200$. The ℓ^2 error norm was used, as in [100], and is defined as $Err = \|I_{filt} - I_{ref}\|_2 / \|I_{ref}\|_2$, where I_{filt} and I_{ref} are the filtered and reference images, and $\|.\|_2$ represents the ℓ^2 norm.

4.3.4.2 Various Filter Implementation Considerations

Each iteration of a scalar anisotropic diffusion method, using semi-implicit discretisation, requires the construction of matrices by (2.28) and also that (2.30) be solved. This requires a total of $12N$ multiplications and $12N - 8$ additions, where N is the number of pixels in the image.

As noted by [100], convolution with a Gaussian can be performed using isotropic diffusion, i.e. anisotropic diffusion with a unity diffusion function. The effort required for this in the 2D case is $6N - 4$ multiplications and $4N - 4$ additions.

Solution of the semi-implicit discretisation for matrix schemes as per (2.35), including the calculation of all A matrices, can be accomplished with $17N - 9$ multiplications and $24N - 14$ additions.

The 2D convolutions required in the evaluated filters (e.g. windowed means in (2.17) and (2.52), and both types of wavelet transform) are performed using separable convolution kernels. Time domain convolution requires $2NM$ multiplications and $2N(M-1)$ additions, for a separable kernel of size $M \times M$. A commonly used alternative is frequency domain convolution using the FFT. Here the bit reversed radix 2 FFT is considered, which can perform an N point FFT of a real signal in $N \log_2 N$ real multiplies and $\frac{3}{2}N \log_2 N + \frac{1}{2}N$ real additions [121]. Separable 2D convolution using this FFT, of a square image of size $N = N_1^2$ with a separable $M \times M$ kernel will thus take $2N^{'}(1+2N_1)\log_2 N^{'} + 2N^{'}(1+6N_1)$ real multiplies and $3N^{'}(1+2N_1)\log_2 N^{'} + N^{'}(1+6N_1)$ real additions, where $N^{'} = N_1 + M - 1$. For a 512×512 image, frequency domain convolution is more efficient for $M = 25$ and greater. As the kernels used here are smaller than this, time domain convolution is used.

The wavelet transforms can be efficiently applied with the 'Algorithme à trous'. The wavelet transform used in the GLM method requires $546N$ multiplications and $504N$ additions for three levels of decomposition and reconstruction. This figure is based on convolution as above. Mallat and Zhong's DWT, as used in the NMWD and Zong filters, requires $248N$ multiplications and $240N$ additions for an image with N pixels, for convolution as above. This figure includes four levels of analysis and reconstruction (as in both the NMWD and Zong methods).

Calculation of the median of N values can be implemented efficiently using a 'divide and conquer' algorithm using $O(N)$ sorting operations [122]. These sorting operations are not quantified when assessing the number of multiplications and additions, although the reader should be aware that a computational effort is still required.

Functions which need to be evaluated, such as exponential, logarithmic and diffusion functions are assumed to be implemented in look-up tables for greatest efficiency, although this is at the expense of added memory utilisation.

This section has presented the filter evaluation framework used in this chapter. The speckled images, both clinical and simulated, to which the filters are applied are presented. The image quality metrics used to quantify the performance of these filters are presented, and the different methods in which they are applied for simulated and clinical images is presented. The computational analysis performed is also detailed, and some important implementational aspects are reviewed. The next section in this chapter details the results of echo of the tests presented here.

4.4 Results

This section presents the results of the tests outlined in Section 4.3 above. For both simulated and clinical ultrasound images, example speckle reduced images are displayed for each filter. The image quality metrics detailed in Section 4.3.3 are applied in each case, and the results presented. The results of the analysis of filters computational requirements is also presented.

4.4.1 Simulated Images

The speckle filters are applied to the simulated amplitude images as described in Section 4.3.1. Ten simulated images with different speckle patterns are produced for each template. Typical filtered images from each template are displayed here, and these are representative of the results from each template. Figs. 4.7-4.8 display results of application of the speckle filters to simulated echocardiograms of Figs. 4.3(c)-(d), respectively.

Each of the filters described in Section 2.4 have user definable parameters, the values of which can determine to differing degrees the nature of the filtered output. The approach taken in the filter evaluation strategy in this thesis is to apply the filters, each with parameters chosen experimentally to produce to greatest degree of speckle reduction. This approach ensures that the strongest possible speckle suppression is achieved in each case. Each filter is applied to the set of simulated images, with various parameter values. The average SNR_A metric over the simulated test set is used to quantify speckle level. The filter parameters achieving these maximum levels of speckle suppression are used in the rest of the evaluation. These parameters are displayed in Table 4.1.

The images filtered by the SAR local statistics filters can be grouped into two groups of differing speckle reduction performance. The first of these contains the output of the Lee filter (shown in Figs. 4.7(a) and 4.8(a)), the Frost filter (shown in Figs. 4.7(c) and 4.8(c)), and the enhanced Frost filter (shown in Figs. 4.7(e) and 4.8(e)). The output of these filters can be seen to have most of the speckle pattern removed. Some speckle does remain however. In all cases a degree of blurring is evident in the edges of the image

Table 4.1: Filter Parameters.

Filter	**Filtering Parameters**
PMAD	$t = 7.5s, \sigma = 4$
SRAD	$t = 150s$
DPAD	$t = 150\ s, Z = 5$
CED	$t = 12.5s, \sigma = 2, \rho = 2.5, C_1 = 0.001, C_2 = 1$
NCD	$t = 25s, \rho = 6$
OSRAD	$t = 25s, \sigma = 2, \rho = 2.5$
Zong	$\alpha = 0.5, T_1 = 0.2, T_2 = 0.25, T_3 = 0.5, c = 4, b = 0.5$
NMWD	$Q = 8, Z = 5$
GLM	$\alpha = 15$
Lee	$Z = 7$
Frost	$Z = 9$
Kuan	$Z = 9$
EnhLee	$Z = 9, K = 10$
EnhFrost	$Z = 9, K = 10$
Geo	$Q = 7$

structures. The other SAR filters, namely the Kuan filter (the output of which is shown in Figs. 4.7(b) and 4.8(b)), and the enhanced Lee filter (shown in Figs. 4.7(d) and 4.8(d)), are seen to be remove considerably less of the speckle pattern from the images.

The anisotropic diffusion filters exhibit a range of quality of their filtered images. The output of the PMAD filter, shown in Figs. 4.7(f) and 4.8(f), displays a high level of blurring. This is due to a large number of diffusion iterations, and it is noted experimentally that this is necessary to remove the speckle.

The SRAD and DPAD filter output (shown in Figs. 4.7(g) and 4.8(g), and in 4.7(h) and 4.8(h)) both display strong speckle suppression, with most of the speckle texture removed. The DPAD filter appears to remove more of the speckle however, producing images which are more uniform. This is due to the use of larger window sizes possible with this method.

The output of the CED method (shown in Figs. 4.7(i) and 4.8(i)), and that of the

NCD filer (shown in Figs. 4.7(j) and 4.8(j)), both display artefacts. These are introduced by the enhancement of contours by these methods, leading to the enhancement of the contours of the speckle texture. This effect is more pronounced for the NCD filter (i.e. finer scale speckle texture has its contour enhanced), due to the lack of an initial smoothing step with this method.

The OSRAD filter (the output of which is displayed in Figs. 4.7(k) and 4.8(k)) displays a strong degree of speckle suppression, as most of the speckle texture can be seen to be removed. In addition, the borders of the image structure are seen to be preserved by the filtering method, i.e. image edges are not blurred.

The outputs of the Zong filter (shown in Figs. 4.7(l) and 4.8(l)), and the output of the GLM filter (displayed in Figs. 4.7(m) and 4.8(m)), are both seen to have a degree of speckle present after filtering. Thus it is observed that these methods due not perform as well as the others considered at removing speckle.

The NMWD filter is seen to produce output with very little of the speckle pattern remaining, as seen in Figs. 4.7(n) and 4.8(n) and Finally, the Geometric filters output (shown in Figs. 4.7(o) and 4.8(o)) displays a good level of speckle reduction, however some of the speckle pattern does still remain.

Table 4.2 displays the SNR_A metric values measured for the simulated images of Fig. 4.3. In both cases, these SNR_A values are averaged over the set of twenty images. The average SNR_A within the each of the three regions of different scatterer density are displayed for all of the filters. The overall average for each filter is also reported in these tables, as is the improvement in this value due to filtering, relative to the noisy amplitude image.

The average SNR_A values for each image region shows how speckle suppression varies with scatterer density. The filters are arranged from highest to lowest overall average SNR_A improvement. In the speckled input images, an SNR_A value of 1.91 is expected in regions of fully developed speckle [9]. The values observed for the input images here are similar to this value, and are closer for individual regions with higher scatterer density values. This is the expected result, which is also reported in the work of [5], wherein a similar simulation approach is taken. As described in Chapter 2, for speckle to be fully developed, a sufficient number of scatterers must be present in each resolution cell of the imaging system. Regions in these simulated images with lower scatterer densities are more likely to contain resolution cells with a small number of scatterers, and thus speckle will be

Table 4.2: Average SNR_A for the simulated echocardiographic images of Figs. 4.3(c)-(d).

Filter	R1(25%)	R2(1%)	R3(5%)	Average	Improvement
MW	5.94	3.02	4.06	4.34	142.46%
NMWD	3.51	3.49	3.99	3.66	104.47%
OSRAD	3.53	3.13	4.06	3.57	99.44%
SRAD	3.94	2.64	3.81	3.46	93.30%
NCD	3.15	2.77	3.36	3.09	72.63%
DPAD	3.21	2.73	3.24	3.06	70.95%
PMAD	3.38	2.30	3.27	2.98	66.48%
Lee	3.11	2.68	3.16	2.98	65.31%
Geo	2.84	2.26	3.14	2.75	53.63%
EnhFrost	2.89	2.42	2.73	2.68	49.72%
CED	2.67	2.39	2.69	2.58	44.13%
GLM	2.49	2.05	2.76	2.43	35.75%
EnhLee	2.57	2.21	2.41	2.40	34.08%
Zong	2.45	2.15	2.43	2.34	30.73%
Kuan	2.45	2.15	2.34	2.32	29.61%
Frost	2.22	1.93	2.11	2.09	16.76%
Input	1.88	1.69	1.81	1.79	0%

less than fully developed within these regions.

All of the speckle reduction filters are seen to improve the SNR_A, in all regions of scatterer density. The MW reference has a higher average SNR_A than the post processing filters as expected, although some post processing filters exceed or match the MW SNR_A in individual regions. The NMWD and OSRAD filters achieve the highest average SNR_A values of the speckle reduction filters. The SRAD filter also performs well in this measure of speckle reduction. The OSRAD filter exceeds the average SNR_A of the MW reference in one image region in Table 4.2, and matches it in another. SRAD exceeds the MW reference average SNR_A in a single image region, also in Table 4.2. The rest of the anisotropic diffusion methods have varying results as quantified by SNR_A. The NCD, DPAD and

PMAD filters are rated highly, however the CED method performs poorly. Of the SAR filters, the Lee filter achieves the highest SNR_A score. Apart from the NMWD filter, the wavelet based methods perform quite poorly in this test.

Table 4.3: Average image quality metrics for speckle filtered simulated echomap images of Figs. 4.3(c)-(d), measured relative to MW reference images. The CNR value is the relative difference with respect to the MW reference. Each column of metric values is arranged from best to worst.

Filter	FoM	Filter	SSIM	Filter	CNR	Filter	EdgeMSE
Zong	0.69	CED	0.20	PMAD	0.21	OSRAD	1.624×10^{-4}
Kuan	0.63	NCD	0.19	OSRAD	0.15	NCD	1.691×10^{-4}
Frost	0.63	OSRAD	0.14	NMWD	0.14	CED	1.707×10^{-4}
EnhLee	0.58	SRAD	7.90×10^{-3}	SRAD	0.11	SRAD	3.897×10^{-4}
NCD	0.53	PMAD	4.94×10^{-3}	GLM	0.06	PMAD	4.043×10^{-4}
Lee	0.53	Frost	4.79×10^{-3}	DPAD	0.02	DPAD	4.169×10^{-4}
EnhFrost	0.49	Kuan	4.30×10^{-3}	Geo	-0.08	Frost	4.192×10^{-4}
CED	0.43	EnhLee	4.14×10^{-3}	NCD	-0.10	Kuan	4.193×10^{-4}
SRAD	0.38	Zong	3.90×10^{-3}	Lee	-0.19	Lee	4.194×10^{-4}
DPAD	0.31	DPAD	3.76×10^{-3}	EnhFrost	-0.32	EnhLee	4.194×10^{-4}
Geo	0.29	EnhFrost	3.63×10^{-3}	CED	-0.33	EnhFrost	4.195×10^{-4}
OSRAD	0.27	Lee	2.98×10^{-3}	Zong	-0.43	NMWD	4.196×10^{-4}
NMWD	0.25	GLM	2.61×10^{-3}	EnhLee	-0.43	Geo	4.210×10^{-4}
PMAD	0.24	NMWD	2.36×10^{-3}	Kuan	-0.45	Zong	4.211×10^{-4}
GLM	0.19	Geo	1.67×10^{-3}	Frost	-0.51	GLM	4.220×10^{-4}

Application of the other image quality metrics to the simulated image set results in the values of Table 4.3 for the filtered simulated echocardiographic images of Fig. 4.3. These metrics compare each filters output to the speckle free result from the MW technique. In each case, these metrics are averaged over the twenty simulated images.

The OSRAD, CED and NCD filters have the lowest edge region MSE, and so have on average the smallest difference in the intensity of pixels close to image edges. By contrast, here the FoM metric quantifies the average distortion in edge pixel locations between each filtered image and the MW reference image. The filters which perform best here are the Zong wavelet filter and the SAR filters. The filters with output most similar to the MW edges, as measured by the edge region MSE, perform poorly in the FoM. The SSIM metric compares the average structural similarity of filtered output with the MW reference. Calculated values are quite low, with the CED filter having the highest average of 0.22. Thus the speckle filters output are not structurally similar to the MW reference images. The CNR metric quantifies the average difference in contrast between each filtered output and the corresponding MW reference. Negative CNR values here indicate a lower contrast value than the MW reference. Over one third of the filters show improved average contrast relative to the MW image, the greatest of which are for the PMAD and OSRAD filters.

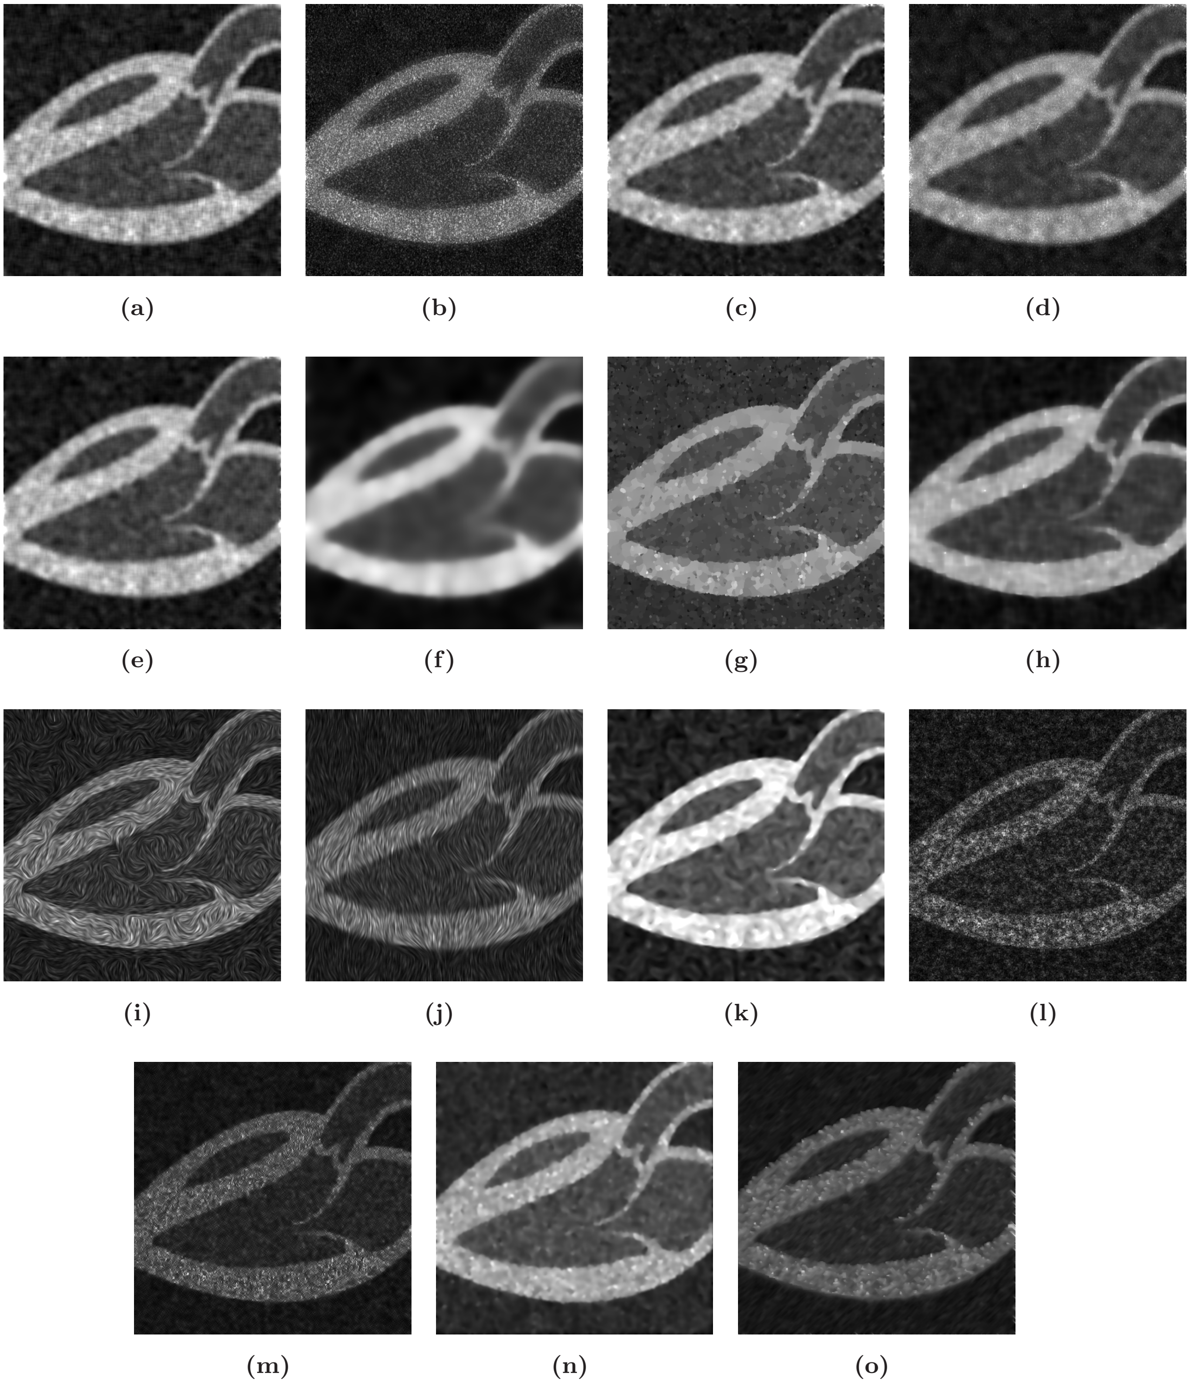

(a) **(b)** **(c)** **(d)**

(e) **(f)** **(g)** **(h)**

(i) **(j)** **(k)** **(l)**

(m) **(n)** **(o)**

Figure 4.7: Sample speckle filter output for simulated images using the template of Fig. 4.3(c). (a) Lee, (b) Kuan, (c) Frost, (d) EnhLee, (e) EnhFrost, (f) PMAD, (g) SRAD, (h) DPAD, (i) CED, (j) NCD, (k) OSRAD, (l) Zong, (m) GLM, (n) NMWD, (o) Geo.

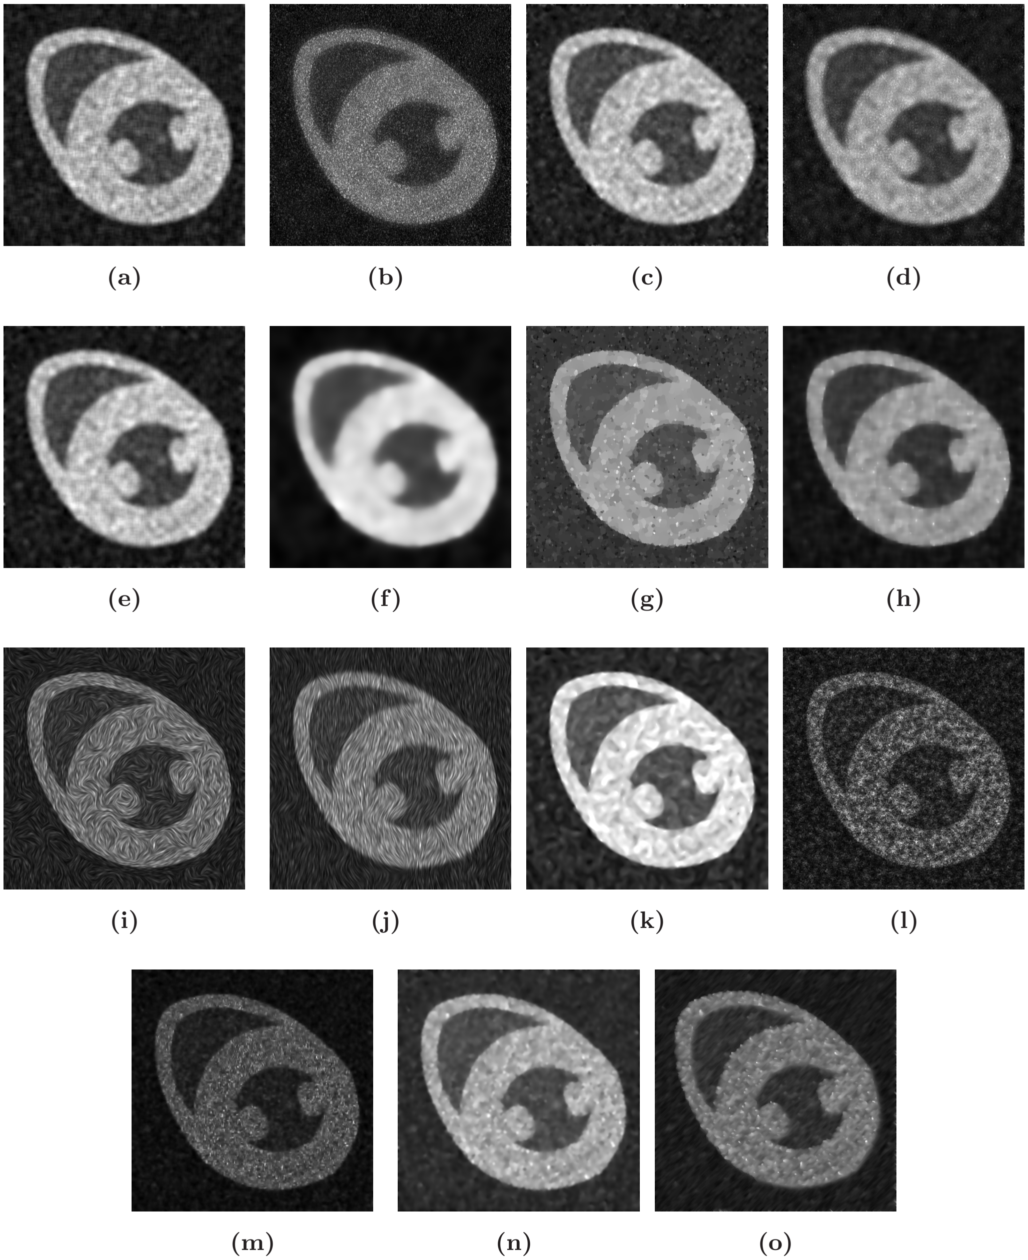

(a) **(b)** **(c)** **(d)**

(e) **(f)** **(g)** **(h)**

(i) **(j)** **(k)** **(l)**

(m) **(n)** **(o)**

Figure 4.8: Sample speckle filter output for simulated images using the template of Fig. 4.3(d). (a) Lee, (b) Kuan, (c) Frost, (d) EnhLee, (e) EnhFrost, (f) PMAD, (g) SRAD, (h) DPAD, (i) CED, (j) NCD, (k) OSRAD, (l) Zong, (m) GLM, (n) NMWD, (o) Geo.

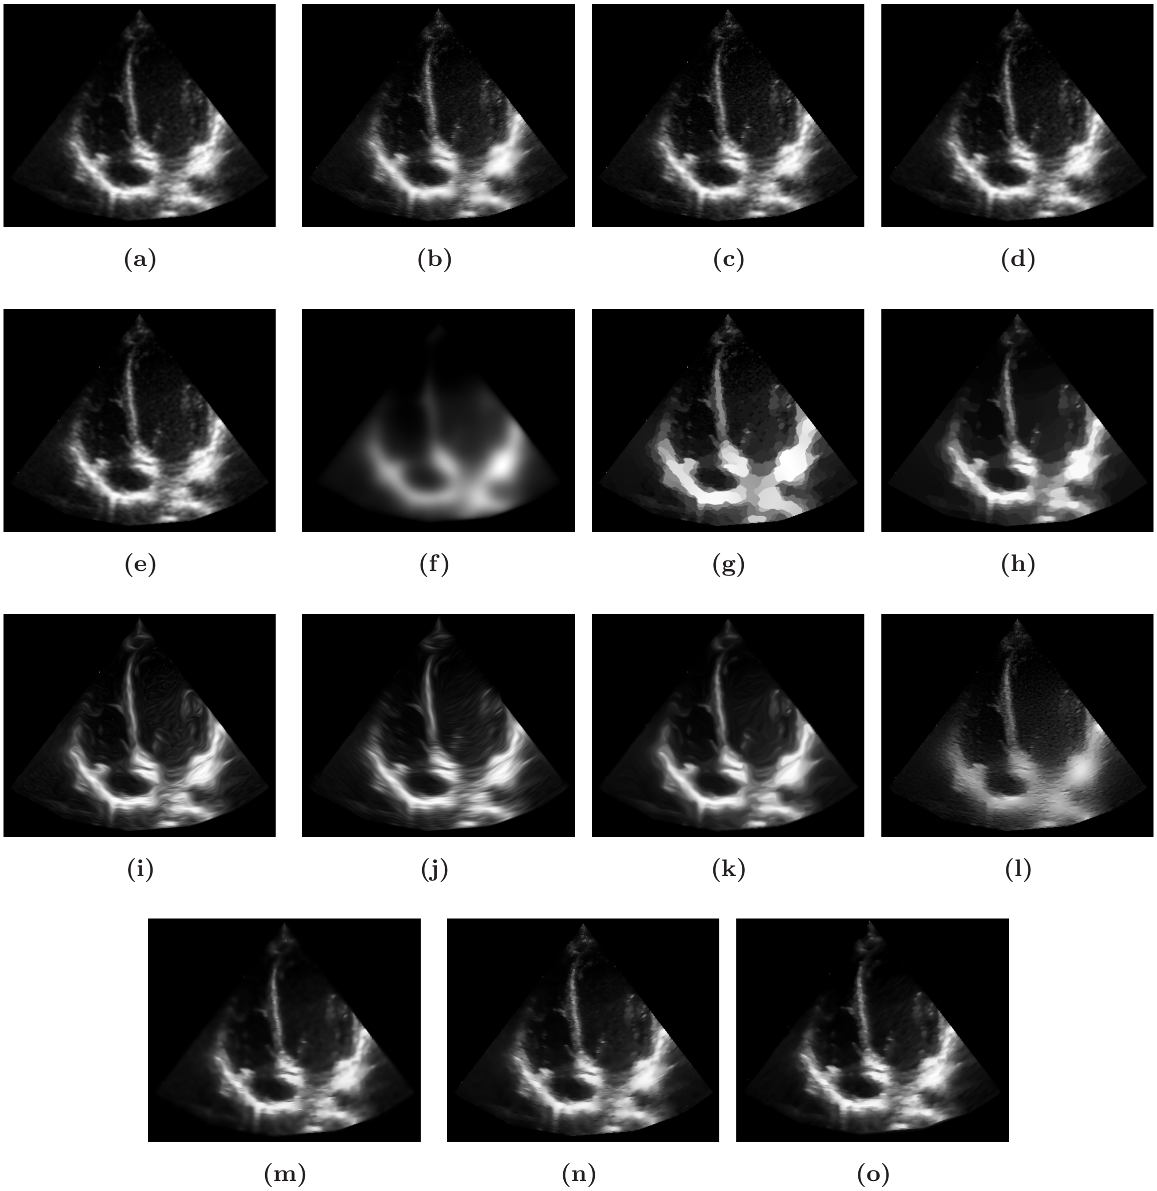

Figure 4.9: Speckle filter output for the clinical image of Fig. 4.5(a). (a) Lee, (b) Kuan, (c) Frost, (d) EnhLee, (e) EnhFrost, (f) PMAD, (g) SRAD, (h) DPAD, (i) CED, (j) NCD, (k) OSRAD, (l) Zong, (m) GLM, (n) NMWD, (o) Geo.

4.4.2 Clinical Images

Filtering of the clinical echocardiographic video frames is performed in the same manner as for the simulated images, after each frame is converted to amplitude form as

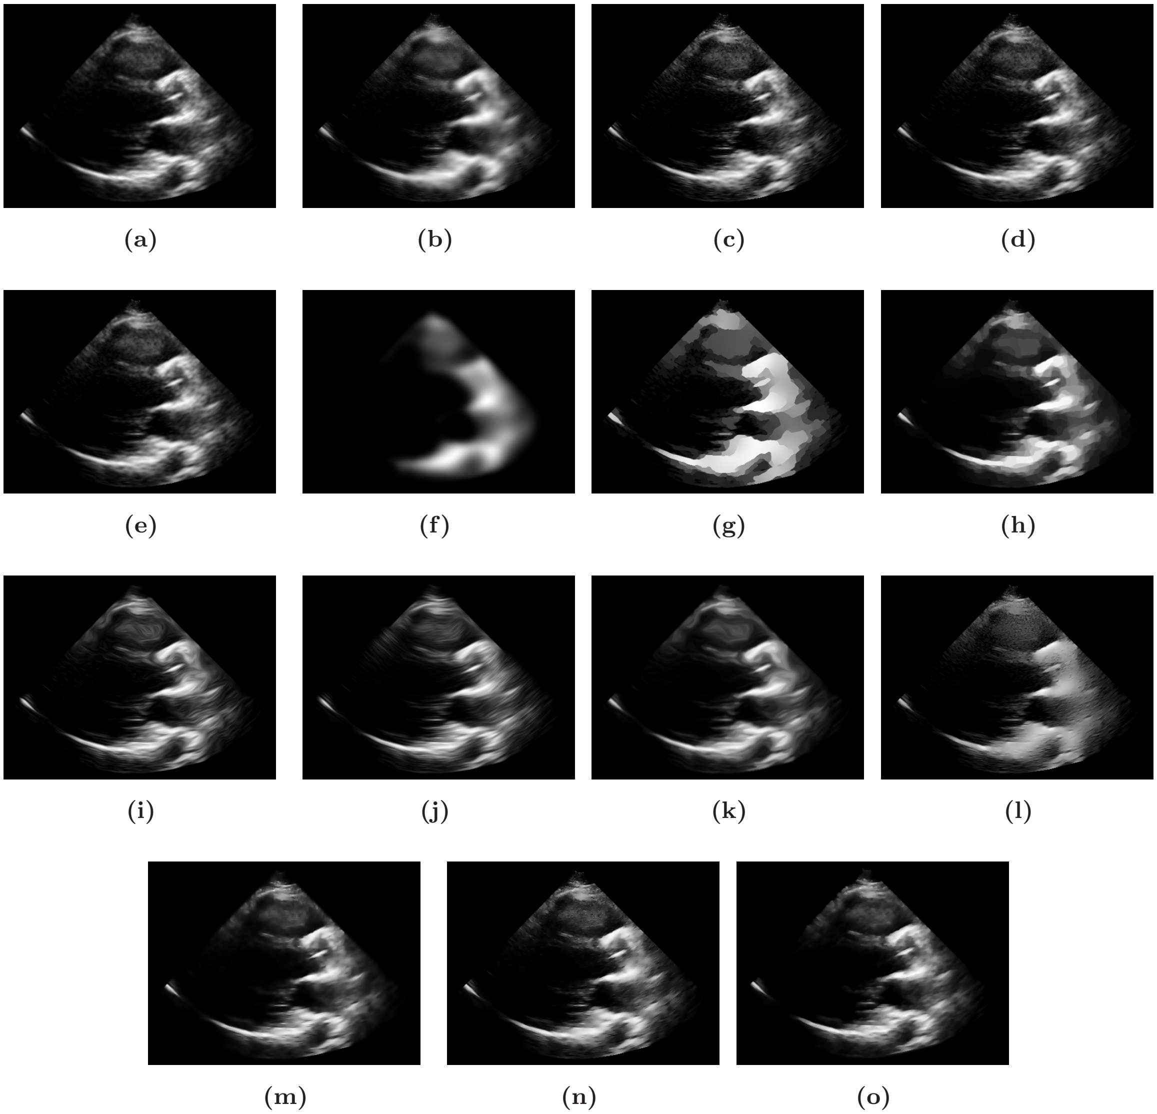

Figure 4.10: Speckle filter output for the clinical image of Fig. 4.5(b). (a) Lee, (b) Kuan, (c) Frost, (d) EnhLee, (e) EnhFrost, (f) PMAD, (g) SRAD, (h) DPAD, (i) CED, (j) NCD, (k) OSRAD, (l) Zong, (m) GLM, (n) NMWD, (o) Geo.

described in Section 4.3.2 above. The filtered versions of the example frames of Fig. 4.5 are displayed in Figs. 4.9-4.12. These are representative of the output from the entire clinical test set.

In visual inspection of the images, it is noted that the SAR filters output (shown in Figs.4.9(a)-(e), 4.10(a)-(e), 4.11(a)-(e) and 4.12(a)-(e)) are quite similar in appearance,

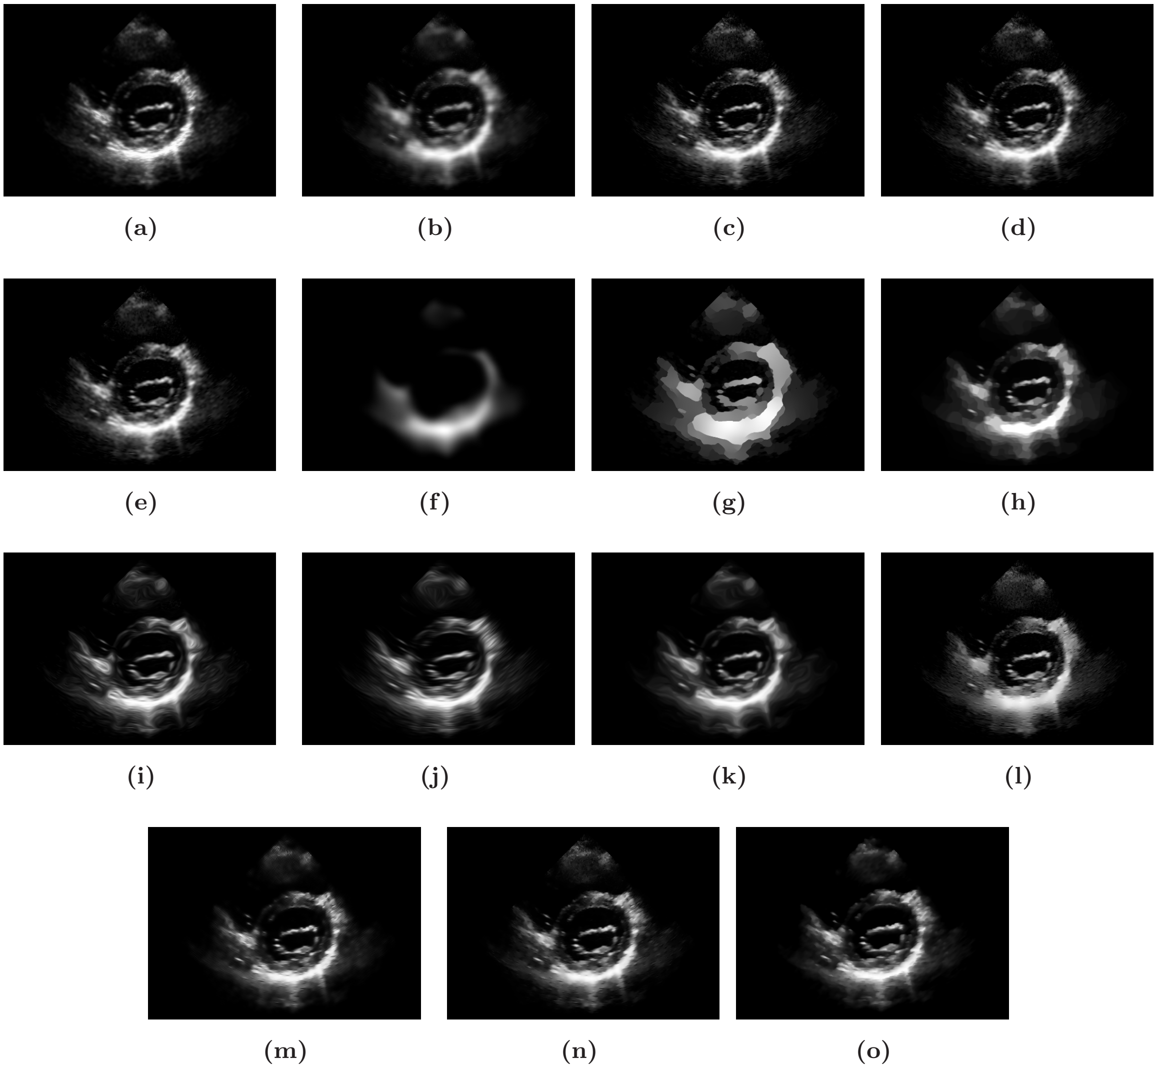

Figure 4.11: Speckle filter output for the clinical image of Fig. 4.5(c). (a) Lee, (b) Kuan, (c) Frost, (d) EnhLee, (e) EnhFrost, (f) PMAD, (g) SRAD, (h) DPAD, (i) CED, (j) NCD, (k) OSRAD, (l) Zong, (m) GLM, (n) NMWD, (o) Geo.

and still contain some speckle.

As with the simulated images discussed above, the anisotropic diffusion filters produce output with a wide range of characteristics. The PMAD filter produces images which are extremely blurred, as can be seen in Fig. 4.9(f), 4.10(f), 4.11(f) and 4.12(f). The SRAD and DPAD filters both show strong speckle suppression, with no trace of speckle remaining in the frames shown in Figs. 4.9(g)-(h), 4.10(g)-(h), 4.11(g)-(h) and 4.12(g)-(h).

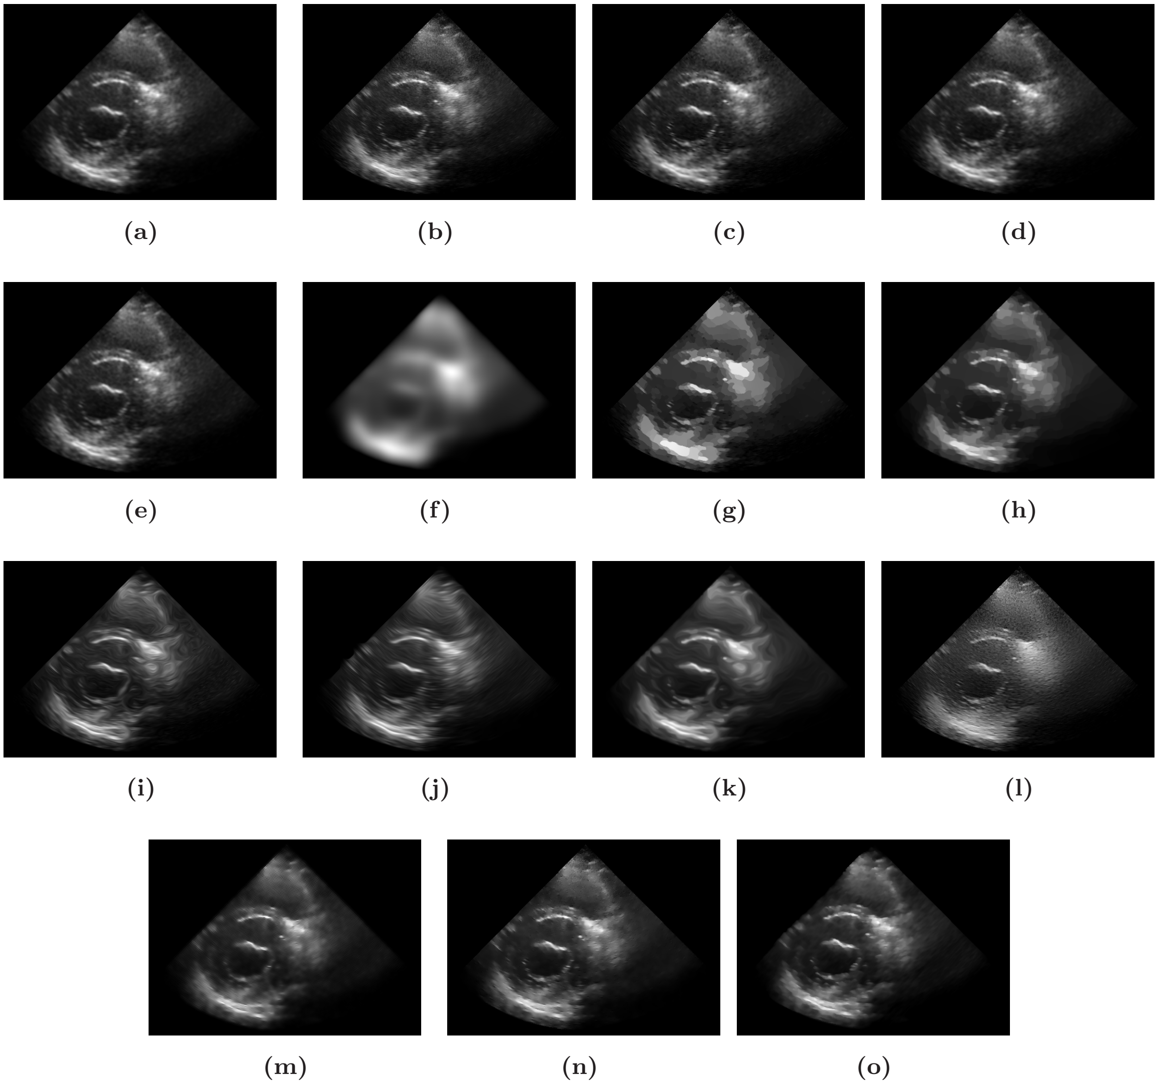

Figure 4.12: Speckle filter output for the clinical image of Fig. 4.5(d). (a) Lee, (b) Kuan, (c) Frost, (d) EnhLee, (e) EnhFrost, (f) PMAD, (g) SRAD, (h) DPAD, (i) CED, (j) NCD, (k) OSRAD, (l) Zong, (m) GLM, (n) NMWD, (o) Geo.

The SRAD output appears more distorted however.

As with the simulated images, the CED and NCD filters (shown in Figs. 4.9(i)-(j), 4.10(i)-(j), 4.11(i)-(j) and 4.12(i)-(j)) introduce small scale artefacts to the images. Again this is due to unwanted enhancement of the contours of the speckle pattern. Similar to the simulated image set, the NCD filter introduces a greater level of this artefact due to the lack of an initial smoothing step. The OSRAD filter shows strong speckle suppression, as can be seen in Figs. 4.9(k), 4.10(k), 4.11(k) and 4.12(k), and the output from this filter appears to have preserved more of the image details than the SRAD and DPAD methods.

Frames processed with the Zong wavelet filter have a somewhat washed-out appearance, and not all of the speckle is removed (as can be seen in Figs. 4.9(l), 4.10(l), 4.11(l) and 4.12(l)).

The GLM, NMWD and Geometric filtered frames (Figs. 4.9(m)-(o), 4.10(m)-(o), 4.11(m)-(o) and 4.12(m)-(o)) have most of the speckle removed, however the GLM output appears slightly blurred.

Table 4.4: Average image quality metrics for speckle filtered clinical amplitude videos, measured relative to speckled input images. The CNR value is the relative difference with respect to the input image. Each column of metric values is arranged from best to worst.

Filter	**FoM**	**Filter**	**SSIM**	**Filter**	**CNR**	**Filter**	**EdgeMSE**	**Filter**	SNR_A **Change**
OSRAD	0.76	Frost	0.92	OSRAD	0.408	NMWD	6.16×10^{-4}	NMWD	70.44%
EnhLee	0.73	EnhLee	0.90	SRAD	0.353	OSRAD	6.21×10^{-4}	OSRAD	62.59%
EnhFrost	0.68	Zong	0.89	NMWD	0.280	Frost	7.80×10^{-4}	SRAD	54.04%
CED	0.65	Kuan	0.88	DPAD	0.276	CED	8.93×10^{-4}	DPAD	39.34%
NCD	0.62	NMWD	0.86	Geo	0.254	EnhFrost	9.54×10^{-4}	Kuan	34.81%
Frost	0.62	CED	0.86	Zong	0.228	Kuan	1.03×10^{-3}	Zong	27.07%
Lee	0.6	EnhFrost	0.84	PMAD	0.221	NCD	1.11×10^{-3}	NCD	25.83%
Kuan	0.57	OSRAD	0.82	NCD	0.191	Lee	1.21×10^{-3}	Lee	23.57%
NMWD	0.56	GLM	0.81	Kuan	0.167	EnhLee	1.34×10^{-3}	EnhFrost	21.45%
GLM	0.54	NCD	0.81	EnhFrost	0.110	GLM	1.51×10^{-3}	GLM	20.66%
DPAD	0.53	Lee	0.79	GLM	0.107	DPAD	1.70×10^{-3}	EnhLee	19.28%
Zong	0.48	DPAD	0.72	Lee	0.105	SRAD	1.80×10^{-3}	CED	18.68%
SRAD	0.44	SRAD	0.69	EnhLee	0.097	Zong	1.84×10^{-3}	PMAD	18.44%
Geo	0.39	Geo	0.59	CED	0.086	PMAD	7.88×10^{-3}	Frost	15.68%
PMAD	0.13	PMAD	0.46	Frost	0.066	Geo	1.20×10^{-2}	Geo	13.54%

Table 4.4 displays the average metrics measured for each filter when applied to the clinical image set. Displayed values are those averaged over the entire clinical test set. Due to the lack of a noise free reference, the metrics in the clinical test measure change relative to the speckled filter input. As expected due to the differing nature of measurement, the metric values differ from that of the simulated test. The columns of this table are arranged from best to worst for each metric. The SNR_A metric shows the amount of speckle suppression

each filter performs. The filters removing the most speckle are the same as in the simulated case, although the level of improvement in the SNR_A metric is lower than for the simulated images. However, visual inspection of the clinical and simulated output indicates that the amount of speckle suppression is no less for the clinical images. The smaller SNR_A increase can be explained by the differences in image content between the simulated and clinical images. Unlike the simulated images, the clinical images contain specular as well as scattered reflections. In addition they contain deviations from Rayleigh statistics [123].

The FoM metric measures the average distortion in edge pixel locations relative to that of the unfiltered noisy image. This is the case shown to correlate with expert opinion on overall quality in [120]. The OSRAD filter achieves the highest average FoM value. The diffusion filters in general achieve mixed FoM scores, while some of the SAR filters achieve quite high scores. Application of the SSIM metrics quantifies the average structural similarity between the filtered and unfiltered images. The Enhanced Lee, Frost and Zong filters are on average the most similar to the speckled input according to this metric. While a high degree of similarity with input images may be generally desirable, the SNR_A values for the Enhanced Lee and Frost filters indicates that this may be due to a low level of speckle reduction filtering in these cases. The average increase in contrast relative to the unfiltered speckled input is measured by the improvement in CNR measure. This increases for all post processing filters, with the OSRAD, SRAD and NMWD attaining the greatest improvement. The lowest average difference in edge region pixel intensity due to filtering is observed for the NMWD and OSRAD filters. The Frost, CED and Kuan filters also preserve the content of edge region pixels quite well as measured by this metric.

4.4.3 Computational Requirements

4.4.3.1 Diffusion Discretisation

Table 4.5: Discretisation Scheme Accuracy.

Discretization	τ	**PMAD Error**	**DPAD Error**
Explicit	0.03	0.01%	0.02%
	0.05	0.02%	0.06%
	0.1	0.06%	0.18%
	0.15	0.10%	0.37%
	0.21	0.14%	0.75%
	0.25	0.32%	1.66%
Semi-implicit	0.03	0.08%	0.20%
	0.05	0.11%	0.31%
	0.1	0.20%	0.61%
	0.15	0.29%	0.89%
	0.21	0.38%	1.14%
	0.25	0.46%	1.31%
	0.5	0.88%	2.14%
	1.25	2.17%	3.72%
	1.67	2.86%	4.31%
	2.5	4.08%	5.18%
	5	6.68%	6.68%
Jacobi	0.03	0.92%	0.37%
	0.05	1.29%	0.55%
	0.1	2.32%	1.04%
	0.15	3.21%	1.53%
	0.21	4.04%	1.98%
	0.25	4.56%	2.26%
	0.5	6.85%	3.45%
	1.25	10.17%	4.91%
	1.67	11.19%	5.27%
	2.5	12.59%	5.71%
	5	14.78%	6.76%

The evaluation of discretisation methods for the PMAD and DPAD diffusivity functions is performed as detailed in Section 4.3.4.1. The results of this analysis are presented in Table 4.5.

The error for all discretisation schemes increases with larger τ, and is in general larger than the figures quoted in [100]. This can be explained by the fact that the reference image used here is a more accurate discretisation, calculated for $\tau = 0.025$, as opposed to $\tau = 0.1$ in [100].

The error measured for the explicit scheme is small as expected, but the valid range of τ is constrained to small values. The error for the semi-implicit scheme and the Jacobi method are higher than the explicit method. It is observed that the semi-implicit scheme performs similarly for both the PMAD and DPAD diffusion methods. In both case the measured error is less than 5% for $\tau = 3$, and less than 7% for $\tau = 5$. The Jacobi discretisation scheme has a slightly higher error for small τ in the DPAD filter, exceeding 5% error for $\tau < 1.5$. The error is similar to the semi-implicit scheme for higher values of τ. For the PMAD method Jacobi discretisation results in a significantly higher error compared to its semi-implicit counterpart, i.e. the accuracy of this discretisation method is observed to vary with the choice of diffusion function.

As noted in [100], the effect of these errors on the resultant images is a weakening of the filtering effect. For application of the anisotropic diffusion filters to the simulated and clinical images in this study, a maximum acceptable error of 5% is chosen. This results in adoption of the AOS scheme in this study, as described by (2.30) for scalar-valued, and (2.35) for matrix-valued diffusion methods, with a maximum value of $\tau = 3$.

4.4.3.2 Overall Filter Requirements

The evaluation of filter complexity is performed taking into account the considerations detailed in Section 4.3.4.2. Results are displayed in Table 4.6, in which the computational requirements are quantified as the number of additions, multiplications, and look up operations required in terms of the parameters of each filter. Computational requirements are independent of all filter parameters except (where applicable) the number of iterations, the window size, and the number of pixels in the image. The NMWD filter uses EM parameter estimation, which can be slow to converge. In this study an average of approximately 500 iterations of are required, and this is reflected in the complexity in Table 4.6. The

Table 4.6: Speckle filter computational requirements. N = number of pixels in image, Q = number of iterations, Z = size of window.

Filter	**Multiplications/Divisions**	**Additions/Subtractions**	**Look Up Operations**
Lee	$5N + NZ(Z+2)$	$N + NZ(Z+2)$	-
Kuan	$6N + NZ(Z+2)$	$2N + NZ(Z+2)$	-
Frost	$5N + 2NZ(2Z+1)$	$2NZ(2Z+1) - 4N$	NZ^2
EnhLee	$8N + NZ(Z+2)$	$2N + 2NZ(Z-1)$	N
EnhFrost	$7N + 2NZ(2Z+1)$	$2NZ(2Z+1) - 2N$	NZ^2
PMAD	$Q(22N-4)$	$Q(19N-12)$	$QN + N$
SRAD	$Q(27N+1)$	$Q(26N-7)$	-
DPAD	$Q(17N + NZ(Z+1) + 1)$	$Q(11N + 2NZ(Z+1) - 7)$	-
CED	$Q(66N-25)$	$Q(56N-29)$	QN
NCD	$Q(60N-21)$	$Q(52N-26)$	-
OSRAD	$Q(64N + NZ(Z+2) - 24)$	$Q(48N + NZ(Z+2) - 29)$	-
NMWD	$Q(261N + 14NZ + 4)$ $+8,004N + 2,530$	$Q(234N + 14Nz - 5)$ $+4,502N - 1,996$	$3QN + 1,003N$
Zong	$288N + 44$	$284N + 16$	$8N + 8$
GLM	$1,160N + 31$	$1,052N - 42$	$30N$
Geo	-	$Q(48N)$	-

computational overhead of the geometric filter is mostly comparative operations, and is not reflected in a quantification of multiplications and additions.

Table 4.7 lists the number of operations required for an image of 512 × 512 pixels, and the parameters of Table 4.1. The most computationally intensive method is the NMWD filter, requiring almost five times as many multiplications as the next most demanding. This high complexity is due to the calculation of the wavelet transform in each iteration, and use of EM parameter estimation. The DPAD and SRAD filters have higher requirements than the other diffusion methods, but this is due to the large number of iterations required. The DWT used in the GLM filter requires much more computation than the DWT of the Zong filter. The efficiency of the semi-implicit scheme for the anisotropic diffusion filters is demonstrated by the similarity between their requirements and the SAR filters.

This section has presented the results of the three individual tests of the filter evaluation framework presented in this chapter:

- The application of the speckle filters to the simulated images, and the use of image

Table 4.7: Filter Computational Requirements, for $N = 512 \times 512$ pixels. Filters are ordered in decreasing complexity, as measured by relative number of multiplications.

Filter	**Mult/Div**	**Add/Sub**	**LUT**	**Relative Mult/Div**
NMWD	2.79x10^{9}	6.3810^{8}	1.18x10^{9}	1
DPAD	6.16x10^{8}	9.31x10^{8}	-	0.22
SRAD	3.54x10^{8}	3.41x10^{8}	-	0.13
GLM	3.04x10^{8}	2.76x10^{8}	7.86x10^{6}	0.11
OSRAD	1.30x10^{8}	1.09x10^{8}	0	4.65x10^{-2}
EnhFrost	9.15x10^{7}	8.91x10^{7}	2.12x10^{7}	3.28x10^{-2}
Frost	9.10x10^{7}	8.86x10^{7}	2.12x10^{7}	3.26x10^{-2}
CED	8.65x10^{7}	7.34x10^{7}	1.31x10^{6}	3.10x10^{-2}
NCD	7.86x10^{7}	6.82x10^{7}	-	2.82x10^{-2}
Zong	7.55x10^{7}	7.44x10^{7}	2.10x10^{6}	2.07x10^{-2}
PMAD	5.77x10^{7}	4.98x10^{7}	2.88x10^{6}	2.01x10^{-2}
EnhLee	2.80x10^{7}	3.83x10^{7}	2.62x10^{5}	1.01x10^{-2}
Kuan	2.75x10^{7}	2.65x10^{7}	-	9.86x10^{-3}
Lee	1.78x10^{7}	1.68x10^{7}	-	6.38x10^{-3}
Geo	0	8.81x10^{7}	-	-

quality metrics to compare the filtered output to the maximally speckle free MW reference image

- Application of the speckle filters to the clinical image set. In this case the image quality metrics are used to quantify the effects of speckle filtering relative to the speckled input values
- The detailed analysis of the computational requirements of each of the speckle reduction filters. In the case of the anisotropic diffusion filters, in which the choice of discretisation method would affect these requirements, an analysis of the accuracy of the two most efficient approaches is presented to justify the choice.

The following section concludes this chapter by discussing these results and using them to determine the optimal speckle reduction approach of those considered.

4.5 Conclusions

This chapter has presented a framework for evaluating the performance of speckle reduction filters when applied to echocardiography. A range of speckle reduction filters are considered, many of which are relatively recent. These methods are applied to both realistically simulated and clinical echocardiographic images, separate tests which highlight different aspects of the filters performance. Evaluation incorporates the use of objective quality metrics, some of which are shown to have a relationship with expert assessment. A detailed examination of the computational requirements of each method is also included in the evaluation.

A large number of numerical results were presented in the previous section. Table 4.8 summarises graphically the conclusions drawn from these numerical results, in categories of speckle reduction, edge preservation, improvement in contrast, and computational complexity.

The use of simulated images permits comparison of speckle reduced filtered output with a maximally noise free reference. Quantification of speckle reduction capabilities using the SNR_A has shown that anisotropic diffusion based methods have in general the strongest suppression of speckle. Some of these filters, in particular the OSRAD and NMWD methods, are able to match the SNR_A of the MW reference for certain regions of scatterer density. This suggests that these filters can approach the maximal amount of speckle suppression. Removal of speckle noise is not the only criterion for successful filtering, particularly for images which are to be used diagnostically. The application of objective metrics quantifies other aspects of the filtering process. In this study two metrics, the FoM and the edge region MSE, quantify the similarity in image edges between the filtered and MW reference images. It is found that the matrix diffusion filters (OSRAD,CED and NCD) have the closest average edge region intensities to the MW images. However, the similarity in location of edge pixels, as measured by the FoM metric, is higher for other methods. The improvement in CNR values of the PMAD, OSRAD and NMWD filters shows that these methods can achieve greater contrast than the MW reference. Low overall values in the SSIM metric here indicates that although some of the filters are capable of suppressing speckle to a similar level as that of the MW reference, their output is not structurally similar.

Table 4.8: Summary of relative filter performance. Green in a particular category indicates relatively good performance for that filter, amber indicates a medium performance in that category, while red indicates poor performance.

Filter	**Speckle Reduction**	**Edge Preservation**	**Contrast Improvement**	**Complexity**
Lee	🟧	🟧	🟥	🟩
Kuan	🟧	🟩	🟥	🟩
Frost	🟥	🟩	🟥	🟧
EnhLee	🟥	🟩	🟥	🟧
EnhFrost	🟧	🟧	🟥	🟩
PMAD	🟧	🟥	🟩	🟩
SRAD	🟩	🟧	🟩	🟥
DPAD	🟩	🟧	🟩	🟥
CED	🟧	🟩	🟥	🟧
NCD	🟩	🟩	🟧	🟧
OSRAD	🟩	🟩	🟩	🟧
Zong	🟥	🟩	🟥	🟧
GLM	🟥	🟥	🟧	🟥
NMWD	🟩	🟧	🟩	🟥
Geo	🟧	🟥	🟧	-

Some of the filters are capable of removing most (and in some cases all) of the speckle in simulated images. In the case of the SRAD and OSRAD filters this is done with similar edge region pixel values to the maximally speckle reduced MW image, but with a higher level of contrast.

Application of speckle reduction to the set of clinical echocardiographic videos, and the subsequent calculation of objective metrics, quantifies how filtered clinical images compare to the unfiltered images. Average edge pixel distortion due to filtering was lowest in the matrix diffusion and SAR filters, as seen by both the high FoM and low edge region MSE values. This indicates that these methods distort image boundaries the least amount. In the case of the SAR filters however, this is due to a low overall level of filtering. Evidence for this can be found in the high SSIM values for these filters, representing a large similarity between the SAR filter output and the speckle corrupted input, and observation of Figs. 4.9-

4.12. The OSRAD filter also shows the greatest average improvement in contrast relative to the unfiltered input.

Based on analysis of computational complexity, it is clear that there is a large disparity in the requirements of the speckle reduction methods considered here. The SAR and geometric filters have the lowest computational overhead, but this comes at the expense of lower speckle reduction capability. The wavelet based approaches are hindered from a performance perspective by the requirements of implementing wavelet analysis and reconstruction. In particular, the NMWD filter performs wavelet analysis and reconstruction for each iteration, leading to the largest requirement of all considered methods. The anisotropic diffusion methods all have similar processing needs, and these fall between those of the SAR methods and the wavelet based filters. Efficient implementation of these methods is only possible by the use of a discretisation method which allows a large timestep. It is shown in Section 4.4.3 that the accuracy of the Jacobi discretisation scheme varies with the choice of diffusion function. The AOS discretisation scheme is shown to have a greater accuracy, independence of diffusion function, and in addition can be implemented in a very efficient manner as described in [100].

In general the anisotropic diffusion (in particular OSRAD) and NMWD filters exhibit the strongest speckle suppression. The other wavelet based methods are not capable of a comparable level of speckle removal in the study reported here. This also applies to the geometric and SAR filters. The matrix based diffusion methods perform well in terms of how they affect image edges in both simulated and clinical applications. To determine which filtering method is optimal for application to echocardiography, various criteria should be taken into consideration. If the main concern is a constraint on available processing capability, the SAR filters are the best due to their low requirements. In particular the Lee filter is a reasonable choice given its speckle suppression ability and low overhead. If however the main objective is to remove as much speckle as possible, the NMWD filter has the strongest speckle suppression capabilities. This comes at the expense of the highest computational complexity however, and the preservation of edges is not optimal. The OSRAD method represents the best trade-off between both of these situations. The level of speckle suppression achievable using this approach is very close to that of the NMWD, with the advantage of a much smaller processing overhead. In addition, image edges experience a minimum of blurring or other distortion. The OSRAD method is therefore considered the best suitable of those considered for clinical application, based on consideration of its performance on both simulated and clinical data, and also evaluation of its computational requirements.

This chapter has presented a framework for evaluating the performance of speckle reduction filters for speckled ultrasonography, with a focus on echocardiographic imaging. This evaluation framework is applied to a range of such filters, and the OSRAD is deemed the optimal approach. The next chapter presents a method of accurately detecting edges in speckled ultrasound imagery, which can also be used as the basis for a method of ultrasonic speckle removal.

Chapter 5

Edge Detection and Speckle Reduction using Gabor Zero Crossings

5.1 Introduction

This chapter details the development of a method for edge detection in the presence of speckle noise. In ultrasonography, accurate detection of tissue boundaries is important for higher level tasks such as image segmentation and analysis. The task of automated edge detection is however hindered by speckle. The widely used Laplacian of Gaussian (LoG) [124] method has been shown to result in a large number of erroneous edges in speckled images [109]. Gradient based edge detection, using both simple differences and the Sobel operator [125] display intensity dependent false edges due to speckle [112]. Edge detection is also often performed as a step in speckle reduction methods. The speckle reducing anisotropic diffusion of [3] (expanded upon in [60]), uses an edge detection function suited to speckle, which is evolved iteratively along with diffusion based smoothing of the image. The effectiveness of this edge detector is further analysed in [99]. In [80], which presented a wavelet based approach, an intermediate step involves empirical estimation of a mask which determines the location of the signal in each wavelet subband, effectively an edge detector. In the hybrid wavelet-diffusion method proposed in [81], the normalised modulus of wavelet coefficients is used to locate image edges.

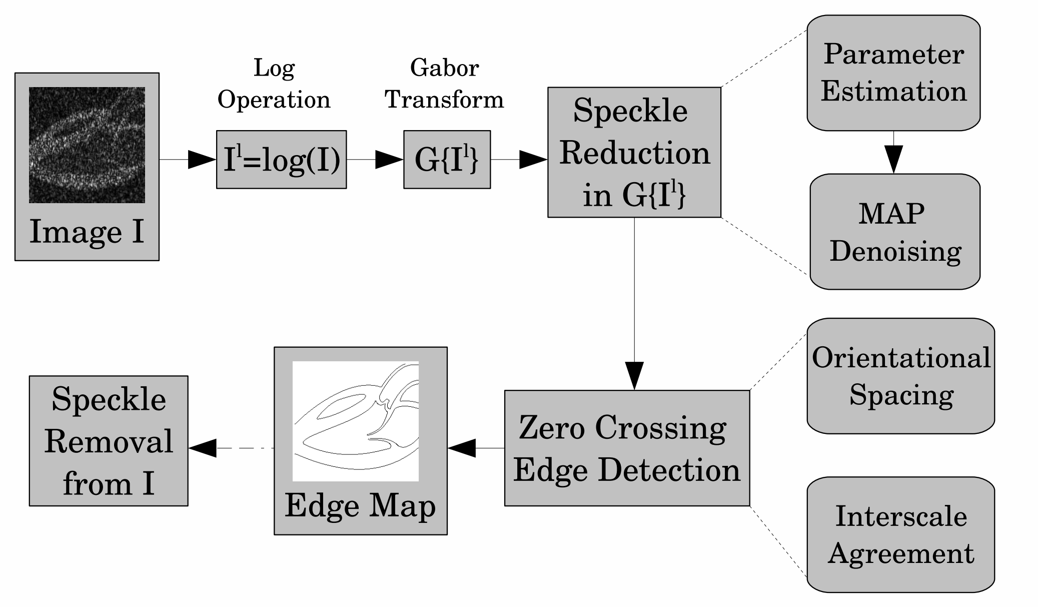
Image I
Log Operation
I^l=log(I)
Gabor Transform
G{I^l}
Speckle Reduction in G{I^l}
Parameter Estimation
MAP Denoising
Orientational Spacing
Interscale Agreement
Zero Crossing Edge Detection
Edge Map
Speckle Removal from I

Figure 5.1: Flowchart of the edge detection scheme presented in this chapter.

The edge detection method presented in this chapter is comprised of a number of steps, summarised graphically in Fig. 5.1. Taking a speckled image as input, the first step performed is a logarithmic operation, to allow treatment of the speckle as additive noise. This is followed by decomposition of the speckled ultrasound images into scale and orientation specific subbands using Gabor filtering. The zero crossings of the Gabor coefficients are used to find the image edges. To avoid detecting spurious edges due to speckle, statistical estimation is used to remove the Gabor coefficients resulting from image speckle. This step includes accurate modelling of the Gabor coefficients as a combined symmetric α stable (SαS) and Gaussian process, and the estimation of the parameters of this model. These parameters are then used to perform MAP estimation of the speckle free coefficients. Edge detection is performed using the zero crossings of the Gabor transform. In order to determine which zero crossings coincide with image edges, the behaviour of zero crossings in the vicinity of image edges is observed. Two properties are observed and investigated: the spacing of zero crossings in the direction of Gabor orientation (referred to as *orientational spacing* here) and the agreement between zero crossings in adjacent scales (referred to as *interscale agreement*). These properties are quantified and combined to generate the edge detection measure presented in this chapter. An accurate edge map is produced by these steps, which can be used for a number of tasks including image enhancement, segmentation or aiding in speckle removal. This latter use for edge maps is proposed in this chapter, by their inclusion in an anisotropic diffusion method for speckle reduction.

In the proposed scheme, speckle is removed from the Gabor transform of the image as an intermediate step. As discussed in Chapter 2, a number of approaches to speckle reduction proceed by removing the coefficients due to speckle from the wavelet transform of an image, followed by an inverse transform to reconstruct the speckle free image. Reconstruction of the image from the speckle reduced Gabor coefficients here could in theory yield a speckle free image also. However, reconstruction is usually not possible from Gabor coefficients due to their non-orthogonal nature. Lee [126] presented an analysis of this issue, and determined the specific cases in which stable reconstruction is possible. Image edges are found here within the Gabor domain without the need for reconstruction. One possible use for these edges is a second speckle removal step, which removes the speckle component from the original image. This application is explored here, by incorporation of the edge maps into a diffusion process for speckle removal.

In Chapter 6, the proposed edge detection scheme is compared to a number of other methods found in the literature in a series of tests involving simulated, phantom and clinical ultrasound images. The primary contributions of the research presented in this chapter are:

- The proposed method of parameter estimation for the Gabor transform model
- The method of edge detection via Gabor zero crossings
- Proposal of a speckle reducing diffusion method, based on a static edge map

The organisation of this chapter is as follows: Background information on relevant topics is given in Section 5.2, including Gabor filtering, stable statistical distributions, and details of some existing methods of edge detection. Section 5.3 details the modelling of the ultrasound Gabor coefficients. This includes a demonstration of the accuracy of this model, details of the proposed method of estimating the model parameters, and statistical removal of speckle using MAP estimation. The method of edge detection proposed in this chapter is developed in Section 5.4. The two properties of interest displayed by Gabor zero crossings in image edge locations, orientational spacing and interscale agreement, are discussed here. Methods of quantifying these properties are presented, and their use in constructing an edge detection measure is detailed. The use of the proposed zero crossing edge detector for speckle removal, by its inclusion in a tensor valued diffusion scheme, is presented in Section 5.5. Concluding remarks are presented in Section 5.6.

5.2 Background

This section presents information in three areas of relevance to the topics discussed in this chapter:

- Gabor filtering
- Stable statistical distributions
- Existing methods of edge detection

5.2.1 Gabor Filters

Gabor [127] presented a seminal proof that a fundamental limit exists to the level of specificity in which a signal can be simultaneously represented in both time and frequency. A class of functions is also presented for the representation of time varying signals, providing maximal simultaneous localisation in both time and frequency. In the time domain, these functions are constructed as the product of a complex sinusoid with a Gaussian envelope:

$$g(t) = \exp\left(-\frac{t^2}{2\sigma^2} + 2\pi j f_0 t\right) \tag{5.1}$$

Where f_0 is the frequency of the complex sinusoid, $j = \sqrt{-1}$, and σ determines the shape of the Gaussian envelope, and therefore the spatial extent of the function. Scaling constants have been omitted here for simplicity. The real and imaginary components from a quadrature pair, and have a phase offset of $\frac{\pi}{2}$. The frequency response of these functions has a Gaussian profile, positioned on the frequency plane centred at f_0. An area in which these Gabor functions have found widespread use is in the early vision models, Unidimensional Gabor function are used to model the profiles of cortical cells in [128, 129]. Daugman [130] extended the use of Gabor functions in visual modelling to two-dimensional Gabor functions. Two dimensional Gabor functions are directional in nature, and when applied to a signal they encode information in a preferred orientation. They can be constructed by extending (5.1) to the product of a two dimensional Gaussian envelope $e(x, y)$ and a complex sinusoid carrier $c(x, y)$:

$$\begin{aligned} g(x, y) =& e(x, y)c(x, y) \\ c(x, y) =& \exp\{2\pi j \left(u_0 x' + v_0 y'\right)\} \\ e(x, y) =& \exp\{-\frac{x'^2 + \gamma y'^2}{2\sigma^2}\} \\ x' =& x \cos\theta + y \sin\theta \\ y' =& -x \sin\theta + y \cos\theta \end{aligned} \tag{5.2}$$

where the centre frequency of the function is defined by (u_0, v_0), and the angle of orientation by θ. The parameters of the Gaussian envelope σ and γ define its spread and aspect ratio (level of ellipticity), respectively. A set of these 2D Gabor functions can be applied to an image in a filter bank, to encode image information from a range of frequencies and orientations. Such filter banks have found widespread use in image processing, perhaps inspired by their success in modelling aspects of mammalian visual systems produced by evolution. Common applications include texture classification and texture edge detection, including applications to ultrasonography found in references [131–134].

A modification to the 2D Gabor functions proposed by Field [135] is employed here. The motivation in this work is again the modelling of visual cortical elements. It is noted that the efficiency of an image representation by a particular image coding scheme is best determined by considering the properties of those images likely to be encountered. Images typical of the 'natural environment' are considered, such as those representing trees, rocks, and other common outdoor scenes. The efficiency of encoding such images is best achieved by Gabor filters which have a Gaussian frequency distribution in the log frequency domain, known as log Gabor filters. In the linear frequency domain, a log Gabor filter has the transfer function:

$$G(\omega) = \exp\left(-\frac{\log\left(\omega/\omega_0\right)^2}{2\log\left(\sigma/\omega_0\right)^2}\right)\exp\left(\frac{-(\Theta - \Theta_0)^2}{2T\Delta\Theta^2}\right) \tag{5.3}$$

where $\omega_0 = (u_0, v_0)$ is the filters centre frequency, σ controls radial bandwidth, Θ_0 is the orientation angle of the filter, $\Delta\Theta$ is the spacing between filters of different orientations, and T is a scaling factor.

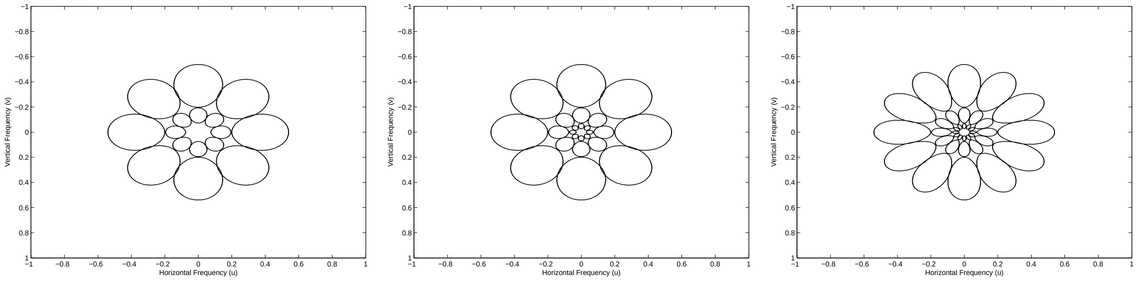

(a) **(b)** **(c)**

Figure 5.2: Half magnitude plot of log Gabor filter banks, with various scales and orientations. Spectra are mirrored in the negative half of the frequency plane. Radial bandwidth is two octaves in all cases. (a) 2 scales, 4 orientations, (b) 3 scales, 4 orientations, (c) 3 scales, 6 orientations.

These allow a wider bandwidth than is possible with conventional Gabor filters by (5.2), which develop an unwanted DC component for higher bandwidth values. The linear-phase nature of these filters has led to their use in extraction of local frequency information, for phase based edge detection [136–138]. In [5], adapted log Gabor filters are applied to ultrasound RF (radio frequency) data, to reduce speckle noise. These filters are usually applied in a group, or bank, containing filters of differing centre frequency and orientation.

Banks of log Gabor filters can be generated by varying the parameters of (5.3). As neither the original 2D Gabor functions of [130] or the log Gabor functions described above form orthonormal bases, there is no optimal arrangement of filters. Figs. 5.2-5.3 display

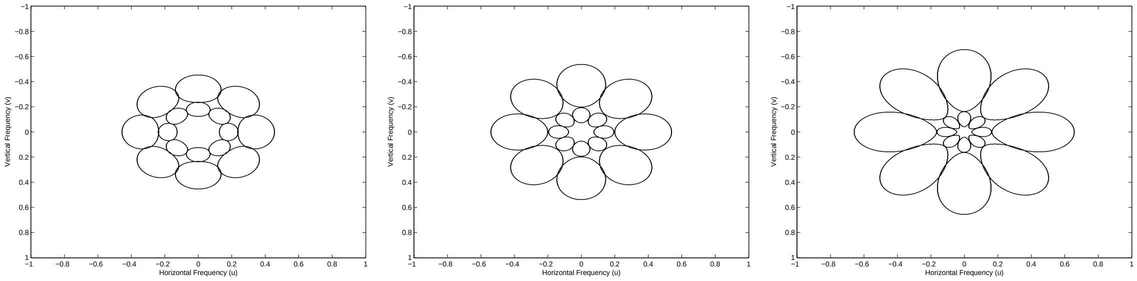

(a) **(b)** **(c)**

Figure 5.3: Half magnitude plot of log Gabor filter banks, with two scales and four orientations, showing the effect of various radial bandwidth values. (a) ~ 1 octaves, (b) ~ 2 octaves, (c) ~ 3 octaves.

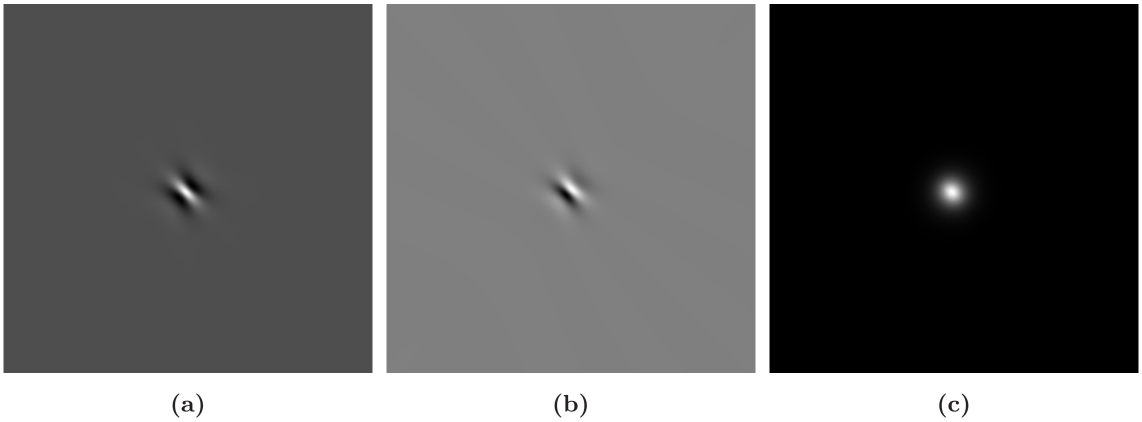
(a) (b) (c)

Figure 5.4: Time domain representation of a single log Gabor filter, orientated at $\theta = \frac{\pi}{3}$ rad. (a) Even symmetric, (b) Odd symmetric, (c) Magnitude.

the spectra of different arrangements of Gabor filter banks, as half magnitude profiles on the frequency plane. The filter banks in Fig. 5.2 have the same radial bandwidth, and the number of scales and orientations is varied. To maintain a constant radial bandwidth over different scales, the ratio σ/ω_0 must be held constant as the centre frequency ω_0 is changed. Fig. 5.3 displays the effect of varying radial bandwidth, from approximately one octave in Fig. 5.3(a) to three octaves in 5.3(c). In Figs. 5.2-5.3, the radial and orientational spacing of the filters is chosen so the half magnitude profiles touch without overlap. This approach ensures a high level of independence in output of the filters. A competing filter bank design criterion is the desire for even spectral coverage, which requires a degree of overlap in the regions of the spectrum covered by each filter. This reduces the independence of the filter outputs. The effect of filter bank parameters on the edge detection scheme of this chapter is explored in Section 5.4 below.

While the log Gabor filters are designed and applied in the frequency domain, a time domain representation may be acquired by application of the inverse Fourier transform. Time domain log Gabor filters are complex valued, with the real and imaginary parts containing the even and odd symmetric components of a quadrature pair. Fig. 5.4 displays an example time domain representation of a single log Gabor filter, including both quadrature components and their magnitude.

A log Gabor filter bank with N_{scale} scales and N_{ori} orientations is applied in the frequency domain by

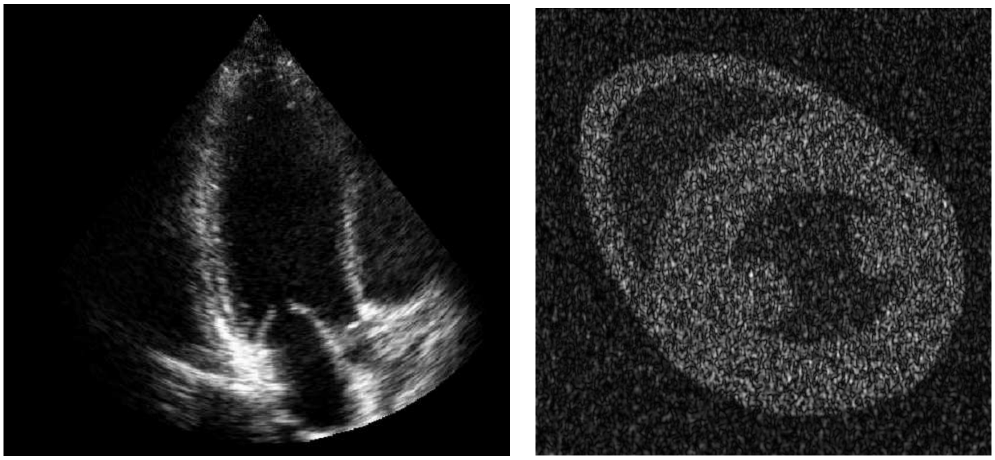

(a) **(b)**

Figure 5.5: Sample simulated and clinical images, used to demonstrate the effect of log Gabor transformation. (a) Clinical image, (b) Simulated image.

$$G_{s,o}\{x\} = \mathcal{F}^{-1}\{G_{s,o}.\mathcal{F}\{x\}\} \tag{5.4}$$

where $\mathcal{F}$ represents the Fourier transform, and $(o = 1, 2, .., N_{ori},\ s = 1, 2, .., N_{scale})$ are the indices of scale and orientation, respectively. To demonstrate the effect of log Gabor filtering on example ultrasound image, a log Gabor filter bank with spectra as in Fig. 5.2(a) is applied to the simulated and clinical images of Fig. 5.5. The magnitude of the Gabor response to the clinical image of Fig. 5.5(a) is shown in Fig. 5.6, while the magnitude of the Gabor output of the simulated image of Fig. 5.5(a) is displayed in Fig. 5.7.

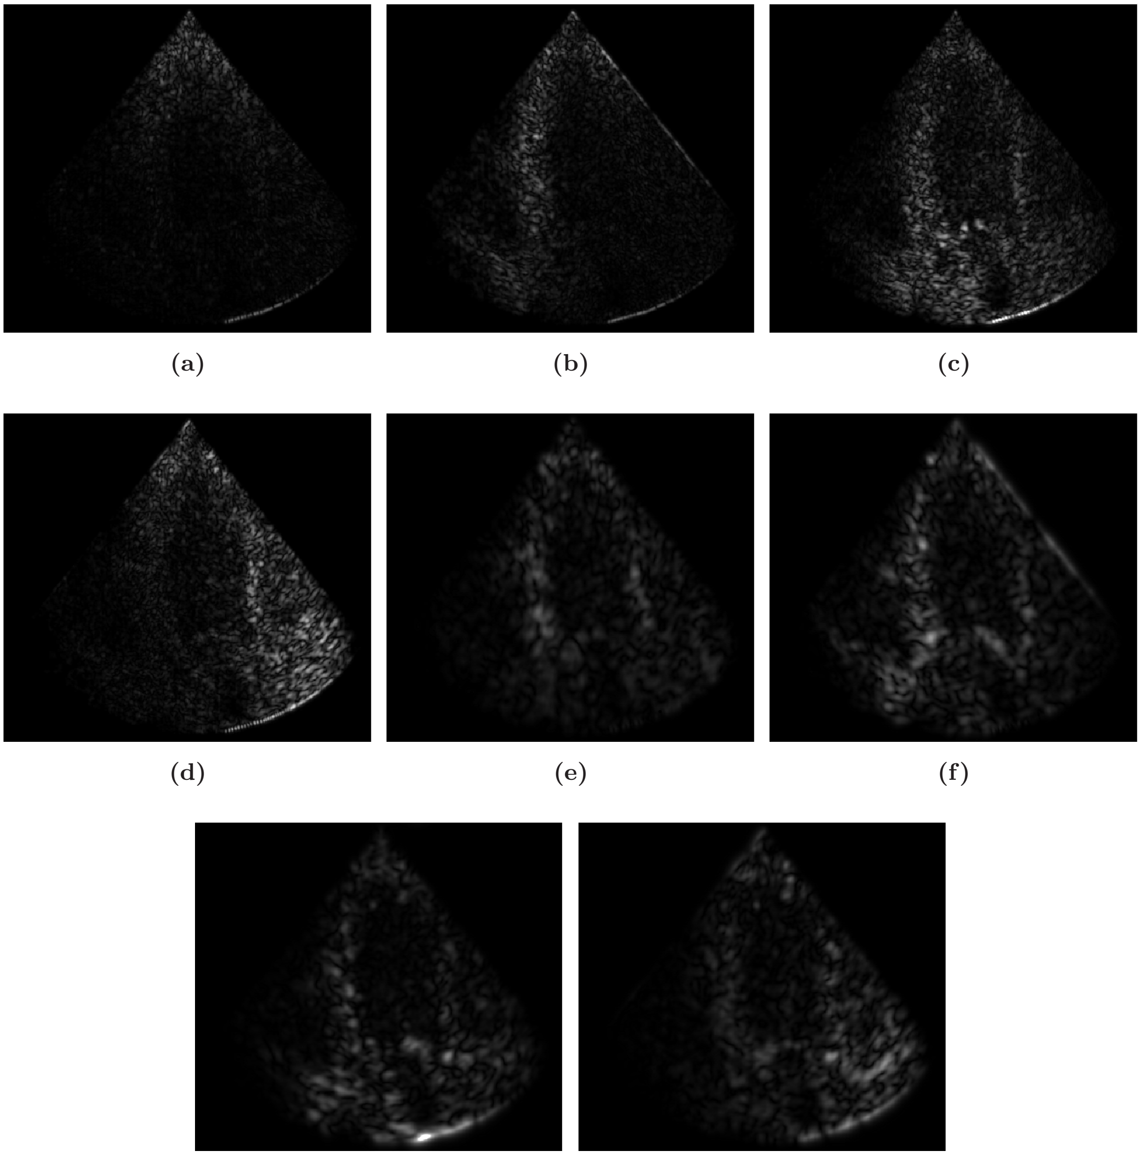

(a) (b) (c)

(d) (e) (f)

(g) (h)

Figure 5.6: Magnitude of Gabor transform of the clinical image of Fig. 5.5(a). (a) Scale 1, $\theta = 0$ rad, (b) Scale 1, $\theta = \frac{\pi}{4}$ rad, (c) Scale 1, $\theta = \frac{\pi}{2}$ rad, (d) Scale 1, $\theta = \frac{3\pi}{4}$ rad, (e) Scale 2, $\theta = 0$ rad, (f) Scale 2, $\theta = \frac{\pi}{4}$ rad, (g) Scale 2, $\theta = \frac{\pi}{2}$ rad, (h) Scale 2, $\theta = \frac{3\pi}{4}$ rad.

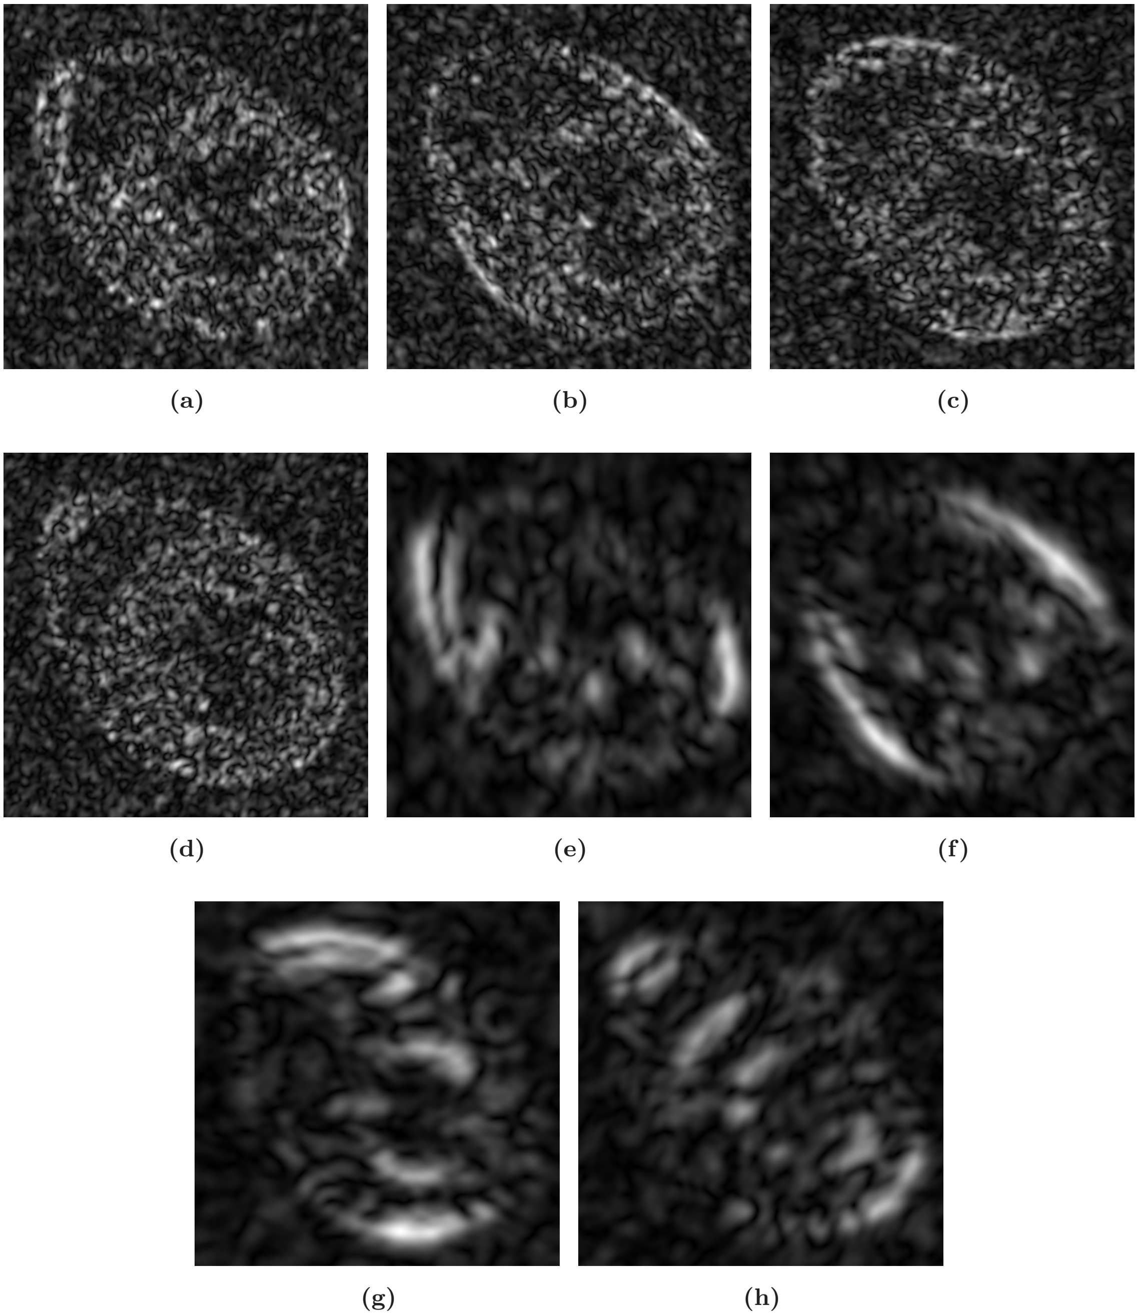

(a) (b) (c)

(d) (e) (f)

(g) (h)

Figure 5.7: Magnitude of Gabor transform of the simulated image of Fig. 5.5(b). (a) Scale 1, $\theta = 0$ rad, (b) Scale 1, $\theta = \frac{\pi}{4}$ rad, (c) Scale 1, $\theta = \frac{\pi}{2}$ rad, (d) Scale 1, $\theta = \frac{3\pi}{4}$ rad, (e) Scale 2, $\theta = 0$ rad, (f) Scale 2, $\theta = \frac{\pi}{4}$ rad, (g) Scale 2, $\theta = \frac{\pi}{2}$ rad, (h) Scale 2, $\theta = \frac{3\pi}{4}$ rad.

5.2.2 Stable Processes

This section provides a brief introduction to α stable statistical models. For a thorough review, the reader is referred to the literature, particularly [139–143]. Stability in the statistical sense refers to a property of some statistical distributions whereby they remain unchanged in shape under addition. Specifically, a random variable X is stable if the sum of two or more independent variables, each with distribution X, has itself that distribution, i.e.:

$$a_1X_1 + a_2X_2 + \cdots + a_nX_n \doteq cX + d \tag{5.5}$$

for $c \in \mathbb{R}_{>0}$, $a_n, d \in \mathbb{R}$, and $\doteq$ represents equality of distribution. There are three cases of stable random variables for which closed forms are expressible for the probability density function (PDF); these are the Gaussian, Cauchy and Lévy distributions [141].

The general case of stable variables are expressed in terms of their characteristic functions. Characteristic functions are an alternate method of describing a random variable. For a random variable X with a PDF $p_x(x)$, the characteristic function is given by:

$$\phi_x(t) = \int_{-\infty}^{+\infty} \exp\{jtx\} p_x(x)\, dx \tag{5.6}$$

where $t \in \mathbb{R}$ is the characteristic function argument. The characteristic function is the dual of the PDF, in the same way that the Fourier transform of a variable is the dual of its time domain representation. A number of parameterisations are used to describe the general case of stable distributions, a common approach is of the form:

$$\phi_x(t) = \exp\{j\delta t - \gamma|t|^{\alpha}\left(1 - \beta sign(t)\omega(t;\alpha)\right)\} \tag{5.7}$$

where $0 < \alpha \leq 2$ is the characteristic exponent, $\gamma \in \mathbb{R}_{>0}$ is the dispersion parameter, $\beta \in [-1, 1]$ determines the skewness of the distribution, and $\delta \in \mathbb{R}$ is the location parameter. The expression $sign(t)$ refers to the sign of t, while the expression $\omega(t;\alpha)$ is given by:

$$\omega(t;\alpha) = \begin{cases} \tan\frac{\pi\alpha}{2} & \text{if } \alpha \neq 1 \\ -\frac{2}{\pi}\log|t| & \text{if } \alpha = 1 \end{cases} \tag{5.8}$$

The closed form cases are found for parameter values of $\alpha = 2$ (where (5.7) reduces to a Gaussian distribution), $\alpha = 1$, $\beta = 0$ (reduces to a Cauchy distribution) and $\alpha = 0.5$,

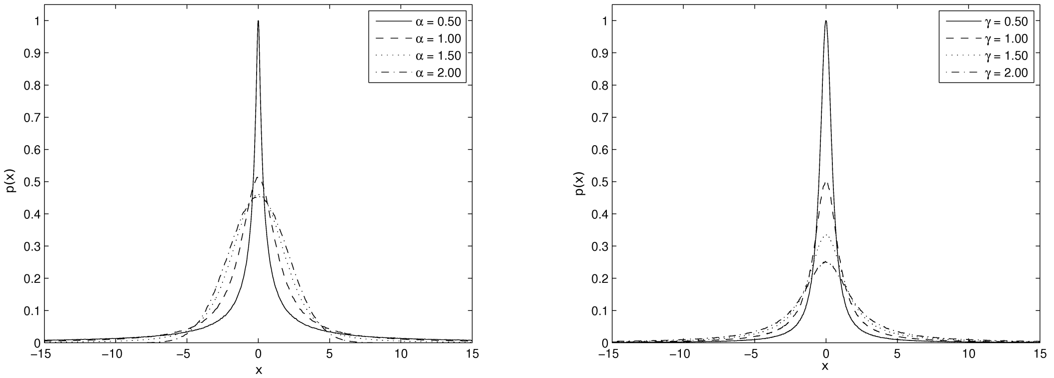

(a) **(b)**

Figure 5.8: Example symmetric α stable probability density functions. (a) Varying α, for $\gamma = 1.5$, (b) Varying γ, for $\alpha = 1.0$.

$\beta = 1$ (reduces to a Lévy distribution). Stable distributions have been used to model a wide range of heavy tailed stochastic processes [144–147]. For the modelling of Gabor coefficients presented in this chapter, the symmetric subclass of these distributions, known as Symmetric α Stable (SαS), are employed. SαS distributions are found by (5.7), as those distributions without any skewness, and centred at the origin ($\beta = \delta = 0$). The characteristic function is thus greatly simplified:

$$\phi_s(t) = \exp\{-\gamma|t|^\alpha\} \tag{5.9}$$

A set of empirically calculated SαS PDFs are shown in Fig. 5.8. It can be seen that for small α values, these distribution have heavy tails and pronounced peaks. As α increases this effect is lessened, until Gaussian behaviour is seen for $\alpha = 2$ The effect of increasing γ for a given α value is to make the density more dispersed, with a less pronounced peak.

5.2.3 Existing Edge Detection Methods

To evaluate the performance of the edge detection method proposed in this chapter, its performance is compared against a number of existing edge detection approaches from the literature:

- CV (coefficient of variation) based anisotropic diffusion
- The two step RoA (ratio of averages) method
- PC (phase congruency) edge detection
- The FA (feature asymmetry) method

Some well known approaches to edge detection have been shown to perform poorly in speckled imagery, notably the Laplacian of Gaussian (or LoG, discussed in [109]) and gradient based methods [112]. These are therefore not implemented here. Instead the methods compared to the proposed edge detector are those which have been presented specifically for edge detection in the presence of speckle (CV, RoA, FA). The exception to this is the PC method, but as this approach is phase based (and so intensity independent), in theory it may be capable of good performance in the case of speckled images. In addition, a quantification of the PC methods edge detection capabilities for speckled images has not been published, to the best of the author's knowledge.

5.2.3.1 Coefficient of Variation Edge Detection

CV based filtering has been used for the removal of speckle in SAR imaging [42, 44–46]. Use of the coefficient of variation (the ratio of local deviation to mean) is extended to anisotropic diffusion based speckle removal in [3]. The use of an arbitrary sized window for calculation of the local statistics is explored in [60]. As detailed in Section 2.4.1.3 above, the CV measure is calculated at each location (x, y) and diffusion iteration t as:

$$CV(x,y;t) = \sqrt{\frac{\sigma_I^2(x,y;t)}{\bar{I}(x,y;t)^2}} = \sqrt{\frac{\frac{1}{|\eta_{x,y}|-1}\sum_{p\in\eta_{x,y}}\left(I_p - \bar{I}(x,y;t)\right)^2}{\bar{I}(x,y;t)^2}} \tag{5.10}$$

where $I(x, y; t)$ is the image at diffusion iteration t, $\eta_{x,y}$ is a square $N \times N$ neighbourhood, and $\bar{I}(x, y; t) = (1/|\eta_{x,y}|) \sum_{p\in\eta_{x,y}} I_p$. Similar to the evaluation of [99], which considers

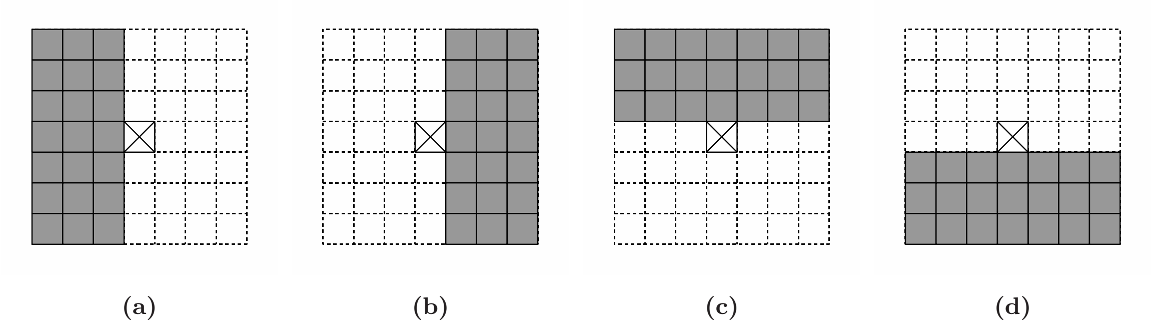

Figure 5.9: Local masks used in the RoA edge detection method, from a 7 × 7 window. The shaded square represents the centre pixel. (a) Left, (b) Right, (c) Up, (d) Down.

CV based edge detection for a fixed neighbourhood, diffusion is performed until the result becomes stable (i.e. no further diffusion is taking place) and the final value of CV used for edge detection. The explicit method of discretisation is used, as detailed in Section 2.4.1.7.

5.2.3.2 Ratio of Averages Edge Detection

In [109], the shortcomings of the LoG method for edge detection in speckled imagery are noted. Specifically, while the LoG is capable of discriminating image edges, a large amount of spurious edges are generated by the speckle pattern, having no correlation with image structure. To counter this effect, a ratio of averages edge detector is developed. Estimates of the horizontal and vertical components of edge magnitude are computed as:

$$\begin{aligned} H(x,y) &= \max\left(\frac{R(x,y)}{L(x,y)}, \quad \frac{L(x,y)}{R(x,y)}\right) \\ V(x,y) &= \max\left(\frac{U(x,y)}{D(x,y)}, \quad \frac{D(x,y)}{U(x,y)}\right) \end{aligned} \tag{5.11}$$

The quantities $R(x,y)$,$L(x,y)$,$U(x,y)$ and $D(x,y)$ are the average values of image intensity in neighbourhoods to the right, left, above and below each pixel (x,y). These local averages can be computed by convolving the image with averaging kernels of similar aspects to those in Fig. 5.9. Different sizes of neighbourhood N can be considered for the calculation of these properties, and the example kernels shown here are drawn from a 7 × 7 mask. The estimates of edge magnitude components $H(x,y)$ and $V(x,y)$ are independent of image intensity. The final RoA measure is found as the magnitude of these quantities:

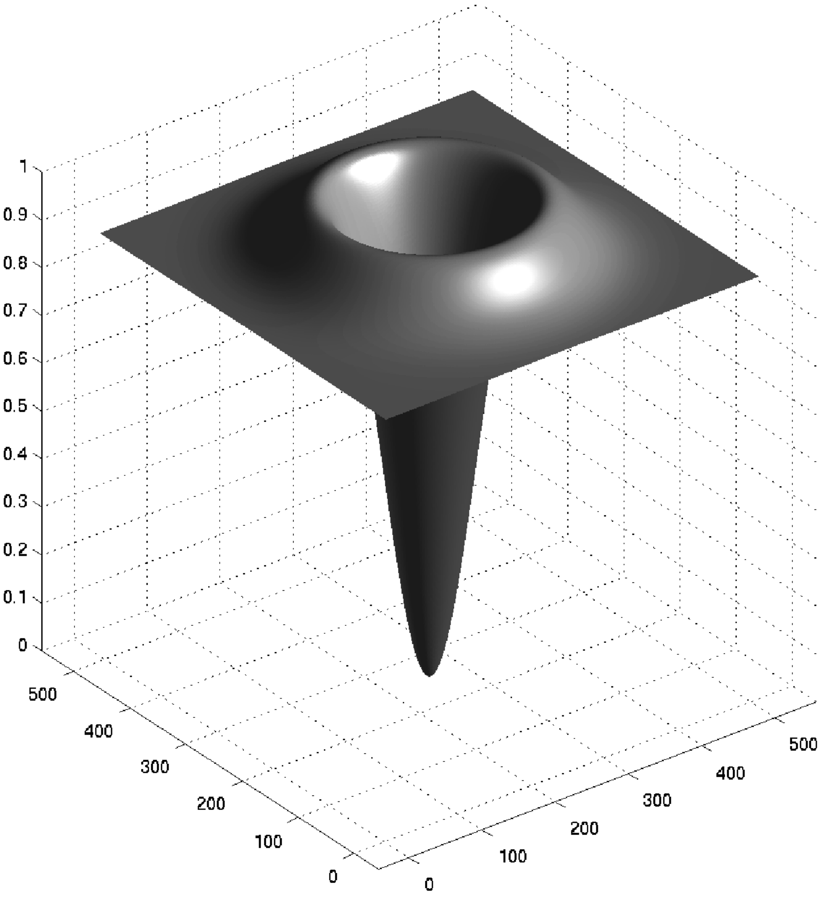

Figure 5.10: Example of a Laplacian of Gaussian filter kernel.

$$RoA(x,y) = \sqrt{H^2(x,y) + V^2(x,y)} \tag{5.12}$$

The RoA measure is thresholded to produce a binary edge map. RoA edges can be wide in extent and so poorly localised, especially for larger values of N. Thus the RoA step can be said to determine the approximate locations of image edges.

By combining the RoA step with the more localised edges found from the zero crossings of an LoG filter, a more precise edge estimate is found. The RoA step can be seen as suppressing the spurious edges produced by the LoG filter, by constraining these edges to regions containing image edges. Fig. 5.10 displays an example of a 2D Laplacian of Gaussian filter kernel, known for its appearance as the ‘Mexican hat’ function.

5.2.3.3 Phase Congruency Edge Detection

A wide range of image features including step edges and lines, are found to correspond with locations of maximal phase agreement of local Fourier components [148, 149]. The PC method of Kovesi [136] uses a log Gabor filter bank with multiple scales to isolate local frequency information. Areas where the phase is in agreement can then be found by calculation of a phase deviation function.

Kovesi develops a phase congruency measure which is shown to be highly localised. Central to this method is the following function of local phase deviation:

$$\zeta(x,y) = \sum_s \{F_{s,o}\bar{\phi}_{F,o} + H_{s,o}\bar{\phi}_{H,o} - |F_{s,o}\bar{\phi}_{H,o} - H_{s,o}\bar{\phi}_{F,o}|\} \tag{5.13}$$

where s, o indicate the Gabor scale orientation, respectively. Here the quantities ($\bar{\phi}_{F,o}$ and $\bar{\phi}_{H,o}$) are the components of the mean weighted phase vector, given by:

$$(\bar{\phi}_{F,o}, \bar{\phi}_{H,o}) = \frac{(\sum_s F_{s,o}, \quad \sum_s H_{s,o})}{\sqrt{(\sum_s F_{s,o})^2 + (\sum_s H_{s,o})^2}} \tag{5.14}$$

The real and imaginary (even and odd symmetric) components of the Gabor subbands are denoted by $F_{s,o}(x,y) = \Re(G_{s,o}\{I(x,y)\})$ and $H_{s,o}(x,y) = \Im(G_{s,o}\{I(x,y)\})$, where $I(x,y)$ is the input image. A weighting function is defined to suppress output when the local frequency spread is low. A measure of frequency spread is defined for each orientation as:

$$S_o(x,y) = \frac{1}{N_{scale}} \left(\frac{A_o(x,y)}{A_{o,[max]}(x,y) + \varepsilon} \right) \tag{5.15}$$

where $A_{s,o}(x,y) = |G_{s,o}\{I(x,y)\}|$, $A_o(x,y) = \sum_s A_{s,o}(x,y)$, and $A_{o,[max]}(x,y) = max_s(A_{s,o}(x,y))$. The weighting function is then found by applying a sigmoid to $S_o(x,y)$ as:

$$W_o(x,y) = \frac{1}{1 + \exp\{g\,(c - S_o(x,y))\}} \tag{5.16}$$

Where g and c are gain and cut-off parameters. Finally, the phase congruency is calculated from the above quantities as:

$$PC(x,y) = \sum_o \frac{W_o(x,y)\lfloor \zeta(x,y) - T_o \rfloor}{A_o(x,y) + \varepsilon} \tag{5.17}$$

Where the operation $\lfloor \cdot \rfloor$ represents zeroing of negative values. Compensation for image noise is performed in [136] by removing small values by subtraction of T_o above.

5.2.3.4 Feature Asymmetry Edge Detection

The FA method of [138] is another phase based method employing log Gabor filtering to separate local phase components. The Gabor filter bank used here has a single scale, as the authors state that phase based approaches do not need fine scale information for feature localization. Two-dimensional feature asymmetry is defined as the sum of the asymmetry over the orientations:

$$FA(i,j) = \sum_{o=1}^{N_{ori}} FA_o = \sum_{o=1}^{N_{ori}} \frac{\lfloor D_o(i,j) - T_o \rfloor}{A_o(i,j) + \varepsilon} \tag{5.18}$$

Where D_o is the difference of the magnitude of the odd and even symmetric components in orientation o: $D_o(i,j) = |H_o(i,j)| - |F_o(i,j)|$. As in the description of the PC method above, the real and imaginary (even and odd symmetric) components of the Gabor coefficients are denoted by $F_o(i,j)$ and $H_o(i,j)$, respectively, while $A_o(i,j)$ represents their magnitude. T_o is a noise threshold found as a fraction of the deviation of D_o: $T_o = k \cdot std\{D_o(i,j)\}$ [150]. Parameter k determines the level of detail rejected as noise, ε is a small constant to avoid zero division, and $\lfloor \cdot \rfloor$ represents zeroing of negative values. FA edge detection has been incorporated into echocardiographic enhancement schemes [150], and has been extended to use temporal information for video sequences [137, 138]. Six orientations are used here, as in [138].

5.3 Removal of Speckle Noise from Gabor Coefficients

In the development of the edge detector presented in this chapter, the commonly used multiplicative noise model of [76] for a speckle corrupted image is used. The equally common approach of discarding the effect of additive noise is taken here [2, 60, 75]. A logarithmic operation is applied to transform the multiplicative noise to an additive model. A log Gabor filter bank, as described in Section 5.2.1 above, is then applied. These steps are represented as:

$$\begin{aligned} y(i,j) =& s(i,j)n(i,j) \\ y^l(i,j) =& s^l(i,j) + n^l(i,j) \\ Y^l(i,j) =& S^l(i,j) + N^l(i,j) \end{aligned} \tag{5.19}$$

where y is the observed speckled image, s is the noise free signal, n is the speckle noise, and (i,j) are the spatial indices. The log transform is represented as $y^l = log(y)$, and the application of the log Gabor filter bank is denoted by $Y = G_{s,o}\{y\}, \forall s, o$. The following sections describe the steps involved in the estimation of the speckle free Gabor coefficients S^l from the observed speckled coefficients Y^l above. Firstly, a statistical model is proposed for (5.19) in Section 5.3.1, using a stable process to describe S^l, and a Gaussian process for N^l. Next, a method for simultaneously estimating the parameters of this model is then presented, in Section 5.3.2. Using this model and parameter estimation method, the effect of speckle noise is removed from the Gabor coefficients using MAP estimation as described in Section 5.3.3

5.3.1 Modelling Gabor Coefficients

5.3.1.1 SαS Modelling of Signal Gabor Coefficients

The speckle free Gabor coefficients (S^l in (5.19) above) are modelled here as SαS processes. This is motivated by observation of the distributions of log Gabor coefficients of speckle free ultrasound images. Two types of speckle free images are considered; maximally speckle free MW simulated images (generated using the technique described in Section 4.2), and also clinical images which have their speckle content removed by speckle filtering. Fig. 5.11 presents two example ultrasound images, one of which is a maximally speckle free simulated MW image (Fig. 5.11(a)), and one of which is a clinical image treated with the

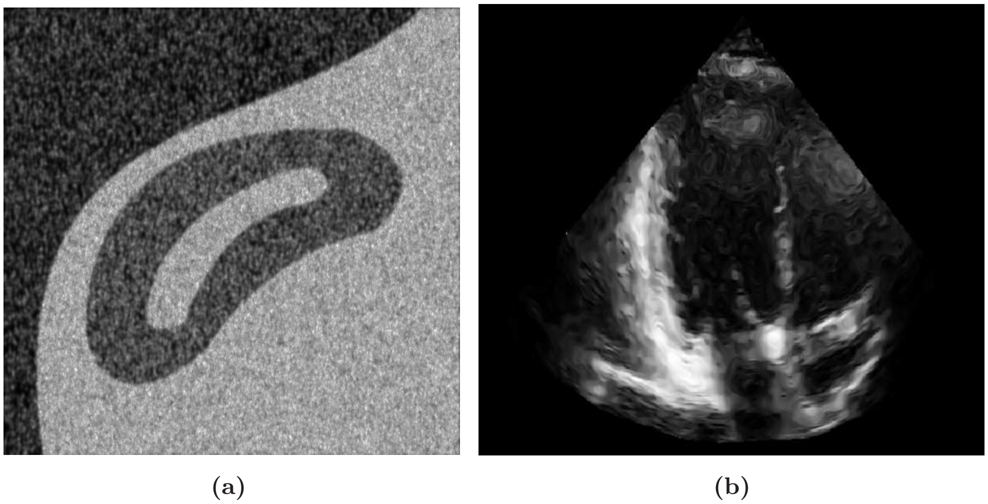
(a) (b)

Figure 5.11: Speckle free images to demonstrate SαS fitting. (a) simulated MW image, (b) clinical image treated with the OSRAD filter.

OSRAD filter of Section 2.4.1.6 to suppress its speckle content (Fig. 5.11(b)). Histograms of the real components Gabor coefficients of each of these images are presented in Figs. 5.12-5.13. These distributions of Gabor coefficients of both types of speckle free image are noted to exhibit heavy-tailed behaviour and pronounced peaks to varying degrees. The imaginary components of these Gabor coefficients exhibit also these characteristics, and can be found in Figs. A.1-A.2. A notable observation is that the histograms of the Gabor coefficients of the clinical image (Fig. 5.13) have more pronounced peaks at the origin than their simulated counterpart (Fig. 5.12). This is due to the differing intensity profiles of the respective images, i.e. in Fig. 5.11 it can be seen that the clinical image has more dark regions. These empirical densities resemble symmetric stable distributions in both cases, and a further investigation of the suitability of this model is undertaken.

Previous work which uses stable models for speckle free multiscale coefficients includes that of Achim *et al.* In [151, 152], these authors use SαS processes to model the wavelet coefficients of log transformed SAR imagery, after speckle has been suppressed with the gamma MAP filter of [153]. The same authors use SαS distributions to accurately model wavelet coefficients of log transformed ultrasound data in [75, 154], after speckle has been suppressed using the homomorphic Weiner filter of [76]. In [155–157], bivariate isotropic stable distributions are employed for image processing tasks such as denoising and image

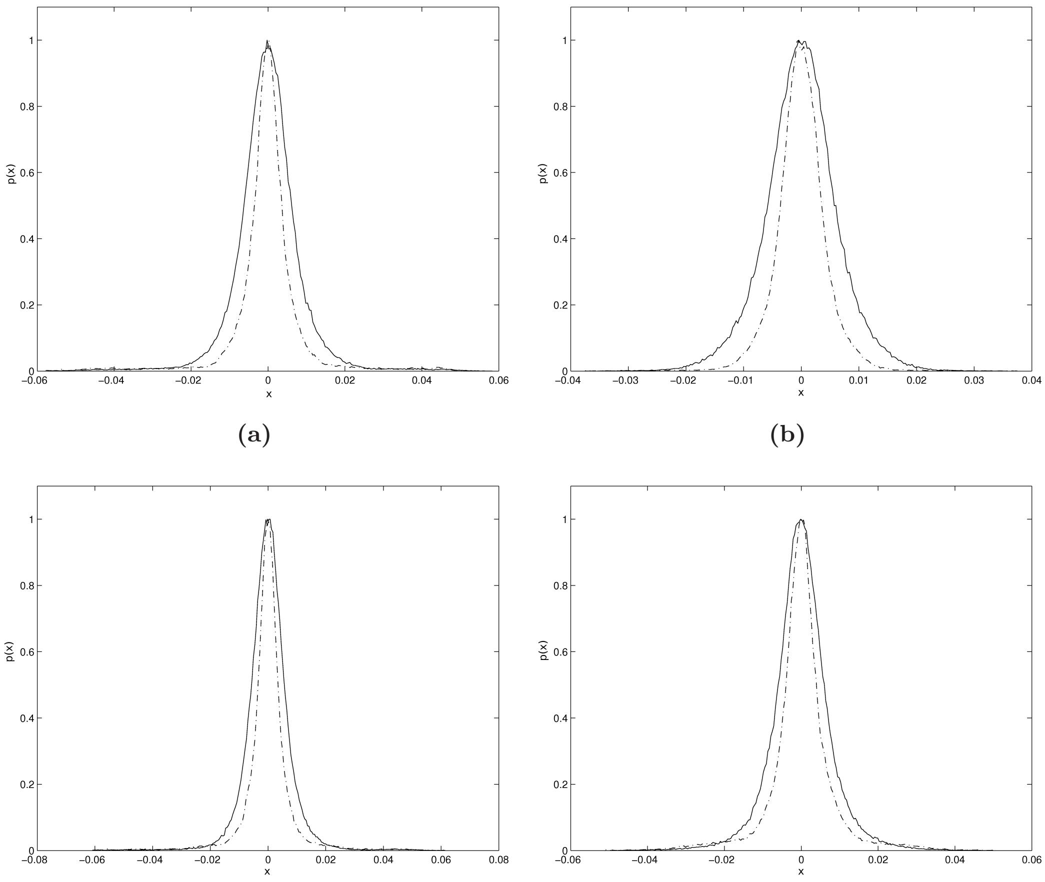

(a) **(b)**

(c) **(d)**

Figure 5.12: Histograms of real-valued Gabor coefficients of Fig. 5.11(a), plotted per orientation. Solid lines are coefficients from scale 1, dash dot lines from scale 2. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

fusion.

To verify the suitability of a SαS model for speckle free Gabor coefficients, the approach of [75, 152] is followed and expanded upon. The approach presented here differs from prior work in a number of ways:

- Multiscale decomposition is accomplished with orientation specific log Gabor filtering, rather than Daubechies Symmlet 8 wavelet used in [152], or the Daubechies Symmlet 4 wavelet used in [75].

- The use of maximally speckle free simulated MW images allows verification of this

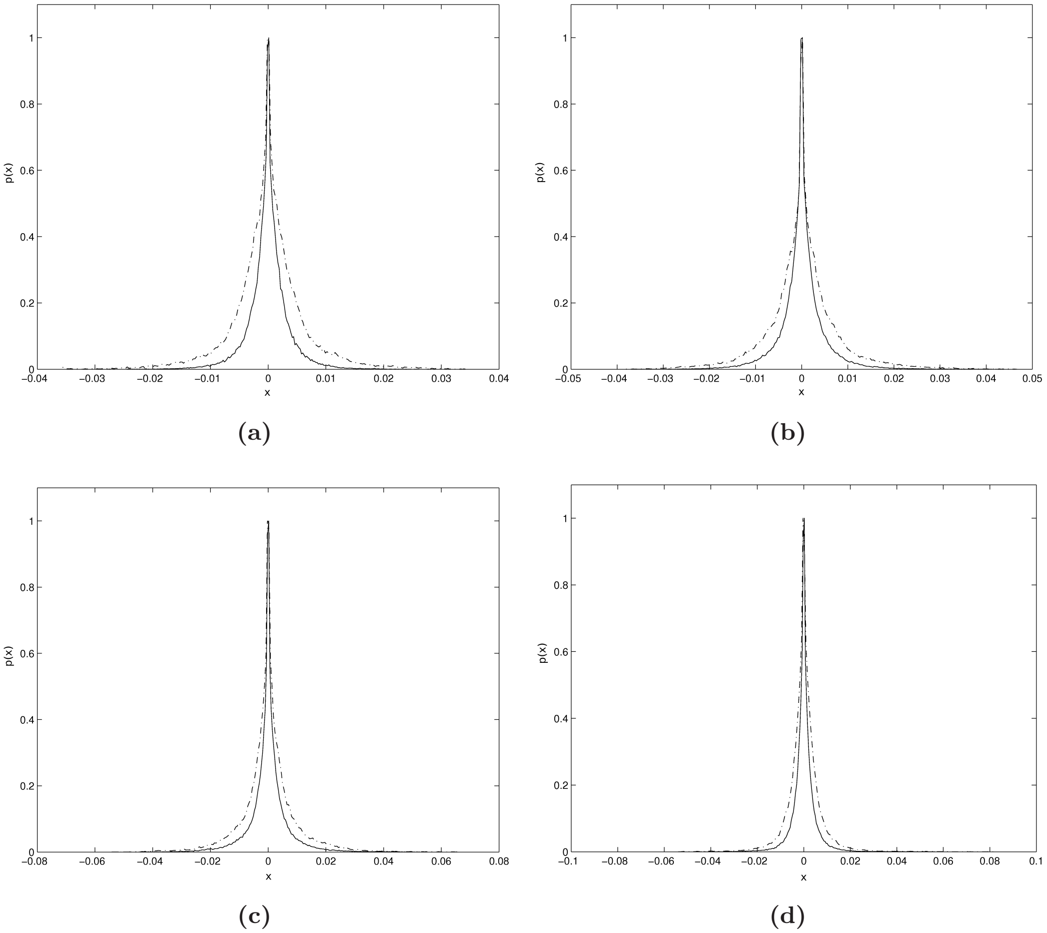

(a) **(b)**

(c) **(d)**

Figure 5.13: Histograms of real-valued Gabor coefficients of Fig. 5.11(b), plotted per orientation. Solid lines are coefficients from scale 1, dash dot lines from scale 2. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

model on statistically speckle free data. This is in comparison to the practice of estimating speckle free data from real world images, using a speckle reduction technique such as gamma MAP filtering of SAR imagery [152], or a homomorphic Weiner filter applied to clinical ultrasound [75]. Approximating speckle free data in this way has proven sufficient in these works however, and is employed here in addition to simulated MW images to demonstrate model suitability in clinical as well as simulated images.

The suitability of the SαS model is demonstrated here using the images of Fig. 5.11. Firstly, maximum likelihood (ML) estimates are found for the parameters of the stable distributions best describing the real and imaginary log Gabor coefficients. These are found using the method of [158], for which the author has made an executable program available online[1]. These parameter estimates, with their corresponding 95% confidence intervals, are presented in Tables 5.1 (containing α estimates) and Tables 5.2 (presenting γ estimates). The confidence intervals for both parameters are narrow relative to the estimate itself in all cases. This is evidence that the stable distributions which best fit the data do so in an accurate manner. The confidence intervals are observed to be very similar for both the simulated MW and speckle reduced clinical images. In addition, the intervals do not vary significantly with either decomposition scale or orientation. Estimates for the other parameters of stable distributions (β, δ) are also found. In order for the best fitting stable distribution to be symmetric (and therefore SαS), both (β, δ) must be zero. In all cases, the absolute value of the estimates of both β and δ are less than the corresponding 95% confidence interval. This provides strong evidence that the value of these parameters was very close to zero, and so that the stable processes best fitting the data can be reasonably described as SαS.

In addition to calculating and observing the confidence intervals of ML parameter estimates, the suitability of the SαS model for speckle free Gabor coefficients is analysed by the use of probability plots. These compare the observed distributions with those expected from the ML parameter estimates, examples of which are shown in Figs. 5.14-5.15. These images contain the probability plots for the first scale of the real Gabor coefficients, over each orientation, for the images of Fig. 5.11. These are observed to be representative of the full set of probability plots, i.e. both real and imaginary components, over both the scales considered. The full set of probability plots can be seen in Figs. A.3-A.6.

These plots can be described as stabilised P-P (probability-probability) [159], and as such they display the cumulative distribution functions (CDF) of the observed and expected data against each other. A straight diagonal reference line is included in each plot, which represents the expected CDF from the ML estimated parameters. The + marks are points from the observed CDF. Close proximity of these point to the diagonal reference line indicates similarity between the expected and observed distributions.

Deviations of the observed CDF line from the expected diagonal represent differences in the Both sets of data are subject to a variance stabilising transform, so deviations from the diagonal should be the same at each location.

---

[1] `http://academic2.american.edu/~jpnolan/stable/stable.html`

Table 5.1: Maximum likelihood estimates for SαS parameter α, for the speckle free images of Fig. 5.11.

$\hat{\alpha}_{ML}$		Real		Imaginary	
		Scale 1	Scale 2	Scale 1	Scale 2
Image 1 Fig. 5.11(a)	$\theta = 0$ rad	$1.62 \pm 3.13\times10^{-2}$	$1.31 \pm 3.04\times10^{-2}$	$1.61 \pm 3.14\times10^{-2}$	$1.27 \pm 3.00\times10^{-2}$
	$\theta = \frac{\pi}{4}$ rad	$1.86 \pm 2.47\times10^{-2}$	$1.71 \pm 3.02\times10^{-2}$	$1.88 \pm 2.38\times10^{-2}$	$1.72 \pm 2.99\times10^{-2}$
	$\theta = \frac{\pi}{2}$ rad	$1.66 \pm 3.08\times10^{-2}$	$1.31 \pm 3.05\times10^{-2}$	$1.66 \pm 3.08\times10^{-2}$	$1.27 \pm 3.00\times10^{-2}$
	$\theta = \frac{3\pi}{4}$ rad	$1.68 \pm 3.06\times10^{-2}$	$1.33 \pm 3.06\times10^{-2}$	$1.69 \pm 3.05\times10^{-2}$	$1.32 \pm 3.04\times10^{-2}$
Image 2 5.11(b)	$\theta = 0$ rad	$1.38 \pm 3.24\times10^{-2}$	$1.36 \pm 3.22\times10^{-2}$	$1.36 \pm 3.23\times10^{-2}$	$1.40 \pm 3.26\times10^{-2}$
	$\theta = \frac{\pi}{4}$ rad	$1.25 \pm 3.11\times10^{-2}$	$1.35 \pm 3.21\times10^{-2}$	$1.23 \pm 3.07\times10^{-2}$	$1.31 \pm 3.17\times10^{-2}$
	$\theta = \frac{\pi}{2}$ rad	$1.12 \pm 2.92\times10^{-2}$	$1.22 \pm 3.06\times10^{-2}$	$1.12 \pm 2.92\times10^{-2}$	$1.21 \pm 3.05\times10^{-2}$
	$\theta = \frac{3\pi}{4}$ rad	$1.27 \pm 3.13\times10^{-2}$	$1.30 \pm 3.17\times10^{-2}$	$1.22 \pm 3.06\times10^{-2}$	$1.32 \pm 3.19\times10^{-2}$

From these probability plots, it can be seen that the observed data deviates very little from that expected from the ML parameter estimates. For the probability plots of the speckle free simulated MW image (Figs. 5.14), the deviation (and therefore the difference between the observed and expected distributions) is particularly small. A greater degree of deviation is observed for the clinical image (Figs. 5.15), however this is still quite small and the expected and observed lines are in general agreement. It is noted that the deviations here occur around the (0.5,0.5) point, indicating differences on each side of the median of the data sets. These deviations from the expected values are explained by the nature of the clinical image considered; specifically that it is only approximately speckle free. As well as presenting the probability plots graphically, the correlation between the observed and expected P-P lines is quantified in each case. These are each calculated with 512

Table 5.2: Maximum likelihood estimates for SαS parameter γ, for the speckle free images of Fig. 5.11.

$\hat{\gamma}_{ML}$		Real		Imaginary	
		Scale 1	Scale 2	Scale 1	Scale 2
Image 1 Fig. 5.11(a)	$\theta = 0$ rad	4.30×10^{-3} $\pm8.66\times10^{-5}$	2.84×10^{-3} $\pm6.77\times10^{-5}$	4.42×10^{-3} $\pm8.96\times10^{-5}$	2.79×10^{-3} $\pm6.79\times10^{-5}$
	$\theta = \frac{\pi}{4}$ rad	4.12×10^{-3} $\pm7.08\times10^{-5}$	2.58×10^{-3} $\pm4.94\times10^{-5}$	4.20×10^{-3} $\pm7.10\times10^{-5}$	2.56×10^{-3} $\pm4.86\times10^{-5}$
	$\theta = \frac{\pi}{2}$ rad	4.00×10^{-3} $\pm7.87\times10^{-5}$	2.81×10^{-3} $\pm6.71\times10^{-5}$	4.04×10^{-3} $\pm7.95\times10^{-5}$	2.70×10^{-3} $\pm6.58\times10^{-5}$
	$\theta = \frac{3\pi}{4}$ rad	4.44×10^{-3} $\pm8.62\times10^{-5}$	3.29×10^{-3} $\pm7.77\times10^{-5}$	4.36×10^{-3} $\pm8.42\times10^{-5}$	3.33×10^{-3} $\pm7.87\times10^{-5}$
Image 2 5.11(b)	$\theta = 0$ rad	1.55×10^{-3} $\pm3.71\times10^{-5}$	2.70×10^{-3} $\pm6.54\times10^{-5}$	1.52×10^{-3} $\pm3.67\times10^{-5}$	2.80×10^{-3} $\pm6.62\times10^{-5}$
	$\theta = \frac{\pi}{4}$ rad	1.86×10^{-3} $\pm4.77\times10^{-5}$	3.30×10^{-3} $\pm8.04\times10^{-5}$	1.81×10^{-3} $\pm4.72\times10^{-5}$	3.25×10^{-3} $\pm8.11\times10^{-5}$
	$\theta = \frac{\pi}{2}$ rad	2.17×10^{-3} $\pm6.03\times10^{-5}$	3.42×10^{-3} $\pm8.97\times10^{-5}$	2.10×10^{-3} $\pm5.84\times10^{-5}$	3.41×10^{-3} $\pm8.97\times10^{-5}$
	$\theta = \frac{3\pi}{4}$ rad	1.76×10^{-3} $\pm4.49\times10^{-5}$	2.78×10^{-3} $\pm6.95\times10^{-5}$	1.71×10^{-3} $\pm4.49\times10^{-5}$	2.89×10^{-3} $\pm7.15\times10^{-5}$

points, for the real and imaginary Gabor components of the scales and orientations of each image, and Pearson's correlation coefficient is calculated. These are displayed in Tables 5.3. It can be seen that all probability plots have correlation measures very close to unity, indicating a very strong correlation between the observed distributions and those from the ML parameter estimates. This is the case even for the probability plots of the clinical image, in which some deviations from the expected values are observed. It is therefore determined that these deviations are not large relative to the overall distributions. Both the narrow confidence intervals and the close agreement observed in the probability plots indicate that a SαS model for the speckle free Gabor coefficients is a suitable and accurate one.

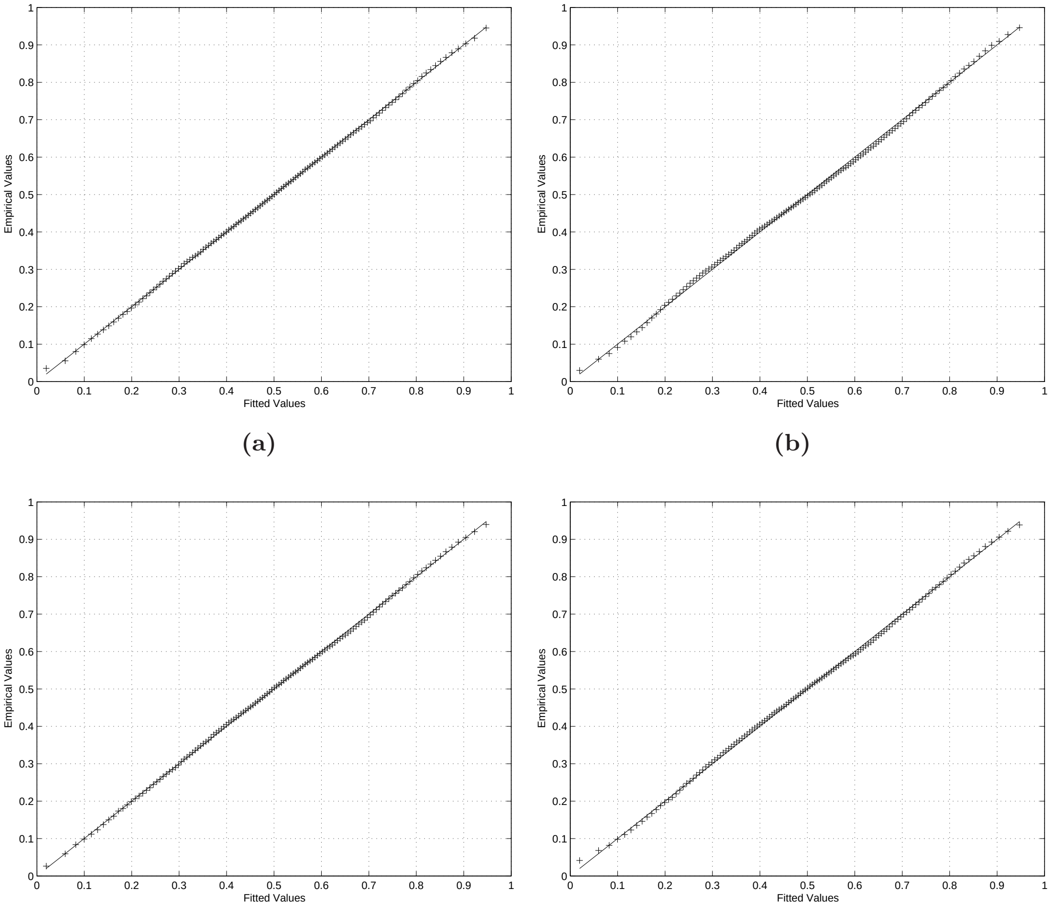

(a) **(b)**

(c) **(d)**

Figure 5.14: Stabilised probability plots of the real-valued Gabor coefficients from the first scale of image 5.11(a), plotted per orientation. + marks are points from the observed distribution, and the thin diagonal reference lines are those expected from the ML parameter estimates. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

5.3.1.2 Combined SαS and Gaussian Model for Speckled Gabor Coefficients

Having established the accuracy of the SαS model for speckle free Gabor coefficients, the inclusion of the speckle noise component is now addressed. As in [152], it is expected that the Gabor coefficients of the log transformed speckle noise (N^l in (5.19)) will approach a Gaussian distribution through the central limit theorem. The characteristic function of a Gaussian centred at the origin is:

$$\phi_n(t) = \exp\{-\frac{\sigma^2}{2}|t|^2\} \tag{5.20}$$

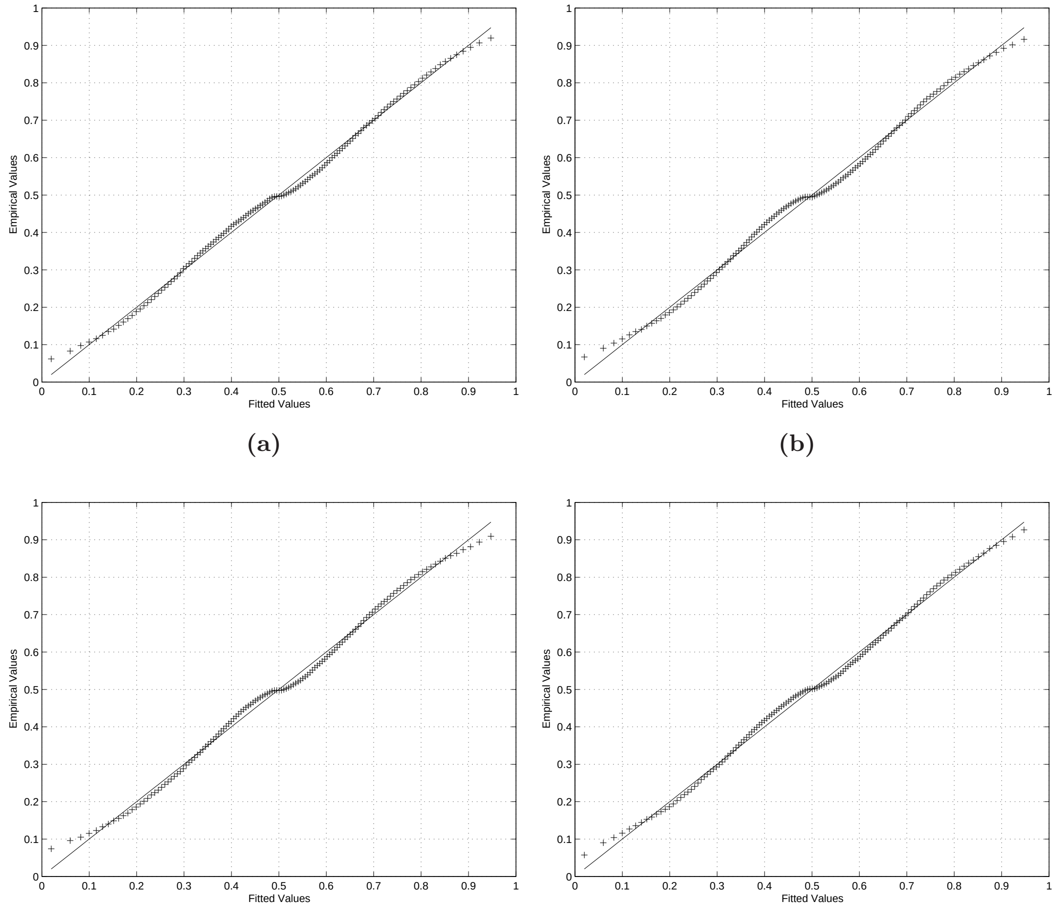

(a) **(b)**

(c) **(d)**

Figure 5.15: Stabilised probability plots of the real-valued Gabor coefficients from the first scale of image 5.11(b), plotted per orientation. + marks are points from the observed distribution, and the thin diagonal reference lines are those expected from the ML parameter estimates. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

Using the SαS model for speckle free Gabor coefficients S^l and the above zero mean Gaussian for noise coefficients N^l, the complete model for the observed Gabor coefficients Y^l can be expressed by its characteristic function, $\phi_y(t)$. This is given by the product of the characteristic functions of its components:

$$\phi_y(t) = \phi_s(t)\phi_n(t) = \exp\{-\gamma|t|^\alpha - \frac{\sigma^2}{2}|t|^2\} \tag{5.21}$$

To demonstrate the suitability of the model (5.21), it is fitted to the Gabor coeffi-

Table 5.3: Correlation values for the probability plots of Gabor coefficients of Fig. 5.11, quantifying similarity between observed and expected densities. Probability plots can be found in Section A.1.

$\hat{\alpha}_{ML}$		Real		Imaginary	
		Scale 1	Scale 2	Scale 1	Scale 2
Gabor coefficients of Fig. 5.11(a). Probabilty plots in Fig. 5.14 and Figs. A.3-A.4	$\theta = 0$ rad	0.9999	0.9999	0.9999	0.9998
	$\theta = \frac{\pi}{4}$ rad	0.9996	0.9997	0.9997	0.9998
	$\theta = \frac{\pi}{2}$ rad	0.9999	0.9999	0.9999	0.9998
	$\theta = \frac{3\pi}{4}$ rad	0.9997	0.9997	0.9995	0.9995
Gabor coefficients of Fig. 5.11(b). Probabilty plots in Fig. 5.15 and Figs. A.5-A.6	$\theta = 0$ rad	0.9985	0.9990	0.9984	0.9990
	$\theta = \frac{\pi}{4}$ rad	0.9976	0.9982	0.9978	0.9979
	$\theta = \frac{\pi}{2}$ rad	0.9977	0.9979	0.9979	0.9981
	$\theta = \frac{3\pi}{4}$ rad	0.9984	0.9991	0.9985	0.9989

cients of the speckled images of Fig. 5.16. The parameters of this model (α, γ, σ) are found using the estimation method presented in Section 5.3 below is employed. Figs. 5.17-5.18 compare the observed and fitted densities of the real-valued Gabor coefficients of the first scale of these images, for each orientation. The fitted densities are seen to be in close agreement with the observed data. Accurate fitting is also observed for the imaginary Gabor coefficients, and for the second scale of decomposition used in this section. The full set of fitted densities are presented in Figs. A.7-A.10.

To quantify the accuracy of fitting, the widely used r^2 goodness of fit measure is calculated. These are displayed in Tables 5.4. An r^2 value of unity represents a perfect fit, and it is observed here that the goodness of fit is in all cases very close to this ideal value. Given the high accuracy of the distribution fitting shown in this section, it is concluded that the model of (5.21) is both valid and accurate for simulated and clinical speckled images.

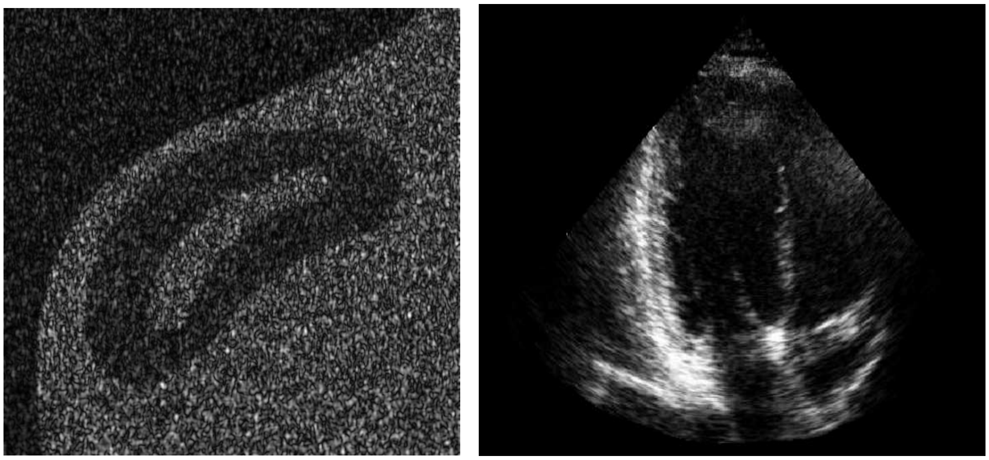
(a) (b)

Figure 5.16: Speckle corrupted versions of the images of Fig. 5.11, used to demonstrate the accuracy of the model of (5.21) for Gabor coefficients. (a) simulated speckled image, (d) speckled clinical images.

Table 5.4: r^2 goodness of fit measures for the densities fitted to the Gabor coefficients of the images of Fig. 5.16.

$\hat{\alpha}_{ML}$		Real		Imaginary	
		Scale 1	Scale 2	Scale 1	Scale 2
Image 1: Fig. 5.16(a). Densities in Figs. A.7-A.8	$\theta = 0$ rad	0.9966	0.9981	0.9977	0.9949
	$\theta = \frac{\pi}{4}$ rad	0.9957	0.9964	0.9979	0.9979
	$\theta = \frac{\pi}{2}$ rad	0.9959	0.9984	0.9973	0.9972
	$\theta = \frac{3\pi}{4}$ rad	0.9946	0.9974	0.9967	0.9970
Image 2: Fig. 5.16(b). Densities in Figs. A.9-A.10	$\theta = 0$ rad	0.9783	0.9885	0.9801	0.9891
	$\theta = \frac{\pi}{4}$ rad	0.9779	0.9811	0.9771	0.9822
	$\theta = \frac{\pi}{2}$ rad	0.9776	0.9772	0.9827	0.9865
	$\theta = \frac{3\pi}{4}$ rad	0.9762	0.9815	0.9826	0.9818

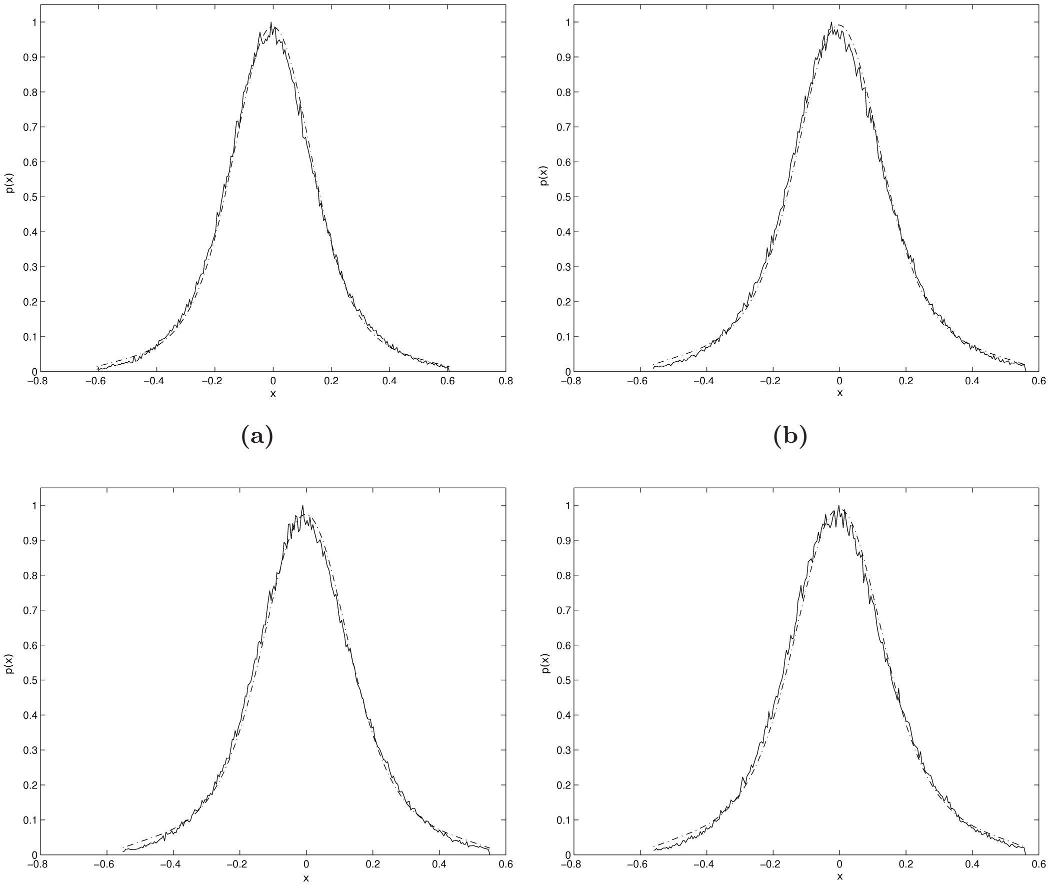

(a) **(b)**

(c) **(d)**

Figure 5.17: Observed densities of the real-valued Gabor coefficients from the first scale of the speckled image of Fig. 5.16(a), and those fitted using the model of (5.21), for each orientation. Solid lines are the observed densities, dash-dot lines are the fitted densities. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

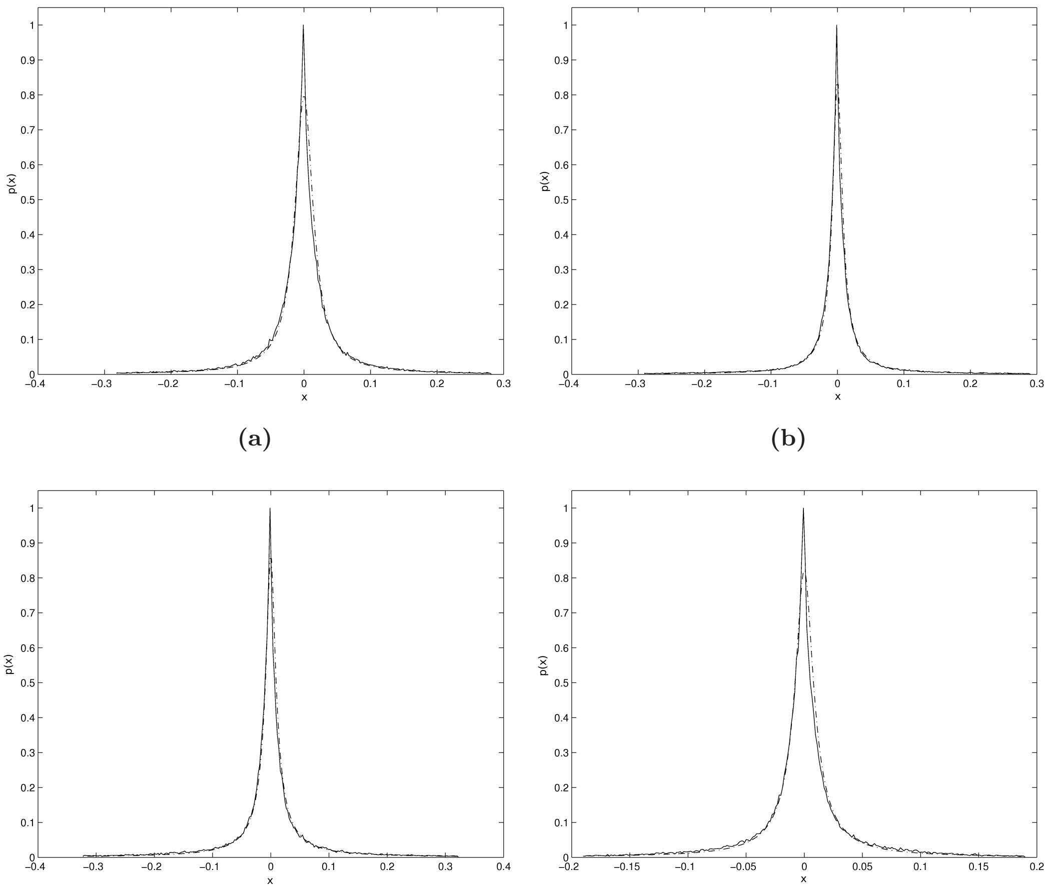

(a) **(b)**

(c) **(d)**

Figure 5.18: Observed densities of the real-valued Gabor coefficients from the first scale of the speckled image of Fig. 5.16(b), and those fitted using the model of (5.21), for each orientation. Solid lines are the observed densities, dash-dot lines are the fitted densities. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

5.3.2 Parameter Estimation Method

This section presents the development of a method of estimating the parameters of the model (5.21) for the Gabor coefficients of speckled images. This provides simultaneous estimation of the parameters of the SαS signal component (α, γ) and also the parameter (σ) of the Gaussian speckle noise component.

Many methods of parameter estimation for stable processes exist, including approaches based on quantiles [160], and those which approximate the Maximum Likelihood (ML) technique [158, 161, 162]. A large number of parameter estimation methods make use of the empirical characteristic function [163–168]:

$$\hat{\phi}_s(t_k) = \frac{1}{N}\sum_{j=1}^{N} \exp\{it_k x_j\}, \quad k = 1, 2, ..K \tag{5.22}$$

where $i = \sqrt{-1}$ and t_k is a set of K discrete points of the characteristic function argument, t. The parameter estimation method proposed here is based on the approach of Koutrouvelis [164, 165]. This has compared favourably with other methods using the empirical characteristic function [167, 169]. For SαS distributions centred at the origin, Koutrouvelis proceeds by deriving the following from (5.9):

$$\ln[-\ln|\phi_s(t)|^2] = \ln 2\gamma + \alpha \ln|t| \tag{5.23}$$

Then estimates for α and γ are found by performing linear regression on:

$$y_k = \mu + \alpha\omega_k + \epsilon_k, \quad k = 1, 2, ..K \tag{5.24}$$

where $y_k = \ln[-\ln|\hat{\phi}_s(t_k)|^2]$, $\omega_k = \ln|t_k|$, $\mu = \ln 2\gamma$ and ϵ_k is an error term. This method is applied iteratively. For further details on practical issues, such as input standardization and a table of suitable t_k values, the reader is referred to [164, 165].

The model of (5.21) requires estimation of three parameters, (α, γ, σ). Equation (5.21) is also used in [152] and [75]. In these works, σ is estimated first using the method of Donoho [63]. Parameters (α, γ) are then estimated by using least-squares fitting to the Fourier transform of the empirical PDF in [75]. In [152], the regression method of (5.24) is used, with the exception that $y = \ln[-\ln|\hat{\phi}_y(t)|^2 - \sigma^2 t^2]$. The term $-\ln|\hat{\phi}_y(t)|^2$ is positive under normal operation, and issues arise when $\sigma^2 t^2 > -\ln|\hat{\phi}_y(t)|^2$, leading to complex y.

This can be shown to occur for $\sigma > \sqrt{-\ln|\hat{\phi}_y(t)|^2/t^2}$, placing an upper limit on the range of σ that can be used in this method.

The method presented here estimates the three parameters (α, γ, σ) simultaneously, and also removes the above limit on the practical range of σ. The proposed method is generated by manipulating (5.21) to give:

$$-\ln|\hat{\phi}_y(t)|^2 = 2\gamma|t|^\alpha + \sigma^2 t^2 \tag{5.25}$$

which can be expressed as the estimation function:

$$y_k = 2b|t_k|^a + c^2 t_k^2, \quad k = 1, 2, ..K \tag{5.26}$$

where $y_k = -\ln|\hat{\phi}_y(t_k)|^2, a = \alpha, b = \gamma, c = \sigma$. Unlike (5.24), this is no longer a linear function, so linear regression cannot be employed. Instead, unconstrained nonlinear minimization is required, using the Nelder-Mead method [170, 171]. As with the method of [164,165], the estimation of parameters is performed multiple times in an iterative fashion, with the estimates of each iteration used as the initial guess for the next iteration. In testing, five iterations is found to be sufficient for parameter convergence.

Similar to [164] and [172], the selection of t_k is explored experimentally. It has been noted that the empirical characteristic function provides the most information close to the origin; a typical example is displayed in Fig. 5.19. It can be seen that the empirical and expected characteristic functions are in agreement only for low values of t_k, and this similarity deteriorates for larger t_k. Methods for automatically detecting this useful region have been proposed [173, 174]. However, in [172] it is stated that the method of [173] for selecting this region is not reliable for $\alpha < 1$. It is noted experimentally here that the suitable range of t_k varies with all three distribution parameters. The optimal values for the upper limit of t_k, denoted t_{max}, are determined by Monte Carlo simulation for a range of parameter values. A two stage process is employed: firstly, a candidate range of t_k over which the empirical characteristic function agrees with that expected from the known parameters is determined. Secondly, the value of t_k within this range which gives the lowest combined parameter error is selected as t_{max}.

The distribution parameters are varied in the ranges $(.5 \leq \alpha \leq 2; .5 \leq \gamma \leq 2; .1 \leq \sigma \leq 2)$. Random vectors of 10^4 points are generated for each parameter combination as the sum of appropriate SαS and Gaussian variables, as in (5.19). Each parameter combination

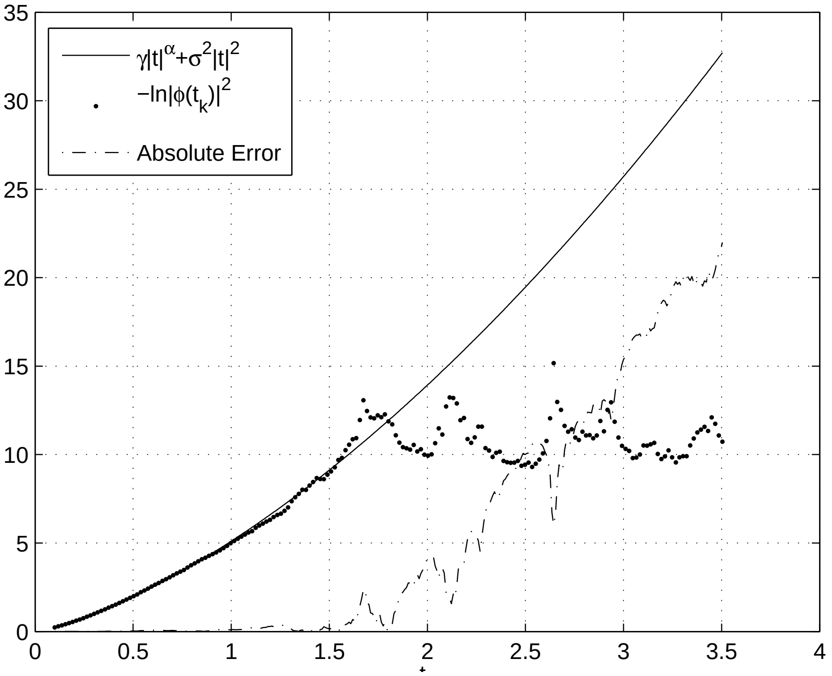

Figure 5.19: Comparison of the empirical and expected characteristic functions, by (5.25). A good match is observed for $t_k \lesssim 1.5$. The values of (α,γ,σ) are $(1.25, 2.00, 1.05)$ here.

is repeated 50 times, with the average result taken. The optimal t_{max} calculation procedure can be summarized as:

1. The candidate range of t_k is found for each parameter combination:
    (a) A t_k vector is defined as 256 linearly spaced points between .1 and the upper limits recommended in [164] for the current value of α.
    (b) The empirical and expected estimation functions are calculated from (5.25). Here y_k is found from the sample random vector, and $-\ln|\hat{\phi}_y(t)|^2$ is found using the known values of (α, γ, σ).
    (c) The error between the empirical and expected functions is calculated, and the candidate range is selected as the t_k with reasonable agreement (RMSE not exceeding 0.5).
2. t_{max} is found as the t_k which gives the optimal parameter estimate within the candidate range found above:
    (a) For each t_k, parameter estimation is performed by (5.26).

(b) t_{max} is found as the value which minimizes and the sample RMSE of the parameter estimates, similar to the method of [164].

A sample of the calculated t_{max} values are displayed in Table A.1.These values are in general lower than those proposed by [164] for a given α value. Higher levels of speckle noise, corresponding to larger σ values, produce lower optimal t_{max} values. The optimal t_{max} values are seen to increase for stable distributions with a large central peak (low α), and to decrease for stable distributions which are more dispersed (larger values of γ). An extensive test of the accuracy of parameter estimation by this method is detailed in Section 6.2.1 using these t_{max} values, and results are presented in Section 6.3.1.

5.3.3 MAP Speckle Removal

The final step in removing the effect of speckle noise from the Gabor coefficients is the use of the maximum a posteriori (MAP) technique to estimate the speckle free coefficients. In terms of (5.19), the aim is to calculate $\hat{S}^l$, an estimate for the speckle free signal Gabor coefficients, given a SαS prior distribution for S^l and a zero mean Gaussian distribution for the speckle nosier N^l. MAP estimation is then performed as:

$$\hat{S}^l(Y^l) = \arg\max_{S^l} P_{S^l|Y^l}(S^l|Y^l) \tag{5.27}$$

where $P_x(x)$ is the PDF of variable x. By Bayes rule, this becomes:

$$\begin{aligned} \hat{S}^l(Y^l) &= \arg\max_{S^l} P_{N^l}(S^l - Y^l)P_{S^l}(S^l) \\ \hat{S}^l(Y^l) &= \arg\max_{S^l} P_{N^l}(N^l)P_{S^l}(S^l) \end{aligned} \tag{5.28}$$

MAP shrinkage functions for values of $(\hat{\alpha}, \hat{\gamma}, \hat{\sigma})$ are calculated numerically. Sample MAP shrinkage curves for various values of these parameters are displayed in Fig. 5.21. Fig. 5.20 displays examples of the effect of MAP estimation on the Gabor coefficients of a simulated speckled ultrasound image. The log Gabor filter bank used here contained a single scale, of wavelength $\lambda = 10$ pixels per cycle, and four orientations.

Table 5.5 displays examples of the reduction in variance in regions of simulated images due to speckle removal by MAP estimation. These are computed from simulated images containing shapes commonly seen in clinical ultrasonography. The images themselves are displayed in Section 6.2.2.1 below, in Fig. 6.3. Each of these simulated images contains three regions of differing scatter density. For each simulated image template, ten simulated images are applied. These are decomposed into orientation specific subbands using a log Gabor filter of a single scale (with wavelength $\lambda = 10\%$ of minimum image support), and four orientations. The parameters $(\hat{\alpha}, \hat{\gamma}, \hat{\sigma})$ of the model (5.21) are estimated using the method proposed in Section 5.3.2. MAP estimation of the speckle free Gabor coefficients is then performed by (5.28). The reduction in variance due to MAP filtering for each image region is measured. These variance reduction values are the averaged over the ten images simulated from each template, and presented in Table 5.5 for each orientation and scatter region.

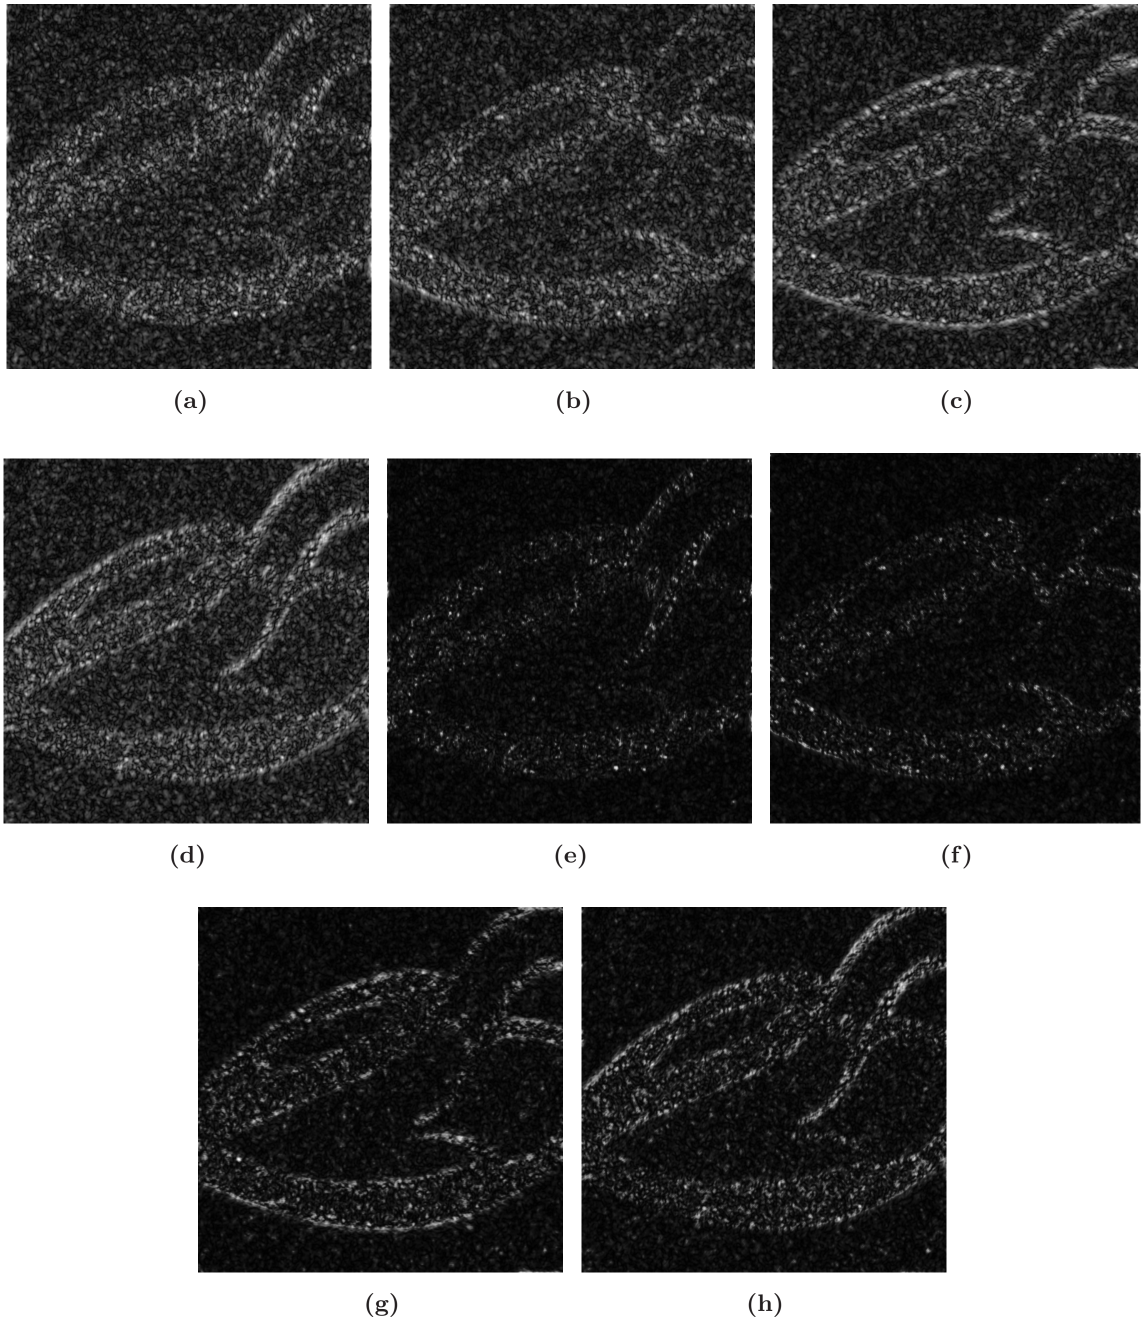
(a) (b) (c) (d) (e) (f) (g) (h)

Figure 5.20: Example of Gabor transform MAP speckle removal. (a)-(d) Magnitudes of Gabor coefficients over four orientations for a simulated speckled ultrasound image ($\theta = 0$ rad, $\theta = \frac{\pi}{4}$ rad, $\theta = \frac{\pi}{2}$ rad, and $\theta = \frac{3\pi}{4}$ rad, respectively). (e)-(h) Corresponding MAP estimates of speckle free values.

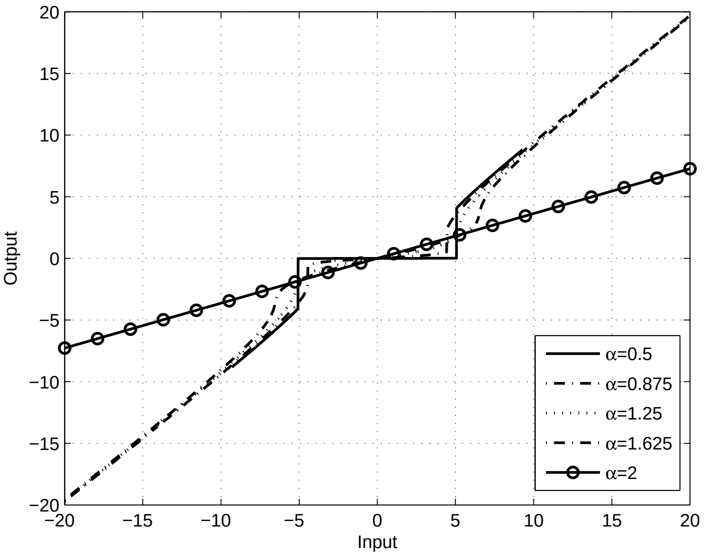

Figure 5.21: Sample MAP input/output functions for various values of α. The values of (γ, σ) are $(0.88, 1.75)$.

It can be seen that in all cases the reduction in variance is significant. Within each image, the reduction in variance is larger for denser scatterer regions. Some variation in the level of speckle reduction is observed between images. This effect is greatest in the low scatter density region, with a range of 20% between images. For the denser regions the reduction in variance due to MAP denoising is quite consistent.

Table 5.5: Variance reduction in simulated image regions due to MAP estimation.

		R1 (1%)	R2 (5%)	R3 (25%)
Image 1	$\theta = 0$ rad	35.93%	50.56%	79.59%
	$\theta = \frac{\pi}{4}$ rad	38.62%	51.06%	78.90%
	$\theta = \frac{\pi}{2}$ rad	44.06%	55.95%	78.21%
	$\theta = \frac{3\pi}{4}$ rad	39.99%	53.66%	79.50%
		R1 (1%)	R2 (5%)	R3 (25%)
Image 2	$\theta = 0$ rad	53.48%	61.95%	79.56%
	$\theta = \frac{\pi}{4}$ rad	53.56%	67.74%	74.31%
	$\theta = \frac{\pi}{2}$ rad	48.01%	63.03%	81.17%
	$\theta = \frac{3\pi}{4}$ rad	53.09%	67.97%	82.15%
		R1 (1%)	R2 (5%)	R3 (25%)
Image 3	$\theta = 0$ rad	50.90%	69.52%	82.56%
	$\theta = \frac{\pi}{4}$ rad	48.16%	67.61%	80.52%
	$\theta = \frac{\pi}{2}$ rad	44.67%	62.52%	80.58%
	$\theta = \frac{3\pi}{4}$ rad	54.67%	61.02%	82.83%
		R1 (1%)	R2 (5%)	R3 (25%)
Image 4	$\theta = 0$ rad	35.70%	51.74%	75.21%
	$\theta = \frac{\pi}{4}$ rad	34.63%	53.17%	83.95%
	$\theta = \frac{\pi}{2}$ rad	38.38%	56.34%	81.43%
	$\theta = \frac{3\pi}{4}$ rad	37.95%	58.36%	83.22%

5.4 Edge Detection Using Gabor Zero Crossings

This section details the development of an edge detection method based on Gabor zero crossings. A common approach to edge detection is the location of local maxima of an edge strength measure, which is often an approximation of the first derivative or gradient. Examples include the Sobel [125] and Prewitt [175] gradient approximations, and the derivative of Gaussian gradient approximations in the widely used Canny [108] edge detector. An alternate approach is offered by zero crossing edge detectors, which locate edges by determining the zero crossing locations of a second order derivative estimate. The most commonly used second order derivative function is the Laplacian of Gaussian (LoG) [124], which combines image smoothing to suppress noise with isotropic second derivative calculation. A widely known problem with zero crossing based edge detection methods, such as LoG, is the presence of spurious zero crossings which do not correspond to image edges. In the RoA method of [109], detailed in Section 5.2.3.2 above, a local ratio of averages estimate is used to constrain LoG zero crossings to regions likely to contain edges.

Gabor functions share some qualities with second order derivative functions, in particular a zero crossing is produced in the response to a step edge input. Fig. 5.22 displays an example of a 1D LoG kernel, and its response to a step edge input which is seen to contain a zero crossing at the edge location. This is compared to the step edge response of a series of 1D log Gabor filters, which are defined by a reduced form of the 2D log Gabor equation of (5.3):

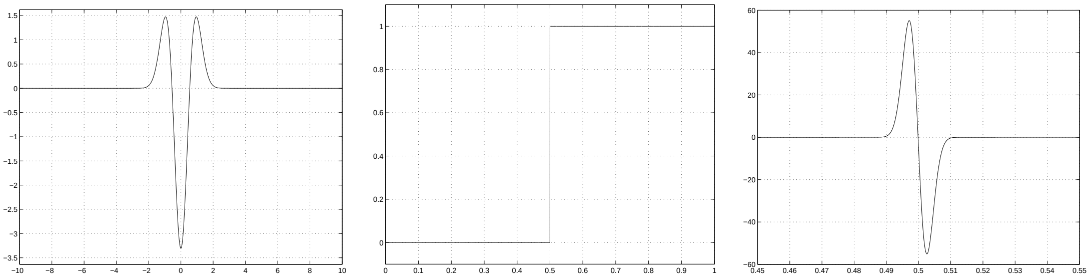

(a) **(b)** **(c)**

Figure 5.22: Response of a LoG kernel to a step edge. (a) LoG kernel, (b) Step Edge, (c) LoG Response.

$$G(\omega) = \exp\left(-\frac{\log(\omega/\omega_0)}{2\log(\sigma/\omega_0)}\right) \tag{5.29}$$

As before, ω_0 is the centre frequency and the value of σ/ω_0 defines the bandwidth. Fig. 5.23 displays example frequency magnitudes for 1D log Gabor filters with $\omega_0 = 0.1$ and various values of bandwidth ranging from $\sigma/\omega_0 = 0.45$ (approximately 2.5 octaves) to $\sigma/\omega_0 = 0.75$ (approximately 1 octave). Figs. 5.24(a)-(d) display the real components of the spatial domain representation of each of these log Gabor filters. It can be seen that for wider bandwidth values (low σ/ω_0), the spatial extent of the filter is more compact than in low bandwidth cases. A degree of ringing is observed for low bandwidth values. Figs. 5.24(e)-(h) display the responses of each of these to the step edge illustrated in Fig. 5.22(b). As with the LoG example above, each of the log Gabor filter produces a zero crossing at the edge location. In low bandwidth cases, the ringing noted above produces multiple zero crossings around the edges.

In the two dimensional case, the similarity between log Gabor filters and the LoG filter is less obvious. The clearest difference is the isotropic nature of the LoG kernel, while the Gabor filters are directional. It is noted experimentally however that the zero crossings of real (even symmetric) log Gabor output correspond with image edge locations for the

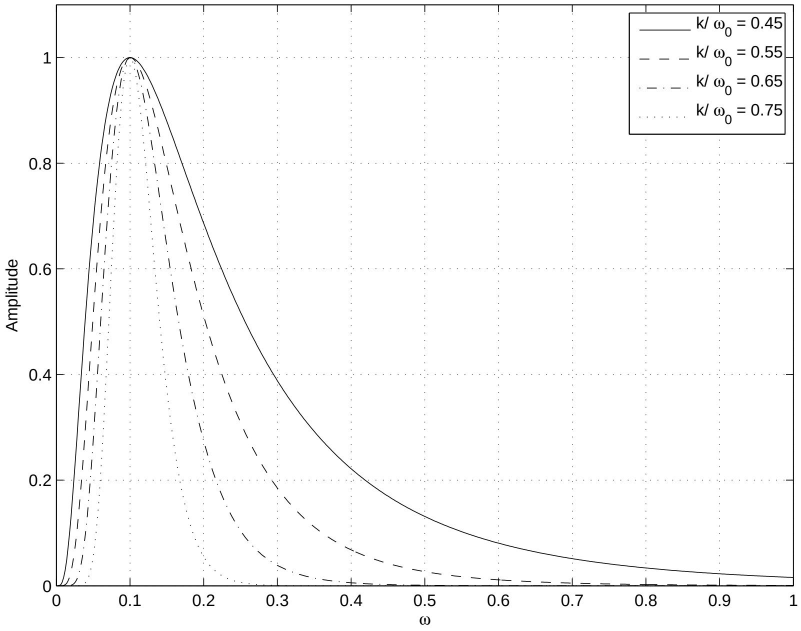

Figure 5.23: Frequency magnitude plots for 1D log Gabor filters with differing bandwidths.

2D case. Examples of real Gabor zero crossings are presented for a simulated speckled test image (displayed in Fig. 5.25(a)) and for a clinical echocardiographic frame (Fig. 5.25(b)). The zero crossing images for these are presented in Figs. 5.26-5.27, respectively. These zero crossing images are representative of what is seen throughout both simulated and clinical data sets. In both of these zero crossing images a grey background region is included, which shows the underlying image structure. The borders of these regions define the ideal edge maps. The edges separating the regions of simulated images are known precisely, however for clinical images, such as Fig. 5.27, these must be approximately delineated by hand. The log Gabor filter bank used in these examples consisted of two scales with four orientations, chosen so that the half magnitude frequency profiles of the filters just touch, similar to Fig. 5.2(a). The centre frequencies of each scale are defined by their wavelengths in pixels per cycle (ppc). The simulated image of Fig. 5.25(a) has support of 600×600 pixels, and the zero crossings in Fig. 5.26, are generated using wavelengths of $\lambda_1 = 50$ ppc, $\lambda_2 = 90$ ppc. The zero crossings of the clinical image of Fig. 5.25(b) (support 376×515), shown in Fig. 5.27, are calculated using wavelengths of $\lambda_1 = 18$ ppc, $\lambda_2 = 34$ ppc. In both cases, the radial bandwidth of the filters is 3 octaves ($\sigma/\omega_0 = 0.41$). Speckle noise is removed from the Gabor transform of these images before zero crossing location, using the method described in Section 5.3.

It is observed that these zero crossings match the image edges with orientation similar to that of the filter, i.e. edges approximately vertical correspond to ZCs in the $\theta = 0$ rad orientation (Figs. 5.26(a),(e) and Figs. 5.27(a),(e)), edges approximately horizontal correspond to ZCs in the $\theta = \frac{\pi}{2}$ rad orientation (Figs. 5.26(c),(g) and Figs. 5.27(c),(g)), while those edge regions close to diagonal are matched by the $\theta = \frac{\pi}{4}$ rad and $\theta = \frac{3\pi}{4}$ rad orientations (Figs. 5.26(b),(d),(f),(h) and Figs. 5.27(b),(d),(f),(h)).

There is an obvious issue to be addressed however, before the zero crossings of Figs. 5.26-5.27 could be used to generate edge maps: there is an abundance of spurious zero crossings present in each Gabor orientation which does not correspond to underlying image edge structure. These unwanted zero crossings are due to small variations in Gabor amplitude. A contributing factor to the amount of these fluctuations is the spatial domain ringing, noted in [176], which arises from the use of a single quadrature component of the Gabor filters.

The solution to this issue proposed here is the development of a measure of zero crossing coincidence with image edges. With the aim of constructing such a measure, two

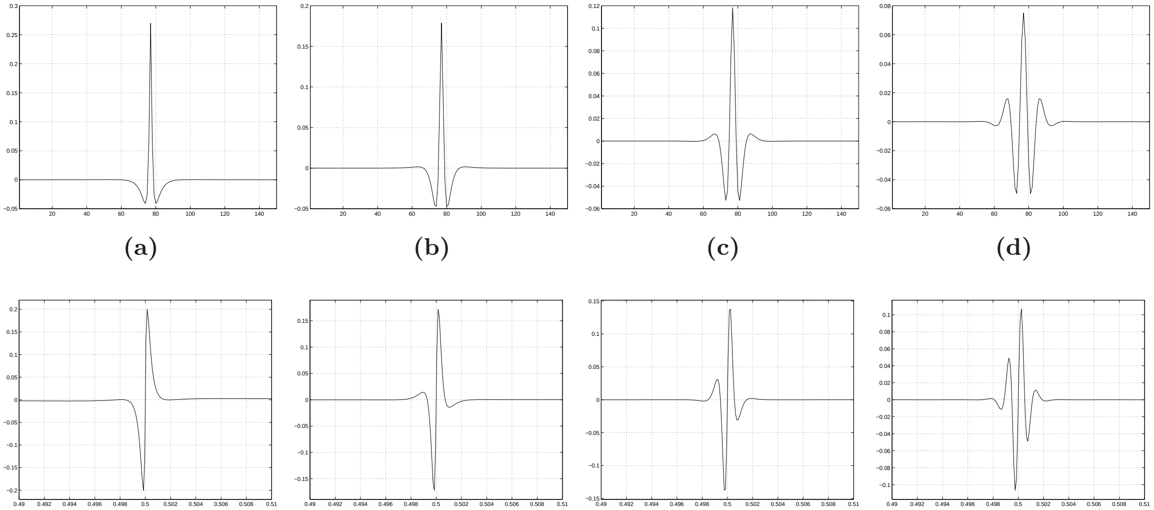

(a) **(b)** **(c)** **(d)**

(e) **(f)** **(g)** **(h)**

Figure 5.24: Response of log Gabor filter to step edge: (a)-(d) Real components of spatial domain representation of 1D log Gabor filters ($\sigma/\omega_0 = 0.45$, $\sigma/\omega_0 = 0.55$, $\sigma/\omega_0 = 0.65$, and $\sigma/\omega_0 = 0.75$, respectively. (e)-(h) Corresponding responses to the step edge of Fig. 5.22(b).

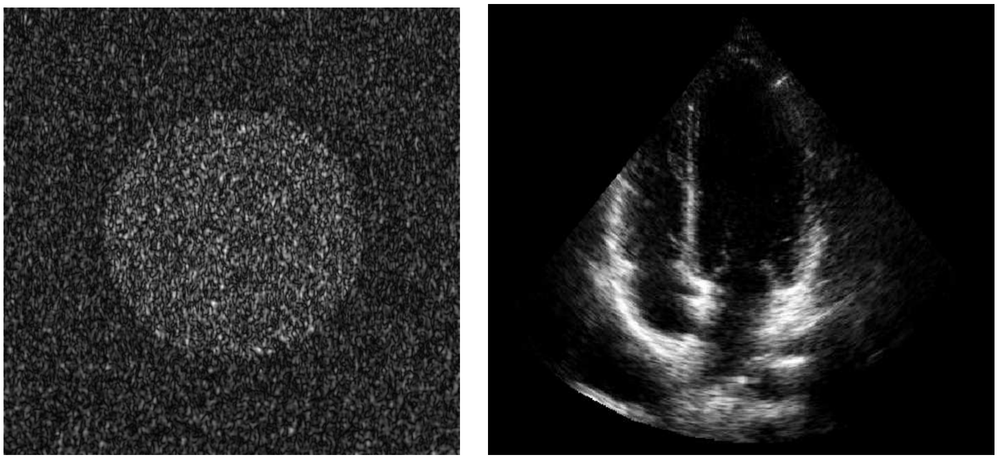

(a) **(b)**

Figure 5.25: Speckled images used in demonstration of Gabor zero crossing behaviour. (a) Simulated, (b) Clinical.

observations are made on the behaviour of the zero crossings at image edge locations:

1. Zero crossings at edge locations are surrounded by a region which is devoid of any other zero crossing. Specifically, this spacing occurs around zero crossings corresponding to image edges in the direction of the given Gabor orientation (i.e. horizontally in the

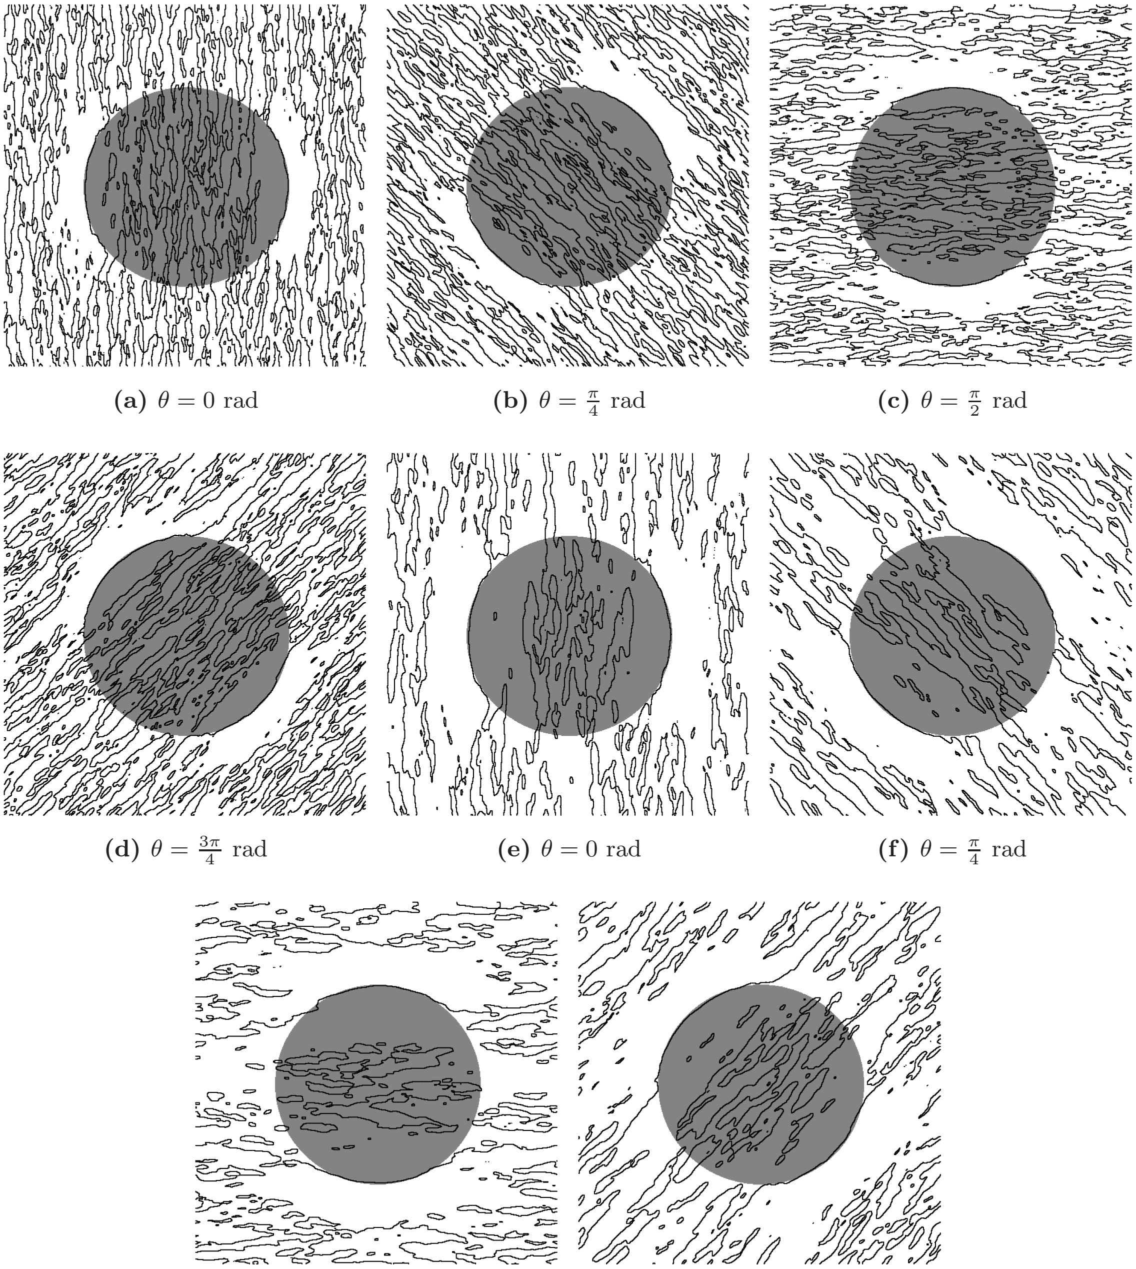

(a) $\theta = 0$ rad

(b) $\theta = \frac{\pi}{4}$ rad

(c) $\theta = \frac{\pi}{2}$ rad

(d) $\theta = \frac{3\pi}{4}$ rad

(e) $\theta = 0$ rad

(f) $\theta = \frac{\pi}{4}$ rad

(g) $\theta = \frac{\pi}{2}$ rad

(h) $\theta = \frac{3\pi}{4}$ rad

Figure 5.26: Example Gabor zero crossings, from the image of Fig. 5.25(a), using two scales and four orientations. The grey region represents underlying image structure, and its borders define the ideal edge. (a)-(d) Scale 1, (e)-(h) Scale 2.

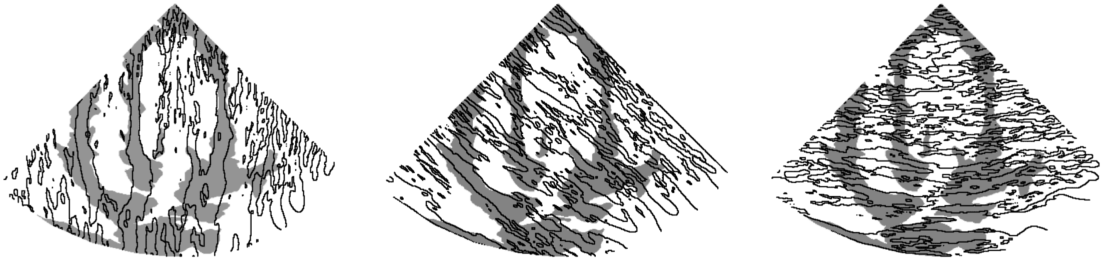

(a) $\theta = 0$ rad **(b)** $\theta = \frac{\pi}{4}$ rad **(c)** $\theta = \frac{\pi}{2}$ rad

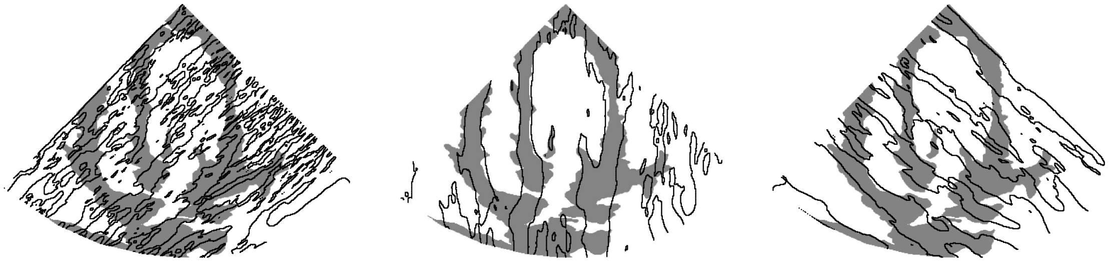

(d) $\theta = \frac{3\pi}{4}$ rad **(e)** $\theta = 0$ rad **(f)** $\theta = \frac{\pi}{4}$ rad

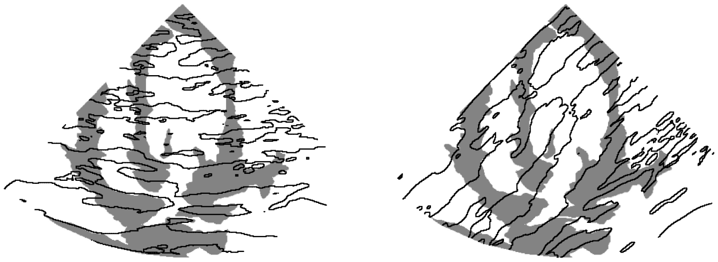

(g) $\theta = \frac{\pi}{2}$ rad **(h)** $\theta = \frac{3\pi}{4}$ rad

Figure 5.27: Example Gabor zero crossings, from the image of Fig. 5.25(b), using two scales and four orientations. The grey region represents underlying image structure, and its borders define the ideal edge. (a)-(d) Scale 1, (e)-(h) Scale 2.

$\theta = 0$ rad orientation, etc).

2. Zero crossings which correspond to image edges in each orientation are in close agreement with the their counterparts in coarser scales of the same orientation. This is a result of the stability of step edges across scales [138].

These properties of zero crossing behaviour are referred to here respectively as *orientational spacing* and *interscale agreement*. The quantification of each of these properties, as well as how they are affected by Gabor frequency and bandwidth, are demonstrated below. Finally, this section concludes by detailing how a measure of zero crossing coincidence with image edges can be constructed from the quantification of the above properties, and how this measure can be used to detect image edges. A convention is followed through out of denoting by $ZC_{s,o}(i,j)$ the zero crossings of the real log Gabor coefficients in scale s and orientation o (i.e. the zero crossings of $\Re[G_{s,o}\{y\}]$).

5.4.1 Zero Crossing Orientational Spacing

5.4.1.1 Quantification

This section details the development of a method of quantifying the orientational spacing of Gabor zero crossings. Firstly, it is noted that the Euclidean distance between two locations in a discrete two dimensional array, such as an image, is given by:

$$d = \sqrt{(i-k)^2 + (j-l)^2} \tag{5.30}$$

where (i,j) and (k,l) are the spatial indices of the image locations. Although these indices are normally integer valued in image processing applications, the distance between locations can be any positive real number ($d \in \mathbb{R}_{\geq 0}$).

Binary images $ZC(i,j)$ are used to represent zero crossing locations here, with $ZC(i,j) = 1$ indicating the presence of a zero crossing at location (i,j). This basic example of (5.30) can be extended to measure the distance from any zero crossing to its nearest neighbour. This is quantified by measuring the distance from (i,j) to each other image location for which $ZC = 1$, and then taking the minimum as the distance to the nearest neighbour:

$$d_{min} = \min \sqrt{(i-k')^2 + (j-l')^2}, \quad \forall (k',l') : ZC(k',l') = 1|_{i \neq k'}^{j \neq l'} \tag{5.31}$$

However, in the present case it is required that distance between zero crossings be measured in the direction of relevant Gabor orientation (e.g. vertically for the $\theta = 0$ orientation, horizontally for the $\theta = \frac{\pi}{2}$ orientation, and so forth). In the case of the horizontal and vertical orientations this is relatively straightforward, requiring only that (5.31) be modified such that either $i = k'$ or $j = l'$. However the problem is non trivial for an arbitrary orientation θ. The solution presented here is to approximate measurement in a given orientation by the use of image rotation, as follows:

- For zero crossing image $ZC_{s,o}(i,j)$, measured at Gabor orientation $\theta = \theta_o$, a rotation is performed through $-\theta_o$. This rotated zero crossing image is denoted $ZC_{s,o}^{-\theta_o}(i,j)$, and the desired direction of measurement is now aligned vertically.
- The orientational spacing for all zero crossing are measured, as the distance from each to its nearest neighbour in the vertical direction. The array of orientational spacing measurements is denoted $S_{s,o}^{-\theta_o}(i,j)$.
- A second inverse rotation through θ_o returns $S_{s,o}(i,j)$ to alignment with the original $ZC_{s,o}(i,j)$ image.

Fig. 5.28 displays an example of this process. The zero crossings in Fig. 5.28(a) are measured in the diagonal orientation of $\theta = \frac{\pi}{4}$ rad. To make measure the distance indicated by the arrow, a rotation through $-\frac{\pi}{4}$ rad is performed, after which the desired measurement can be performed vertically. Calculation of $S_{s,o}^{-\theta_o}(i,j)$ is expressed as:

$$S_{s,o}^{-\theta_o}(i,j) = \begin{cases} 0, & \text{if } ZC_{s,o}^{-\theta_o}(i,j) = 0 \\ min[\sqrt{(i-k)^2}|_{i \neq k}], & \\ \quad \forall k : ZC_{s,o}^{-\theta_o}(k,j) = 1, & \text{otherwise} \end{cases} \tag{5.32}$$

Clockwise rotation through an angle of θ can be computed using the transformation matrix:

$$\begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \tag{5.33}$$

Finally, the orientational spacing measure is normalised within each scale, by dividing each $S_{s,o}$ by the largest value found within each scale. Figs. 5.29-5.30 show the result of the orientational spacing measurement above, when applied to the zero crossings of Figs. 5.26-5.27. The second inverse rotation described above is performed in all cases after calculation of (5.32), to return each $S_{s,o}$ to the original orientation of the zero crossing image.

Comparing these images to the zero crossings of Figs. 5.26-5.27, it can be seen that higher values of zero crossing orientational spacing (represented by darker lines in

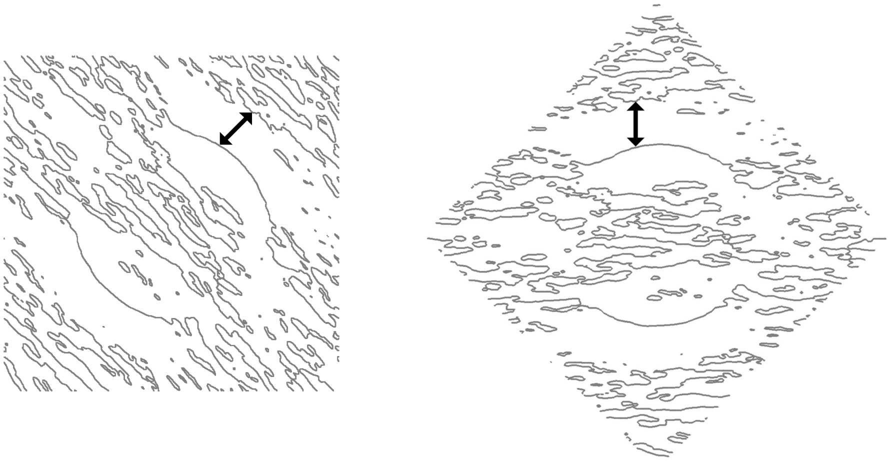
(a) **(b)**

Figure 5.28: Measuring zero crossing orientational spacing using rotation. (a) $ZC_o(i,j)$, (b) $ZC_o^{-\frac{\pi}{4}}(i,j)$.

Figs. 5.29-5.30) are found for zero crossings surrounded by a spacing in the direction of the relevant orientation. Furthermore, these zero crossings are those which are closest to the underlying image edges.

However, an issue can be observed which reduces the degree to which high orientational spacing values are exclusive to image edge locations. High orientational spacing values can be seen in non edge regions, directly opposite zero crossings coinciding with image edges. This is a side effect of measuring edge spacing in this manner: two zero crossing pixels on opposite sides of an empty space are assigned the same value, even if only one corresponds with an image edge. Another way of stating this is that the distance measurement of (5.32) is commutative, i.e. the distance between zero crossings $A \rightarrow B$ necessary equals that of $B \rightarrow A$. As will be shown below that this effect can mitigated when constructing an edge detection measure by combining the orientational spacing measure with a quantification of interscale agreement.

Notwithstanding these spurious high measurements, the values of orientational spacing are seen to be on average higher for zero crossings close to underlying image edges. Next, a method for quantification of how well $S_{s,o}$ can discriminate zero crossings corresponding to image edges is presented.

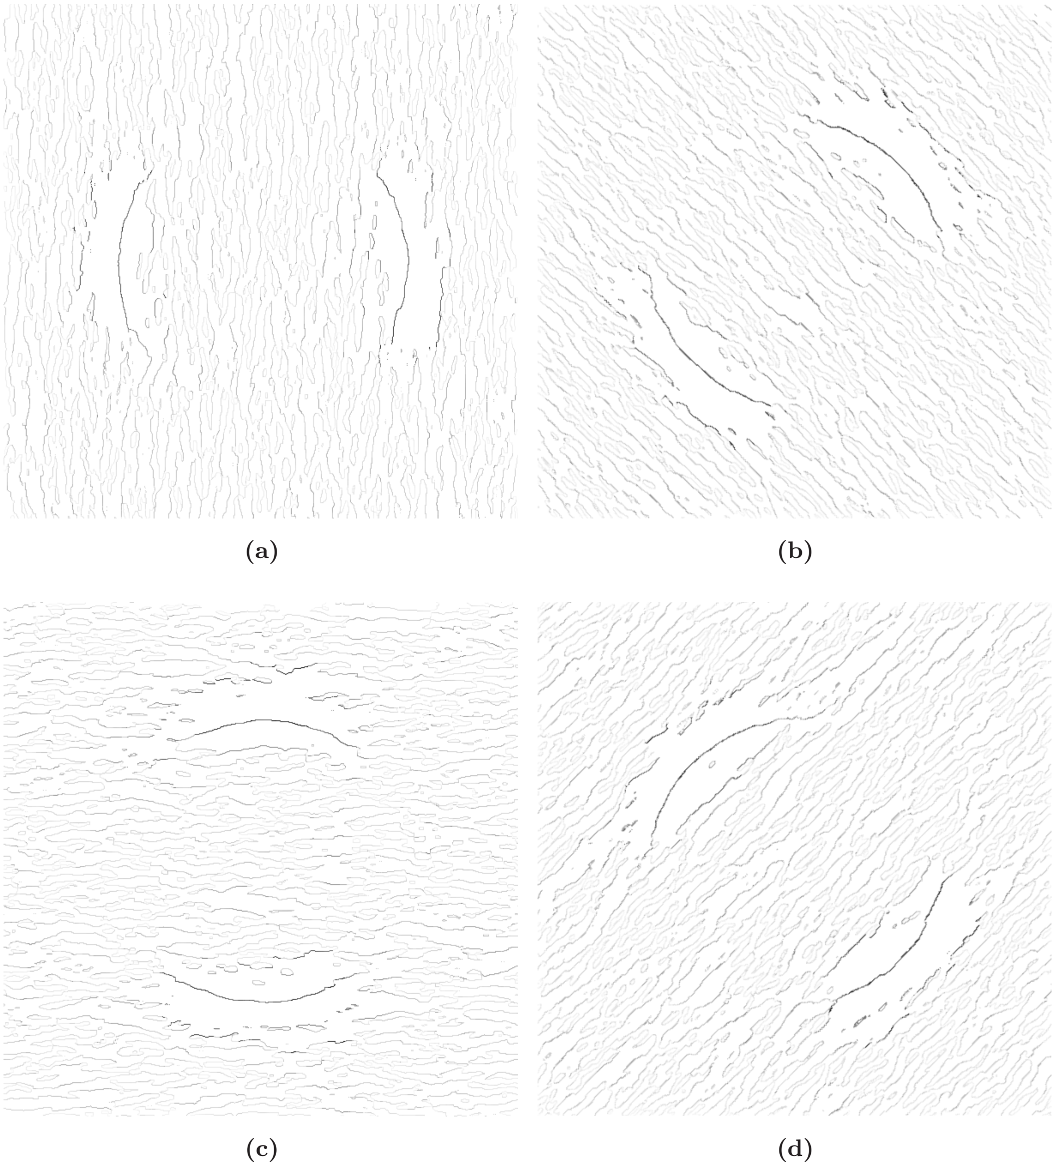

Figure 5.29: Orientational spacing of the first scale of the Gabor zero crossings of Fig. 5.26. Darker regions correspond to larger measured orientational spacing. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

5.4.1.2 Edge Discrimination Properties

Having defined the measure of orientational edge spacing $S_{s,o}$ in the section above, a quantification of how well this measure is able to discriminate zero crossings which correspond with image edges is presented. How this discrimination capability is affected by the

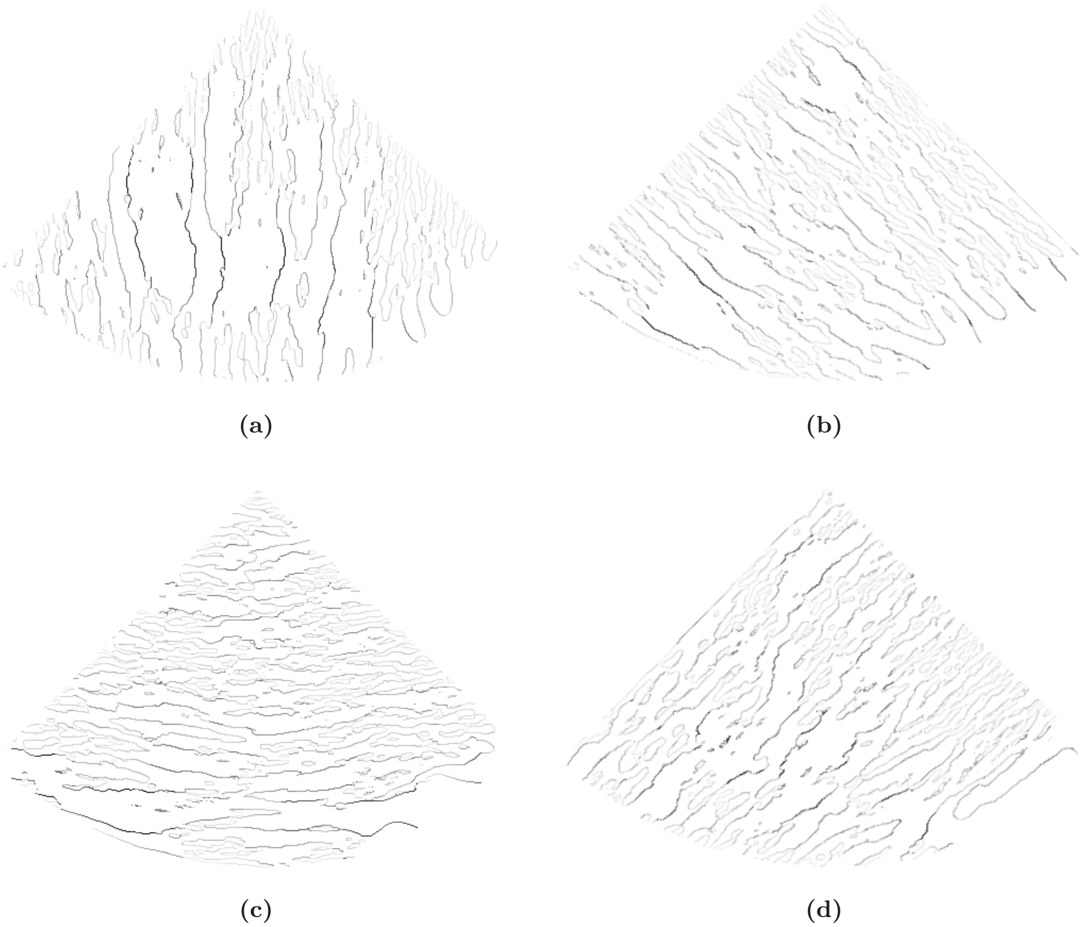

Figure 5.30: Orientational spacing of the first scale of the Gabor zero crossings of Fig. 5.27. Darker regions correspond to larger measured orientational spacing. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

Gabor filter parameters which determine centre frequency and bandwidth is also explored here.

To quantify the ability of $S_{s,o}$ to discriminate zero crossings coinciding with an image edge, a method of finding those coinciding zero crossings must be found. For the simulated and clinical images presented in Fig. 5.25, ideal edge maps are defined. These are used in Figs. 5.26-5.27 above, where they are represented by the borders of grey regions. For simulated images such as Fig. 5.25(a), accurate ideal edge maps are readily available. In the case of clinical images such as Fig. 5.25(b) these edges are manually defined and therefore approximate in nature. However it can be seen that they do represent the underlying image

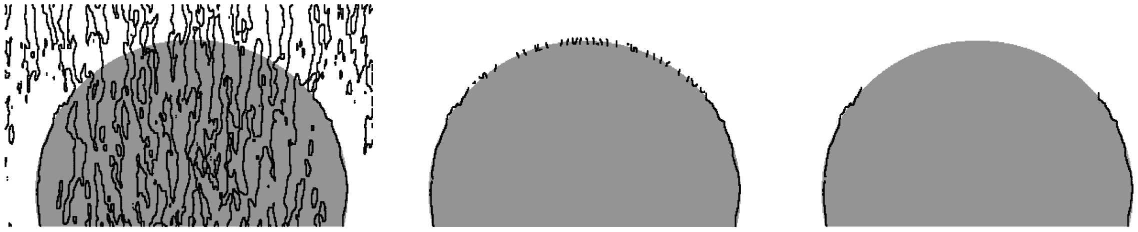
(a) (b) (c)

Figure 5.31: Example Short zero crossing removal, from Fig. 5.26(a) (Scale 1, $\theta = 0$ rad). (a) Zero crossing ZC_o, (b) Thresholded distance $D_{TH,o}$, (c) Short grouping removal $D'_{TH,o}$.

structure of Fig. 5.25(b) quite accurately.

The distance from the zero crossings of each Gabor orientation o, $ZC_o(i,j)$, to the nearest ideal edge pixel is measured as:

$$D_o(i,j) = \begin{cases} 0, & \text{if } ZC_o(i,j) = 0 \\ min[\sqrt{(i-n)^2 + (j-m)^2}], & \\ \quad \forall (n,m) : E_{ideal}(n,m) = 1, & \text{otherwise} \end{cases} \tag{5.34}$$

where E_{ideal} is the ideal edge map, consisting of edge pixels located where $E_{ideal}(n,m) = 1$. A threshold is then applied, so that zero crossings below a given distance to the nearest ideal edge pixel are considered to coincide with an image edge, and those beyond this distance are deemed not to agree with image edges. This is represented as:

$$D_{TH,o}(i,j) = D_o(i,j) \geq T_D \tag{5.35}$$

where T_D is a threshold. After thresholding the distance measure D_o to produce $D_{TH,o}$, a post processing step which discards very short groups of zero crossings is employed. This is necessary to discount zero crossings which intersect ideal image edges which are not aligned with the given Gabor orientation. For example, Fig. 5.31(a) displays a region of the zero crossings of Fig. 5.26(a).

The image edges which are approximately vertical (to the left and right) are aligned with the orientation of this Gabor filter, and it can be seen that zero crossings coincide with these edges. In the upper portion of Fig. 5.31(a) however, the image edges are not aligned with the Gabor orientation, being closer to the horizontal than to the vertical orientation of the filter. The zero crossings do not follow the image edge here, but they are seen to intersect it at numerous locations. Fig. 5.31(b) shows $D_{TH,o}$ calculated by (5.35), in which only those zero crossings close to the ideal image edge remain. In the upper portion of this image, the locations where zero crossings intercept the approximately horizontal edge leaves a number of short groupings of zero crossings. The zero crossing lines which intersect the ideal edges here do so at angles which differ significantly from the local direction of the ideal edge. Thus after thresholding to remove zero crossings far from ideal edges, these lines leave groupings of zero crossings which are always of short length, relative to the length of zero crossing lines coinciding with image edges. They are removed here by measuring the length of each connected region in $D_{TH,o}$, and removing those below a certain length. Connectivity is measured using the 8 neighbours of each pixel. This operation is represented by:

$$D'_{TH,o}(i,j) = \begin{cases} D_{TH,o}(i,j), & \text{if } C_{D,o}(i,j) \geq T_C \\ 0 & \text{otherwise} \end{cases} \tag{5.36}$$

where $C(i,j)$ is the number of `true` binary pixels connected to $D_{TH,o}(i,j)$, and T_C is the desired minimum length of zero crossing grouping. Fig. 5.31(c) displays the effect of removing these short edges. It is observed that only the zero crossings which coincide with image edges remain.

Fig. 5.32 displays the zero crossings which are found to be in agreement with the image edges of the simulated image of Fig. 5.25(a) in each orientation. The zero crossings used here are the first scale of Fig. 5.26, i.e. Figs. 5.26(a)-(d). These are found by calculating the distance to the ideal edges using (5.34), and thresholding by (5.35), using $T_D = 10$ pixels. This threshold is heuristically determined for the simulated images, to select only zero crossings located in close proximity to image edges. As described above short zero crossing regions are removed by (5.36), using $T_C = 15$ pixels. It can be seen that for each orientation, only the zero crossings which are in close proximity to image edges agreeing with that orientation are present.

Similarly, zero crossings which are found to be in agreement with ideal image edges for the clinical image of Fig. 5.25(b) are displayed in Fig. 5.33. In this case a slightly larger threshold is required due to the approximate nature of the ideal edge, and the images of Fig. 5.33 are using $T_D = 20$ pixels. Short zero crossing groupings are also removed by (5.36) using $T_C = 15$ pixels. As in the simulated case, the zero crossings remaining here are those in agreement with ideal image edges matching the orientation in question.

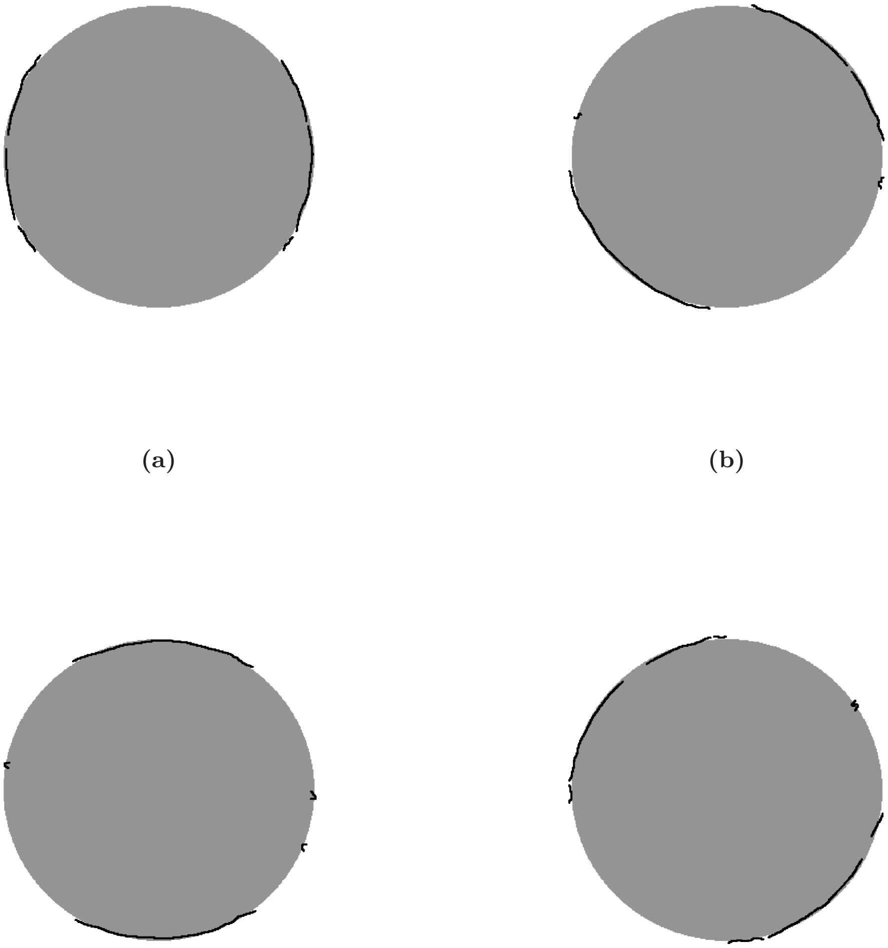

(a) (b) (c) (d)

Figure 5.32: Gabor zero crossings determined to coincide with image edges, for the simulated image of Fig. 5.25(a). Zero crossings are from the first scale, and ideal edges are represented by the borders of the shaded region. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

Now that a method is established to locate the zero crossings in each orientation which coincide with a predefined ideal edge map, it is possible to measure the degree to which the orientational edge spacing measure $S_{s,o}$ is capable of discriminating these zero

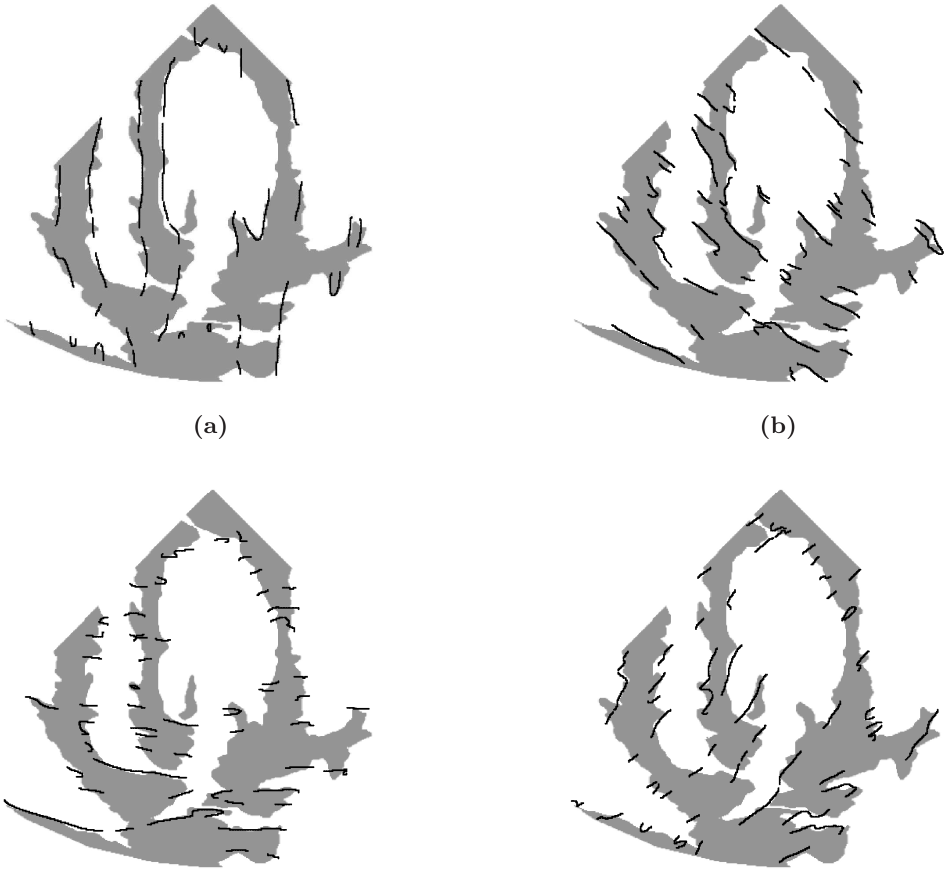

(a)

(b)

(c)

(d)

Figure 5.33: Gabor zero crossings determined to coincide with image edges, for the simulated image of Fig. 5.25(b). Zero crossings are from the first scale, and ideal edges are represented by the borders of the shaded region. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

crossings. A straightforward method of quantifying this is to calculate $r_{D,o}$, the ratio of mean $S_{s,o}$ values for zero crossings coinciding with image edges to the mean value of $S_{s,o}$ for other zero crossings in that orientation:

$$r_{D,o} = \frac{\dfrac{1}{N_1}\displaystyle\sum_{i=1}^{N_1} S_o(m,n)}{\dfrac{1}{N_2}\displaystyle\sum_{i=1}^{N_2} S_o(j,k)}, \tag{5.37}$$
$$\forall(m,n) : D'_{TH,o}(m,n) = 1, \quad \forall(j,k) : D'_{TH,o}(j,k) = 0, ZC_o(j,k) = 1$$

where (m,n) are the indices of the N_1 zero crossings determined to coincide with the ideal edges by the process described by (5.34),(5.35) and (5.36). Here (j,k) are the indices of the N_2 zero crossings which do not coincide with the edge pixels. Values of $r_{D,o} > 1$ indicate that S_o was on average higher for zero crossings which coincided with image edges than for the zero crossings which did not. Higher values of $r_{D,o}$ provide an indication of the strength of the discrimination between these zero crossings possible using the orientational edge spacing measure S_o.

Table 5.6: Discrimination measure $r_{D,o}$ for Gabor zero crossing orientational spacing, of the zero crossings of the simulated ultrasound of Fig. 5.25(a). Values are presented for various frequency and bandwidth values (defined by λ and σ/ω_0, respectively), and are averaged over orientations.

	σ/ω_0					
	0.40	0.47	0.54	0.61	0.68	0.75
$\lambda = 50.0$ ppc	3.019	2.523	1.991	1.587	1.306	1.131
$\lambda = 60.0$ ppc	3.529	2.700	2.070	1.617	1.267	1.111
$\lambda = 70.0$ ppc	3.563	2.782	2.131	1.594	1.212	1.103
$\lambda = 80.0$ ppc	3.532	2.771	2.101	1.536	1.211	1.077
$\lambda = 90.0$ ppc	3.383	2.634	1.991	1.490	1.202	1.083
$\lambda = 100.0$ ppc	3.060	2.440	1.842	1.411	1.185	1.061

The ability of S_o to discriminate those zero crossings which coincide with image edges, measured by $r_{D,o}$, is explored for various values of Gabor frequency and bandwidth, for both simulated and clinical images. For a representative image of each type, 36 different log Gabor filter banks are applied, comprising all combinations of six values of centre frequency (determined by Gabor wavelength λ), and six values of bandwidth. Bandwidth is determined by the ratio of log Gabor parameters σ/ω_0. This is varied linearly from 0.4 to 0.75, corresponding to radial bandwidths decreasing from 3 to 1 octave. In all cases a single Gabor scale and four orientations are used, with orientations spaced so as the half magnitude profiles just touch in the radial direction, as in the examples of Fig. 5.3.

For the simulated image of Fig. 5.25(a), 36 Gabor filter banks are applied as above. In each case, zero crossings ZC_o are found, and then orientational spacings S_o are calculated by (5.34). It is determined which of these zero crossings coincide with ideal edges by (5.35) and (5.36), and the discrimination measure $r_{D,o}$ is found for each orientation. The $r_{D,o}$ values are averaged over Gabor orientations, and summary results are presented in Table 5.6. It is seen that in all cases the average $r_{D,o}$ is greater than one, so S_o here has a higher average value for zero crossings corresponding with image edges than for those which do not. In some cases the average $r_{D,o}$ value is in excess of 3.5, indicating that the orientational spacing measure S_o is quite capable of discriminating those zero crossings which coincide with ideal edges. Higher bandwidths produce larger average $r_{D,o}$ values, indicating that these high bandwidths are better suited to this discrimination task. For lower bandwidth filter banks, the higher centre frequencies appear to be capable of better discrimination by this measure. This is explained by the fact for these low bandwidths, only Gabor filters with higher centre frequencies can capture the high frequency details of the edge contours.

Table 5.7: Discrimination measure for Gabor zero crossing orientational spacing, of zero crossings of the clinical ultrasound of Fig. 5.25(b). Values are presented for various frequency and bandwidth values (defined by λ and σ/ω_0, respectively), and are averaged over orientations.

	σ/ω_0					
	0.40	0.47	0.54	0.61	0.68	0.75
$\lambda = 10.0$ ppc	2.993	2.980	2.955	2.996	3.019	2.975
$\lambda = 18.0$ ppc	3.491	3.223	3.180	3.062	3.048	2.955
$\lambda = 26.0$ ppc	3.579	3.332	3.055	3.000	2.948	3.048
$\lambda = 34.0$ ppc	3.250	3.186	3.060	2.933	2.903	3.022
$\lambda = 42.0$ ppc	3.084	3.160	3.122	3.190	3.119	3.092
$\lambda = 50.0$ ppc	3.100	3.005	3.076	2.992	3.089	2.996

For higher bandwidth values, lower Gabor frequencies achieve higher average $r_{D,o}$ scores. This increase in discrimination capability for lower frequency filters is due to their increased inclusion of higher frequency detail as bandwidth increases.

Similar to the simulated case above, for the clinical image of Fig. 5.25(b), a set of 36 log Gabor filter banks are applied with varying centre frequency and bandwidth. As before, discrimination measures $r_{D,o}$ are calculated for the orientational edge spacing in each Gabor orientation. These values are averaged over orientations, and are displayed in Fig. 5.7. As in the simulated case, the average $r_{D,o}$ values are greater than unity in all cases, indicating that the S_o measure is capable of discriminating zero crossings corresponding to image edges in clinical images. As in the simulated case, average $r_{D,o}$ values are seen to be slightly higher for higher values of bandwidth, although the effect is less pronounced in this case. It is tentatively suggested that this may be due to the stability property of edges in clinical images, as mentioned by [138], i.e. the edges in this clinical image are present across a wide range of frequencies.

In this section, a method of quantifying the ability of the zero crossing orientational measure S_o to discriminate those zero crossings which coincide with image edges is presented. Although it is noted that spurious large values of S_o are measured in some regions which do not correspond to image edges, it is shown that for both simulated and clinical images this measure is capable of discriminating the zero crossings which coincide with image edges. In the subsequent section, a measure of the agreement of Gabor zero crossings between scales is presented.

5.4.2 Zero Crossing Interscale Agreement

5.4.2.1 Quantification

This section details a method of quantifying the interscale agreement between Gabor zero crossings in each orientation. As with the measurement of orientational edge spacing above, this is based on computing the Euclidean distance from one zero crossing to its nearest neighbour. The differences between this measurement and the orientational spacing measurement (5.32) are that:

- The nearest neighbour is permitted to be located in any direction, whereas in (5.32) all distances are measured in the direction of Gabor orientation.
- The measurement is between each zero crossing in the scale of interest $ZC_{s,o}$ and its nearest neighbour in the next coarser scale, $ZC_{s+1,o}$.

A final but crucial consideration is that in this case the aim is to measure the agreement, rather than the distance, between zero crossings. "Agreement" as it is used here can be considered the reciprocal of distance. The approach of inverting distance values mathematically has a number of drawbacks. Firstly, computational issues arise when perfect agreement is found, whereby the distance between pixels is zero. Secondly, this approach would result in a monotonic but non-linear relationship between distance and agreement. The approach taken here avoids these issues as follows:

1. Distances $\varphi_{s,o}$ are measured between each zero crossing in $ZC_{s,o}$ and its nearest neighbour in $ZC_{s+1,o}$.
2. These distances range from 0 (in the case of perfect agreement between scales), to a maximum value representing the worst agreement found.
3. A distinction must be made between locations having perfect agreement $\varphi_{s,o}(i,j) = 0$ and locations where there are no zero crossings in the scale of interest ($ZC_{s,o}(j,k) = 0$). Thus locations devoid of zero crossings in scale s are set to -1.
4. Finally, the distance measure $\varphi_{s,o}$ is transformed to an agreement measure $A_{s,o}$. This is performed by subtracting from $\varphi_{s,o}$ the largest distance value found in scale s, and adding one. The result is an agreement measure $A_{s,o}$ which equals zero where no zero

crossing is present ($A_{s,o}(j,k) = 0$ if $ZC_{s,o}(j,k) = 0$), and varies from unity (for perfect agreement) to a small positive value for the poorest agreement found in scale s.

Mathematically, these steps are represented as:

$$\varphi_{s,o}(i,j) = \begin{cases} -1, & \text{if } ZC_{s,o}(i,j) = 0 \\ min[\sqrt{(i-k)^2 + (j-l)^2}], \quad \forall (k,l) : ZC_{s+1,o}(k,l) = 1, & \text{otherwise} \end{cases} \tag{5.38}$$

$$A_{s,o}(i,j) = |\varphi_{s,o}(i,j) - M_\varphi + 1|, \quad M_\varphi = \max \varphi_{s,o}(i,j)|_{o=1,.,N_{ori}} \tag{5.39}$$

The Gabor zero crossing interscale agreement for the example zero crossings of Figs. 5.26-5.27 are displayed in Fig. 5.34 and Fig. 5.35. These are calculated from simulated and clinical images (Figs. 5.25(a) and 5.25(b), respectively). In both cases it can be seen that the interscale agreement, as measured by $A_{s,o}$, has large values as expected for regions where the zero crossings are in agreement across scales. Zero crossings which align with ideal image edges are seen to have high interscale agreement. Areas of high interscale agreement are not confined just to zero crossings which are coincident with edges however. In Section 5.4.1.1 it is noted that spurious high values of orientational spacing are observed adjacent to true image edges in Figs. 5.29-5.30. An important property of the interscale agreement measured are that they are observed to have particularly low values in regions where spurious effects are present in the orientational spacing measure. This indicates that a combination of interscale agreement and orientational spacing would not suffer from such high values in areas not corresponding to image edges. These regions of low interscale agreement are due to the increased orientational spacing in the coarser scales $ZC_{s+1,o}$, used to calculate the agreement between scales by (5.38).

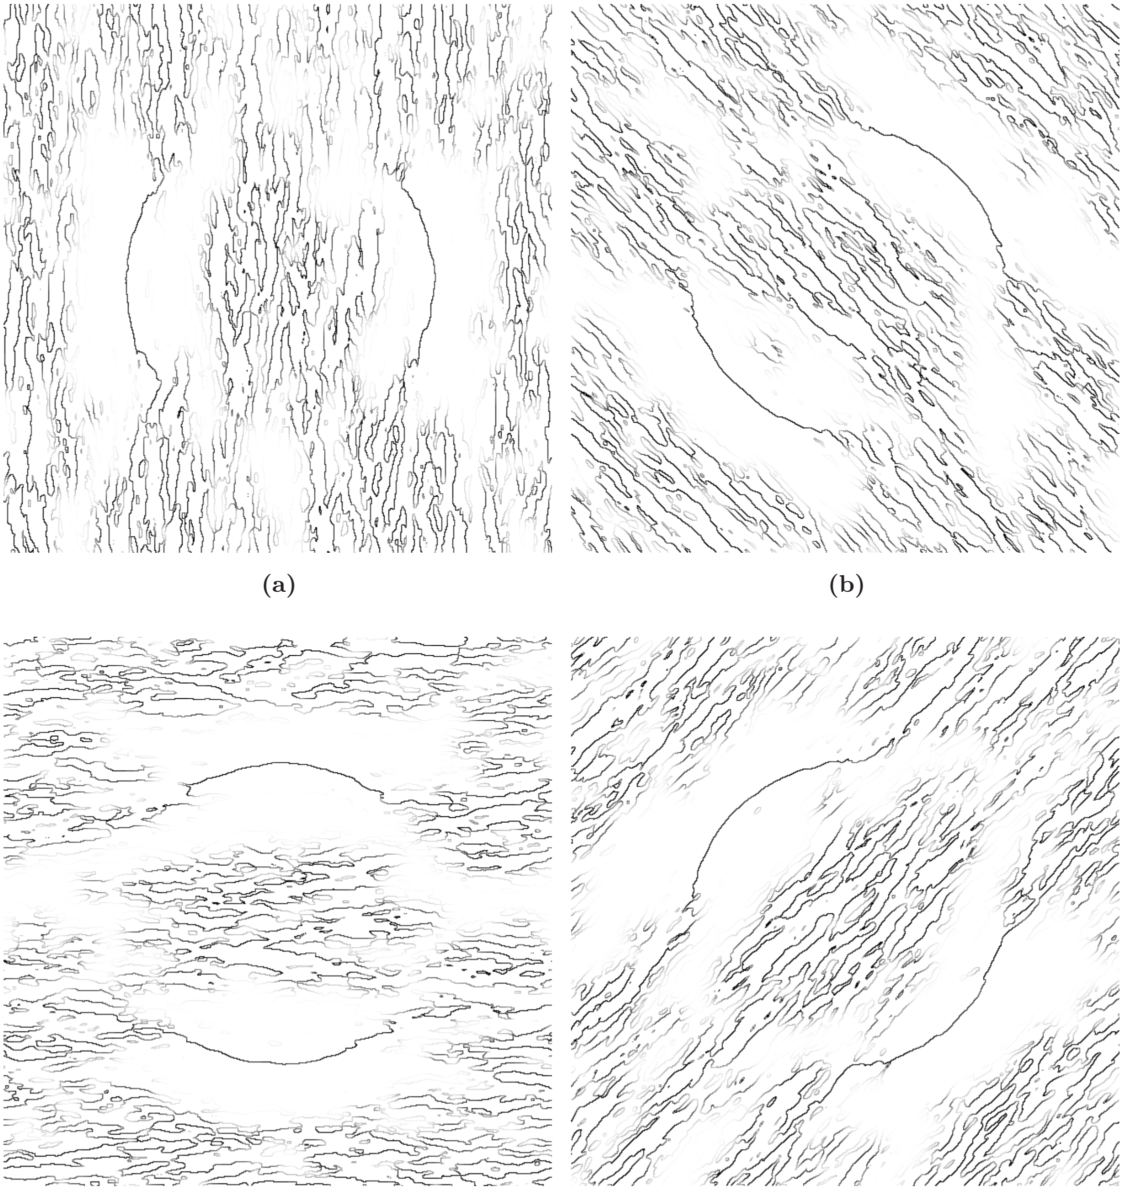

(a) **(b)**

(c) **(d)**

Figure 5.34: Gabor zero crossing interscale agreement, for the zero crossings of Fig. 5.26, from the simulated image of Fig. 5.25(a). Darker regions correspond to higher interscale agreement. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

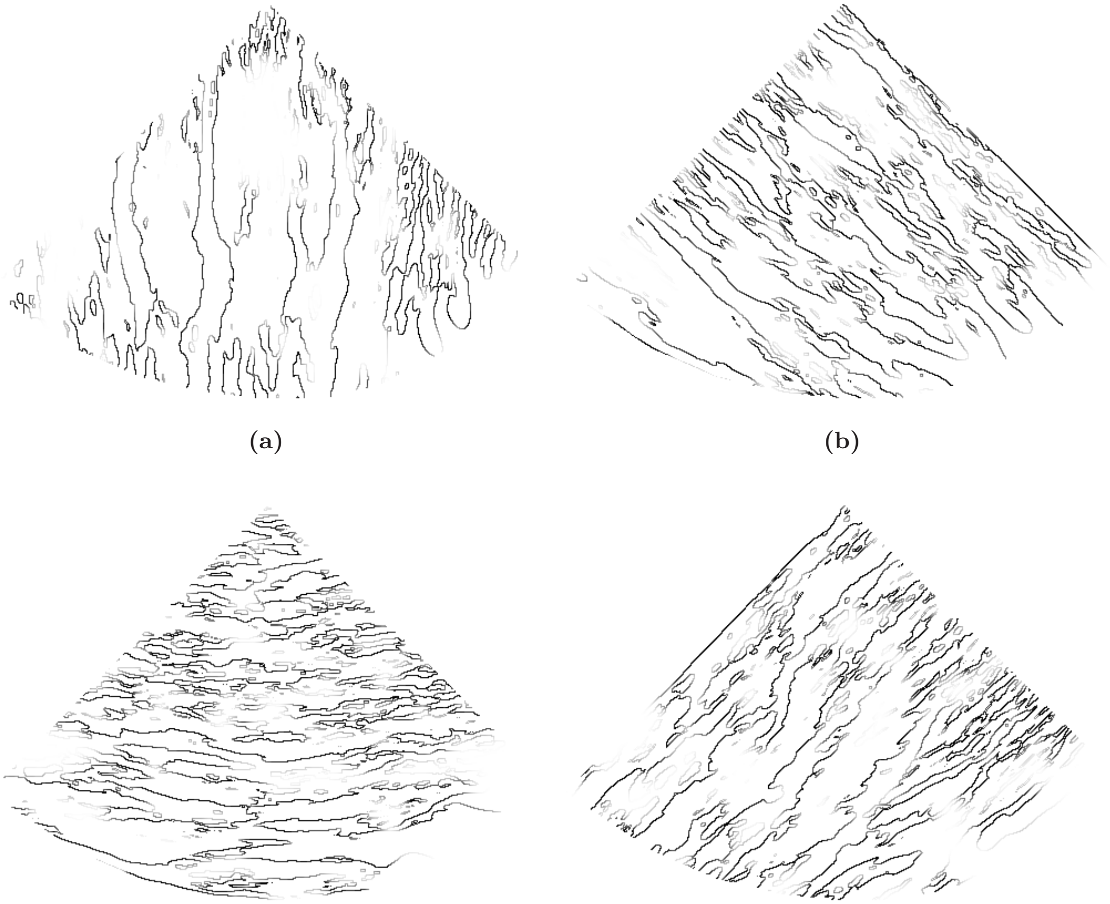

(a) **(b)**

(c) **(d)**

Figure 5.35: Gabor zero crossing interscale agreement, for the zero crossings of Fig. 5.26, from the simulated image of Fig. 5.25(a). Darker regions correspond to higher interscale agreement. (a) $\theta = 0$ rad, (b) $\theta = \frac{\pi}{4}$ rad, (c) $\theta = \frac{\pi}{2}$ rad, (d) $\theta = \frac{3\pi}{4}$ rad.

5.4.2.2 Edge Discrimination Properties

The capability of the interscale agreement measure $A_{s,o}$, presented in Section 5.4.2.1 above, to discriminate zero crossings which coincide with image edges is explored here. This is performed in a similar manner to the investigation of the discrimination properties of the orientational spacing measure $S_{s,o}$ presented in Section 5.4.1.2 above. The same method of determining the zero crossings which coincide with image edges is employed as in Section 5.4.1.2, generating a binary mask $D'_{TH,o}$, which has a `true` value for zero crossings in orientation o which coincide with ideal image edges. Examples of these are shown in Figs. 5.32-5.33 above.

In a manner similar to the investigation of Section 5.4.1.2, 36 log Gabor filter banks are employed, comprised of the combinations of six centre frequencies and six bandwidth values. The values bandwidth (determined by the ratio of log Gabor parameters σ/ω_0) are again varied linearly from 0.4 to 0.75, corresponding to radial bandwidths decreasing from 3 to 1 octave. In contrast to the investigation of the previous section however, here two scales of decomposition are required. A range of wavelength values (λ) are used to define the centre frequencies of the first (finest) scale, and the second scale has a wavelength of 2λ. The scale of interest s in (5.38) is chosen as the first scale in all cases, therefore $A_{s,o}$ measures the agreement between the zero crossings of the first scale with the second. In all cases four orientations are used, spaced so the half magnitude profiles just touch in the radial direction, as in the examples of Fig. 5.3.

Similar to the investigation of Section 5.4.1.2, discrimination capability for each frequency and bandwidth is measured by quantifying the ratio of the mean value of $A_{s,o}$ for zero crossings which coincide with image edges, to the mean value of $A_{s,o}$ for zero crossings which do not. This is expressed as the metric $r_{A,o}$:

$$r_{A,o} = \frac{\frac{1}{N_1}\sum_{i=1}^{N_1} A_{s,o}(m,n)}{\frac{1}{N_2}\sum_{i=1}^{N_2} A_{s,o}(j,k)}, \tag{5.40}$$

$$\forall(m,n) : D'_{TH,o}(m,n) = 1, \quad \forall(j,k) : D'_{TH,o}(j,k) = 0, ZC_{s,o}(j,k) = 1$$

where (n,m) are the indices of the N_1 zero crossings determined to coincide with the ideal edges, as defined by $D'_{TH,o}$. Here (j,k) are the indices of the N_2 zero crossings which do

Table 5.8: Discrimination measure $r_{A,o}$ for Gabor zero crossing interscale agreement, for the zero crossings of the simulated ultrasound of Fig. 5.25(a). Values are presented for various frequency and bandwidth values (defined by λ and σ/ω_0, respectively), and are averaged over orientations. λ is the wavelength of the first scale.

	σ/ω_0					
	0.40	0.47	0.54	0.61	0.68	0.75
$\lambda = 50.0$ ppc	1.733	1.826	1.776	1.920	2.134	2.645
$\lambda = 60.0$ ppc	2.376	2.113	2.230	2.314	2.350	2.534
$\lambda = 70.0$ ppc	3.425	2.677	2.498	2.673	2.653	2.388
$\lambda = 80.0$ ppc	3.964	3.558	2.824	2.411	2.968	2.252
$\lambda = 90.0$ ppc	4.657	4.721	2.866	2.904	3.313	2.420
$\lambda = 100.0$ ppc	4.758	5.316	3.901	2.765	3.107	2.689

not coincide with the edge pixels. Values of $r_{A,o} > 1$ indicate that A_o is on average higher for zero crossings which coincide with image edges than for the zero crossings which do not. Higher values of $r_{A,o}$ provide an indication of the strength of the discrimination between these zero crossings possible using the interscale agreement measure A_o.

Table 5.8 displays discrimination values for the interscale agreement of the zero crossings of the simulated image of Fig. 5.25(a). These are seen to increase with wavelength, indicating that the agreement between scales for zero crossings coinciding with image edges is greater for lower frequencies, relative to the agreement of zero crossings which do not. For lower frequencies, wider bandwidth values are observed to produce larger $r_{A,o}$.

Table 5.9 displays discrimination values for the interscale agreement of the zero crossings of the clinical image of Fig. 5.25(b). Similar to the discrimination properties of the simulated image above, the $r_{A,o}$ value is higher for wider bandwidths, although the effect is smaller, and confined to a region of wavelengths in the middle of the range considered.

A quantification of the edge discrimination properties of the interscale agreement measure A_o is presented here. In case of both simulated and clinical images, the interscale agreement measure introduced here has achieved values of discrimination metric $r_{A,o} > 1$. This indicates that it has values which are on average higher for zero crossings which coincide with image edges, than for those which do not. As noted in Section 5.4.2.1 above however,

Table 5.9: Discrimination measure $r_{A,o}$ for Gabor zero crossing interscale agreement, for the zero crossings of the clinical ultrasound of Fig. 5.25(b). Values are presented for various frequency and bandwidth values (defined by λ and σ/ω_0, respectively), and are averaged over orientations. λ is the wavelength of the first scale.

	σ/ω_0					
	0.40	0.47	0.54	0.61	0.68	0.75
$\lambda = 10.0$ ppc	2.218	2.121	2.128	2.135	2.114	2.117
$\lambda = 18.0$ ppc	2.292	2.179	2.335	2.433	2.176	2.152
$\lambda = 26.0$ ppc	2.236	2.312	2.294	2.229	2.243	2.090
$\lambda = 34.0$ ppc	2.125	2.196	2.419	2.231	2.074	2.033
$\lambda = 42.0$ ppc	2.027	2.145	2.254	2.435	2.052	2.017
$\lambda = 50.0$ ppc	2.043	2.048	2.115	2.321	2.258	2.220

not all high agreement values are located coincident to image edges. The next section details a method of combining this interscale agreement measure with the orientational spacing measure $S_{s,o}$ presented in Section 5.4.1.1, to produce a combined measure of zero crossing edge coincidence.

5.4.3 Zero Crossing Edge Detection

In sections 5.4.1 and 5.4.2 above, metrics quantifying the orientational spacing and interscale agreement of Gabor zero crossings are presented. As detailed in these sections, both of these metrics have on average higher values for zero crossings which coincide with image edges than for those which do not. In the case of both these measures however, spurious large values are observed in regions which do not correspond to image edges. It is also noted that these spurious large values do not occur in the same image locations is these two measurements, while the large values corresponding to image edges do. This section presents a method of combining both of these measures to form an edge detector based on Gabor zero crossings. This edge detector is shown to be capable of discriminating image edges, and does not suffer from the spurious large values which affect its components.

In all cases mentioned below, images are decomposed using a Gabor filter bank of two scales. The finer scale is the scale of interest, and orientational spacing measure $S_{s,o}$ is calculated for this scale alone. The zero crossings of the coarser scale are used only to compute the interscale agreement measure $A_{s,o}$. Thus both $S_{s,o}$ and $A_{s,o}$ consist of a single scale. Both of these are normalised values, having a maximum of unity for the highest value over all orientations.

The edge detector is computed by finding the maximum of both $S_{s,o}$ and $A_{s,o}$ over orientations at each image location. These maxima images coincide at image edge locations, and so the composite edge detector can be found by taking their sum:

$$E_s = \max_o S_{s,o} + \max_o A_{s,o} \tag{5.41}$$

Fig. 5.36 displays the edge detector which results from the simulated and clinical images used as examples throughout this section. Fig. 5.36(a) uses the zero crossings calculated from Gabor decomposition of the simulated image of 5.25(a) (displayed in Fig. 5.26), the orientational spacing values shown in Fig. 5.29, and the interscale agreement values displayed in Fig. 5.34. Fig. 5.36(b) uses the zero crossings calculated from Gabor decomposition of the clinical image of 5.25(b) (displayed in Fig. 5.27), the orientational spacing values shown in Fig. 5.30, and the interscale agreement values displayed in Fig. 5.35.

In both cases the edge detection metric is seen to have high values for regions close to the underlying image edges. As with other real valued functions for edge detection a thresholding step is required to produce a binary edge map. A method of testing the effectiveness of the presented edge detector for binary classification is presented in Section 6.2.2, in which it is compared to the competing edge detection approaches detailed in Section 5.2.3. The results of this performance evaluation are presented in Section 6.3.2. This section has presented the formulation of an edge detector based on the orientational spacing and interscale agreement measures proposed in previous sections. Examples of the output of this edge detector are presented. The next section details how the proposed edge detector may be used as the edge detection step in an anisotropic diffusion process for the suppression of speckle noise.

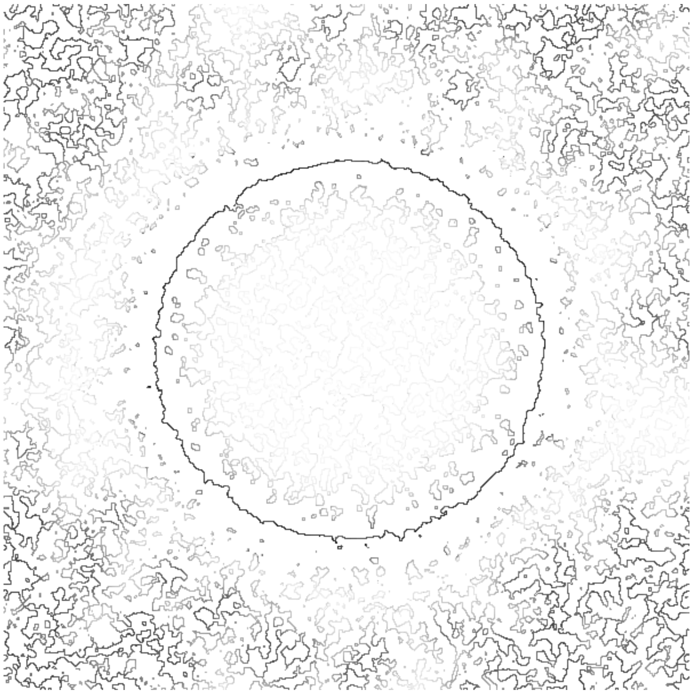

(a)

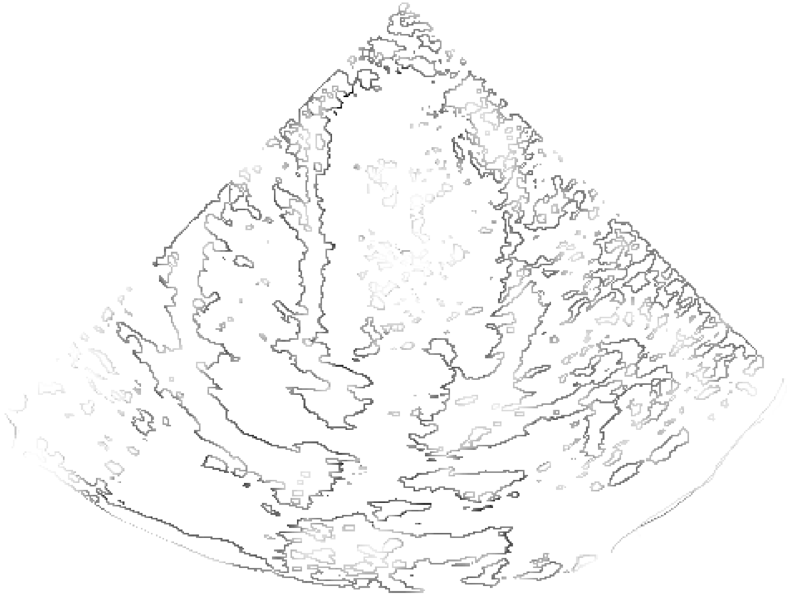

(b)

Figure 5.36: Examples of the final Gabor zero crossing edge detector. (a) For the simulated image of Fig. 5.25(a). (b) For the clinical image of Fig. 5.25(b). Darker regions correspond to higher values.

5.5 Speckle Removal with ZC Edge Detector and Anisotropic Diffusion

This section details a method for incorporating the zero crossing edge detector proposed in Section 5.4 into a tensor diffusion scheme for the suppression of ultrasonic speckle. As detailed previously in Section 2.4.1, noise reduction by anisotropic diffusion proceeds by iteratively smoothing an image. An edge map, also referred to as a diffusivity function, is used to restrict smoothing across image edges. This edge detector is typically calculated iteratively in tandem with the image smoothing step; i.e. after each iteration the edge map is recalculated. In the SRAD diffusion filter [3] of Section 2.4.1.2, an iterative edge detection based on the coefficient of variation is proposed (see Section 5.2.3.1). This edge detection scheme is also used in the DPAD and OSRAD diffusion methods for anisotropic diffusion in speckled imagery. In all of these methods, the edge maps produced early in the diffusion process contain a large amount of false positives from the fine scaled image speckle early in the diffusion process. As diffusion proceeds the image becomes more smoothed, and the edge maps reflect the image structure more accurately. In the coefficient of variation (CV) edge detector of [3, 99], the edge map is obtained after the diffusion process has completed for this reason.

Each iteration of an anisotropic diffusion process involves two steps: the calculation of a diffusivity edge map, $c(x, y)$, and finding a solution to the partial differential equation (PDE). As detailed in Section 2.4.1.7, the solution to anisotropic diffusion PDEs can be performed in a very efficient manner. Computing the edge detector at each diffusion iteration can thus contribute significantly to the processing requirements of these methods. The use of a suitably accurate edge map computed prior to the diffusion process is considered here. With this approach the diffusivity edge map is required to be calculated only once, in contrast to other diffusion methods previously proposed in which the edge map is recalculated upon each iteration. Clearly this approach could potentially reduce the computational effort required to perform diffusion. The zero crossing edge map $E(x, y)$ proposed in Section 5.4 is considered for this task. As the increased local adaptivity of the tensor based OSRAD filter [24] is found to lead to favourable speckle suppression properties in Chapter 4, a tensor based anisotropic diffusion scheme developed here.

The normalised edge map $E(x, y)$ is thresholded and converted to a diffusivity function by:

$$c_{zc}(x,y) = |E_{th}(x,y) - 1| \tag{5.42}$$

where E_{th} is the thresholded edge map. Locations where $c_{zc}(x,y)$ are high will experience greater diffusion. The operation $|E_{th}(x,y) - 1|$ produces an output in which equals zero in edge locations (and so no diffusion is performed) and has a value of unity in all other areas.

This simple approach is noted experimentally to yield instabilities in some cases. A similar case is observed in some situations with the original Perona and Malik formulation, discussed in Section 2.4.1.1 above. As discussed in the work of Kichenassamy [177], the original Perona and Malick PDE is ill-posed, i.e. a unique solution depending on the problem information (coefficients, boundary conditions, and initial conditions) does not exist. Nonetheless, stable discretisations have been proposed and the method has been widely applied to many types of image with success. Instabilities can arise however, and are caused by irregularities of the Perona and Malik diffusion term:

$$\begin{aligned} c(x,y;t) = &\quad f(|\nabla I(x,y;t)|), \\ f(x) = &\quad e^{-\left(\frac{x}{k}\right)^2} \quad \text{or} \quad \left(1 + \left(\frac{x}{k}\right)^2\right)^{-1} \end{aligned} \tag{5.43}$$

where $\nabla I(x,y;t)$ is the gradient of the image under diffusion at time t, and k is a threshold. Kichenassamy shows that as $c(.)$ above is a function of the evolving image $I(x,y;t)$, and not the observed image $I(x,y;0)$, a generalised solution to the PDE does not exist. Catté *et al.* [58] define the problem by noting that if the function $xf(x)$ is nonincreasing, then the PDE becomes ill-posed. These authors address this as described in Section 2.4.1.1 by computing $c(.)$ on a smoothed version of the image gradient.

In the diffusion process proposed here, the diffusion function c_{zc} is a function of the zero crossing edge estimate E, i.e. $c = f(E(x,y))$. Thresholding of the edge estimate is used, and (5.42) can be rewritten as:

$$\begin{aligned} c_{zc}(x,y) &= f(E(x,y)) \\ f(x) &= \begin{cases} 0, & \text{if} \quad x \geq T, \\ 1, & \text{otherwise.} \end{cases} \end{aligned} \tag{5.44}$$

where T is the threshold used to locate edges. In this case it is trivial to show that $xf(x)$ is nondecreasing, and so the observed instabilities are not due to ill-posedness of the PDE.

It is investigated experimentally if the discontinuity around $E = T$ may be responsible for the observed numerical instability. This leads to sharp transitions in the level of diffusion between pixels denoted as edges and their immediate neighbours. The diffusion function is modified by smoothing:

$$c_{zc}(x,y) = |E_{th}(x,y) * G_\sigma - max\{E_{th}(x,y) * G_\sigma\}| \tag{5.45}$$

Similar to the regularisation performed by Catté *et al.*, this G_σ is a Gaussian kernel of variance σ^2. This modification is seen to prevent any instances of numerical instability in the extensive set of tests performed in this thesis. The formulation of (5.45) leads to two specific behaviours: firstly, the $c_{zc}(x,y) = 0$ when $E_{th}(x,y) = 1$, so diffusion is completely inhibited across edges. Conversely, far from edges (i.e. beyond the extent of the neighbourhood used in discretisation) $c_{zc}(x,y) = 1$, leading to isotropic diffusion. Secondly $c_{zc}(x,y)$ decreases from unity as an edge is approached, thus smoothing will be inhibited to a degree on the regions directly adjacent to edges. As with the regularisation Catté *et al.*, σ is automatically chosen based on the diffusion time as $\sigma = \tau N$, where τ and N are respectively the discrete timestep and the current iteration.

The diffusion edge detection function c_{zc} presented here could be applied in a scalar valued approach, as with the PMAD, SRAD and DPAD methods. Given the high performance displayed by the OSRAD tensor diffusion method in the evaluation of Chapter 4 however, a similar tensor valued scheme is pursued. As with the CED, NCD and OSRAD filters, the proposed method is defined by the PDE:

$$\frac{\delta I(x,y;t)}{\delta t} = \nabla.(D\nabla I(x,y;t)) \tag{5.46}$$

where D is the diffusion matrix, based on the structure tensor, T. The structure tensor is a 2×2 array at each image position, and it computation is discussed in Section 2.4.1.4 above. Specifically, D is constructed with the same eigenvectors as T, which define the direction of local gradient and curvature. The corresponding eigenvalues (λ_1, λ_2) determine the strength of diffusion in these directions. Here they are chosen by the zero crossing diffusivity function to be:

$$\begin{aligned} \lambda_1 &= c_1 \\ \lambda_2 &= c_{zc}(x,y) \end{aligned} \tag{5.47}$$

similar to the CED method, c_1 is a small constant $0 < c_1 \ll 1$. Given the basis of this anisotropic diffusion method on image edges estimated from Gabor zero crossings, it is referred to as ZCAD where abbreviation is necessary.

5.6 Conclusions

This chapter has detailed the development of a novel edge detector for operation in the presence of ultrasound speckle noise. Three main contributions are made in the process of edge detector development:

- A method of estimating the parameters of the stable and Gaussian combined model for Gabor coefficients is proposed;
- A method of edge detection using the properties of Gabor zero crossing is presented;
- The inclusion of an accurate static edge map in a tensor diffusion process for speckle reduction is proposed.

The edge detection method operates by performing a Gabor transform of the speckled image. The effect of speckle is removed from the Gabor coefficients by the statistical estimation of the speckle free values. This includes the use of a suitable model for these coefficients. A combined SαS and Gaussian model is suggested for this purpose, and the accuracy of this model is verified by fitting to maximally speckle free simulated images, which are generated using Burckhardt's maximum writing technique. The proposed method for simultaneous estimation of all parameters of this combined model is then utilised, to perform MAP estimation of the speckle free Gabor coefficients.

The location of edges in these images is based on two observed properties for the zero crossings of real Gabor coefficients in close proximity to known image edges. Methods of quantifying these properties of orientational spacing and interscale agreement are presented, and are demonstrated to have higher values in locations corresponding to image edges. This ability to discriminate zero crossing which coincide with image edges is investigated for a range of Gabor filter frequency and bandwidth values. It is concluded that both of these properties are capable of discriminating image edges, however each contain spurious values. These would lead to false edges if the zero crossing properties were used individually for edge detection. The edge detection measure proposed in this chapter instead combines these two quantities, so that the effect of the spurious values in each are much reduced. A method of using the proposed edge detection method for the reduction of speckle is presented, in which the edge detector is included in a tensor valued anisotropic diffusion scheme.

In summary, this chapter has presented the development of a novel edge detection method for speckled ultrasound images, and its utilisation for speckle reduction. The next chapter presents an evaluation of the performance of this technique.

Chapter 6

Performance Evaluation

6.1 Introduction

This chapter presents the performance evaluation of the technique presented in Chapter 5. There are three steps proposed in the technique of Chapter 5 above, and each of these are tested in turn:

1. The method of parameter estimation presented in Section 5.3.2 is tested to determine the achievable level of accuracy. Monte Carlo simulation is performed, and errors in parameter estimates are compared to other approaches from the literature.

2. The zero crossing edge detection method proposed in Section 5.4 is compared to the other edge detection methods detailed in Section 5.2.3. The performance of these edge detectors is evaluated by application to simulated, phantom, and clinical ultrasound images.

3. The use of the proposed edge detection method for speckle reduction as part of a tensor diffusion process is evaluated by applying the evaluation framework proposed in Chapter 4. The proposed ZCAD filter is applied to the same simulated and clinical images as the filters compared in Chapter 4, and image quality metrics are computed. An analysis of the computational requirement of the proposed ZCAD method is performed. For all tests, results are compared with the speckle reduction methods considered in Chapter 4.

Section 6.2 details the performance evaluation tests described above. The results of these tests, for each of the three proposed steps are presented in Section 6.3. Conclusions from these results are drawn in Section 6.4.

6.2 Test Procedures

6.2.1 Stable Model Parameter Estimation

The proposed method of parameter estimation presented in Section 5.3.2 simultaneously estimates the three parameters (α, γ, σ) of the model (5.21) for speckled Gabor coefficients. Monte Carlo simulation is performed to determine the accuracy of this method. The parameters (α, γ, σ) within the ranges $(.5 \leq \alpha \leq 2; .5 \leq \gamma \leq 2; .1 \leq \sigma \leq 2)$. Eight values of each parameter, linearly spaced within these ranges, are used here. This leads to a total of 512 different parameter combinations.

Random vectors of 10^4 data points are produced for each of the 512 parameter combinations. Each random vector follows the distribution of model (5.21), i.e. the product of a SαS and a Gaussian characteristic function. By the duality between characteristic functions and probability density functions (PDFs), the combined SαS-Gaussian density is the result of convolution between SαS and Gaussian PDFs:

$$\phi_y(t) = \phi_s(t)\phi_n(t), \quad \Rightarrow \quad p_y(y) = p_s(s) * p_s(n) \tag{6.1}$$

where $\phi_x(t)$ and $p_x(x)$ denote the characteristic function and PDF of variable x respectively, t is the characteristic function argument and $*$ represents convolution. Subscripts (s, n, y) respectively indicate quantities which belong to the SαS, Gaussian, and combined distributions. Thus for each parameter combination, a suitable random variable with density $p_y(y)$ is generated.

For each parameter combination, 50 individual random vectors of 10^4 points are generated, and the estimated parameters $(\hat{\alpha}, \hat{\gamma}, \hat{\sigma})$ are computed using the proposed method in each case. Errors are then calculated for each parameter combination as the root mean squared error (RMSE) between the known true parameter value and the set of 50 estimates. These errors are denoted $(\xi(\hat{\alpha}), \xi(\hat{\gamma}), \xi(\hat{\sigma}))$. Results of this parameter estimation test are reported in Section 6.3 below, and are compared with other approaches from the literature.

6.2.2 Edge Detection

To evaluate the performance of the edge detection methods, they are applied to a number of simulated, phantom, and clinical ultrasound images. The images used in evaluation from each of these three types are discussed in this section. Details of how the edge detectors methods are applied are also presented.

6.2.2.1 Simulated Images

The set of simulated ultrasound images are generated using the technique described in Section 4.2. The generation of these simulated images uses templates to define regions of differing scatter density in each image. These templates thus define the shape of the structure in each image, and ideal edge maps can be found as the borders between template regions. While the shape and density of scatterers within each image region are predetermined, the magnitude, phase and positioning of the scattering elements are randomly selected. Thus a number of similar images can be produced for each template, each having the same basic image structure, but with differing speckle properties.

Eight templates are used to test the performance of the edge detection methods (both the proposed zero crossing approach, and the methods from the literature detailed in Section 5.2.3). Four of these templates define images of simple shapes, with two distinct regions of differing scatterer density. Examples from each of these templates are displayed in Fig. 6.1, and their corresponding ideal edge maps are shown in Fig. 6.2. These images contain shapes which are not seen in the simulated images of clinical shapes below, such as non-smooth curves and sharp corners. The densities of scattering elements in the regions of these image are 10% and 40%, respectively for the darker and lighter areas.

The remaining four simulation templates define images containing simplified versions of shapes commonly seen in ultrasonography. These contain three different regions of scatterers, with densities of $70\%, 20\%$ and 10%. Examples from each of these 'clinical' templates are shown in Fig. 6.3, with their corresponding ideal edge maps shown in Fig. 6.4. In these images brighter regions are those with higher scatterer density.

As in the simulations of Chapter 4, and the work of [5], the phase values of the scattering elements have uniformly random values from $0 \rightarrow 2\pi$, and Gaussian distributed magnitudes of unity mean and $\sigma = 0.1$. For each simulation template, ten individual images with differing speckle content are produced, giving a total data set of eighty simulated images.

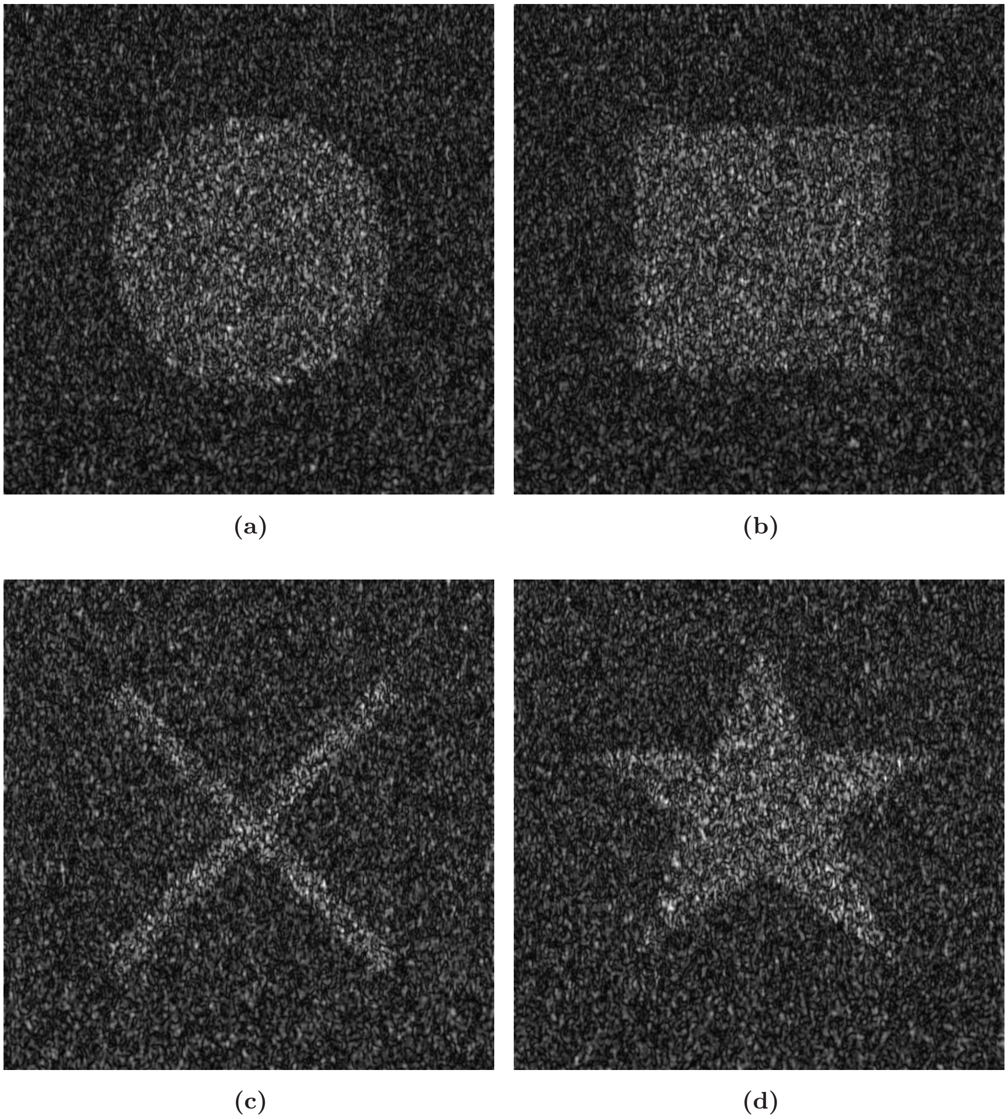

Figure 6.1: Examples of simulated ultrasound images from each of the simulated shape image templates. (a) Circle, (b) Square, (c) Cross, (d) Star.

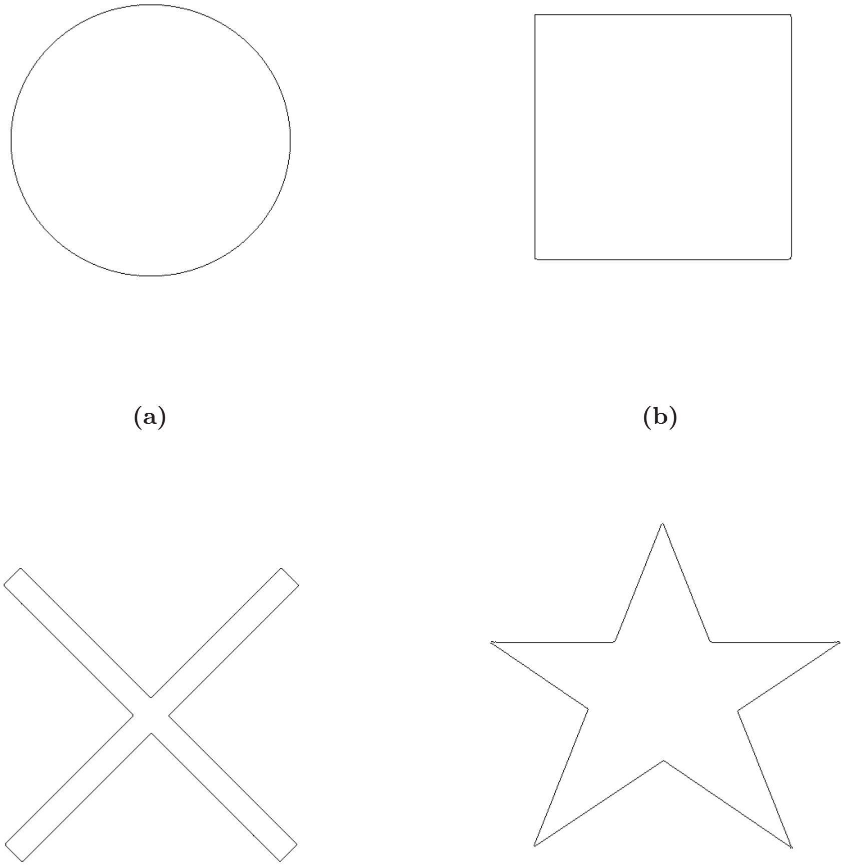

(a) (b) (c) (d)

Figure 6.2: Ideal edge maps for the simulated shape images of Fig. 6.1. (a) Circle, (b) Square, (c) Cross, (d) Star.

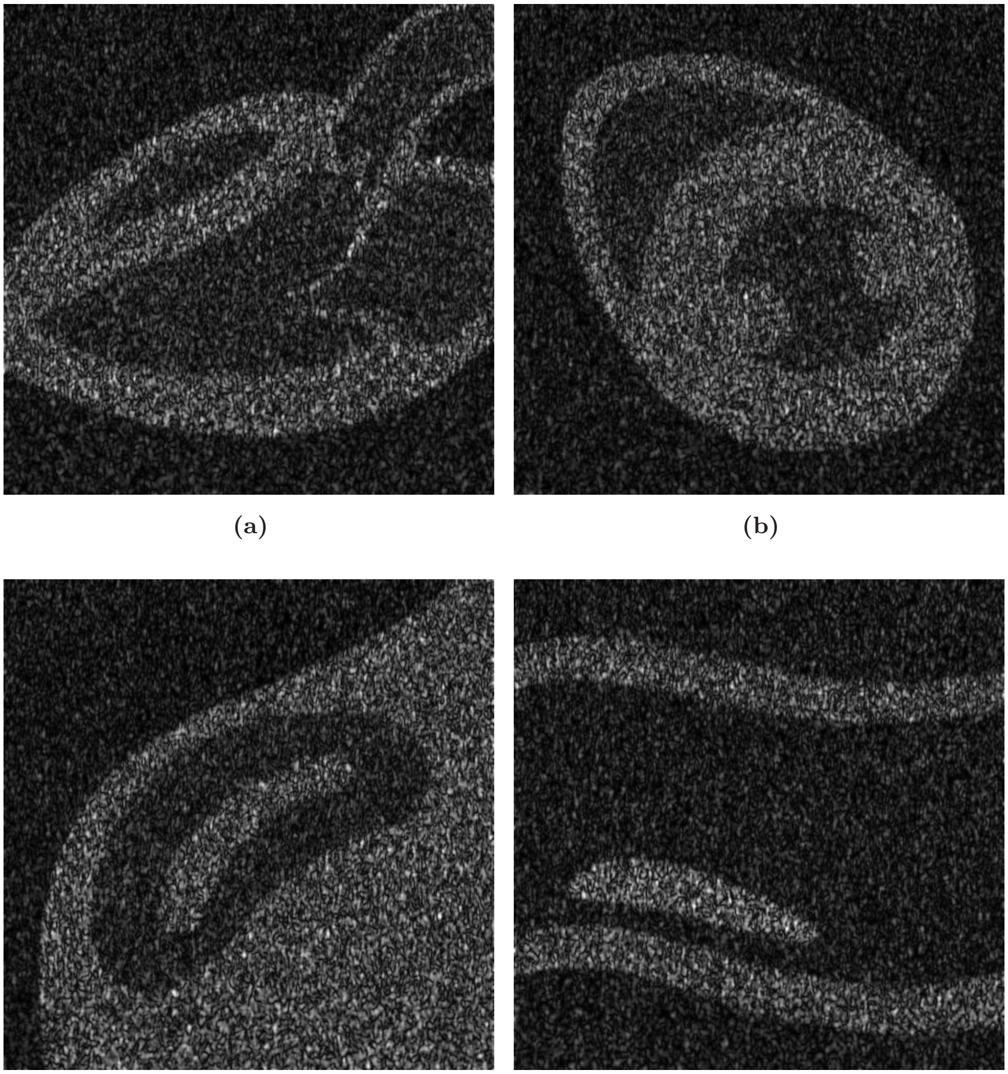

(a) **(b)** **(c)** **(d)**

Figure 6.3: Examples of simulated ultrasound images from each of the clinical ultrasound image templates. (a) Long axis echocardiogram, (b) Short axis echocardiogram, (c) Renal Ultrasound, (c) Carotid Ultrasound.

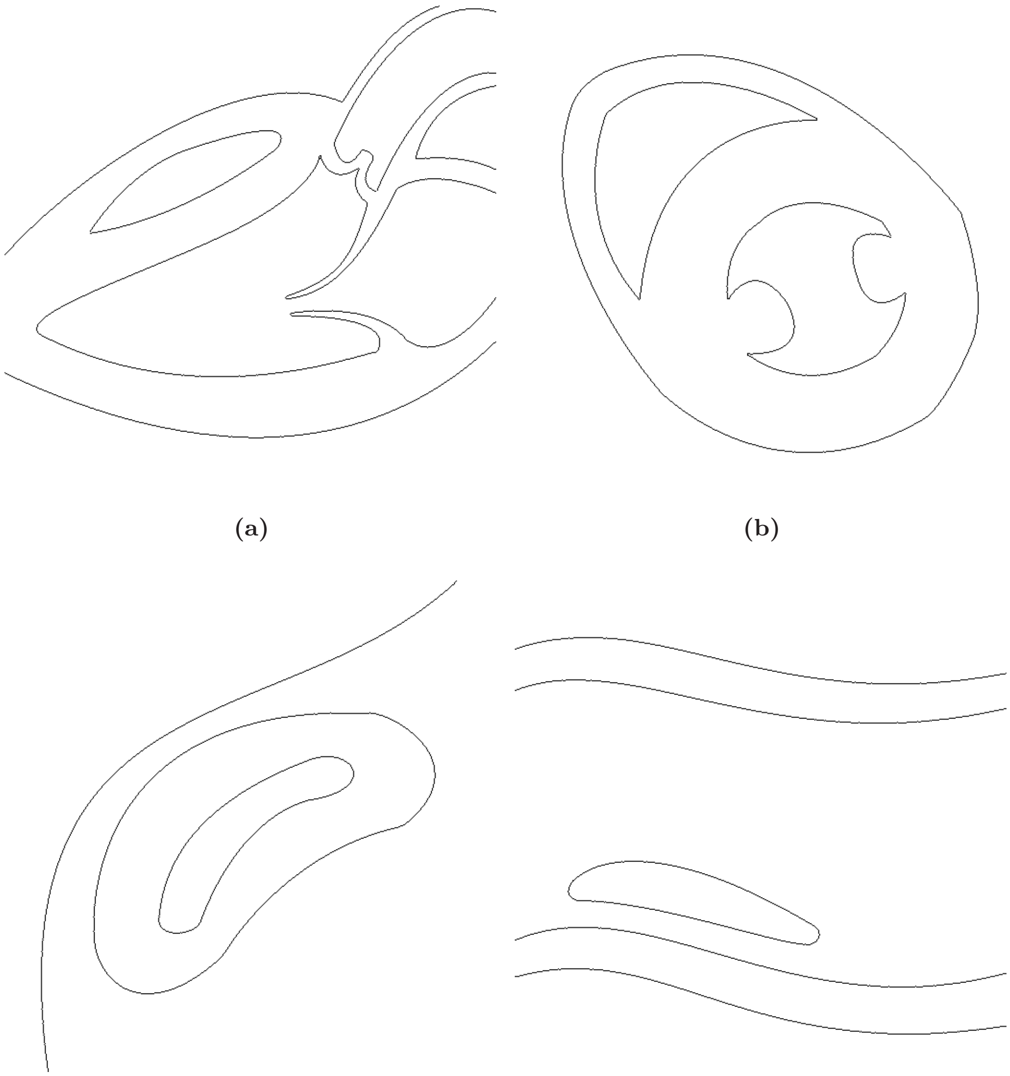

(a) **(b)**

(c) **(d)**

Figure 6.4: Ideal edge maps for the simulated clinical ultrasound images of Fig. 6.3. (a) Long axis echocardiogram, (b) Short axis echocardiogram, (c) Renal Ultrasound, (c) Carotid Ultrasound.

6.2.2.2 Phantom Images

A set of images from a tissue mimicking ultrasound phantom are used to evaluate edge detection performance. The phantom used was constructed for the purposes of these tests using the technique developed by Bude and Adler [178]. Ultrasound phantoms can be grouped into those which attempt to mimic the acoustic properties of tissue, and those which attempt to mimic the appearance of ultrasonographic images. The latter type (referred to as 'biopsy' phantoms in [178]) are appealing in many cases due to their ease of construction. In the present work, the aim of employing a phantom is to generate ultrasound images similar in nature to those obtained clinically. These must contain a known underlying structure to allow evaluation of edge detection performance. As is observed in the images below, the biopsy phantom constructed using the method of [178] is capable of achieving these goals.

A substance which mimics the echogenicity was produced, using a commercially available dried powder of a fibrous plant material (*psyllium hydrophilic mucilloid* fibre). This provided the echogenic scattering material, the particles of which were suspended in a gelatine mixture. Gelatine is sonolucent, so all reflections were provided by the psyllium powder. Detailed mixing instructions, including ingredient proportions, can be found in [178]. This mixture produced a speckled pattern when insonified, and regions of differing echogenicity are introduced to provide boundaries for edge detection. Dark regions are produced here by inserting the water filled tips of common latex gloves.

It is concluded in [178] (the authors of which are both physicians) that this technique produces images similar to those observed clinically. The method of construction has the advantage of requiring less precision to produce realistic phantoms than other methods. It was noted experimentally here that images similar to those reported by Bude and Adler can be generated readily using this method. A set of thirteen of these phantom images were used for edge detection evaluation here, which are displayed in Fig. 6.5. Insonification was performed using a General Electric LOGIQe ultrasound system (GE Healthcare, Piscataway, NJ, USA), with a linear array transducer. The gain and focal length settings were varied to generate the differing appearance of these phantom images. No speckle reduction techniques were used in acquisition. It can be seen this set contains images of varying quality, with some images suffering from missing or obscured borders. These were deliberately included to test edge detection in a range of image quality situations.

As the shapes in these images result from known physical structure, the desired edges can be manually delineated for quantitative evaluation of edge detection performance. The manually defined ideal edge maps for the phantom images of Fig. 6.5 are displayed in Fig. 6.6.

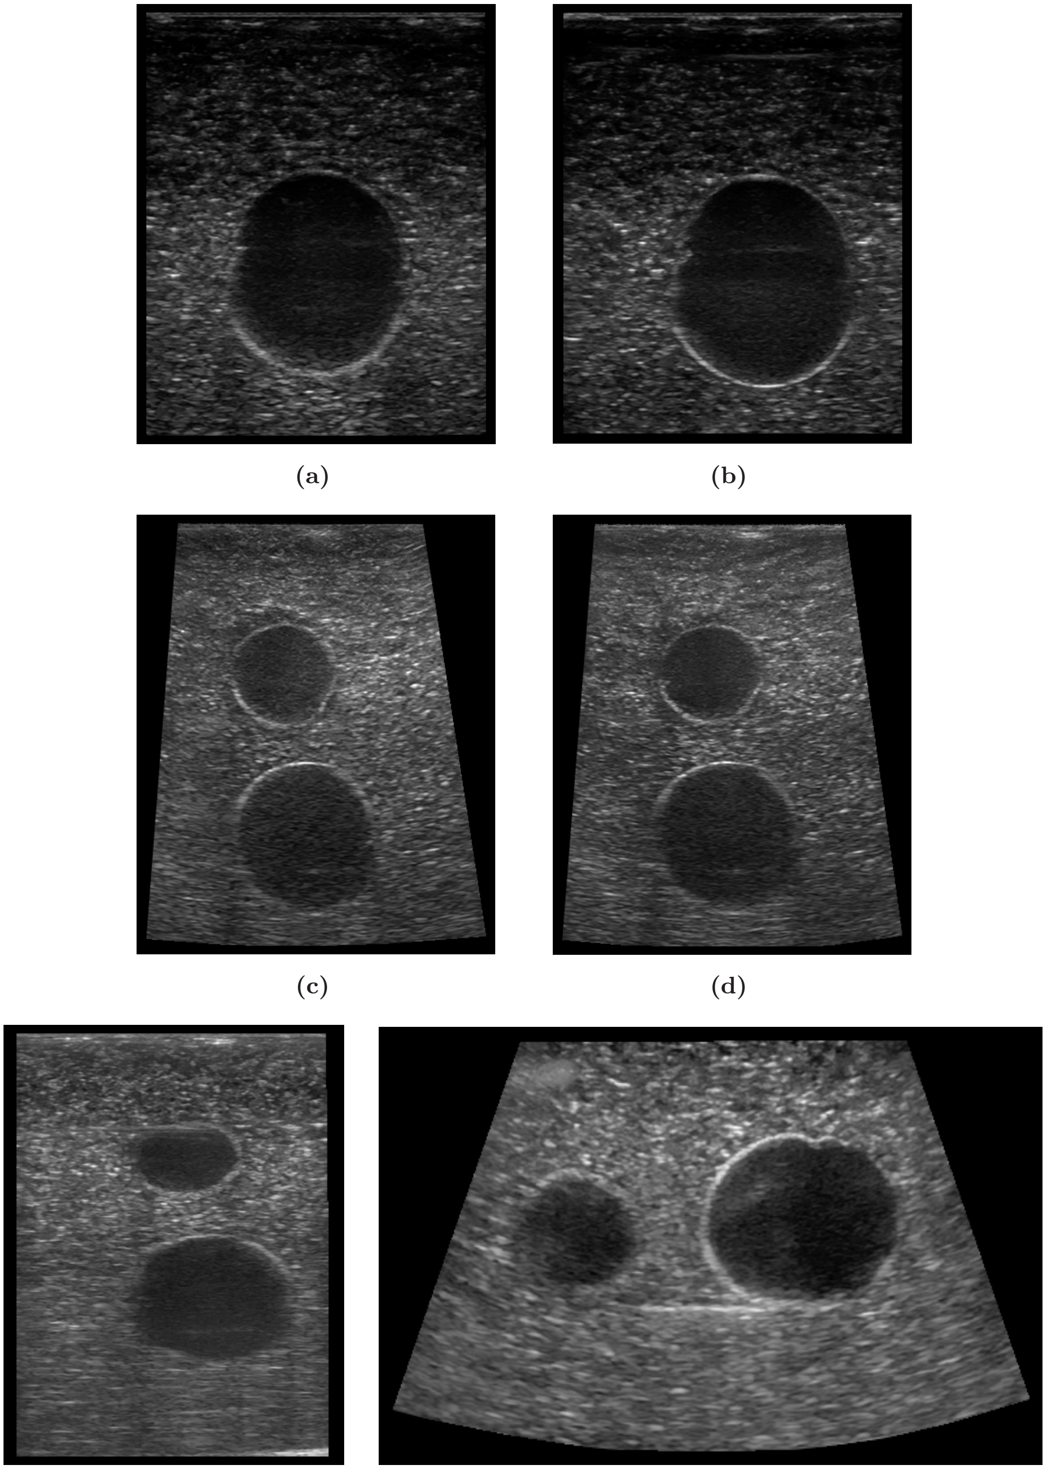

(a) **(b)**

(c) **(d)**

(e) **(f)**

Figure 6.5: Phantom ultrasound images used in evaluating edge detector performance. (a)–(f) Phantom Images 1–6, respectively.

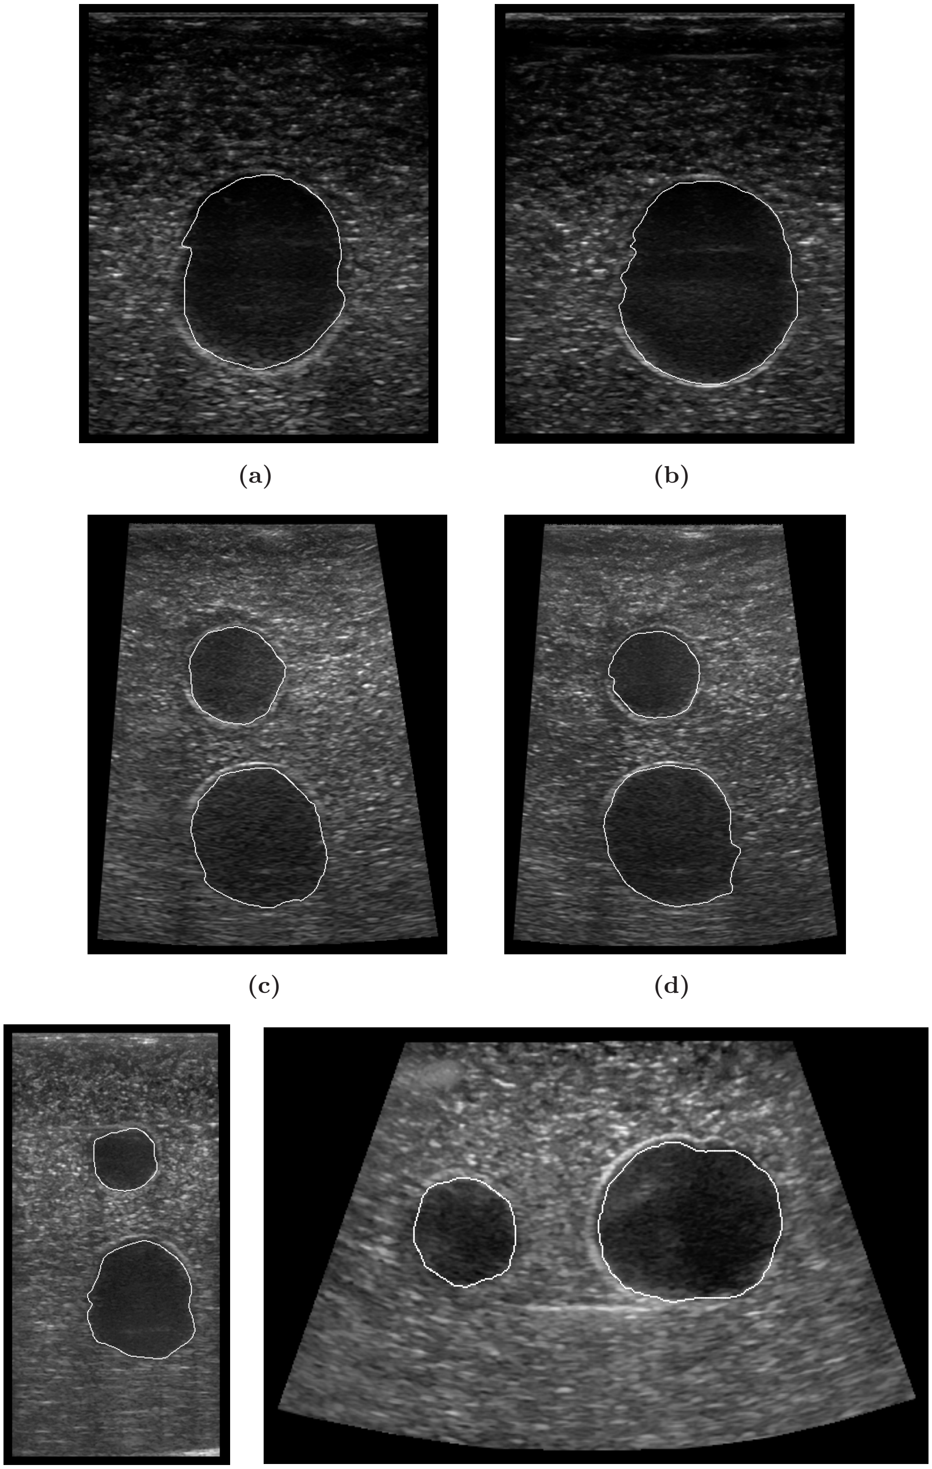

(a) (b)

(c) (d)

(e)

Figure 6.6: Ideal edge maps for the phantom ultrasound images of Fig. 6.5. (a)–(f) Phantom Images 1–6, respectively.

6.2.2.3 Clinical Images

The edge detection methods were also applied to a set of clinical echocardiographic images to demonstrate their performance in practical applications. A set of twelve of these images were used here, shown in Figs. 6.7-6.8 They were scanned using a General Electric Vivid 7 Series scanner (GE Healthcare, Piscataway, NJ, USA). These images were subject to log compression, however no image optimization or speckle reduction (such as compounding from multiple angles) was performed.

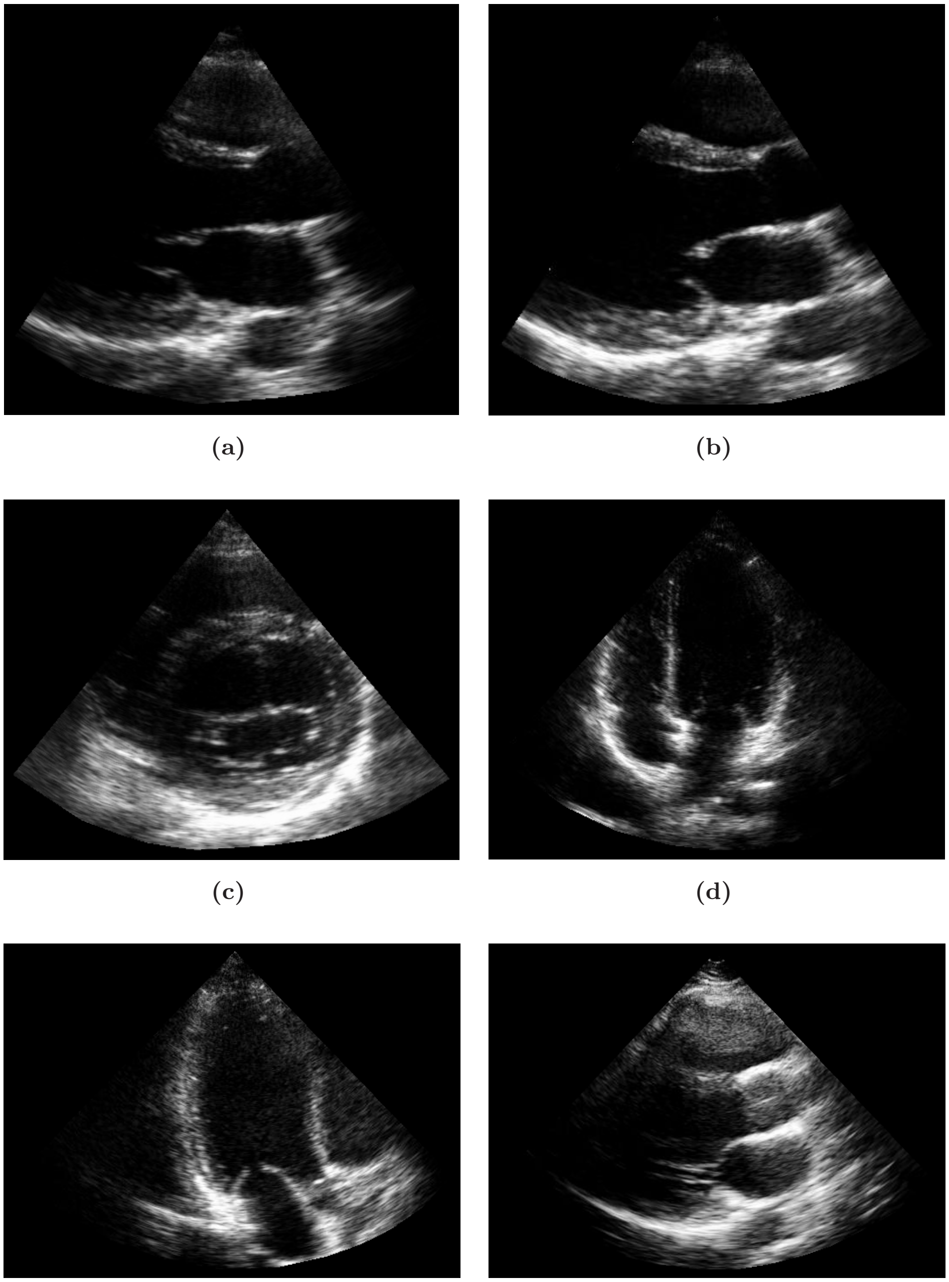

(a) **(b)**

(c) **(d)**

(e) **(f)**

Figure 6.7: Clinical ultrasound images used in evaluating edge detector performance. (a)–(f) Clinical images 1–6, respectively.

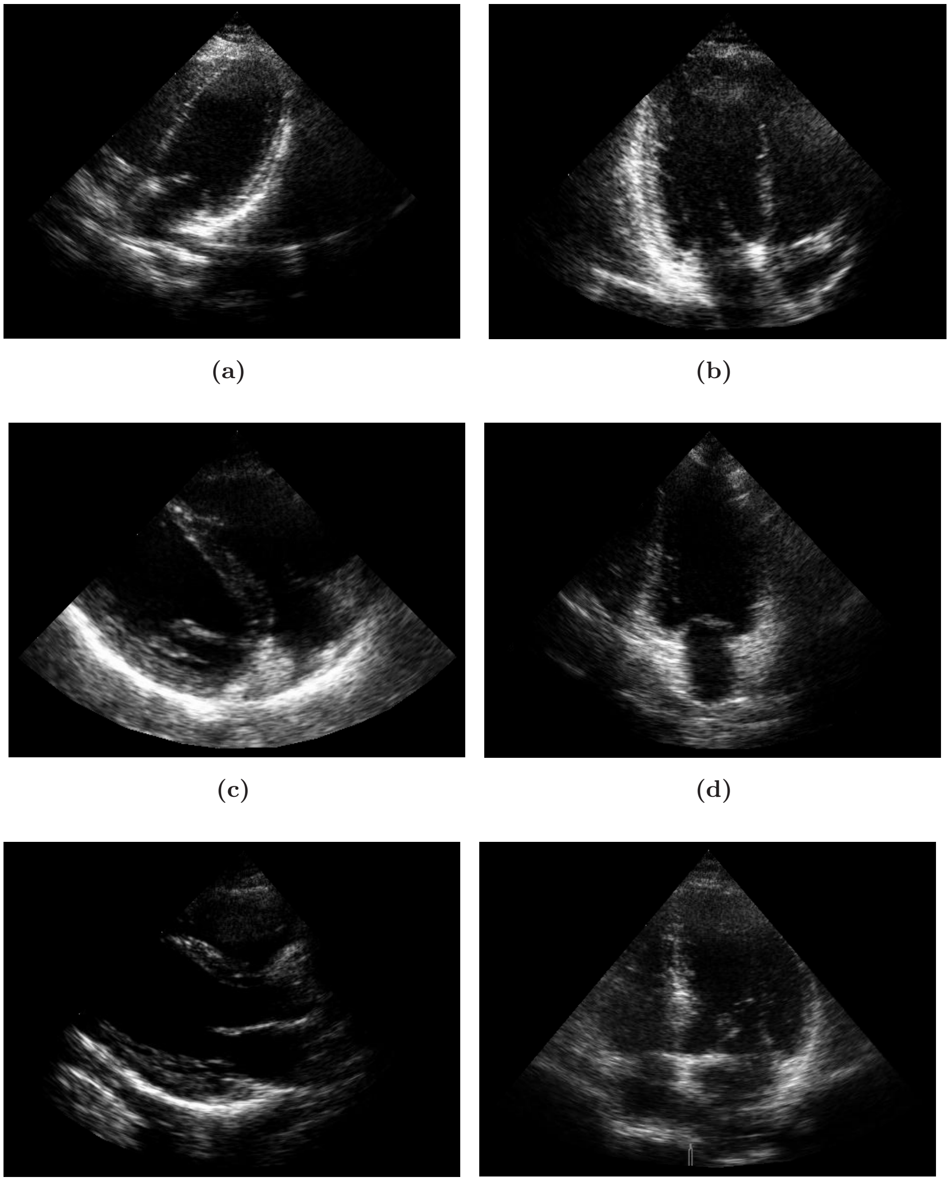

(a) (b)

(c) (d)

(e) (f)

Figure 6.8: Clinical ultrasound images used in evaluating edge detector performance. (a)–(f) Clinical images 7–12, respectively.

6.2.2.4 Edge Detector Performance Evaluation

The edge detector proposed in Section 5.4 is compared to the other edge detection methods detailed in Section 5.2.3 (CV, RoA, PC, FA), by application to the simulated, phantom and clinical images. In the case of the simulated and phantom image sets, quantitative evaluation by comparison with ideal edge maps is possible. This is accomplished using the figure of merit (FoM) of Pratt [59]. This measure is described previously in Section 3.3. For the set of clinical images however, no such ideal edge maps exist. Evaluation of edge detection performance is limited to qualitative visual analysis in this case.

All of the edge detection methods have parameters which can take on a range of values, and thresholding must be performed to produce binary edge maps. The choice of parameter and threshold values is clearly an important contributing factor to the accuracy of edge detection. The edge detector evaluation undertaken here aims to avoid the additional degree(s) of freedom introduced by the selection of parameter and threshold values, and to determine the best possible performance of each method for the images considered. This is accomplished by performing edge detection with each method repeatedly, using a range of parameter and threshold values. The FoM is measured for each repetition, and the optimal result is taken, producing the best possible edge detection performance for each method under this criterion. The parameters and thresholds which produced the optimal FoM for the phantom images are used in performing edge detection on the clinical images.

The CV method is applied as in [99], as part of an iterative diffusion scheme. Diffusion is performed until the solution becomes stable. The extension of [60] is incorporated, allowing the use of varying window sizes for local statistics calculations. Similar to the test employed in [99], a number of thresholds are applied to generate binary edge maps, and the FoM measured in each case. Here fifty threshold values are applied, linearly spaced between .75 and .995 of the maximum value. Window sizes from 3×3 to 33×33 are tested. The RoA method is applied with multiple window sizes from 9×9 to 45×45. The standard deviation of the LoG kernel is chosen so as to fill the current window size, i.e. $\sigma_{LoG} = (N-1)/6$, where N is the window size. Thresholding of the ratio components applied as with the CV method above. The FA method is applied using a single scale log Gabor filter bank with six orientations and a wide bandwidth of three octaves. Multiple values of centre frequency (specified in terms of λ, the period in pixels per cycle) are tested, ranging from 5% to 20% of the minimum support of the image. The noise scaling parameter, k, is varied from 0.1 to

2.5. The PC method is applied using a log Gabor filter bank with three scales and six orientations. These are arranged so as to cover the frequency plane as in Fig. 5.2, with minimal overlap in the half magnitude profiles. The wavelength of the middle scale is chosen as with the FA method above, and the filter bandwidth is two octaves. The proposed ZC method has two parameters to select: the frequency of the Gabor filters, and the threshold values used. As with the parameters of the other edge detection methods, a range of values are used, and the best by the FoM metric is selected. It is found in most cases that the optimal thresholds are approximately 80% of the maximum value of $E_{s,o}(x, y)$, while the optimal frequency for the first scale has values between 5-10% of the minimum image support.

6.2.3 Speckle Reduction

This section details how the performance of the ZCAD speckle removal method described in Section 5.5 above is evaluated. The comprehensive evaluation strategy proposed in Chapter 4 is applied to this filter. This includes the application of the ZCAD filter to the test set of 500 clinical echocardiographic video frames and the simulated image set comprised of 40 images. The quantification of the filters performance in both cases is achieved using the FoM, SSIM, EdgeMSE, CNR and SNR_A image quality metrics. These metrics, as well as the simulated and clinical image sets are detailed in Chapter 4 above. The parameters chosen for the proposed ZCAD method are mostly similar to those used for the OSRAD filter in Chapter 4, i.e. $\sigma = 2, \rho = 2.5$ for the construction of the diffusion tensor. The edge maps are generated using log Gabor filter banks with two scales, the finest of which has a wavelength equal to 10% of the minimum image support. As Gabor filtering can be computationally expensive, only two orientations are used. The thresholds used to generate the binary edge maps are chosen as the 80% point in cumulative histogram of the real valued edge detection function E in all cases. The level of diffusion is chosen as $t = 50$s.

An analysis of the computational requirements of the ZCAD filter is carried out in the same fashion as the analyses of Chapter 4. The filtering approach is analysed by quantifying the number of multiplications or divisions, additions or subtractions, and look up table operations required for each constituent step:

1. *Log Gabor filtering:* A bank of log Gabor filters are designed in the frequency domain. The efficient approach suggested by Kovesi is used here, demonstrated in a tutorial

available on the authors website[1]. Each log Gabor filter, defined by (5.3), consists of the product of an axial and a radial term. These are constructed separately, using a two dimensional frequency axis of the same size as the input image (i.e. N pixels). For a filter bank of N_s scales and N_o orientations, the construction of the filter bank can be performed with $5N+(5N+2)N_s+8NN_o$ multiplications, $N+2NN_o$ additions and $3N+(N+1)N_s+(2N+4)N_o$ look up operations. The application of the filter bank by (5.4) requires that the 2D Fourier transform of the input image I is found. This N point FFT is multiplied with each of the N_sN_o log Gabor frequency responses. An inverse FFT is performed on each of the resulting N_sN_o log Gabor output spectra. As detailed in Section 4.3.4, the bit reversed radix 2 FFT is used, which requires $N\log_2 N$ multiplications and $\frac{3}{2}N\log_2 N+\frac{1}{2}N$ additions for an N point image.

2. *Statistical Model Parameter Estimation:* In order to perform MAP noise removal from the speckled Gabor responses, the parameters (α,γ,σ) for the model of (5.21) are estimated, using the technique proposed in Section 5.3.2. The parameter estimation method begins by computing the quantify $y(t_k)$ from the observed characteristic function (5.22), by $y(t_k)=-\ln|\phi_s(t_k)|^2$. This is a vector of length K. Nonlinear minimisation using the Nelder-Mead method [170, 171] is employed to estimate the three parameters. This is achieved by minimising the error between the observed function $y(t_k)$, and the equivalent computed from the parameter estimates $(\hat{\alpha},\hat{\gamma},\hat{\sigma})$, denoted $\hat{y}(t_k)$:

$$\hat{y}(t_k;\hat{\alpha},\hat{\gamma},\hat{\sigma})=2\hat{\gamma}|t_k|^{\hat{\alpha}}+\hat{\sigma}^2t_k^2,\quad k=1,2,\ldots,K \tag{6.2}$$

The error function used is the sum of squared error (SSE):

$$\xi(\hat{\alpha},\hat{\gamma},\hat{\sigma})=\sum_{k=1}^{K}\left(y(t_k)-\hat{y}(t_k;\hat{\alpha},\hat{\gamma},\hat{\sigma})\right)^2 \tag{6.3}$$

Minimising this function of three variables by the Nelder mead method involved the calculation of $\xi(\hat{\alpha},\hat{\gamma},\hat{\sigma})$ at a number of points, each of which costs $5K+2$ multiplications and $3K-2$ additions, along with K look up operations. Each iteration of the Nelder-Mead method applied to a three-dimensional function operates by computation of that function at five points (four upon the vertices of the three dimensional simplex shape, and one at its centroid). This is followed by two more possible steps, giving a worst case maximum of seven function evaluations at each iteration [171]. The remainder of the algorithm requires far less computation than these function evaluations [122], especially when the function is relatively complex as in this case. It noted experimentally that the algorithm converges in an average of 150 iterations. As recommended by Press *et al.* [122], the algorithm is reapplied multiple times (5 here) to ensure the convergence criteria were not falsely triggered. Thus the minimisation algorithm operates for an average of $Z = 750$ iterations.

---

[1] `http://www.csse.uwa.edu.au/~pk/Research/MatlabFns/PhaseCongruency/Docs/convexpl.html`

3. *MAP Shrinkage:* Once the parameters of the combined SαS and Gaussian model are found, the MAP estimation step is performed. This operates by generating a shrinkage function for each of the $N_s N_o$ log Gabor subbands. The number of elements in the MAP shrinkage function is denoted by N_m. It is computed numerically as an analytic solution does not exist, due to the use of the SαS distribution. A Gaussian and a SαS density function are computed, both of length N_p. The SαS PDF is found by computing an N_p length characteristic function by (5.9), and then taking the FFT. The SαS prior distribution is then multiplied with a shifted version of the Gaussian to find each point in the N_m length MAP shrinkage function. The application of the MAP function can be then performed using a look up table, created to match each input value to its MAP estimate.

4. *Zero Crossing Detection:* The approach taken to find zero crossings in the two dimensional image arrays is to first apply the *sign*(.) function to each Gabor subband:

$$sign(x) = \begin{cases} 1 & \text{if} \quad x > 0 \\ 0 & \text{if} \quad x = 0 \\ -1 & \text{if} \quad x < 0 \end{cases} \tag{6.4}$$

   A zero crossing is deemed to have occurred at a pixel if there is a change of sign in its immediate neighbours in the vertical, horizontal or either diagonal direction. For each of these four pixel pairs, a change of sign is found if the absolute value of their sum is neither zero nor two. Thus the zero crossings of an image with N pixels can be found $4N$ addition operations, plus some comparisons.

5. *Orientational Spacing Calculation:* The computation of orientation spacing measure

$S_{s,o}$ by (5.32) depends on the number of zero crossing pixels in the image, denoted N_z. The image is assumed to be square, so that it contains $\sqrt{N}$ rows and columns. Each of the $\sqrt{N}$ columns will then contain an average of $\frac{N_z}{\sqrt{N}}$ zero crossings, and so the orientation spacing can be calculated using $\frac{N_z^2}{\sqrt{N}}$ multiplications and additions. For the case of a Gabor filter bank with two orientations as considered here, the rotation of the image is accomplished with a transposition operation.

6. *Interscale Agreement Calculation:*

   The calculation of the interscale agreement $A_{s,o}$ by (5.38) is evaluated using the same assumptions as above. This then requires a total of $3N_z^2$ multiplications and $3N_z^2 + N + 1$ additions

7. *Edge Measure Calculation:*

   The zero crossing edge measure E_s by 5.41 uses a number of sorting operations and N additions.

6.3 Results

This section presents the results of the three performance evaluation tests detailed in Section 6.2 above. Firstly, the results of the Monte Carlo test to determine the accuracy of the proposed method of model parameter estimation are presented and compared with those from the literature. Next, the evaluation of the proposed edge detector, compared against a number of other methods by application to simulated, phantom and clinical ultrasound is reported on. Finally, the effectiveness of the proposed edge detector as a part of a tensor anisotropic diffusion scheme for speckle reduction is reported upon.

6.3.1 Model Parameter Estimation

The Monte Carlo test detailed in Section 6.2.1 above, is performed to evaluate the level of accuracy of the parameter estimation proposed in Section 5.3.2. The errors for each of the three parameters, denoted $(\xi(\hat{\alpha}), \xi(\hat{\gamma}), \xi(\hat{\sigma}))$, are presented in Tables 6.1-6.3 below, respectively. In each case, the error of the parameter in question is presented as the RMSE of the 50 individual parameter estimates, expressed as a ratio of the true parameter. As described in Section 6.2.1, 8 values of each parameter are used. For brevity, Tables 6.1-6.3

display parameter estimation errors for three values of σ. The full set of values are presented in Tables B.1-B.3. Thus each table is divided into eight sections, each which reports the error of the given parameter for a particular value of σ, while parameters α and γ vary. As one progresses downward through the different sections of each table the parameter σ becomes larger, corresponding to the density of the Gaussian component becoming wider, leading to a greater level of 'smearing' of the stable component (i.e. the Gaussian component becomes more dominant). After each table is presented below, the error of the parameter in question is discussed.

In Table 6.1, it can be seen that the error in α estimation is larger for higher values of σ. For values of $\sigma \geq 1$, the largest errors are seen when for α is closer to 2, Thus when the Gaussian component is more pronounced it becomes more difficult to estimate α, particularly when the SαS component is itself close to Gaussian. The value of γ does not seem to affect these values significantly, indicating the accuracy of the estimates of parameter α using the proposed method does not depend on γ to a significant degree.

Table 6.2 displays the errors of γ estimation. Similar to the α parameter estimate errors, the error for γ estimates are seen to increase with larger σ, i.e. it becomes more difficult to estimate the parameters of the stable component as the Gaussian component becomes wider in extent. Larger γ estimate errors are observed for larger α values, i.e. as the stable component becomes more Gaussian. When the SαS is completely Gaussian ($\alpha = 2$) γ estimate errors are relatively low. These estimation errors here do not appear to vary with the true value of γ.

In Table 6.3 the errors of σ estimation are presented. In general, the error in σ estimate does not change significantly with the true value of σ, or that of γ. Similar to the case of γ errors above, the highest values are seen for large α, but when the stable component is fully Gaussian ($\alpha = 2$), relatively low errors are observed. The proposed method estimates the parameters of both the stable and Gaussian components simultaneously, while other techniques estimate the stable parameters (α, γ) alone. The error in estimation of the stable parameters $(\xi(\hat{\alpha}), \xi(\hat{\gamma}))$ can be compared with those reported in [172]. This work includes a comprehensive comparison of a number parameter estimation methods for stable distributions, including the SαS case. Table 6.4 reproduces the errors reported in [172], for SαS distributions. The abbreviations used in this table to refer to the different estimation methods are detailed in Table 6.5. For $\alpha = 0.5$, the 'mc' method does not produce an estimate for α, as it only operates for $\alpha > 0.5$. These estimation errors are obtained using

SαS distributions of 500 data points, for $\gamma = 1$ varying values of α, and are transformed to the same unites as those of Tables 6.1-6.3.

When the errors $(\xi(\hat{\alpha}), \xi(\hat{\gamma}))$ of Table 6.4 are compared with of the proposed method in Tables 6.1-6.2 it is observed that the proposed method is of similar accuracy to the best performing methods reported in [172]. This is true even for large values of σ, which make estimation of α and γ more difficult due to increased contribution of the Gaussian component to the combined distribution. In some cases the parameter errors of the proposed method is lower than the best reported in [172].

However, a contributing factor in this may be the differing size of data sets (10^4 points for the proposed method versus 500 points for the values quoted from [172]). In [172] a general increase in estimation accuracy for increasing data length is also observed, indicating that the errors for the methods of 6.4 could possibly be improved by use of more data points.

Given the accuracy of the proposed method for the simulations reported here, and that larger datasets are routinely encountered in image processing applications, it is reasonable to infer that the proposed method of parameter estimation is suitable for multiscale image denoising, in cases where the model of (5.22) is applicable.

Table 6.1: α parameter estimation errors, for various values of α, γ and σ.

$\xi(\hat{\alpha})$, for $\sigma = 0.10$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	7.10×10^{-3}	8.80×10^{-3}	9.11×10^{-3}	6.25×10^{-3}	5.48×10^{-3}	4.93×10^{-3}	2.55×10^{-3}	1.50×10^{-2}
0.71	7.74×10^{-3}	9.23×10^{-3}	5.30×10^{-3}	4.65×10^{-3}	4.40×10^{-3}	3.49×10^{-3}	3.08×10^{-3}	1.68×10^{-2}
0.93	8.36×10^{-3}	6.16×10^{-3}	8.03×10^{-3}	6.00×10^{-3}	4.29×10^{-3}	4.96×10^{-3}	3.07×10^{-3}	1.78×10^{-2}
1.14	8.87×10^{-3}	9.31×10^{-3}	7.63×10^{-3}	6.43×10^{-3}	3.25×10^{-3}	4.93×10^{-3}	2.05×10^{-3}	1.66×10^{-2}
1.36	9.11×10^{-3}	7.89×10^{-3}	9.03×10^{-3}	9.35×10^{-3}	4.70×10^{-3}	3.81×10^{-3}	2.72×10^{-3}	1.81×10^{-2}
1.57	1.10×10^{-2}	8.84×10^{-3}	6.22×10^{-3}	8.09×10^{-3}	3.45×10^{-3}	3.71×10^{-3}	2.90×10^{-3}	1.66×10^{-2}
1.79	1.10×10^{-2}	9.17×10^{-3}	7.20×10^{-3}	5.78×10^{-3}	5.59×10^{-3}	4.12×10^{-3}	4.05×10^{-3}	1.62×10^{-2}
2.00	1.15×10^{-2}	7.48×10^{-3}	6.84×10^{-3}	5.72×10^{-3}	5.72×10^{-3}	3.85×10^{-3}	3.53×10^{-3}	1.44×10^{-2}

$\xi(\hat{\alpha})$, for $\sigma = 0.91$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	1.29×10^{-2}	1.06×10^{-2}	1.45×10^{-2}	1.59×10^{-2}	2.06×10^{-2}	2.10×10^{-2}	2.99×10^{-2}	1.67×10^{-2}
0.71	1.07×10^{-2}	1.35×10^{-2}	7.92×10^{-3}	1.17×10^{-2}	1.57×10^{-2}	1.87×10^{-2}	1.38×10^{-2}	1.67×10^{-2}
0.93	1.04×10^{-2}	1.42×10^{-2}	1.37×10^{-2}	1.69×10^{-2}	1.57×10^{-2}	1.88×10^{-2}	1.51×10^{-2}	1.52×10^{-2}
1.14	1.07×10^{-2}	9.72×10^{-3}	9.90×10^{-3}	1.22×10^{-2}	1.43×10^{-2}	1.67×10^{-2}	2.74×10^{-2}	1.70×10^{-2}
1.36	9.75×10^{-3}	1.18×10^{-2}	1.45×10^{-2}	9.84×10^{-3}	1.21×10^{-2}	1.70×10^{-2}	1.36×10^{-2}	1.45×10^{-2}
1.57	1.35×10^{-2}	1.18×10^{-2}	1.33×10^{-2}	1.39×10^{-2}	1.40×10^{-2}	1.57×10^{-2}	1.46×10^{-2}	1.40×10^{-2}
1.79	1.24×10^{-2}	1.31×10^{-2}	1.34×10^{-2}	1.64×10^{-2}	1.49×10^{-2}	4.48×10^{-2}	1.10×10^{-2}	1.37×10^{-2}
2.00	1.36×10^{-2}	1.03×10^{-2}	1.32×10^{-2}	1.55×10^{-2}	1.44×10^{-2}	4.32×10^{-2}	6.70×10^{-3}	1.39×10^{-2}

$\xi(\hat{\alpha})$, for $\sigma = 2.00$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	2.44×10^{-2}	3.18×10^{-2}	2.00×10^{-2}	3.38×10^{-2}	3.29×10^{-2}	6.09×10^{-2}	2.08×10^{-2}	1.29×10^{-2}
0.71	1.47×10^{-2}	1.53×10^{-2}	2.32×10^{-2}	2.22×10^{-2}	2.73×10^{-2}	4.22×10^{-2}	5.38×10^{-2}	1.34×10^{-2}
0.93	2.17×10^{-2}	1.91×10^{-2}	1.19×10^{-2}	1.91×10^{-2}	2.29×10^{-2}	4.71×10^{-2}	3.46×10^{-2}	1.16×10^{-2}
1.14	1.21×10^{-2}	1.42×10^{-2}	1.41×10^{-2}	1.73×10^{-2}	3.35×10^{-2}	3.67×10^{-2}	3.05×10^{-2}	1.18×10^{-2}
1.36	1.53×10^{-2}	1.19×10^{-2}	1.77×10^{-2}	1.45×10^{-2}	1.96×10^{-2}	2.29×10^{-2}	4.10×10^{-2}	1.23×10^{-2}
1.57	1.20×10^{-2}	1.29×10^{-2}	1.23×10^{-2}	2.83×10^{-2}	1.98×10^{-2}	6.09×10^{-2}	4.58×10^{-2}	1.37×10^{-2}
1.79	1.15×10^{-2}	1.41×10^{-2}	1.21×10^{-2}	1.41×10^{-2}	1.98×10^{-2}	4.53×10^{-2}	3.05×10^{-2}	1.20×10^{-2}
2.00	1.14×10^{-2}	1.43×10^{-2}	9.57×10^{-3}	1.59×10^{-2}	2.29×10^{-2}	2.66×10^{-2}	2.27×10^{-2}	1.51×10^{-2}

Table 6.2: γ parameter estimation errors, for various values of α, γ and σ.

$\xi(\hat{\gamma})$, for $\sigma = 0.10$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	5.69×10^{-3}	7.51×10^{-3}	9.63×10^{-3}	6.79×10^{-3}	6.28×10^{-3}	6.10×10^{-3}	8.10×10^{-3}	2.16×10^{-2}
0.71	4.19×10^{-3}	6.50×10^{-3}	5.12×10^{-3}	5.79×10^{-3}	6.45×10^{-3}	5.54×10^{-3}	5.89×10^{-3}	1.26×10^{-2}
0.93	4.80×10^{-3}	4.94×10^{-3}	5.45×10^{-3}	5.97×10^{-3}	6.59×10^{-3}	6.33×10^{-3}	6.47×10^{-3}	1.46×10^{-2}
1.14	2.95×10^{-3}	5.10×10^{-3}	4.90×10^{-3}	5.25×10^{-3}	5.65×10^{-3}	7.11×10^{-3}	4.96×10^{-3}	1.05×10^{-2}
1.36	2.84×10^{-3}	3.89×10^{-3}	6.08×10^{-3}	9.76×10^{-3}	5.19×10^{-3}	5.87×10^{-3}	5.03×10^{-3}	7.60×10^{-3}
1.57	4.18×10^{-3}	4.48×10^{-3}	5.11×10^{-3}	5.92×10^{-3}	6.16×10^{-3}	5.93×10^{-3}	6.03×10^{-3}	8.43×10^{-3}
1.79	3.97×10^{-3}	4.97×10^{-3}	5.15×10^{-3}	5.32×10^{-3}	6.44×10^{-3}	6.44×10^{-3}	7.06×10^{-3}	8.17×10^{-3}
2.00	4.82×10^{-3}	6.01×10^{-3}	6.32×10^{-3}	5.50×10^{-3}	6.59×10^{-3}	6.09×10^{-3}	7.28×10^{-3}	7.32×10^{-3}

$\xi(\hat{\gamma})$, for $\sigma = 0.91$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	1.04×10^{-2}	1.20×10^{-2}	2.07×10^{-2}	3.04×10^{-2}	5.25×10^{-2}	8.47×10^{-2}	2.75×10^{-1}	1.55×10^{-2}
0.71	7.45×10^{-3}	1.49×10^{-2}	1.21×10^{-2}	1.72×10^{-2}	4.01×10^{-2}	8.76×10^{-2}	1.86×10^{-1}	1.30×10^{-2}
0.93	7.30×10^{-3}	1.13×10^{-2}	2.00×10^{-2}	2.26×10^{-2}	3.35×10^{-2}	6.74×10^{-2}	1.58×10^{-1}	1.54×10^{-2}
1.14	6.31×10^{-3}	9.25×10^{-3}	1.33×10^{-2}	1.77×10^{-2}	3.03×10^{-2}	7.97×10^{-2}	8.83×10^{-2}	9.79×10^{-3}
1.36	4.39×10^{-3}	7.23×10^{-3}	9.85×10^{-3}	1.49×10^{-2}	2.66×10^{-2}	7.61×10^{-2}	6.64×10^{-2}	9.18×10^{-3}
1.57	5.95×10^{-3}	1.02×10^{-2}	1.77×10^{-2}	2.36×10^{-2}	2.68×10^{-2}	6.99×10^{-2}	7.84×10^{-2}	7.42×10^{-3}
1.79	7.20×10^{-3}	8.09×10^{-3}	1.09×10^{-2}	2.32×10^{-2}	4.01×10^{-2}	5.88×10^{-2}	4.33×10^{-2}	7.82×10^{-3}
2.00	5.53×10^{-3}	7.65×10^{-3}	1.20×10^{-2}	1.90×10^{-2}	3.83×10^{-2}	6.59×10^{-2}	4.41×10^{-2}	7.52×10^{-3}

$\xi(\hat{\gamma})$, for $\sigma = 2.00$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	2.25×10^{-2}	4.29×10^{-2}	4.29×10^{-2}	8.87×10^{-2}	1.40×10^{-1}	6.52×10^{-1}	9.36×10^{-1}	2.35×10^{-2}
0.71	1.40×10^{-2}	1.83×10^{-2}	4.36×10^{-2}	5.22×10^{-2}	1.04×10^{-1}	2.64×10^{-1}	9.95×10^{-1}	1.63×10^{-2}
0.93	1.95×10^{-2}	2.49×10^{-2}	2.33×10^{-2}	5.04×10^{-2}	7.19×10^{-2}	2.66×10^{-1}	3.19×10^{-1}	1.42×10^{-2}
1.14	1.07×10^{-2}	1.79×10^{-2}	2.33×10^{-2}	4.26×10^{-2}	9.82×10^{-2}	2.03×10^{-1}	4.80×10^{-1}	1.32×10^{-2}
1.36	1.19×10^{-2}	1.23×10^{-2}	3.15×10^{-2}	3.92×10^{-2}	6.64×10^{-2}	1.50×10^{-1}	3.21×10^{-1}	1.00×10^{-2}
1.57	1.02×10^{-2}	1.16×10^{-2}	1.77×10^{-2}	4.84×10^{-2}	6.78×10^{-2}	2.13×10^{-1}	2.23×10^{-1}	8.36×10^{-3}
1.79	6.66×10^{-3}	1.60×10^{-2}	2.20×10^{-2}	3.53×10^{-2}	7.46×10^{-2}	1.55×10^{-1}	1.70×10^{-1}	7.27×10^{-3}
2.00	7.76×10^{-3}	1.60×10^{-2}	1.41×10^{-2}	3.61×10^{-2}	9.24×10^{-2}	1.91×10^{-1}	1.93×10^{-1}	7.60×10^{-3}

Table 6.3: σ parameter estimation errors, for various values of α, γ and σ.

$\xi(\hat{\sigma})$, for $\sigma = 0.10$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	7.58×10^{-3}	1.19×10^{-2}	6.91×10^{-2}	1.11×10^{-2}	1.01×10^{-1}	2.35×10^{-2}	2.36×10^{-2}	3.24×10^{-2}
0.71	1.59×10^{-2}	1.12×10^{-2}	2.18×10^{-2}	4.03×10^{-2}	2.09×10^{-2}	2.19×10^{-2}	2.62×10^{-2}	3.53×10^{-2}
0.93	1.65×10^{-2}	3.61×10^{-2}	1.20×10^{-2}	1.07×10^{-1}	1.99×10^{-2}	2.54×10^{-2}	2.48×10^{-2}	3.29×10^{-2}
1.14	1.46×10^{-2}	9.37×10^{-3}	1.26×10^{-1}	1.97×10^{-2}	1.99×10^{-2}	2.45×10^{-2}	2.51×10^{-2}	4.53×10^{-2}
1.36	1.38×10^{-2}	1.03×10^{-2}	1.60×10^{-2}	3.94×10^{-1}	2.21×10^{-2}	2.52×10^{-2}	2.69×10^{-2}	4.03×10^{-2}
1.57	3.65×10^{-2}	3.97×10^{-2}	1.67×10^{-2}	2.06×10^{-2}	2.17×10^{-2}	2.55×10^{-2}	2.82×10^{-2}	2.89×10^{-2}
1.79	1.80×10^{-2}	1.65×10^{-2}	2.33×10^{-2}	2.58×10^{-2}	2.58×10^{-2}	2.67×10^{-2}	2.78×10^{-2}	3.07×10^{-2}
2.00	8.78×10^{-2}	2.07×10^{-1}	2.06×10^{-2}	2.35×10^{-2}	2.79×10^{-2}	2.76×10^{-2}	2.75×10^{-2}	4.06×10^{-2}

$\xi(\hat{\sigma})$, for $\sigma = 0.91$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	6.97×10^{-3}	8.35×10^{-3}	1.21×10^{-2}	1.84×10^{-2}	3.08×10^{-2}	5.20×10^{-2}	2.85×10^{-1}	9.97×10^{-3}
0.71	6.79×10^{-3}	1.20×10^{-2}	1.12×10^{-2}	1.61×10^{-2}	3.60×10^{-2}	8.68×10^{-2}	2.79×10^{-1}	1.25×10^{-2}
0.93	8.30×10^{-3}	1.40×10^{-2}	2.19×10^{-2}	2.64×10^{-2}	4.08×10^{-2}	6.99×10^{-2}	2.91×10^{-1}	1.23×10^{-2}
1.14	1.09×10^{-2}	1.28×10^{-2}	2.10×10^{-2}	2.57×10^{-2}	4.34×10^{-2}	1.24×10^{-1}	2.20×10^{-1}	1.11×10^{-2}
1.36	7.19×10^{-3}	1.49×10^{-2}	2.02×10^{-2}	2.35×10^{-2}	4.75×10^{-2}	1.62×10^{-1}	1.05×10^{-1}	1.07×10^{-2}
1.57	1.45×10^{-2}	2.12×10^{-2}	3.15×10^{-2}	5.12×10^{-2}	5.85×10^{-2}	2.00×10^{-1}	2.82×10^{-1}	1.20×10^{-2}
1.79	1.61×10^{-2}	1.66×10^{-2}	2.60×10^{-2}	5.53×10^{-2}	9.26×10^{-2}	1.32×10^{-1}	1.79×10^{-1}	1.14×10^{-2}
2.00	1.27×10^{-2}	1.97×10^{-2}	2.71×10^{-2}	5.53×10^{-2}	9.82×10^{-2}	2.16×10^{-1}	1.79×10^{-1}	1.36×10^{-2}

$\xi(\hat{\sigma})$, for $\sigma = 2.00$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	5.42×10^{-3}	9.95×10^{-3}	1.03×10^{-2}	1.60×10^{-2}	2.05×10^{-2}	1.44×10^{-1}	2.33×10^{-1}	2.83×10^{-3}
0.71	5.64×10^{-3}	6.09×10^{-3}	1.16×10^{-2}	1.38×10^{-2}	2.23×10^{-2}	5.33×10^{-2}	3.38×10^{-1}	4.00×10^{-3}
0.93	9.24×10^{-3}	1.02×10^{-2}	9.11×10^{-3}	1.53×10^{-2}	2.12×10^{-2}	7.91×10^{-2}	1.38×10^{-1}	4.40×10^{-3}
1.14	7.23×10^{-3}	8.59×10^{-3}	8.63×10^{-3}	1.87×10^{-2}	3.22×10^{-2}	6.73×10^{-2}	2.10×10^{-1}	4.47×10^{-3}
1.36	8.98×10^{-3}	8.03×10^{-3}	1.80×10^{-2}	1.92×10^{-2}	2.85×10^{-2}	6.26×10^{-2}	1.46×10^{-1}	5.77×10^{-3}
1.57	1.07×10^{-2}	8.57×10^{-3}	1.07×10^{-2}	2.17×10^{-2}	3.14×10^{-2}	1.09×10^{-1}	1.16×10^{-1}	5.46×10^{-3}
1.79	8.50×10^{-3}	1.10×10^{-2}	1.63×10^{-2}	2.23×10^{-2}	4.26×10^{-2}	7.98×10^{-2}	9.32×10^{-2}	7.06×10^{-3}
2.00	9.90×10^{-3}	1.70×10^{-2}	1.44×10^{-2}	2.49×10^{-2}	5.69×10^{-2}	1.21×10^{-1}	1.11×10^{-1}	6.52×10^{-3}

Table 6.4: α and γ parameter estimate errors, of various methods of estimation, as reported in [172]. Errors are for SαS distributions with $\gamma = 1$ and various values of α, and have been transformed to the same units as Tables 6.1-6.3. Abbreviations describing estimation methods are explained in Table 6.5.

$\xi(\hat{\alpha})$	Etimation Method						
α	ml	k	kar	ekl	p	mc	kw
0.5	5.74×10^{-2}	3.13×10^{-2}	2.34×10^{-2}	2.52×10^{-2}	2.70×10^{-2}	-	4.79×10^{-2}
0.9	4.38×10^{-2}	4.89×10^{-2}	4.58×10^{-2}	4.44×10^{-2}	4.70×10^{-2}	6.06×10^{-2}	6.16×10^{-2}
1.3	6.59×10^{-2}	6.93×10^{-2}	6.72×10^{-2}	6.63×10^{-2}	7.23×10^{-2}	7.65×10^{-2}	7.99×10^{-2}
1.7	6.91×10^{-2}	6.96×10^{-2}	7.00×10^{-2}	6.91×10^{-2}	7.44×10^{-2}	1.22×10^{-1}	8.14×10^{-2}

$\xi(\hat{\gamma})$	Etimation Method						
α	ml	k	kar	ekl	p	mc	kw
0.5	1.45×10^{-1}	1.18×10^{-1}	1.15×10^{-1}	1.17×10^{-1}	1.16×10^{-1}	1.67×10^{-1}	1.46×10^{-1}
0.9	6.85×10^{-2}	7.00×10^{-2}	7.00×10^{-2}	6.91×10^{-2}	6.89×10^{-2}	7.99×10^{-2}	8.48×10^{-2}
1.3	5.31×10^{-2}	5.33×10^{-2}	5.34×10^{-2}	5.32×10^{-2}	5.32×10^{-2}	5.74×10^{-2}	5.88×10^{-2}
1.7	4.42×10^{-2}	4.43×10^{-2}	4.46×10^{-2}	4.42×10^{-2}	4.47×10^{-2}	5.68×10^{-2}	4.55×10^{-2}

Table 6.5: List of abbreviations used for parameter estimation methods in Table 6.4.

abbreviation	method
ml	Maximum Likelihood [158]
k	Koutrouvelis [164, 165]
kar	Koutrouvelis with nonlinear arithmetic spacing
ekl	Emperical k-L method [172]
p	Pourahmadi [179]
mc	McCulloch [160]
kw	Kogon Williams [167]

6.3.2 ZC Edge Detection Performance

The results of application of edge detection methods to simulated, phantom, and clinical edge maps are presented in this section. The visual properties of these edge maps are discussed. Favourable edge detection performance is signified by accurate location of all image edges, with a minimum of false edges. The degree to which these edges are localised is also considered, and this is qualitatively assessed for the binary edge maps reproduced below by noting the nature of these edges. Edges which are seen to be narrow in extent offer a more localised estimate of edge location than is provided by wider edge maps. For the simulated and phantom image sets ideal edge maps are available. FoM scores can be computed in these cases, and tables of these values are included.

6.3.2.1 Simulated Images

Figs. 6.9–6.10 display typical examples of the edge maps produced by each of the edge detectors when applied to the simulated images. These are representative of the results for the simulated images as a whole, and the remaining edge maps for the simulated images can be found in Appendix B, in Figs. B.1–B.6. The edges displayed are those using the parameter values for each filter which produce the highest FoM.

The edges produced by the CV method are displayed in Figs. 6.9(a) and 6.10(a). Most of the edges in the simulated images are located successfully, however some edges are incomplete with missing sections (particularly Fig. 6.9(a)). Some of spurious edges are observed in Fig. 6.10(a). It is also seen that the edges are poorly localised, i.e. the edges produced are quite thick in extent. This is a result of the relatively large window size N which produces the optimal FoM value. It is noted experimentally that smaller values of N produce higher levels of spurious edges, which reduce the resulting FoM.

The RoA method produces edge maps which can be seen in Figs. 6.9(b) and 6.10(b). All of the ideal simulated edges are located with highly localised (narrow) edges. A high level of erroneous edges are observed, indicating that the RoA step in this edge detection method is unable to constrain the LoG edges to correct regions. This effect is particularly noticeable in Fig. 6.10(b).

The FA method edges for the simulated images are displayed in Figs. 6.9(c) and 6.10(c). This edge detection method locates the image edges successfully in a large number of cases. The edge maps can be poorly localised however. In Fig. 6.10(c), some images

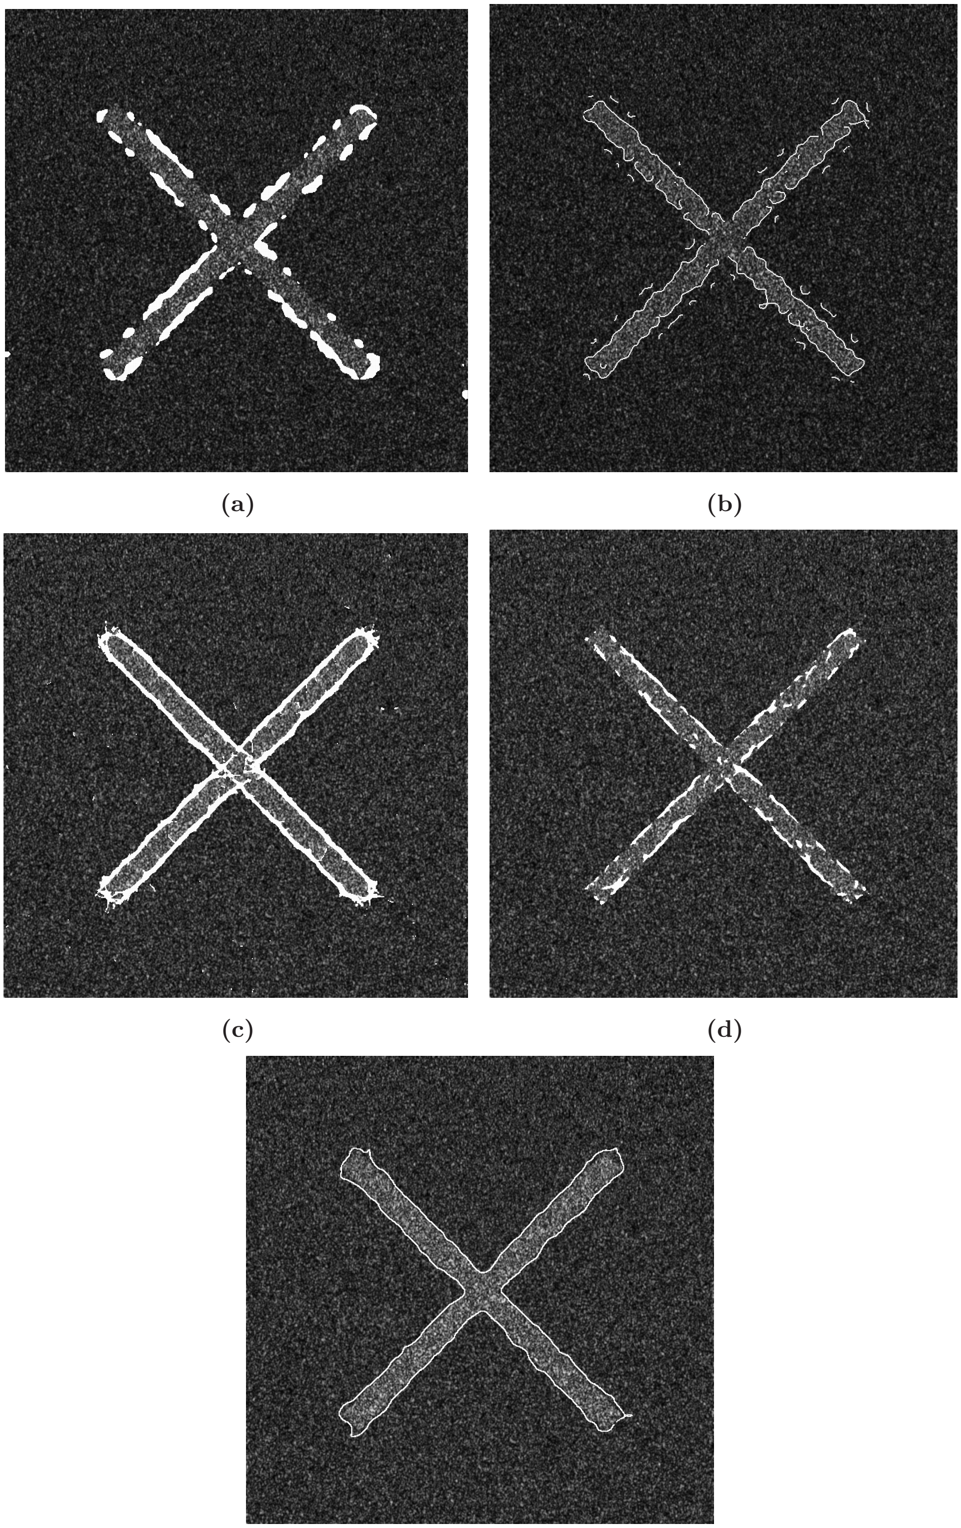

(a) **(b)** **(c)** **(d)** **(e)**

Figure 6.9: Edge detector output for the simulated edge images of Fig. 6.1(c). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

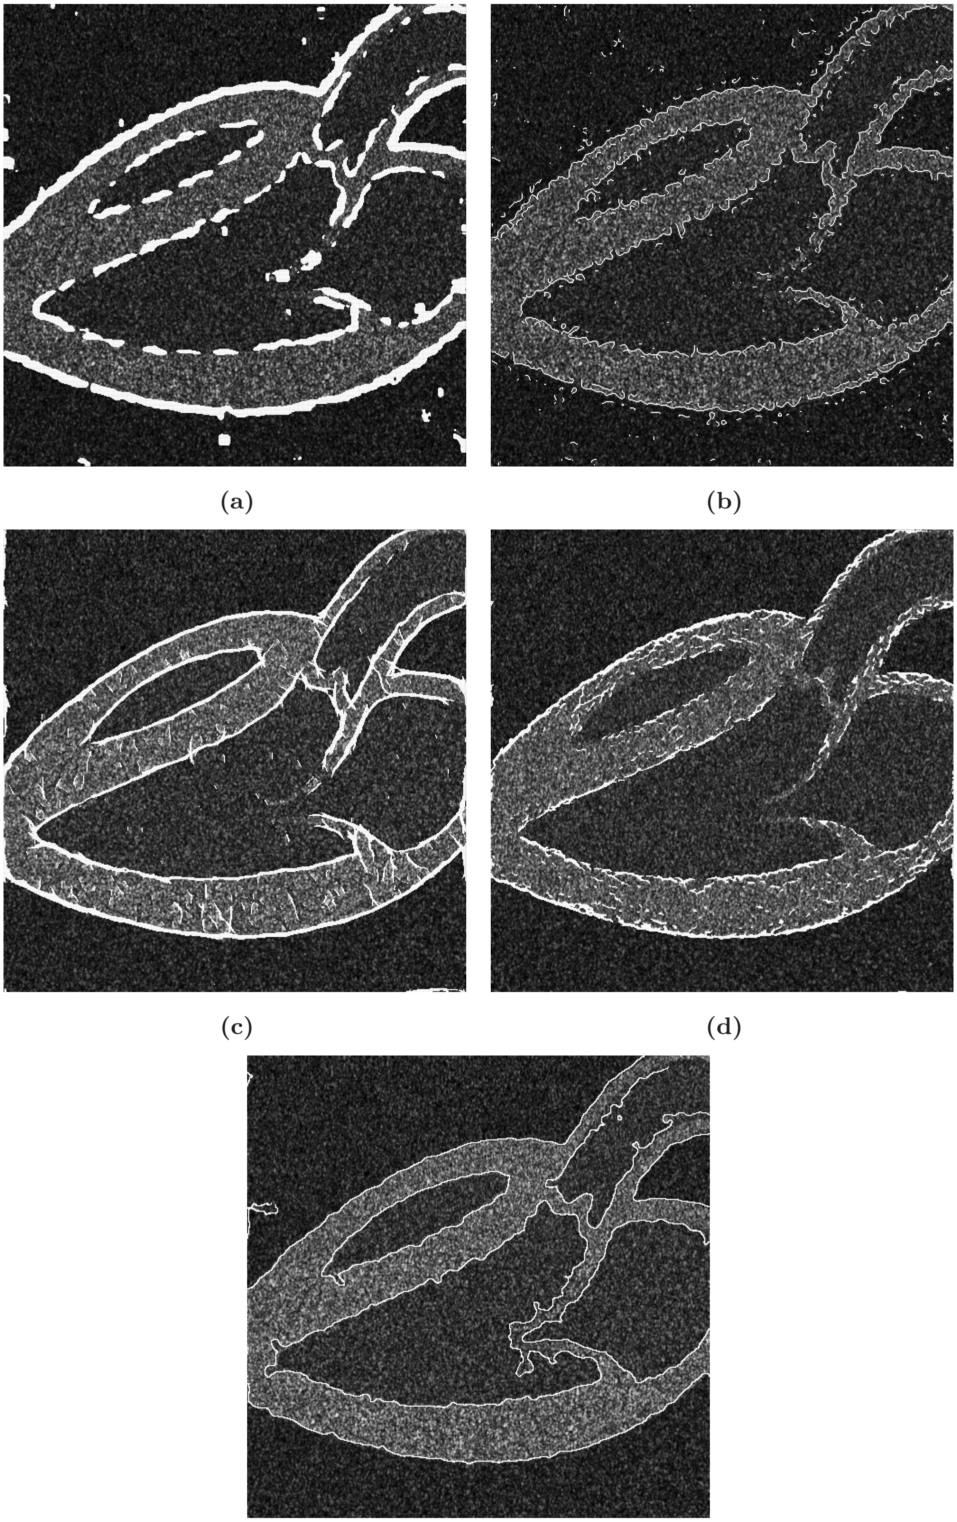

(a) **(b)**

(c) **(d)**

(e)

Figure 6.10: Edge detector output for the simulated edge images of Fig. 6.3(a). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

edges are not located by this method, while in Fig. 6.9(c) some spurious edges are observed.

When applied to the simulated images, the PC method produces edge maps as in Figs. 6.9(d) and 6.10(d). The performance of this method varies from image to image. The edges of the PC method shown in this section contain both erroneous edge segments, and also missing edge regions. Some of the edges displayed in Appendix B are more accurate however, particularly Figs. B.1(d) and B.2(d).

The edge maps generated by the proposed ZC edge detector for simulated images are displayed in Figs. 6.9(e) and 6.10(e). This is generally seen to perform well, detecting most image edges. Some slight localisation errors are seen in Figs. 6.9(e). In general however, these edges are very clearly defined, with narrow edge maps observed in all cases.

The optimal FoM values for the simulated shape templates are reported in Tables 6.6. These values compare the ideal edge maps of Fig. 6.2 to the corresponding edge maps produced by the edge detectors. As noted in Section 6.2.2, ten simulated images are produced for each template, and each edge detection method is applied using a range of parameter values. This table thus presents the largest FoM achieved by each edge detector for each of the ten images produced from the shape image templates. The optimal mean FoM values for each of the simulated shape templates are presented in Tables 6.7. These values are found for each filter by averaging the FoM values of each parameter combination, over the ten images for each template. The largest average FoM for each template are then reported here.

Similarly the FoM values for images produced from simulated clinical templates are displayed in Tables 6.6 and 6.9. These compare the ideal edge maps of Fig. 6.4 to the corresponding edge detector outputs. Table 6.6 displays the best FoM values for each of the ten images simulated from each template, while Tables 6.9 displays the optimal average FoM values for each template.

In all cases, the FoM for the proposed ZC method is highest, reflecting the accuracy of this method in locating image edges. The FoM values attained by the CV method are relatively low, corresponding to the lower edge detection performance observed in these simulated images. The FoM values for the PC method are highest for the simpler shapes of templates 1 and 2 in Tables 6.7. The FA methods FoM is high in some cases, but not consistently so. This confirms what is noted by visual analysis above, i.e. performance is seen to vary with image. The FoM scores of the RoA method are surprisingly high, considering the level of erroneous edges noted in the discussion above. This is explained by noting the location of these incorrect edges, which is usually in close proximity to the ideal edge. Thus even though such false edges are visually quite noticeable, they do not adversely effect the FoM scores significantly.

Table 6.6: Optimal edge detector FoM values, for each of the 10 images produced using the simulated shape templates. Templates 1-4 refer those illustrated in Figs. 6.1(a)-(d).

	Image	ZC	FA	RoA	PC	CV
Template 1	1	0.885	0.750	0.739	0.760	0.638
	2	0.862	0.694	0.713	0.765	0.614
	3	0.864	0.750	0.724	0.760	0.588
	4	0.862	0.674	0.701	0.758	0.678
	5	0.872	0.761	0.717	0.768	0.591
	6	0.864	0.758	0.729	0.752	0.616
	7	0.876	0.758	0.688	0.768	0.624
	8	0.844	0.729	0.687	0.754	0.649
	9	0.896	0.753	0.706	0.767	0.657
	10	0.863	0.755	0.729	0.773	0.656
Template 2	1	0.775	0.648	0.697	0.723	0.593
	2	0.836	0.640	0.714	0.726	0.610
	3	0.800	0.635	0.706	0.709	0.636
	4	0.820	0.629	0.737	0.716	0.644
	5	0.835	0.638	0.767	0.731	0.601
	6	0.799	0.633	0.736	0.719	0.598
	7	0.799	0.636	0.730	0.721	0.638
	8	0.825	0.644	0.748	0.725	0.649
	9	0.799	0.634	0.736	0.729	0.670
	10	0.831	0.633	0.734	0.721	0.619
Template 3	1	0.786	0.664	0.642	0.611	0.660
	2	0.779	0.669	0.638	0.630	0.621
	3	0.787	0.669	0.629	0.638	0.561
	4	0.811	0.669	0.656	0.628	0.635
	5	0.786	0.668	0.618	0.596	0.605
	6	0.789	0.669	0.616	0.579	0.616
	7	0.793	0.670	0.671	0.624	0.651
	8	0.788	0.672	0.635	0.605	0.561
	9	0.773	0.676	0.628	0.581	0.632
	10	0.772	0.669	0.633	0.585	0.648
Template 4	1	0.737	0.684	0.672	0.673	0.670
	2	0.661	0.676	0.669	0.630	0.610
	3	0.711	0.690	0.641	0.677	0.605
	4	0.708	0.683	0.670	0.693	0.630
	5	0.714	0.678	0.659	0.678	0.638
	6	0.720	0.682	0.667	0.669	0.658
	7	0.696	0.687	0.689	0.690	0.628
	8	0.726	0.675	0.655	0.667	0.606
	9	0.694	0.685	0.678	0.671	0.587
	10	0.757	0.697	0.682	0.669	0.590

Table 6.7: Optimal mean edge detector FoM values, averaged over 10 images produced using the simulated shape templates. Templates 1-4 refer those illustrated in Figs. 6.1(a)-(d), respectively.

	ZC	FA	RoA	PC	CV
Template 1	0.866	0.738	0.713	0.760	0.631
Template 2	0.808	0.636	0.731	0.720	0.626
Template 3	0.765	0.670	0.636	0.607	0.619
Template 4	0.692	0.676	0.668	0.671	0.622

Table 6.8: Optimal edge detector FoM values, for each of the 10 images produced using the simulated clinical templates. Templates 1-4 refer those illustrated in Figs. 6.3(a)-(d).

	Image	ZC	FA	RoA	PC	CV
Template 1	1	0.739	0.663	0.630	0.577	0.614
	2	0.724	0.658	0.623	0.600	0.619
	3	0.746	0.663	0.647	0.608	0.638
	4	0.707	0.660	0.619	0.608	0.605
	5	0.738	0.666	0.630	0.590	0.627
	6	0.741	0.661	0.621	0.603	0.628
	7	0.726	0.667	0.624	0.598	0.633
	8	0.738	0.661	0.627	0.600	0.626
	9	0.739	0.671	0.633	0.606	0.649
	10	0.726	0.669	0.634	0.587	0.626
Template 2	1	0.830	0.695	0.639	0.691	0.660
	2	0.789	0.696	0.639	0.699	0.643
	3	0.804	0.691	0.629	0.693	0.645
	4	0.787	0.702	0.659	0.702	0.643
	5	0.794	0.706	0.645	0.708	0.659
	6	0.788	0.695	0.642	0.697	0.638
	7	0.791	0.700	0.639	0.694	0.654
	8	0.804	0.690	0.624	0.686	0.659
	9	0.821	0.688	0.639	0.696	0.648
	10	0.790	0.703	0.634	0.695	0.660
Template 3	1	0.772	0.568	0.650	0.480	0.660
	2	0.785	0.550	0.654	0.469	0.635
	3	0.724	0.544	0.649	0.488	0.639
	4	0.752	0.562	0.647	0.455	0.644
	5	0.759	0.545	0.647	0.481	0.648
	6	0.759	0.546	0.656	0.472	0.664
	7	0.761	0.546	0.639	0.473	0.630
	8	0.755	0.559	0.648	0.504	0.661
	9	0.753	0.554	0.640	0.490	0.641
	10	0.742	0.554	0.632	0.502	0.640
Template 4	1	0.837	0.701	0.702	0.621	0.654
	2	0.840	0.695	0.692	0.639	0.664
	3	0.849	0.692	0.699	0.633	0.665
	4	0.828	0.691	0.703	0.629	0.669
	5	0.833	0.697	0.687	0.639	0.670
	6	0.818	0.689	0.695	0.638	0.657
	7	0.841	0.698	0.697	0.643	0.670
	8	0.827	0.707	0.682	0.624	0.663
	9	0.827	0.695	0.688	0.635	0.665
	10	0.843	0.701	0.671	0.633	0.666

Table 6.9: Optimal mean edge detector FoM values, averaged over 10 images produced using the simulated clinical templates. Templates 1-4 refer those illustrated in Figs.6.3(a)-(d), respectively.

	ZC	FA	RoA	PC	CV
Template 1	0.728	0.661	0.629	0.596	0.626
Template 2	0.799	0.687	0.639	0.694	0.651
Template 3	0.746	0.551	0.646	0.480	0.646
Template 4	0.826	0.695	0.692	0.629	0.664

6.3.2.2 Phantom Images

A selection of the edge maps produced by each filter when applied to the phantom images of Fig. 6.5 are displayed in Figs. 6.11–6.12. These are the edges produced using the optimal parameter and threshold values for each method, as determined by the FoM criterion. The edge maps shown in Figs. 6.11–6.12 represent the results that are seen generally for the phantom images. The complete set of edge maps for the phantom images of Fig. 6.5 can be found in Appendix B, in Figs. B.7–B.10. The phantom images provide a greater challenge to the edge detectors than do the simulated images, having in general less clearly defined borders. This results in edge maps which can be seen to be of lower accuracy than in the simulated case.

The edge maps produced by the CV method can be seen in Figs. 6.11(a) and B.10(a). These are seen to contain poorly localised edges of wide extent. In many cases a large amount of false edges are seen, and in some cases (e.g. Fig. 6.11(a)), a large amount of image edges are missed.

The RoA method produces edges which are displayed in Figs. 6.11(b) and B.10(b). This edge detection approach is capable of delineating almost all of the phantom image edges in most images. As is observed in the simulated case above, a large amount of spurious edges are present. In contrast to the simulated case however, these are not restricted to regions close to the ideal edges, and are instead spread throughout the images.

When the FA method is applied to the phantom images, the edges yielding the highest FoM values are shown in Figs. 6.11(c) and B.10(c). These can be seen to contain nearly all phantom image edges. However, some false edges are observed, and often the edge maps are poorly localised.

The PC method produces edges for the phantom images which are shown in Figs. 6.11(d) and B.10(d). This edge detection method does not appear to perform well when applied to these images. A large amount of image edges are missed, and those that detected are not well localised.

The proposed ZC edge detection method produces edges displayed in Figs. 6.11(e) and B.10(e). This method is seen to locate nearly all of the phantom image edges, with very localised (i.e. narrow) edge maps. In most of these images the amount of false edges present is very small. In cases where the phantom image borders are missing or obscured, the edges produced by this method are seen to follow the local image texture in missing

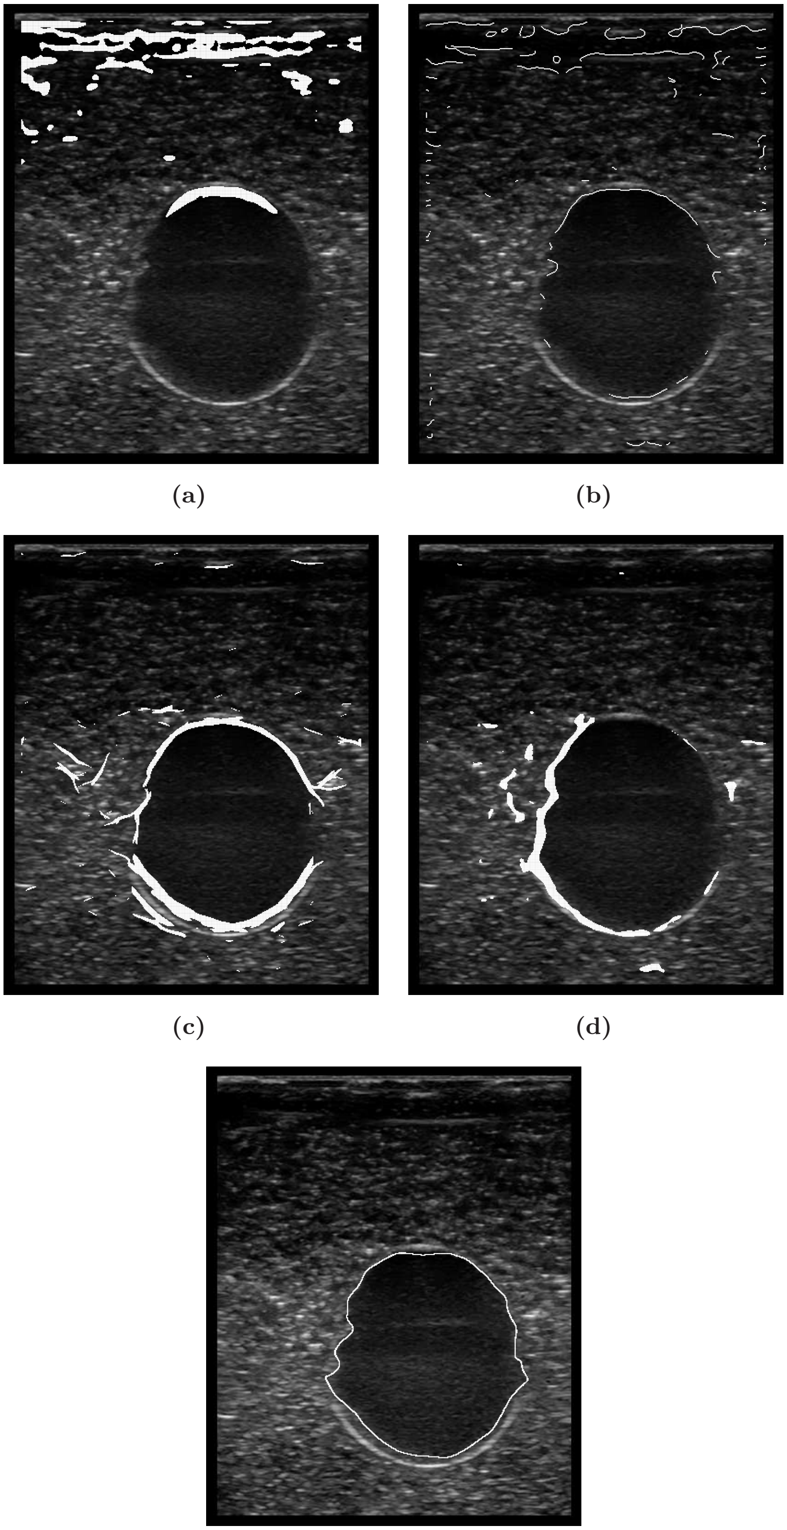

(a) **(b)**

(c) **(d)**

(e)

Figure 6.11: Edge detector output for the phantom ultrasound images of Figs. 6.5(b). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

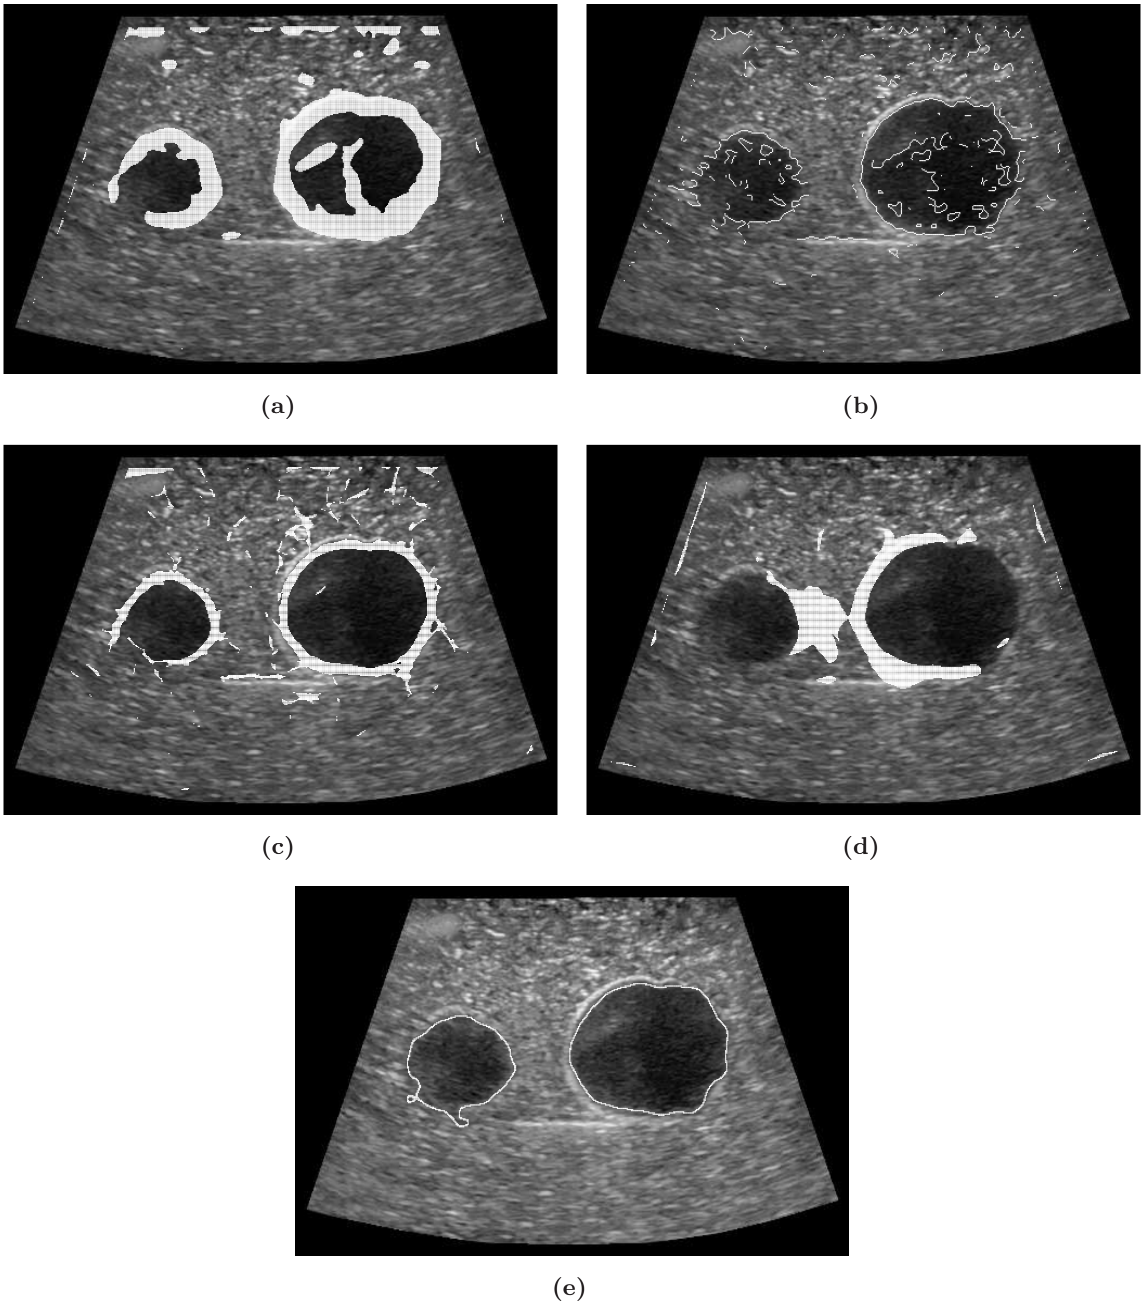

(a) (b) (c) (d) (e)

Figure 6.12: Edge detector output for the phantom ultrasound images of Figs. 6.5a(f). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

boundary regions. Examples of this behaviour can be seen in Figs. 6.11(e) and B.10(e).

Table 6.10 displays the optimal FoM values measured for each edge detection method when they are applied to the phantom images. These are in general lower than

Table 6.10: Optimal edge detector FoM values for the phantom images of Fig. 6.5. Images 1-7 refer to Figs. 6.5(a)–6.5(f), respectively.

	ZC	FA	RoA	PC	CV
Image 1	0.584	0.445	0.180	0.391	0.104
Image 2	0.662	0.480	0.197	0.292	0.092
Image 3	0.518	0.393	0.297	0.298	0.319
Image 4	0.540	0.405	0.267	0.337	0.289
Image 5	0.722	0.436	0.348	0.385	0.360
Image 6	0.722	0.460	0.379	0.369	0.346

the FoM values achieved for the simulated images. This reflects the greater difficulty in accurately locating edges in these images, as noted above. The FoM values are highest in all phantom images for the proposed ZC edge detector. This corresponds with the favourable performance observed in Figs. 6.11–B.10. The CV and PC methods are seen to have poor edge detection performance in these images, and this is reflected in the their relatively low FoM scores. The false edges produced by the RoA method appear to adversely affect this methods FoM values, due to their location in the images (which is often far removed from the true edge location). This is in contrast to the simulated case above, in which the false edges present do not have such an adverse affect on FoM scores, due to close proximity to the ideal edges. The FoM scores achieved by the FA edge detector are relatively high here, which corresponds with the good performance observed for this method.

6.3.2.3 Clinical Images

The edge detection methods are applied to the clinical images of Figs. 6.7–6.8. A selection of the resulting edge maps are displayed in Figs. 6.13–6.14, which are representative of the full results of the clinical images. The entire set of edge maps are produced in Appendix B, in Figs. B.11-B.20.

As no ideal edge map exists for these images, it is not possible to select optimal parameter and threshold values for each filter by optimising the FoM. The parameter and threshold values used here are those found to give be optimal result when applied to the phantom images above, and are indicative of the general performance of each filter. The edge detectors appear to perform somewhat better when applied to these images, successfully locating a higher level of image edges when compared to phantom image set discussed above.

The CV method edge maps are displayed in Figs. 6.13(a)–6.14(a). These are seen to locate most (but not all of) the image edges. However, as is the case in the simulated and phantom image sets, these edges are poorly localised. The edge maps produced for the clinical images by the RoA method are displayed in Figs. 6.13(b)–6.14(b). These maps contain highly localised narrow edges. As seen previously, spurious edge are observed in images locations where no edge is discernible. The amount of these false edges is seen to be less than in for the case of the phantom images above. The FA method produces edge for the clinical images as seen in Figs. 6.13(c)–6.14(c). This method can be seen to perform well when applied to these images, locating nearly all image edges in the majority of images presented. A small level of false edges can be seen, and in some cases the edges found are poorly localised. Neither of these effects are as pronounced as in the case of the phantom image set. The PC method clinical edge maps are displayed in Figs. 6.13(d)–6.14(d). These edge maps do not appear to reflect the underlying image edge structure accurately. In most of these images a large amount of false edges are seen, while highly visible edge regions are missed.

The proposed ZC method produces clinical edge maps as shown in Figs. 6.13(e)–6.14(e). As in the case of the simulated in phantom image sets, this method produces edge maps which are seen to locate most visible image edges with highly localised edge maps. Both high and low contrast boundaries are located, and small scale features such as the valve leaflets of Fig. 6.14(e) are successfully delineated.

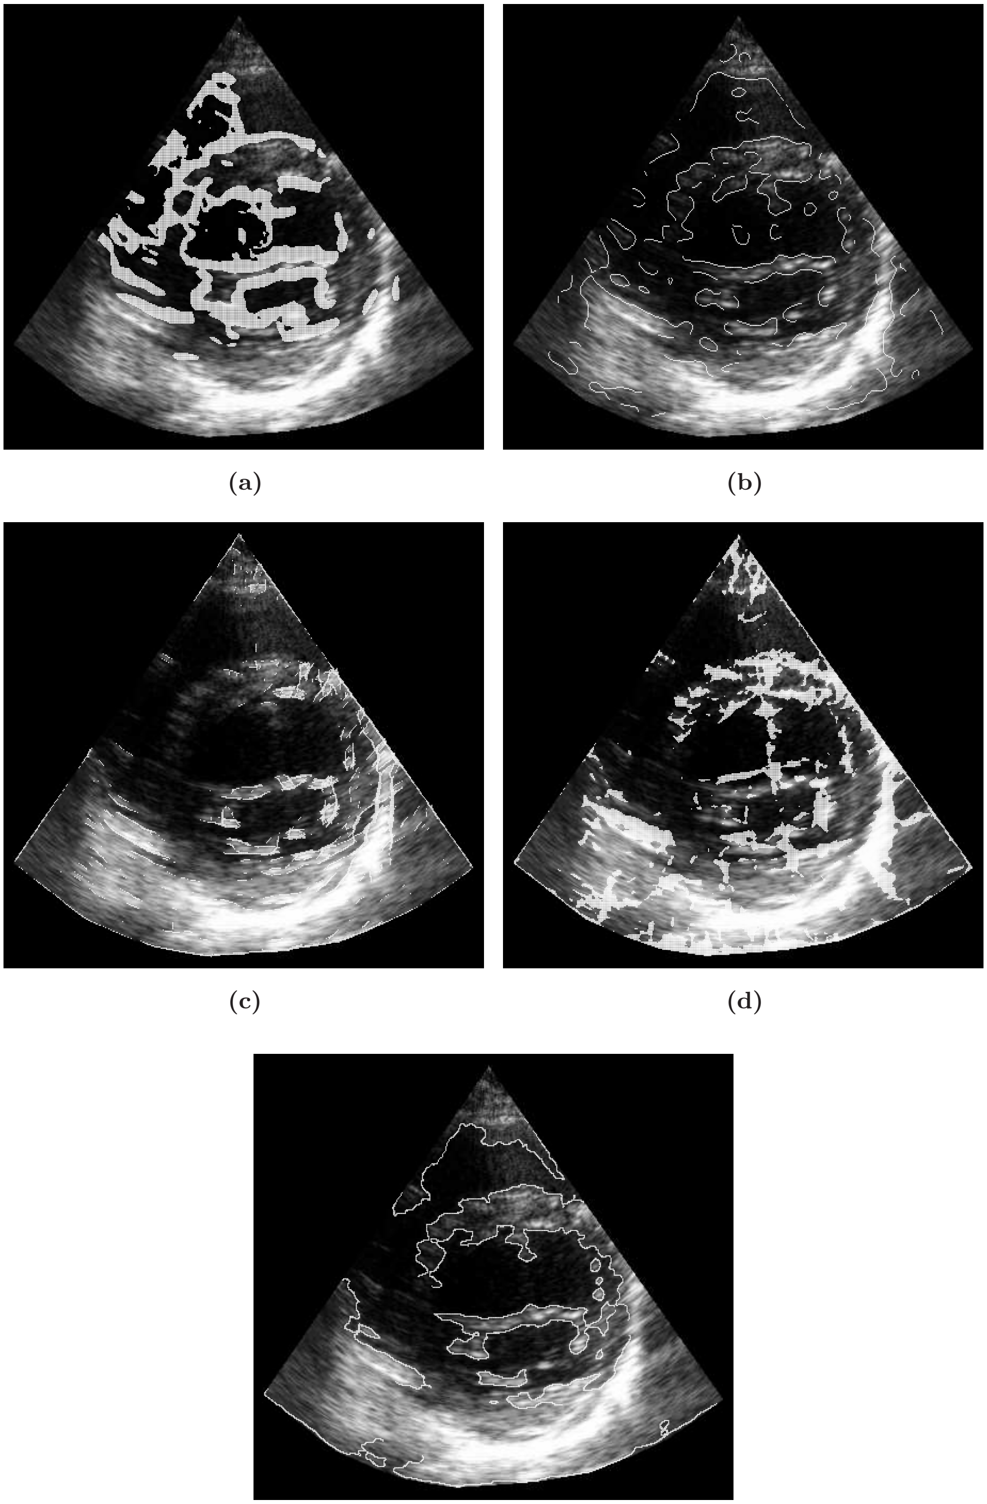

(a)

(b)

(c)

(d)

(e)

Figure 6.13: Edge detector output for the clinical ultrasound images of Fig. 6.7(c). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

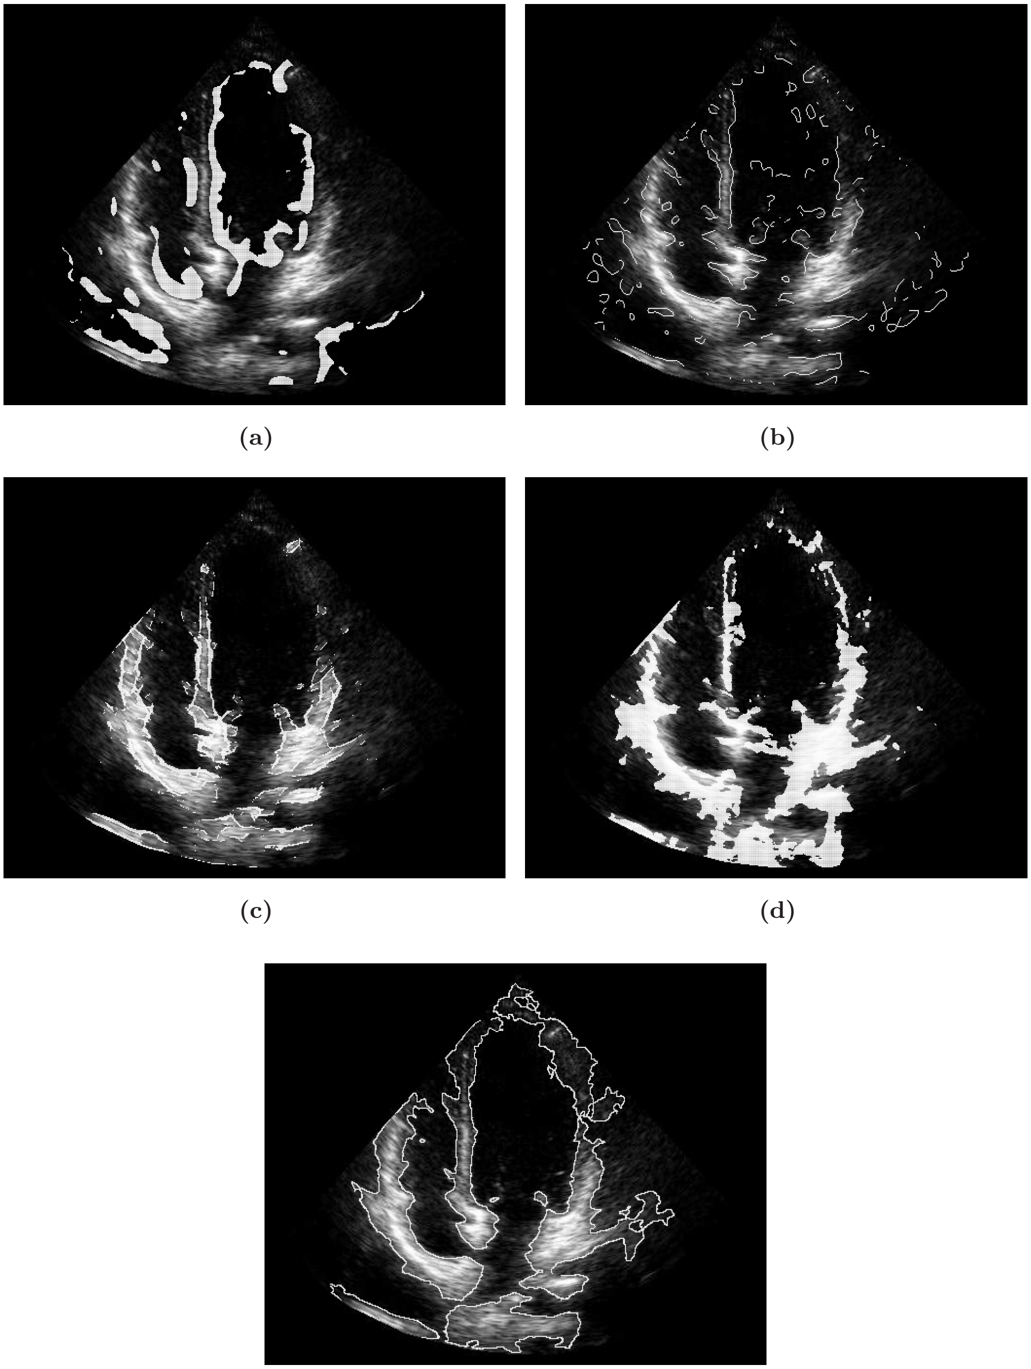

(a) **(b)** **(c)** **(d)** **(e)**

Figure 6.14: Edge detector output for the clinical ultrasound images of Fig. 6.7(e). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

6.3.3 Speckle Reduction using ZC Edges

This section presents the results of the tests detailed in Section 6.2.3, which asses the performance of the zero crossing edge detector based speckle reduction scheme (ZCAD) proposed in Section 5.5.

6.3.4 Simulated Images

Fig. 6.15 displays the result of application of the proposed ZCAD filter to simulated images. These images are also used in the evaluation of Chapter 4. The output of the best performing filters from Chapter 4 (the OSRAD and NMWD filters) is reproduced here as a comparison.

The images filtered by the ZCAD method are seen to have a greater level of smoothing in the different image regions than the other methods, and the edges of the ZCAD filtered images are clearly defined. The use of a thresholded edge function in the proposed ZCAD method results in sharply inhibited diffusion across image borders. Thus a greater level of smoothing can be performed without blurring the image edges. The accuracy of the input edge is seen to be crucial, however. When these detected edges are seen to deviate from the true edge, the diffused image is seen to reflect this deviation. This is to be expected, and as the detected edges do not deviate far from the true edges this effect is considered to be a minor issue. A more serious issue is seen where the detected edge is missing a segment, as in the top right corner of Fig. 6.15(g). In this region the lack of an edge has resulted in blurring due to uninhibited diffusion. In general however the edge maps produced by the Gabor zero crossing technique proposed in this thesis are quite accurate, as demonstrated in the results of Section 6.3.2. The results of diffusion as demonstrated here are observed to produce images with clean well defined borders matching those of the input image, due to the accuracy of the edge detection method.

Table 6.11 displays the SNR_A measures calculated from the simulated ultrasound images. The values produced by the entire filter set examined in Chapter 4 are reproduced in this table, with the added inclusion of the values calculated for the proposed ZCAD filter. These are values averaged over the entire test set of simulated images, totalling twenty images (ten are produced for each template). The average SNR_A in both cases is seen to exceed that of the other speckle filtering techniques, corresponding to a higher level of filtering. The ZCAD filter comes close to matching the SNR_A of the maximally speckle

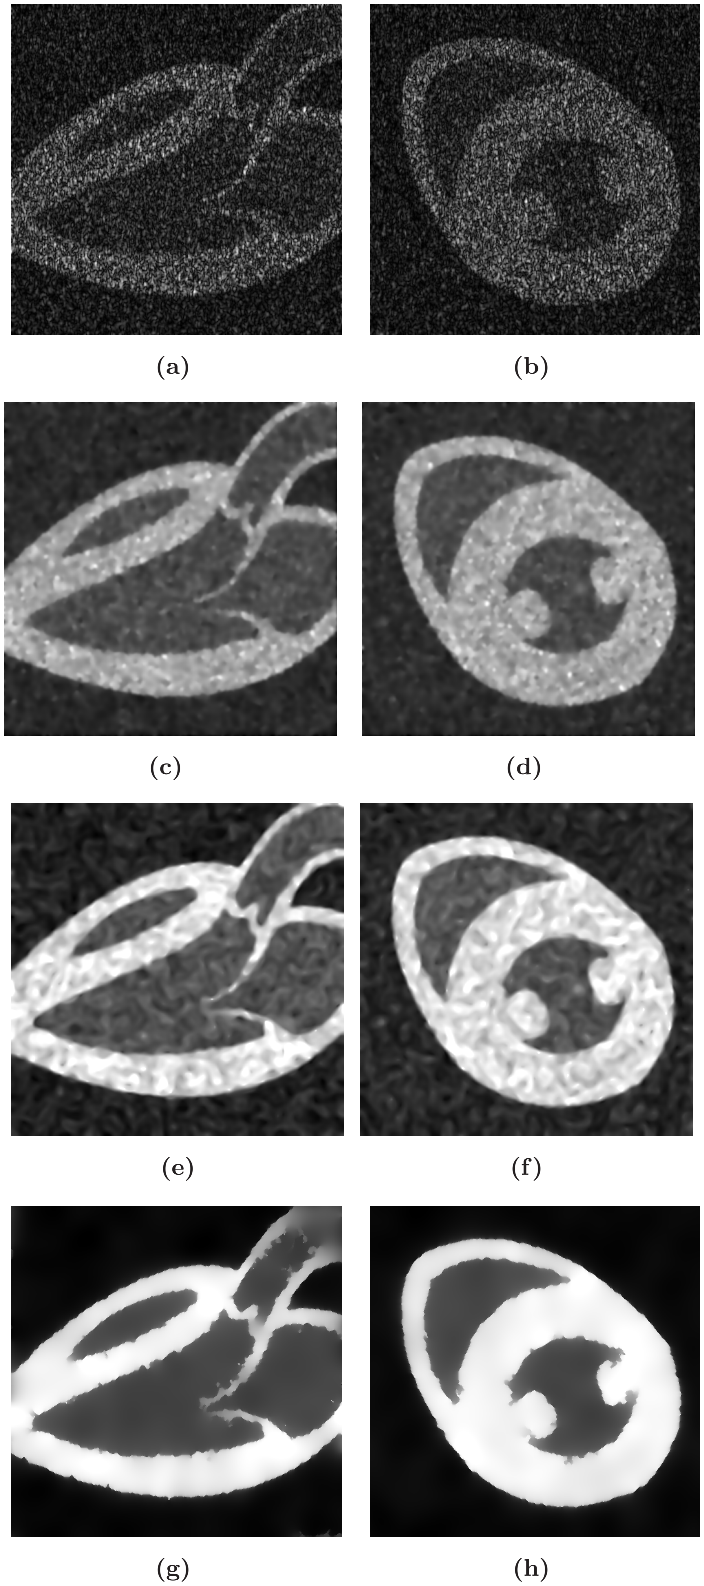

Figure 6.15: Speckle filter output for simulated images, including the proposed ZCAD method. (a)–(b) Original, (c)–(d) NMWD, (e)–(f) OSRAD, (g)–(h) ZCAD.

Table 6.11: Average SNR_A for the simulated echocardiographic images of Figs. 4.3(c)-(d).

Filter	R1(25%)	R2(1%)	R3(5%)	Average	Improvement
MW	5.94	3.02	4.06	4.34	142.46%
ZCAD	**5.27**	**3.15**	**4.02**	**4.15**	**131.84%**
NMWD	3.51	3.49	3.99	3.66	104.47%
OSRAD	3.53	3.13	4.06	3.57	99.44%
SRAD	3.94	2.64	3.81	3.46	93.30%
NCD	3.15	2.77	3.36	3.09	72.63%
DPAD	3.21	2.73	3.24	3.06	70.95%
PMAD	3.38	2.30	3.27	2.98	66.48%
Lee	3.11	2.68	3.16	2.98	65.31%
Geo	2.84	2.26	3.14	2.75	53.63%
EnhFrost	2.89	2.42	2.73	2.68	49.72%
CED	2.67	2.39	2.69	2.58	44.13%
GLM	2.49	2.05	2.76	2.43	35.75%
EnhLee	2.57	2.21	2.41	2.40	34.08%
Zong	2.45	2.15	2.43	2.34	30.73%
Kuan	2.45	2.15	2.34	2.32	29.61%
Frost	2.22	1.93	2.11	2.09	16.76%
Input	1.88	1.69	1.81	1.79	0%

free MW reference image. The high level of speckle suppression as measured by this metric confirms what is observed visually in Fig. 6.15.

Table 6.12 display the FoM, SSIM, CNR and EdgeMSE image quality metrics for the simulated images. Similar to the evaluation performed in Chapter 4, these metrics are measured comparing the images filtered by each method with the MW reference image. These tables reproduce the metrics for the filtering methods, with the ZCAD metrics added for comparison, and highlighted in boldface.

Edge similarity, as measured by the FoM and EdgeMSE values is seen to be favourable in comparison to the other methods. The large FoM values are due to the well defined borders seen in the ZCAD filtered images. These value would likely be even higher if the minor fluctuations in image borders observed were not present. Similarly, the low EdgeMSE values indicate that the output of the proposed ZCAD filter is generally quite similar to the speckle free MW image in regions surrounding image edges. As expected

Table 6.12: Average image quality metrics for speckle filtered simulated echomap images of Figs. 4.3(c)-(d), measured relative to MW reference images. The CNR value is the relative difference with respect to the MW reference. Each column of metric values is arranged from best to worst.

Filter	FoM	Filter	SSIM	Filter	CNR	Filter	EdgeMSE
Zong	0.69	CED	0.20	**ZCAD**	**0.23**	OSRAD	1.624×10^{-4}
ZCAD	**0.67**	**ZCAD**	**0.19**	PMAD	0.21	NCD	1.691×10^{-4}
Kuan	0.63	NCD	0.19	OSRAD	0.15	CED	1.707×10^{-4}
Frost	0.63	OSRAD	0.14	NMWD	0.14	SRAD	3.897×10^{-4}
EnhLee	0.58	SRAD	7.90×10^{-3}	SRAD	0.11	PMAD	4.043×10^{-4}
NCD	0.53	PMAD	4.94×10^{-3}	GLM	0.06	DPAD	4.169×10^{-4}
Lee	0.53	Frost	4.79×10^{-3}	DPAD	0.02	Frost	4.192×10^{-4}
EnhFrost	0.49	Kuan	4.30×10^{-3}	Geo	-0.08	Kuan	4.193×10^{-4}
CED	0.43	EnhLee	4.14×10^{-3}	NCD	-0.10	Lee	4.194×10^{-4}
SRAD	0.38	Zong	3.90×10^{-3}	Lee	-0.19	EnhLee	4.194×10^{-4}
DPAD	0.31	DPAD	3.76×10^{-3}	EnhFrost	-0.32	EnhFrost	4.195×10^{-4}
Geo	0.29	EnhFrost	3.63×10^{-3}	CED	-0.33	NMWD	4.196×10^{-4}
OSRAD	0.27	Lee	2.98×10^{-3}	Zong	-0.43	**ZCAD**	$\mathbf{4.204}\times10^{-4}$
NMWD	0.25	GLM	2.61×10^{-3}	EnhLee	-0.43	Geo	4.210×10^{-4}
PMAD	0.24	NMWD	2.36×10^{-3}	Kuan	-0.45	Zong	4.211×10^{-4}
GLM	0.19	Geo	1.67×10^{-3}	Frost	-0.51	GLM	4.220×10^{-4}

given the large degree of smoothing observed inside the regions of the simulated images, the proposed ZCAD filter displays a large increase in contrast as measured by the CNR metric. This represents an increase in contrast relative to the MW reference image. The structural similarity to the MW reference image as quantified by the SSIM metric is also quite large. Thus it is concluded that in terms of speckle reduction, edge similarity and structural similarity, the proposed ZCAD speckle filter is comparable to the speckle free MW reference image. The contrast of the ZCAD filtered output is however a good deal greater than for the MW reference image.

6.3.5 Clinical Images

Fig. 6.16 displays the result of applying the proposed ZCAD filter to clinical images. These are the same clinical images used to demonstrate the performance of speckle filters in Chapter 4, and as with the simulated image test reported above, the best performing OSRAD and NMWD filters have their output reproduced here for comparison. Similar to the simulated case, the proposed ZCAD filter is observed to produce images with clearly defined edges and a high level of smoothing in the flat image regions. The ZCAD filtered clinical images have experienced a greater level of speckle suppression than the other methods considered. This has led to a greater loss of the false fine detail added to the image by the presence of speckle. However, observation of the borders of the imaged structure in these images displays that the edges of anatomical structures are not blurred by this process.

Table 6.13: Average image quality metrics for speckle filtered clinical amplitude videos, measured relative to speckled input images. The CNR value is the relative difference with respect to the input image. Each column of metric values is arranged from best to worst.

Filter	FoM	Filter	SSIM	Filter	CNR	Filter	EdgeMSE	Filter	SNR_A Change
OSRAD	0.76	Frost	0.92	**ZCAD**	**0.457**	NMWD	6.16×10^{-4}	**ZCAD**	**73.52%**
ZCAD	**0.74**	EnhLee	0.90	OSRAD	0.408	OSRAD	6.21×10^{-4}	NMWD	70.44%
EnhLee	0.73	Zong	0.89	SRAD	0.353	**ZCAD**	$\mathbf{6.93}\times10^{-2}$	OSRAD	62.59%
EnhFrost	0.68	Kuan	0.88	NMWD	0.280	Frost	7.80×10^{-4}	SRAD	54.04%
CED	0.65	**ZCAD**	**0.87**	DPAD	0.276	CED	8.93×10^{-4}	DPAD	39.34%
NCD	0.62	NMWD	0.86	Geo	0.254	EnhFrost	9.54×10^{-4}	Kuan	34.81%
Frost	0.62	CED	0.86	Zong	0.228	Kuan	1.03×10^{-3}	Zong	27.07%
Lee	0.60	EnhFrost	0.84	PMAD	0.221	NCD	1.11×10^{-3}	NCD	25.83%
Kuan	0.57	OSRAD	0.82	NCD	0.191	Lee	1.21×10^{-3}	Lee	23.57%
NMWD	0.56	GLM	0.81	Kuan	0.167	EnhLee	1.34×10^{-3}	EnhFrost	21.45%
GLM	0.54	NCD	0.81	EnhFrost	0.110	GLM	1.51×10^{-3}	GLM	20.66%
DPAD	0.53	Lee	0.79	GLM	0.107	DPAD	1.70×10^{-3}	EnhLee	19.28%
Zong	0.48	DPAD	0.72	Lee	0.105	SRAD	1.80×10^{-3}	CED	18.68%
SRAD	0.44	SRAD	0.69	EnhLee	0.097	Zong	1.84×10^{-3}	PMAD	18.44%
Geo	0.39	Geo	0.59	CED	0.086	PMAD	7.88×10^{-3}	Frost	15.68%
PMAD	0.13	PMAD	0.46	Frost	0.066	Geo	1.20×10^{-2}	Geo	13.54%

The calculation of the image quality metrics for the proposed ZCAD on the clinical images produced the values highlighted in boldface in Table 6.13. This is the same test as that of Chapter 4, in which a total of 500 clinical frames are processed. The metrics are measured relative to the speckle input images, and are averaged over the entire clinical test set. The values of the same metrics computed for the range of speckle filters compared in Chapter 4 are reproduced here for comparison.

The level of speckle suppression as measured by the SNR_A metric is seen to be highest for the proposed ZCAD filter. This confirms what is noted visually. Edge preservation, as measured by the FoM metric is seen to be high. This indicates that the proposed speckle reduction method does not significantly alter the edges of the clinical image through filtering. The EdgeMSE results are relatively quite high however, indicating that filtering alters the intensities of pixels close to the edges. Therefore although the proposed method preserves image edges, it does perform filtering in the vicinity (i.e. on either side of) image edges. The proposed filter is seen to have relatively high similarity in its output to the speckled input, as measured by the SSIM metric. As with the simulated images, the proposed method is achieves a larger degree of improvement in contrast as measured by the CNR metric than the other filtering approaches considered.

Thus the proposed filtering technique is capable of removing the speckle in clinical images to a high degree. This does not come at the cost of decreased edge region preservation. In addition, the images filtered by this method are reasonably similar in a structural sense to the input images (as quantified by the SSIM metric) and also display an improved level of contrast.

6.3.6 Computational Requirements

The computational requirement analysis of the ZCAD filter results in the values presented in Table 6.14, in terms of the parameters of the filter. It is observed that the computational requirements of the filter are a complex function of the parameters used to determine its computation. This is due to the different steps involved in the process. If reasonable average values are taken for these parameters, absolute values for a given situation can be determined. This is performed by examining application to an image of size 512×512 pixels, similar to the evaluation of Chapter 4. The other parameters are chose with values as follows: The number of scales and orientations in the log Gabor filter bank are both set to two ($N_s = N_o = 2$), as applied here. The number of diffusion iterations,

Table 6.14: Computational requirements for the ZCAD filter. N = number of pixels in image, Q = number of diffusion iterations. N_s and N_o are the number of Gabor scales and orientations. K is the length of the characteristic function used in parameter estimation, and Z is the number of iterations of the minimisation function. N_m and N_p are the number of points used to calculate the MAP shrinkage curves, and the PDFs used in their computation.

Multiplications/Divisions	**Additions/Subtractions**	**Look Up Operations**
$5N + (5N+2)N_s + 8NN_o + N\log_2 N(1 + N_sN_o) + NsN_o\{(2N+2)K + (35K+14)Z + N_mN_p + N_p\log_2 N_p + 6N_p + N_z^2(\frac{1}{\sqrt{N}} + 3) + 10\} + 12NQ$	$N + 2NN_o + (1 + N_sN_o)\times(\frac{3}{2}N\log_2 N + \frac{1}{2}N) + NsN_o\{(21K - 14)Z + \frac{3}{2}N_p\log_2 N_p + \frac{1}{2}N_p + 5N + N_z^2(\frac{1}{\sqrt{N}} + 3) + 1\} + (12N - 8)Q$	$3N + (N+1)N_s + (2N+4)N_o + NsN_o\{7KZ + 2N_p + N\}$

Q, is selected as 20. Estimation of the parameters of the Gabor model are performed using characteristic functions of length $K = 512$, while 5 applications of the minimisation algorithm, each lasting 150 iterations, give a total of $Z = 750$. MAP estimation is performed using PDFs of length $N_p = 1024$ to generate MAP curves with $N_m = N/20$ points. As a worst case situation, the number of zero crossings (for which the orientational spacing and interscale agreement are computed) are assumed to be $N/25$. This assumes that every 25th pixel contains a zero crossing, which exceed what is observed experimentally.

The results of this absolute complexity are presented in Table 6.15. It can be seen that the proposed method has a higher requirement in terms of both multiplications and additions then most of the filters considered in Chapter 4. However, the NMWD method requires greater computational effort then the proposed method. It is noted that the largest contribution to the complexity of the ZCAD filter is the computation of the interscale agreement and orientational spacing measures, which introduce $O(N^2)$ terms to the computational overhead.

Table 6.15: Filter Computational Requirements, for $N = 512 \times 512$ pixels. Filters are ordered in decreasing complexity, as measured by relative number of multiplications.

Filter	**Mult/Div**	**Add/Sub**	**LUT**	**Relative**
NMWD	2.80×10^9	1.80×10^9	2.70×10^8	1
ZCAD	2.58×10^9	1.46×10^9	1.42×10^7	0.92
DPAD	6.20×10^8	9.30×10^8	-	0.22
SRAD	3.50×10^8	3.40×10^8	-	0.13
GLM	3.00×10^8	2.80×10^8	7.90×10^6	0.11
OSRAD	1.30×10^8	1.10×10^8	0	4.60×10^{-2}
Frost	9.10×10^7	8.90×10^7	2.10×10^7	3.30×10^{-2}
EnhFrost	9.10×10^7	8.90×10^7	2.10×10^7	3.30×10^{-2}
CED	8.70×10^7	7.30×10^7	1.30×10^6	3.10×10^{-2}
NCD	7.90×10^7	6.80×10^7	-	2.80×10^{-2}
Zong	7.50×10^7	7.40×10^7	2.10×10^6	2.70×10^{-2}
PMAD	6.00×10^7	5.00×10^7	2.90×10^6	2.20×10^{-2}
EnhLee	2.80×10^7	3.80×10^7	2.60×10^5	1.00×10^{-2}
Kuan	2.80×10^7	2.60×10^7	-	9.90×10^{-3}
Lee	1.80×10^7	1.70×10^7	-	6.40×10^{-3}
Geo	0	8.80×10^7	-	-

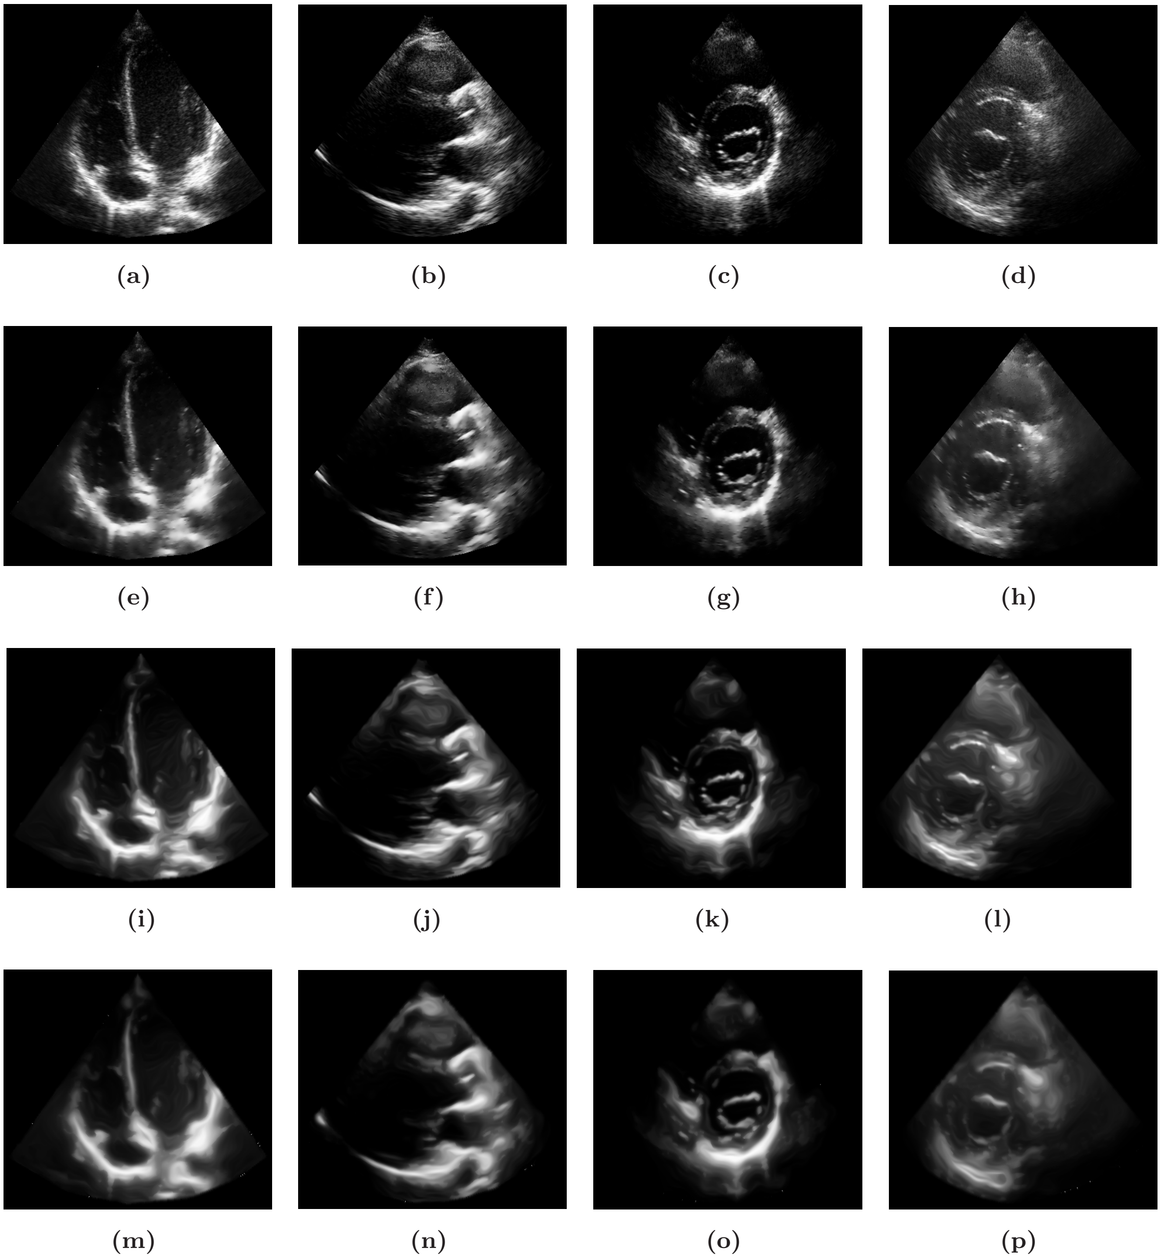
(a)
(b)
(c)
(d)
(e)
(f)
(g)
(h)
(i)
(j)
(k)
(l)
(m)
(n)
(o)
(p)

Figure 6.16: Speckle filter output for clinical images, including the proposed ZCAD method. (a)–(d) Original, (e)–(h) NMWD, (i)–(l) OSRAD, (m)–(p) ZCAD.

6.4 Conclusions

This chapter has detailed the evaluation of the performance of the technique proposed in Chapter 5, for edge detection and speckle removal in speckled imagery.

Evaluation is performed on the accuracy of the proposed method of parameter estimation, using an extensive Monte Carlo simulation. Error values for each of the parameters are presented, and these are seen to compare favourably with reported upon in the literature for the related task of estimating SαS parameters.

The Gabor zero crossing based method of edge detection is also evaluated, by comparing it to a number of other methods from the literature. Sets of simulated, phantom, and clinical ultrasound images are used in this test, and the results in each case are presented. For the simulated and phantom image sets ideal edge maps are available, allowing quantification of performance with the widely used FoM metric. In evaluation by both visual assessment and FoM quantification, the edge maps of the proposed method are seen to offer better performance than the other considered. Of the other edge detection methods considered, the FA and PC method also employ Gabor filtering. While the FA method is seen to perform relatively well, the PC method is not seen to produce accurate edge maps for these speckled images. It is concluded that the proposed method of edge detection is superior to the approaches presented previously in the literature, for use in speckled ultrasound imagery.

When the edge detection method presented in this chapter is used as part of a diffusion scheme for speckle reduction, it is shown to compare favourably with the OSRAD and NMWD filters (deemed to be the best performing in the study of Chapter 4). This demonstrates that the proposed edge detector, once applied to an image, can also serve as the basis for an effective speckle removal algorithm. Unlike the edge detection schemes used in other anisotropic diffusion approaches, this edge detector is sufficiently accurate to not require recalculation of the edge map as speckle is iteratively removed. Thus edge detection is performed only once, and speckle removal can be performed as an optional additional task which requires no further modification to the edge map. While calculation of the edge map is required for diffusion, the edge map could be used instead for other tasks such as edge enhancement or segmentation. Separating the calculation of the image edge map and the performance of diffusion computationally as in the proposed ZCAD method mirrors the functional separation which can exist between these two tasks.

Chapter 7

Conclusions and Future Work

7.1 Summary of Thesis

This thesis presents the development and evaluation of image processing techniques for the reduction of speckle in ultrasonic imagery. The research presented is motivated by the widespread use of ultrasound imagery in many field of medicine, and the open nature of the problem of speckle removal.

A comprehensive strategy for the evaluation of the performance of speckle filtering techniques has been presented. This involves testing filters by application to a large test set of clinical video frames, as well as a set of realistically simulated images. A set of objective image quality metrics are applied to the filtered clinical and simulated images. The relationship between objective image quality metrics and subjective evaluation by clinical experts is investigated using statistical analysis. As well testing functional performance, filtering methods are compared in terms of their computational requirements. The evaluation strategy presented is applied to a wide range of speckle reducing filters from the literature.

The development of edge detection and speckle reduction methods employs orientation specific multiscale decomposition by Gabor filtering. The effect of speckle in the Gabor domain is removed by MAP estimation. This method includes the use of an accurate model for the speckled Gabor coefficients. A method is proposed for the estimation of the parameters of this model, and the accuracy of this parameter estimation is verified by Monte Carlo simulation.

An accurate edge detection method for speckle corrupted ultrasound is developed, which is based on the quantification of observed properties of the zero crossings of the

Gabor transform coefficient. These properties are demonstrated using simulated and clinical images, and methods for their measurement are presented. The accuracy of the proposed edge detection method is quantified by comparison to a number of other edge detection methods from the literature.

The edge detection method presented in this thesis is then used to construct a method of speckle suppression. Edge maps generated using the proposed method are incorporated into a tensor valued anisotropic diffusion scheme. This filter is then compared to those in the literature by the comprehensive evaluation strategy proposed in this thesis.

7.2 Contributions

Based on the research presented in this thesis, the following conclusions have been drawn and contributions made:

1. A comprehensive study which correlates objective quality measures with subjective clinical opinion has been carried out. In this study, it is found that the experts generally associate speckle reduction filtering with a reduction in “overall quality”. It is unclear however if this is due to a true decrease in diagnostic usefulness, or due to the perceived reduction in sharpness noted by Dantas *et al.* [5] which results from the removal of false detail due to speckle.

2. Several objective image quality metrics (FoM, MSE and EdgeMSE) are shown to be good indicators of subjective expert assessment, and therefore may be appropriate for use where subjective opinion is not available. This work has resulted in one international conference paper:

   S. Finn, E. Jones, and M. Glavin, “Objective and subjective evaluations of quality for speckle reduced echocardiography”, in *Proc. IEEE Intl. Annu. Conf. Engineering Medicine and Biology Society*, Sep. 2009, pp. 503-506. doi: 10.1109/IEMBS.2009.5333572.

3. The comparative review of speckle filters in Chapter 4 presents a number of a results relating to the evaluation of filters from the literature. Functional performance in terms of the level of speckle suppression, edge preservation, contrast enhancement and structural similarity varies throughout the filters. This is the first time such a comprehensive evaluation and comparison of different methods has been carried

out. Anisotropic diffusion based methods are seen to have in general the strongest suppression of speckle. In particular the OSRAD diffusion method and the NMWD multiscale diffusion approach are seen to have the best performance, combining strong speckle suppression with favourable detail preservation. The analysis of computational requirements shows that local statistics methods are very efficient. Anisotropic diffusion filters can also be implemented in an efficient manner, depending on the choice of discretisation scheme. This work has resulted in one journal paper accepted for publication:

S. Finn, E. Jones, and M. Glavin, "Echocardiographic Speckle Reduction Comparison", *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, 2010.

4. Application of orientation specific multiscale image decomposition using the Gabor transform is shown to be modelled accurately using a combination of Gaussian and SαS distributions.

5. A novel method of parameter estimation is proposed, and is shown to jointly estimate all parameters of this model with an accuracy matching other methods.

6. Edge detection of speckled ultrasound images, with the proposed Gabor zero crossing method is shown to be superior to a range of commonly used approaches in the literature. This is evaluated in an extensive test by application to simulated, clinical and phantom images. Edge detection accuracy is measured using the FoM metric.

7. When the proposed edge detector is incorporated into a anisotropic diffusion for speckle suppression, it is shown that the accurate edge detection step allows a large degree of smoothing to reduce speckle, without blurring of image edges. A further benefit of the method is that it does not require iterative generation of the edge map for speckle reduction. On these criteria the proposed speckle reduction filter is deemed to be superior to those seen so far in the literature.

7.3 Suggestions for Future Work

The research presented in this thesis is comprised of a number of related topics, many of which could be extended further.

The study presented in Chapter 3 compares objective metric and subjective expert analysis for ultrasound imaging, and specifically the echocardiographic application. Similarly the evaluation framework presented in Chapter 4 compares a wide range of speckle reduction techniques using an echocardiographic data set. While the results presented may also be applicable to other clinical ultrasound applications such as breast, liver and fetal imaging, further investigation into these areas is necessary. Additional potential areas in which this investigation could be extended include testing on both simulated and real 3D ultrasound data, as well as the evaluation of recent filtering methods differing in approach to those considered here, such as [55].

In the research presented in Chapter 5, a number of possible extensions could be investigated. The parameter estimation technique could be extended by consideration of t_k values with non-constant spacing for parameter estimation. This has been shown to improve the parameter estimation method of [164] in [172].

As presented in the computational analysis of the ZCAD speckle filter, this proposed method has a high computational overhead. The functional method as presented is not optimised for efficiency of calculation, and the main contribution to its overhead are the steps involved in edge detection. A number of approaches could be considered to improve the efficiency of the proposed edge detection method. The application of Gabor filters is itself a relatively computationally demanding task. The relatively high computational cost of Gabor filtering may be addressed by inclusion of an efficient implementation such as in [180].

The distance calculations in the orientational spacing and interscale agreement measurements are quite inefficient in their current form. These are variations on a class of image processing algorithms known as distance transforms. Efficient approximations to the Euclidean distances computed in this thesis could be used instead, including that of Bailey [181].

The use of the proposed zero crossing based edge detector for higher level tasks such as segmentation is a possible avenue of future research.

Bibliography

[1] W. D. Middleton and A. B. Kurtz, *Ultrasound: The Requisites (Requisites in Radiology Series)*, 2nd ed. Philadelphia, PA.: Elsevier Mosby, 2004.

[2] X. Zong, A. Laine, and E. Geiser, "Speckle reduction and contrast enhancement of echocardiograms via multiscale nonlinear processing," *IEEE Trans. Med. Imag.*, vol. 17, no. 4, pp. 532–540, Aug. 1998.

[3] Y. Yu and S. Acton, "Speckle reducing anisotropic diffusion," *IEEE Trans. Image Process.*, vol. 11, no. 11, pp. 1260–1270, Nov. 2002.

[4] K. Abd-Elmoniem, A.-B. Youssef, and Y. Kadah, "Real-time speckle reduction and coherence enhancement in ultrasound imaging via nonlinear anisotropic diffusion," *IEEE Trans. Bio-Med. Eng.*, vol. 49, no. 9, pp. 997–1014, Sep. 2002.

[5] R. Dantas and E. Costa, "Ultrasound speckle reduction using modified Gabor filters," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 54, no. 3, pp. 530–538, Mar. 2007.

[6] F. Zhang, Y. M. Yoo, L. M. Koh, and Y. Kim, "Nonlinear diffusion in Laplacian pyramid domain for ultrasonic speckle reduction," *IEEE Trans. Med. Imag.*, vol. 26, no. 2, pp. 200–211, Feb. 2007.

[7] R. G. Dantas, E. T. Costa, and S. Leeman, "Ultrasound speckle and equivalent scatterers," *Ultrasonics*, vol. 43, no. 6, pp. 405–420, 2005.

[8] J. C. Bamber and C. Daft, "Adaptive filtering for reduction of speckle in ultrasonic pulse-echo images." *Ultrasonics*, pp. 41–44, Jan. 1986.

[9] C. Burckhardt, "Speckle in ultrasound B-Mode scans," *IEEE Trans. Sonics Ultrason.*, vol. 25, no. 1, pp. 1–6, Jan. 1978.

[10] C. Rumack, S. Wilson, J. Charboneau, and J. Johnson, *Diagnostic ultrasound*, 3rd ed. Philadelphia, PA.: Elsevier Mosby, 2004.

[11] J. Quistgaard, "Signal acquisition and processing in medical diagnostic ultrasound," *IEEE Signal Process. Mag.*, vol. 14, no. 1, pp. 67–74, Jan. 1997.

[12] J. M. Thijssen, "Ultrasonic speckle formation, analysis and processing applied to tissue characterization," *Pattern Recogn. Lett.*, vol. 24, no. 4-5, pp. 659–675, 2003.

[13] M. Ali, D. Magee, and U. Dasgupta, "Signal processing overview of ultrasound systems for medical imaging," Texas Instruments, Dallas, TX., White paper, Nov. 2008. [Online]. Available: http://focus.ti.com/lit/wp/sprab12/sprab12.pdf

[14] J. Seabra and J. Sanches, "Modeling log-compressed ultrasound images for radio frequency signal recovery," in *Proc. IEEE Intl. Annu. Conf. Engineering Medicine and Biology Society*, Aug. 2008, pp. 426–429.

[15] R. J. Massay, R. B. Logan-Sinclair, J. C. Bamber, and D. G. Gibson, "Quantitative effects of speckle reduction on cross sectional echocardiographic images." *Br. Heart J.*, vol. 62, no. 4, pp. 298–304, 1989.

[16] J. Goodman, "Statistical properties of laser speckle patterns," in *Laser Speckle and Related Phenomena*, ser. Topics in Applied Physics, J. Dainty, Ed. Berlin: Springer, 1975, vol. 9, pp. 9–75.

[17] J. W. Goodman, "Some fundamental properties of speckle," *J. Opt. Soc. Am.*, vol. 66, no. 11, pp. 1145–1150, 1976.

[18] R. Wagner, S. Smith, J. Sandrik, and H. Lopez, "Statistics of Speckle in Ultrasound B-Scans," *IEEE Trans. Sonics Ultrason.*, vol. 30, no. 3, pp. 156–163, May 1983.

[19] R. Wagner, M. Insana, and S. Smith, "Fundamental correlation lengths of coherent speckle in medical ultrasonic images," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 35, no. 1, pp. 34–44, Jan. 1988.

[20] P. Shankar, V. Dumane, J. Reid, V. Genis, F. Forsberg, C. Piccoli, and B. Goldberg, “Classification of ultrasonic b-mode images of breast masses using nakagami distribution,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 48, no. 2, pp. 569–580, Mar. 2001.

[21] P. M. Shankar, “A model for ultrasonic scattering from tissues based on the K distribution,” *Phys. Med. Biol.*, vol. 40, no. 10, pp. 1633–1649, 1995.

[22] R. Smolíovái, M. P. Wachowiak, and J. M. Zurada, “An information-theoretic approach to estimating ultrasound backscatter characteristics,” *Comput. Biol. Med.*, vol. 34, no. 4, pp. 355–370, Jun. 2004.

[23] F. L. Lizzi, M. Astor, T. Liu, C. Deng, D. J. Coleman, and R. H. Silverman, “Ultrasonic spectrum analysis for tissue assays and therapy evaluation,” *Int. J. Imag. Syst. Tech.*, vol. 8, no. 1, pp. 3–10, 1997.

[24] K. Krissian, C.-F. Westin, R. Kikinis, and K. Vosburgh, “Oriented speckle reducing anisotropic diffusion,” *IEEE Trans. Image Process.*, vol. 16, no. 5, pp. 1412–1424, May 2007.

[25] J. Ng, R. Prager, N. Kingsbury, G. Treece, and A. Gee, “Modeling ultrasound imaging as a linear, shift-variant system,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 53, no. 3, pp. 549–563, Mar. 2006.

[26] Y. Kadah, A. Farag, J. Zurada, A. Badawi, and A.-B. Youssef, “Classification algorithms for quantitative tissue characterization of diffuse liver disease from ultrasound images,” *IEEE Trans. Med. Imag.*, vol. 15, no. 4, pp. 466–478, Aug. 1996.

[27] G. E. Trahey, J. W. Allison, and O. T. von Ramm, “Angle independent ultrasonic detection of blood flow,” *IEEE Trans. Biomed. Eng.*, vol. BME-34, no. 12, pp. 965–967, Dec. 1987.

[28] L. De Marchi, N. Testoni, and N. Speciale, “Prostate tissue characterization via ultrasound speckle statistics,” in *IEEE Int. Symp. Signal Processing Information Technology*, Aug. 2006, pp. 208–211.

[29] P.-H. Tsui, S.-H. Wang, and C.-C. Huang, "The effect of logarithmic compression on estimation of the nakagami parameter for ultrasonic tissue characterization: a simulation study," *Phys. Med. Biol.*, vol. 50, no. 14, pp. 3235–3244, 2005.

[30] R. L. Maurice, E. Brusseau, G. Finet, and G. Cloutier, "On the potential of the Lagrangian speckle model estimator to characterize atherosclerotic plaques in endovascular elastography: In vitro experiments using an excised human carotid artery," *Ultrasound Med. Biol.*, vol. 31, no. 1, pp. 85–91, 2005.

[31] G. E. Trahey, J. W. Allison, S. W. Smith, and O. T. von Ramm, "A quantitative approach to speckle reduction via frequency compounding," *Ultrasonic Imaging*, vol. 8, pp. 151–164, 1986.

[32] R. L. Galloway, B. A. McDermott, and F. L. Thurstone, "A frequency diversity process for speckle reduction in real-time ultrasonic images," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, pp. 45–49, 1988.

[33] S. M. Gehlbach and F. G. Sommer, "Frequency diversity speckle processing," *Ultrasonic Imaging*, vol. 9, pp. 92–105, 1987.

[34] P. A. Magnin, O. T. von Ramm, and F. L. Thurstone, "Frequency compounding for speckle contrast reduction in phased array images," *Ultrasonic Imaging*, vol. 4, pp. 267–281, 1982.

[35] V. Newhouse, N. Bilgutay, J. Saniie, and E. Furgason, "Flaw-to-grain echo enhancement by split-spectrum processing," *Ultrasonics*, vol. 20, no. 2, pp. 59–68, 1982.

[36] J. Bamber and J. Phelps, "Real-time implementation of coherent speckle suppression in B-scan images," *Ultrasonics*, vol. 29, no. 3, pp. 218–224, 1991.

[37] P. Stetson, F. Sommer, and A. Macovski, "Lesion contrast enhancement in medical ultrasound imaging," *IEEE Trans. Med. Imag.*, vol. 16, no. 4, pp. 416–425, Aug. 1997.

[38] G. E. Trahey, S. W. Smith, and O. T. von Ramm, "Speckle reduction in medical ultrasound via spatial compounding," in *Proc. SPIE Medicine XIV PACS*, vol. 4, 1986, pp. 629–637.

[39] L. Pai-Chi and M. O'Donnell, "Elevational spatial compounding," *Ultrasonic imaging*, vol. 16, no. 3, pp. 176–189, 1994.

[40] M. O'Donnell and S. Silverstein, "Optimum displacement for compound image generation in medical ultrasound," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 35, no. 4, pp. 470–476, Jul. 1988.

[41] D. Adam, S. Beilin-Nissan, Z. Friedman, and V. Behar, "The combined effect of spatial compounding and nonlinear filtering on the speckle reduction in ultrasound images," *Ultrasonics*, vol. 44, no. 2, pp. 166–181, 2006.

[42] J.-S. Lee, "Digital image enhancement and noise filtering by use of local statistics," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-2, no. 2, pp. 165–168, Mar. 1980.

[43] ——, "Speckle analysis and smoothing of synthetic aperture radar images," *Comput. Vision Graph.*, vol. 17, no. 1, pp. 24–32, 1981.

[44] D. Kuan, A. Sawchuk, T. Strand, and P. Chavel, "Adaptive restoration of images with speckle," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. ASSP-35, no. 3, pp. 373–383, Mar. 1987.

[45] V. S. Frost, J. A. Stiles, K. S. Shanmugan, and J. C. Holtzman, "A model for radar images and its application to adaptive digital filtering of multiplicative noise," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 4, no. 2, pp. 157–166, Mar. 1982.

[46] A. Lopes, R. Touzi, and E. Nezry, "Adaptive speckle filters and scene heterogeneity," *IEEE Trans. Geosci. Remote Sens.*, vol. 28, no. 6, pp. 992–1000, Nov. 1990.

[47] T. Loupas, W. McDicken, and P. Allan, "An adaptive weighted median filter for speckle suppression in medical ultrasonic images," *IEEE Trans. Circuits Syst.*, vol. 36, no. 1, pp. 129 –135, Jan. 1989.

[48] Y. Chen, R. Yin, P. Flynn, and S. Broschat, "Aggressive region growing for speckle reduction in ultrasound images," *Pattern Recogn. Lett.*, vol. 24, no. 4-5, pp. 677–691, 2003.

[49] M. Karaman, M. Kutay, and G. Bozdagi, "An adaptive speckle suppression filter for medical ultrasonic imaging," *IEEE Trans. Med. Imag.*, vol. 14, no. 2, pp. 283–292, Jun. 1995.

[50] J. I. Koo and S. B. Park, "Speckle reduction with edge preservation in medical ultrasonic images using a homogeneous region growing mean filter (HRGMF)," *Ultrasonic Imaging*, vol. 13, no. 3, pp. 211–237, 1991.

[51] H.-C. Huang, J.-Y. Chen, S.-D. Wang, and C.-M. Chen, "Adaptive ultrasonic speckle reduction based on the slope-facet model," *Ultrasound Med. Biol.*, vol. 29, no. 8, pp. 1161–1175, 2003.

[52] P. Tay, S. Acton, and J. Hossack, "A stochastic approach to ultrasound despeckling," in *Proc. 3rd IEEE Int. Symp. Biomedical Imaging: Nano to Macro*, Apr. 2006, pp. 221–224.

[53] ——, "Ultrasound despeckling using an adaptive window stochastic approach," in *Proc. IEEE Int. Conf. Image Processing*, Oct. 2006, pp. 2549–2552.

[54] J. A. Jensen, "FIELD: A program for simulating ultrasound systems," in *Proc. 10th NordicBaltic Conference Biomedical Imaging*, vol. 4, no. 1, 1996, pp. 351–353.

[55] P. Coupé, P. Hellier, C. Kervrann, and C. Barillot, "Nonlocal means-based speckle filtering for ultrasound images," *IEEE Trans. Image Process.*, vol. 18, no. 10, pp. 2221–2229, Oct. 2009.

[56] D. Sakrison, "On the role of the observer and a distortion measure in image transmission," *IEEE Trans. Commun.*, vol. 25, no. 11, pp. 1251–1267, Nov. 1977.

[57] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 7, pp. 629–639, Jul. 1990.

[58] F. Catté, P.-L. Lions, J.-M. Morel, and T. Coll, "Image selective smoothing and edge detection by nonlinear diffusion," *SIAM J. Numer. Anal.*, vol. 29, no. 1, pp. 182–193, 1992.

[59] W. K. Pratt, *Digital Signal Processing.* New York: Wiley, 1977.

[60] S. Aja-Fernandez and C. Alberola-Lopez, "On the estimation of the coefficient of variation for anisotropic diffusion speckle filtering," *IEEE Trans. Image Process.*, vol. 15, no. 9, pp. 2694–2701, Sep. 2006.

[61] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

[62] J. Weickert, *Anisotropic Diffusion in Image Processing.* Stuttgart, Germany: Teubner-Verlag, 1998, out of print.

[63] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Trans. Inf. Theory*, vol. 41, no. 3, pp. 613–627, May 1995.

[64] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D*, vol. 60, no. 1-4, pp. 259–268, 1992.

[65] L. Rudin, P.-L. Lions, and S. Osher, "Multiplicative denoising and de-blurring: Theory and algorithms," in *Geometric Level Set Methods in Imaging, Vision and Graphics.* New York: Springer-Verlag, 2003, ch. 6, pp. 103–120.

[66] K. Krissian, "Flux-based anisotropic diffusion applied to enhancement of 3-d angiogram," *IEEE Trans. Med. Imag.*, vol. 21, no. 11, pp. 1440–1442, Nov. 2002.

[67] P. Burt and E. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Trans. Commun.*, vol. 31, no. 4, pp. 532–540, Apr. 1983.

[68] E. Adelson, C. Anderson, J. Bergen, P. Burt, and J. Ogden, "Pyramid methods in image processing," *RCA Engineer*, vol. 29, no. 6, 1984.

[69] H. Guo, J. Odegard, M. Lang, R. Gopinath, I. Selesnick, and C. Burrus, "Wavelet based speckle reduction with application to SAR based ATD/R," in *Proc. IEEE Int. Conf. Image Processing*, vol. 1, Nov. 1994, pp. 75–79.

[70] P. Moulin, "A wavelet regularization method for diffuse radar-target imaging and speckle-noise reduction," vol. 3, no. 1, pp. 123–134, 1993.

[71] X. Hao, S. Gao, and X. Gao, "A novel multiscale nonlinear thresholding method for ultrasonic speckle suppressing," *IEEE Trans. Med. Imag.*, vol. 18, no. 9, pp. 787–794, Sep. 1999.

[72] S. M. Collins, D. J. Skorton, E. A. Geiser, J. A. Nichols, D. A. Conetta, N. G. Pandian, and R. E. Kerber, "Computer-assisted edge detection in two-dimensional

echocardiography: Comparison with anatomic data," *Am. J. Cardiol.*, vol. 53, no. 9, pp. 1380–1387, 1984.

[73] J. M. Lester, J. F. Brenner, and W. D. Selles, "Local transforms for biomedical image analysis," *Comput. Vision Graph.*, vol. 13, no. 1, pp. 17–30, 1980.

[74] N. Gupta, M. Swamy, and E. Plotkin, "Despeckling of medical ultrasound images using data and rate adaptive lossy compression," *IEEE Trans. Med. Imag.*, vol. 24, no. 6, pp. 743–754, Jun. 2005.

[75] A. Achim, A. Bezerianos, and P. Tsakalides, "Novel Bayesian multiscale method for speckle removal in medical ultrasound images," *IEEE Trans. Med. Imag.*, vol. 20, no. 8, pp. 772–783, Aug. 2001.

[76] A. K. Jain, *Fundamentals of digital image processing.* Upper Saddle River, NJ.: Prentice-Hall, 1989.

[77] F. Sattar, L. Floreby, G. Salomonsson, and B. Lovstrom, "Image enhancement based on a nonlinear multiscale method," *IEEE Trans. Image Process.*, vol. 6, no. 6, pp. 888–895, Jun. 1997.

[78] H. Rabbani, M. Vafadust, P. Abolmaesumi, and S. Gazor, "Speckle noise reduction of medical ultrasound images in complex wavelet domain using mixture priors," *IEEE Trans. Bio-Med. Eng.*, vol. 55, no. 9, pp. 2152–2160, Sep. 2008.

[79] X.-W. Fu, M.-Y. Ding, and C. Cai, "Despeckling of medical ultrasound images based on quantum-inspired adaptive threshold," *Electron. Lett.*, vol. 46, no. 13, pp. 889–891, Jun. 2010.

[80] A. Pižurica, W. Philips, I. Lemahieu, and M. Acheroy, "A versatile wavelet domain noise filtration technique for medical imaging," *IEEE Trans. Med. Imag.*, vol. 22, no. 3, pp. 323–331, Mar. 2003.

[81] Y. Yue, M. Croitoru, A. Bidani, J. Zwischenberger, and J. J. Clark, "Nonlinear multiscale wavelet diffusion for speckle suppression and edge enhancement in ultrasound images," *IEEE Trans. Med. Imag.*, vol. 25, no. 3, pp. 297–311, Mar. 2006.

[82] B. Aiazzi, L. Alparone, and S. Baronti, "Multiresolution local-statistics speckle filtering based on a ratio Laplacian pyramid," *IEEE Trans. Geosci. Remote Sens.*, vol. 36, no. 5, pp. 1466–1476, Sep. 1998.

[83] M. Nagao and T. Matsuyama, "Edge preserving smoothing," *Comput. Vision Graph.*, vol. 9, no. 4, pp. 394–407, 1979.

[84] A. Thakur and R. Anand, "Image quality based comparative evaluation of wavelet filters in ultrasound speckle reduction," *Digit. Signal Process.*, vol. 15, no. 5, pp. 455–465, 2005.

[85] V. Behar, D. Adam, and Z. Friedman, "A new method of spatial compounding imaging," *Ultrasonics*, vol. 41, no. 5, pp. 377–384, 2003.

[86] V. Aurich and J. Weule, "Non-linear Gaussian filters performing edge preserving diffusion," in *Proc. 17th DAGM Symp.* Springer-Verlag, 1995.

[87] Y. Chen, S. L. Broschat, Patrick, and J. Flynn, "Phase insensitive homomorphic image processing for speckle reduction, ultrason," *Ultrasonic Imaging*, vol. 18, pp. 122–139, 1996.

[88] J. L. Mateo and A. Fernández-Caballero, "Finding out general tendencies in speckle noise reduction in ultrasound images," *Expert Syst. Appl.*, vol. 36, no. 4, pp. 7786–7797, 2009.

[89] C. P. Loizou, C. S. Pattichis, C. I. Christodoulou, R. S. H. Istepanian, M. Pantziaris, and A. Nicolaides, "Comparative evaluation of despeckle filtering in ultrasound imaging of the carotid artery," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 52, no. 10, pp. 1653–1669, Oct. 2005.

[90] T. Huang, G. Yang, and G. Tang, "A fast two-dimensional median filtering algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 27, no. 1, pp. 13–18, feb. 1979.

[91] S. Ali and R. Burge, "New automatic techniques for smoothing and segmenting sar images," *Signal Process.*, vol. 14, no. 4, pp. 335–346, 1988.

[92] L. J. Busse, T. R. Crimmins, and J. R. Fienup, "A model based approach to improve the performance of the geometric filtering speckle reduction algorithm," in *Proc. IEEE Ultrason. Symp.*, vol. 2, Nov. 1995, pp. 1353–1356.

[93] J. Saniie, T. Wang, and N. M. Bilgutay, "Analysis of homomorphic processing for ultrasonic grain signal characterization," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 36, pp. 365–375, May 1989.

[94] S. Solbo and T. Eltoft, "Homomorphic wavelet-based statistical despeckling of SAR images," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 4, pp. 711–721, Apr. 2004.

[95] S. Winkler, "Vision models and quality metrics for image processing applications," Ph.D. dissertation, University of Lausanne-Switzerland, Dec. 2000.

[96] Z. Wang and A. C. Bovik, "A universal image quality index," *IEEE Signal Process. Lett.*, vol. 9, no. 3, pp. 81–84, Mar. 2002.

[97] T. Crimmins, "Geometric filter for speckle reduction," *Appl. Opt.*, vol. 24, pp. 1438–1443, 1985.

[98] J. Weickert, "Coherence-enhancing diffusion filtering," *Int. J. Comput. Vision*, vol. 31, no. 2-3, pp. 111–127, 1999.

[99] Y. Yu and S. Acton, "Edge detection in ultrasound imagery using the instantaneous coefficient of variation," *IEEE Trans. Image Process.*, vol. 13, no. 12, pp. 1640–1655, Dec. 2004.

[100] J. Weickert, B. Romeny, and M. Viergever, "Efficient and reliable schemes for nonlinear diffusion filtering," *IEEE Trans. Image Process.*, vol. 7, no. 3, pp. 398–410, Mar. 1998.

[101] J. Weickert, "Anisotropic diffusion in image processing," Ph.D. dissertation, Dept. of Mathematics, University of Kaiserslautern, Germany, 1996.

[102] S. Mallat and S. Zhong, "Characterization of signals from multiscale edges," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 7, pp. 710–732, Jul. 1992.

[103] S. Mallat, *A Wavelet Tour of Signal Processing, 2nd Ed.* New York: Academic Press, 1999.

[104] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. R. Stat. Soc. B*, vol. 39, no. 1, pp. 1–38, 1977.

[105] Z. Shi and K. Fung, “A comparison of digital speckle filters,” in *Proc. Int. Geoscience and Remote Sensing Symp.*, vol. 4, Aug. 1994, pp. 2129–2133.

[106] A. Baraldi and F. Parmiggiani, “An alternative form of the Lee filter for speckle suppression in SAR images,” *Graph. Model Im. Proc.*, vol. 57, no. 1, pp. 75–78, 1995.

[107] F. Ulaby, F. Kouyate, B. Brisco, and T. Williams, “Textural infornation in SAR images,” *IEEE Trans. Geosci. Remote Sens.*, vol. GE-24, no. 2, pp. 235–245, Mar. 1986.

[108] J. Canny, “A computational approach to edge detection,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.

[109] A. C. Bovik, “On detecting edges in speckle imagery,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, pp. 1618–1627, Oct. 1988.

[110] C. Oliver and P. Lombardo, “Simultaneous mean and texture edge detection in SAR clutter,” *Proc. Inst. Elect. Eng. Radar, Sonar, Navig.*, vol. 143, no. 6, pp. 391–399, Dec. 1996.

[111] C. Oliver, D. Blacknell, and R. White, “Optimum edge detection in SAR,” *Proc. Inst. Elect. Eng. Radar, Sonar, Navig.*, vol. 143, no. 1, pp. 31–40, Feb. 1996.

[112] R. Touzi, A. Lopes, and P. Bousquet, “A statistical and geometrical edge detector for SAR images,” *IEEE Trans. Geosci. Remote Sens.*, vol. 26, no. 6, pp. 764–773, Nov. 1988.

[113] R. Czerwinski, D. Jones, and W. J. O’Brien, “Edge detection in ultrasound speckle noise,” in *Proc. IEEE Int. Conf. Image Processing*, vol. 3, Austin, TX., Nov. 1994, pp. 304–308.

[114] W. H. Kruskal and W. A. Wallis, “Errata: Use of ranks in one-criterion variance analysis,” *J. Am. Stat. Assoc.*, vol. 48, no. 264, pp. 907–911, Dec. 1953.

[115] A. Field, *Discovering Statistics Using SPSS*, 3rd ed. Thousand Oaks, CA: SAGE Publications, Jan. 2009.

[116] C. Spearman, “The proof and measurement of association between two things,” *Amer. J. Psychol.*, vol. 15, pp. 72–101, Jan. 1904.

[117] J. Jensen, “A model for the propagation and scattering of ultrasound in tissue,” *J. Acoust. Soc. Am.*, vol. 89, pp. 182–190, Jan. 1991.

[118] J. Ng, R. Prager, N. Kingsbury, G. Treece, and A. Gee, “Wavelet restoration of medical pulse-echo ultrasound images in an em framework,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 54, no. 3, pp. 550–568, Mar. 2007.

[119] O. Michailovich and D. Adam, “A novel approach to the 2-d blind deconvolution problem in medical ultrasound,” *IEEE Trans. Med. Imag.*, vol. 24, no. 1, pp. 86–104, Jan. 2005.

[120] S. Finn, E. Jones, and M. Glavin, “Objective and subjective evaluations of quality for speckle reduced echocardiography,” in *Proc. IEEE Intl. Annu. Conf. Engineering Medicine and Biology Society*, Sep. 2009, pp. 503–506.

[121] D. K. Manolakis and J. G. Proakis, *Digital signal processing: Principles, algorithms, and applications*, 2nd ed. New York: Macmillan, 1992.

[122] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. New York: Cambridge University Press, 2007.

[123] R. C. Molthen, P. M. Shankar, and J. M. Reid, “Characterization of ultrasonic B-scans using non-Rayleigh statistics,” *Ultrasound Med. Biol.*, vol. 21, no. 2, pp. 161–170, 1995.

[124] D. Marr and E. Hildreth, “Theory of Edge Detection,” *Proc. R. Soc. Lond. B*, vol. 207, pp. 187–217, Feb. 1980.

[125] I. Sobel and G. Feldman, “A 3x3 isotropic gradient operator for image processing,” in *Pattern Classification and Scene Analysis*, R. Duda and P. Hart, Eds. John Wiley and Sons, 1973.

[126] T. S. Lee, “Image representation using 2D Gabor wavelets,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 10, pp. 959 –971, Oct. 1996.

[127] D. Gabor, “Theory of communication,” *JIEE*, vol. 93, no. 3, pp. 429–459, 1946.

[128] J. G. Daugman, “Two-dimensional spectral analysis of cortical receptive field profiles,” *Vision Res.*, vol. 20, no. 10, pp. 847–856, 1980.

[129] S. Marčelja, “Mathematical description of the responses of simple cortical cells,” *J. Opt. Soc. Am.*, vol. 70, no. 11, pp. 1297–1300, 1980.

[130] J. G. Daugman, “Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters,” *J. Opt. Soc. Am. A*, vol. 2, no. 7, pp. 1160–1169, 1985.

[131] Y. Zhan and D. Shen, “Deformable segmentation of 3-d ultrasound prostate images using statistical texture matching method,” *IEEE Trans. Med. Imag.*, vol. 25, no. 3, pp. 256–272, Mar. 2006.

[132] D. Shen, Y. Zhan, and C. Davatzikos, “Segmentation of prostate boundaries from ultrasound images using statistical shape model,” *IEEE Trans. Med. Imag.*, vol. 22, no. 4, pp. 539–551, Apr. 2003.

[133] J. Xie, Y. Jiang, and H. tat Tsui, “Segmentation of kidney from ultrasound images based on texture and shape priors,” *IEEE Trans. Med. Imag.*, vol. 24, no. 1, pp. 45–57, Jan. 2005.

[134] C.-M. Chen, H. H.-S. Lu, and K.-C. Han, “A textural approach based on Gabor functions for texture edge detection in ultrasound images,” *Ultrasound Med. Biol.*, vol. 27, no. 4, pp. 515–534, 2001.

[135] D. J. Field, “Relations between the statistics of natural images and the response properties of cortical cells,” *J. Opt. Soc. Am. A*, vol. 4, pp. 2379–2394, 1987.

[136] P. Kovesi, “Image features from phase congruency,” *Videre: A Journal of Computer Vision Research, MIT Press*, vol. 1, no. 3, 1999.

[137] M. Mulet-Parada and J. Noble, “2d+t acoustic boundary detection in echocardiography,” in *Medical Image Computing and Computer-Assisted Interventation: MICCAI 98*. Springer, 1998, vol. 1498, pp. 806–813.

[138] ——, “2d+t acoustic boundary detection in echocardiography,” *Med. Image Anal.*, vol. 4, no. 1, pp. 21–30, 2000.

[139] G. Samorodnitsky and M. S. Taqqu, *Stable Non-Gaussian Random Processes: Stochastic Models with Infinite Variance.* New York: Chapman and Hall, 1994.

[140] C. L. Nikias and M. Shao, *Signal processing with alpha-stable distributions and applications.* New York: Wiley-Interscience, 1995.

[141] J. P. Nolan, *Stable Distributions - Models for Heavy Tailed Data.* Boston, MA: Birkhäuser, 2010, in progress, Chapter 1 online at http://academic2.american.edu/∼jpnolan.

[142] J. Nolan, "Parameterizations and modes of stable distributions," *Stat. Probabil. Lett.*, vol. 38, no. 2, pp. 187–195, 1998.

[143] P. Hall, "A Comedy of Errors: The Canonical Form for a Stable Characteristic Function," *Bull. London Math. Soc.*, vol. 13, no. 1, pp. 23–27, 1981.

[144] J. Nolan, "Modelling financial data with stable distributions," in *Handbook of Heavy Tailed Distributions in Finance*, S. Rachev, Ed. Amsterdam: Elsevier, 2003, pp. 106–129.

[145] E. Kornaropoulos and P. Tsakalides, "A novel kNN classifier for acoustic vehicle classification based on alpha-stable statistical modeling," in *IEEE/SP 15th Workshop Statistical Signal Processing*, Sep. 2009, pp. 1–4.

[146] I. A. Khovanov, N. A. Khovanova, P. V. E. McClintock, and A. Stefanovska, "Intrinsic dynamics of heart regulatory systems on short timescales: from experiment to modelling," *J. Stat. Mech.-Theory E.*, no. 01, p. 01016, 2009.

[147] Q. Zou, D. Feng, Y. Zhu, H. Jiang, X. Ge, and Z. Zhou, "A novel and generic model for synthesizing disk I/O traffic based on the alpha-stable process," in *Proc 16th IEEE Int. Symp. Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, Sep. 2008, pp. 1–10.

[148] M. Morrone and R. Owens, "Feature detection from local energy," *Pattern Recogn. Lett.*, vol. 6, no. 5, pp. 303–313, 1987.

[149] M. C. Morrone, J. R. Ross, D. C. Burr, and R. A. Owens, "Mach bands are phase dependent," *Nature*, vol. 324, no. 6094, pp. 250–253, 1986.

[150] D. Boukerroui, J. A. Noble, and M. Brady, "Feature enhancement in low quality images with applicaiton to echocardiography," *Ultrasound Med. Biol.*, vol. 27, pp. 1583–1594, 2001.

[151] A. Achim, A. Bezerianos, and P. Tsakalides, "SAR image denoising: a multiscale robust statistical approach," in *Proc. 14th Int. Conf. Digital Signal Processing*, vol. 2, 2002, pp. 1235–1238.

[152] A. Achim, P. Tsakalides, and A. Bezerianos, "SAR image denoising via Bayesian wavelet shrinkage based on heavy-tailed modeling," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 8, pp. 1773–1784, Aug. 2003.

[153] A. Baraldi and F. Panniggiani, "A refined gamma MAP SAR speckle filter with improved geometrical adaptivity," *IEEE Trans. Geosci. Remote Sens.*, vol. 33, no. 5, pp. 1245–1257, Sep. 1995.

[154] A. Achim, A. Bezerianos, and P. Tsakalides, "Wavelet-based ultrasound image denoising using an alpha-stable prior probability model," in *Proc. IEEE Image Processing*, 2001, pp. 221–224.

[155] A. M. Achim, "Bivariate wavelet shrinkage using alpha-stable distributions," in *Wavelets XI*, M. Papadakis, A. F. Laine, and M. A. Unser, Eds. SPIE, 2005, vol. 5914, no. 1.

[156] T. Wan, N. Canagarajah, and A. Achim, "Segmentation-driven image fusion based on alpha-stable modeling of wavelet coefficients," *IEEE Trans. Multimedia*, vol. 11, no. 4, pp. 624–633, Jun. 2009.

[157] A. Achim and E. Kuruoglu, "Image denoising using bivariate alpha;-stable distributions in the complex wavelet domain," *IEEE Signal Process. Lett.*, vol. 12, no. 1, pp. 17–20, Jan. 2005.

[158] J. Nolan, "Maximum likelihood estimation and diagnostics for stable distributions," in *Lévy Processes: Theory and Applications*, O. E. Barndorff-Nielsen, T. Mikosch, and S. I. Resnick, Eds. Boston: Birkhäuser, 2001, pp. 379–400.

[159] J. R. Michael, "The stabilized probability plot," *Biometrika*, vol. 70, no. 1, pp. 11–17, 1983.

[160] J. H. McCulloch, “Simple consistent estimators of stable distribution parameters,” *Commun. Stat. B-Simul.*, vol. 15, no. 4, pp. 1109–1136, 1986.

[161] ——, “Numerical approximation of the symmetric stable distribution and density,” in *A practical guide to heavy tails: statistical techniques and applications*, R. Adler, R. Feldman, and M. Taqqu, Eds. Cambridge, MA, USA: Birkhauser Boston Inc., 1998, pp. 489–499.

[162] S. Mittnik, S. Rachev, D. T., and D. Chenyao, “Maximum likelihood estimation of stable paretian models,” *Math. Comput. Model.*, vol. 29, pp. 275–293, May 1999.

[163] S. J. Press, “Estimation in univariate and multivariate stable distributions,” *J. Am. Stat. Assoc.*, vol. 67, no. 340, pp. 842–846, 1972.

[164] I. A. Koutrouvelis, “Regression-type estimation of the parameters of stable laws,” *J. Am. Stat. Assoc.*, vol. 75, no. 372, pp. 918–928, 1980.

[165] ——, “An iterative procedure for the estimation of the parameters of stable laws,” *Commun. Stat. B-Simul.*, vol. 10, no. 1, pp. 17–28, 1981.

[166] A. Feuerverger and P. McDunnough, “On the efficiency of empirical characteristic function procedures,” *J. Roy. Stat. Soc. B Met.*, vol. 43, no. 1, pp. 20–27, 1981.

[167] S. M. Kogon and D. B. Williams, “Characteristic function based estimation of stable distribution parameters,” in *A practical guide to heavy tails: statistical techniques and applications*, R. Adler, R. Feldman, and M. Taqqu, Eds. Boston: Birkhäuser, 1998, pp. 311–335.

[168] A. Antoniadis, A. Feuerverger, and P. Goncalves, “Wavelet-based estimation for univariate stable laws,” *Ann. I. Stat. Math.*, vol. 58, pp. 779–807, Dec 2006.

[169] V. Akgiray and C. G. Lamoureux, “Estimation of stable-law parameters: A comparative study,” *J. Bus. Econ. Stat.*, vol. 7, no. 1, pp. 85–93, 1989.

[170] J. A. Nelder and R. Mead, “A simplex method for function minimization,” *Comp. J.*, vol. 7, no. 4, pp. 308–313, Jan 1965.

[171] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the Nelder–Mead simplex method in low dimensions," *SIAM J. Optimiz.*, vol. 9, no. 1, pp. 112–147, 1998.

[172] P. Besbeas and B. J. T. Morgan, "Improved estimation of the stable laws," *Stat. Comput.*, vol. 18, no. 2, pp. 219–231, Jun. 2008.

[173] A. H. Welsh, "Implementing empirical characteristic function procedures," *Stat. Probabil. Lett.*, vol. 4, no. 2, pp. 65–67, 1986.

[174] C. R. Heathcote and J. Hüsler, "The first zero of an empirical characteristic function," *Stoch. Proc. Appl.*, vol. 35, no. 2, pp. 347–360, 1990.

[175] J. Prewitt, "Object enhancement and extraction," in *Picture Processing and Psychopictorics*, B. Lipkin and A. Rosenfeld, Eds. Academic Press, 1970, pp. 75–149.

[176] F. Pellegrino, W. Vanzella, and V. Torre, "Edge detection revisited," *IEEE Trans. Syst., Man, Cybern. B*, vol. 34, no. 3, pp. 1500–1518, Jun. 2004.

[177] S. Kichenassamy, "The perona-malik paradox," *SIAM J. Appl. Math.*, vol. 57, no. 5, pp. 1328–1342, 1997.

[178] R. O. Bude and R. S. Adler, "An easily made, low-cost, tissue-like ultrasound phantom material," *J. Clin. Ultrasound*, vol. 23, no. 4, pp. 271–273, 1995.

[179] M. Pourahmadi, "Some sampling properties of empirical characteristic functions viewed as harmonizable stochastic processes," *J. Stat. Plan. Inf.*, no. 17, pp. 345–359, 1987.

[180] I. Young, L. van Vliet, and M. van Ginkel, "Recursive Gabor filtering," *IEEE Trans. Signal Process.*, vol. 50, no. 11, pp. 2798–2805, Nov. 2002.

[181] D. Bailey, "An efficient euclidean distance transform," in *Combinatorial Image Analysis*, ser. Lecture Notes in Computer Science, R. Klette and J. Žunic, Eds. Springer Berlin / Heidelberg, 2005, vol. 3322, pp. 394–408.

Appendix A

Additional Data from Chapter 5

A.1 Modelling Gabor Coefficients of Speckled Images

Section 5.3.1 presented a model for the statistics of the Gabor coefficients of log transformed speckled images. This model is a combination of a SαS component for the speckle free Gabor coefficients, and a Gaussian component for the Gabor coefficients resulting from speckle. In Section 5.3.1, the accuracy of this model was verified, and additional data from this verification is included here for completeness.

The histograms of the real and imaginary Gabor coefficients of speckle free simulated and speckle reduced clinical images are presented in Figs. A.1-A.2, which display non-Gaussian characteristics. The fitting of SαS densities to these observed densities was examined using stabilised probability plots. The full set of these probability plots can be seen in Figs. A.3-A.6.

These analysis methods, along with the examination of the confidence intervals of maximum likelihood parameter estimates presented in Section 5.3.1, demonstrate the accuracy of a SαS model for the speckle free Gabor coefficients. Finally, the modelling of images containing speckle, using the combined SαS and Gaussian model, was evaluated by comparing the observed and fitted densities. The full set of these are presented in this appendix in Figs. A.7-A.10. Quantification of the accuracy of these fitted densities using the r^2 value was presented in Section 5.3.1.

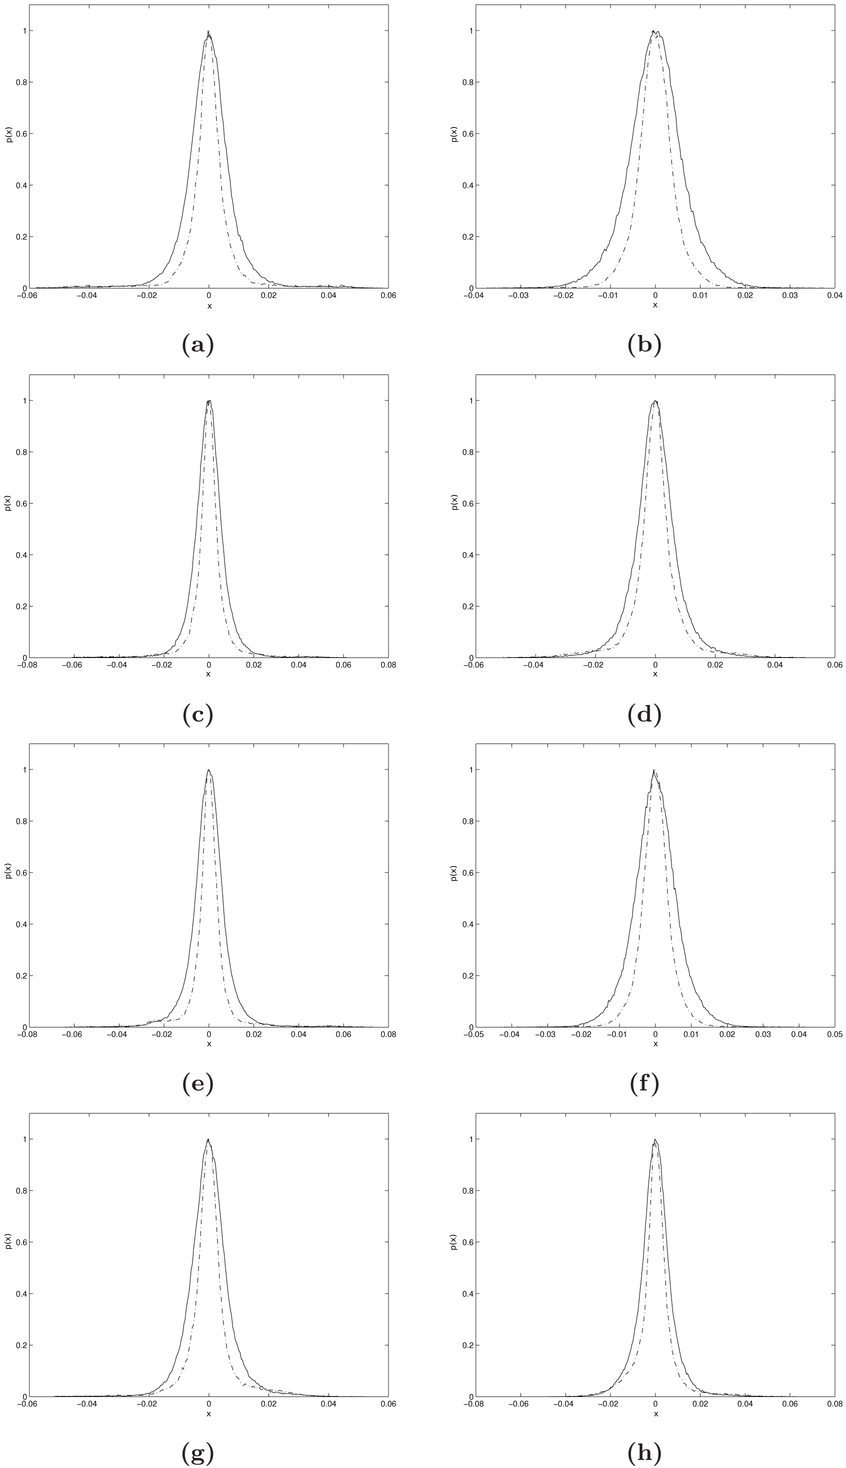

(a) (b) (c) (d) (e) (f) (g) (h)

Figure A.1: Histograms of the real and imaginary Gabor coefficients of each orientation of Fig. 5.11(a). Solid lines are coefficients from scale 1, dash dot lines from scale 2. (a) $\Re\{\theta = 0 \text{ rad}\}$, (b) $\Re\{\theta = \frac{\pi}{4} \text{ rad}\}$, (c) $\Re\{\theta = \frac{\pi}{2} \text{ rad}\}$, (d) $\Re\{\theta = \frac{3\pi}{4} \text{ rad}\}$, (e) $\Im\{\theta = 0 \text{ rad}\}$, (f) $\Im\{\theta = \frac{\pi}{4} \text{ rad}\}$, (g) $\Im\{\theta = \frac{\pi}{2} \text{ rad}\}$, (h) $\Im\{\theta = \frac{3\pi}{4} \text{ rad}\}$.

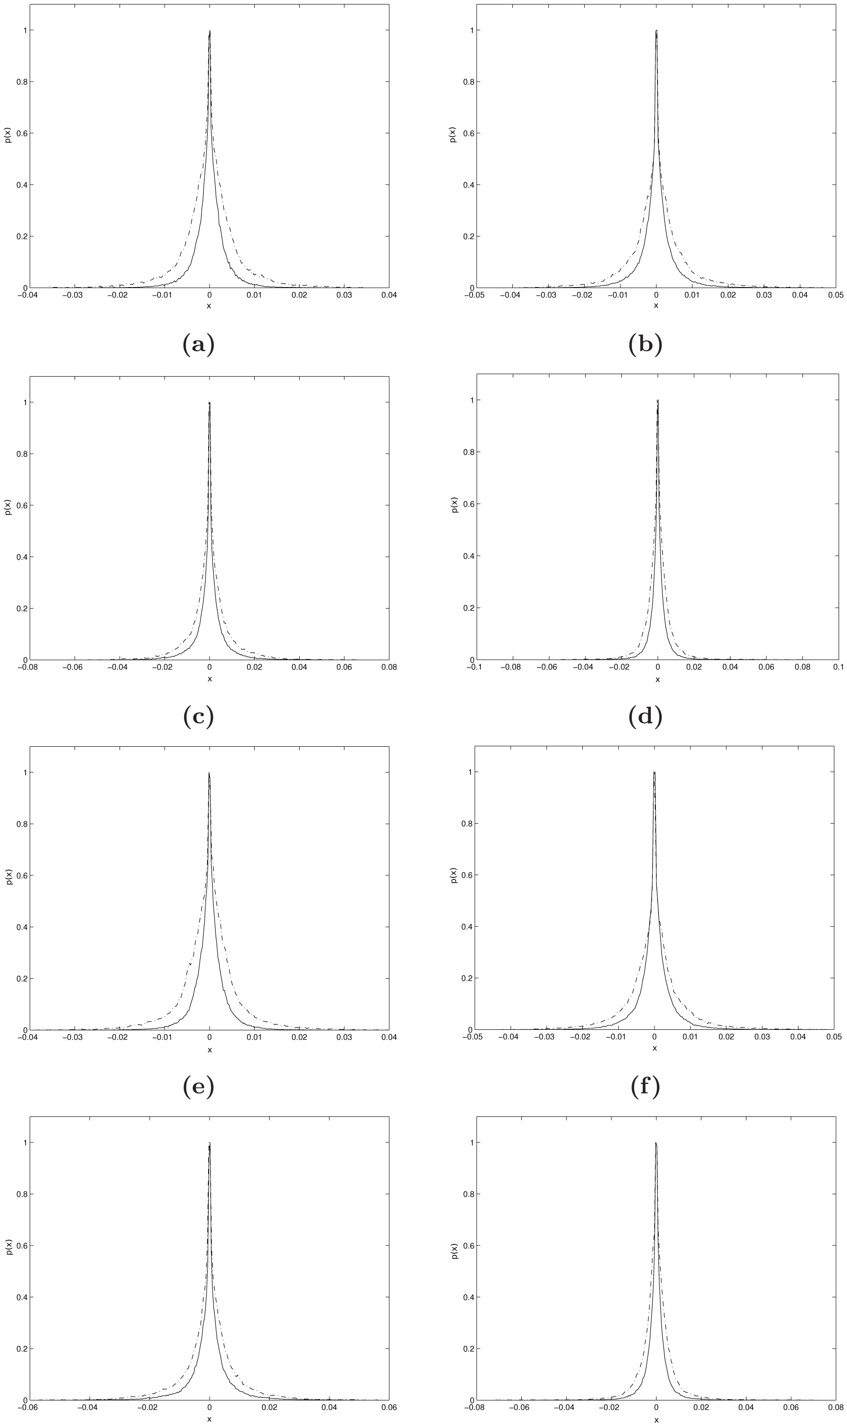

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.2: Histograms of the real and imaginary Gabor coefficients of each orientation of Fig. 5.11(b). Solid lines are coefficients from scale 1, dash dot lines from scale 2. (a) $\Re\{\theta = 0 \text{ rad}\}$, (b) $\Re\{\theta = \frac{\pi}{4} \text{ rad}\}$, (c) $\Re\{\theta = \frac{\pi}{2} \text{ rad}\}$, (d) $\Re\{\theta = \frac{3\pi}{4} \text{ rad}\}$, (e) $\Im\{\theta = 0 \text{ rad}\}$, (f) $\Im\{\theta = \frac{\pi}{4} \text{ rad}\}$, (g) $\Im\{\theta = \frac{\pi}{2} \text{ rad}\}$, (h) $\Im\{\theta = \frac{3\pi}{4} \text{ rad}\}$.

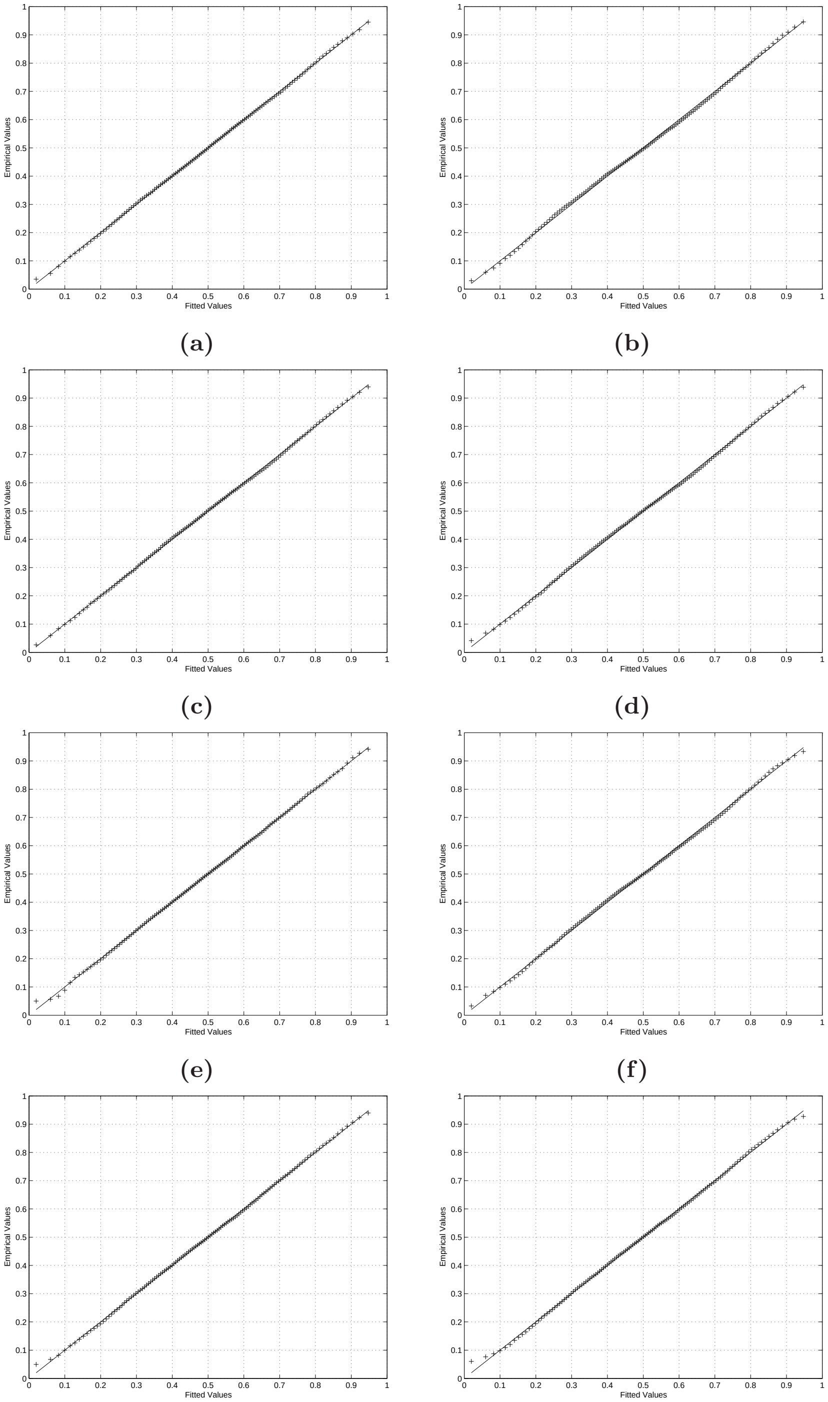

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.3: Stabilised probability plots of the real-valued Gabor coefficients of image 5.11(a), plotted per orientation. Heavy lines are the observed distribution points, thin lines are those expected from the ML parameter estimates. (a) s=1, $\theta = 0$ rad, (b) s=1, $\theta = \frac{\pi}{4}$ rad, (c) s=1, $\theta = \frac{\pi}{2}$ rad, (d) s=1, $\theta = \frac{3\pi}{4}$ rad, (e) s=2, $\theta = 0$ rad, (f) s=2, $\theta = \frac{\pi}{4}$ rad, (g) s=2, $\theta = \frac{\pi}{2}$ rad, (h) s=2, $\theta = \frac{3\pi}{4}$ rad.

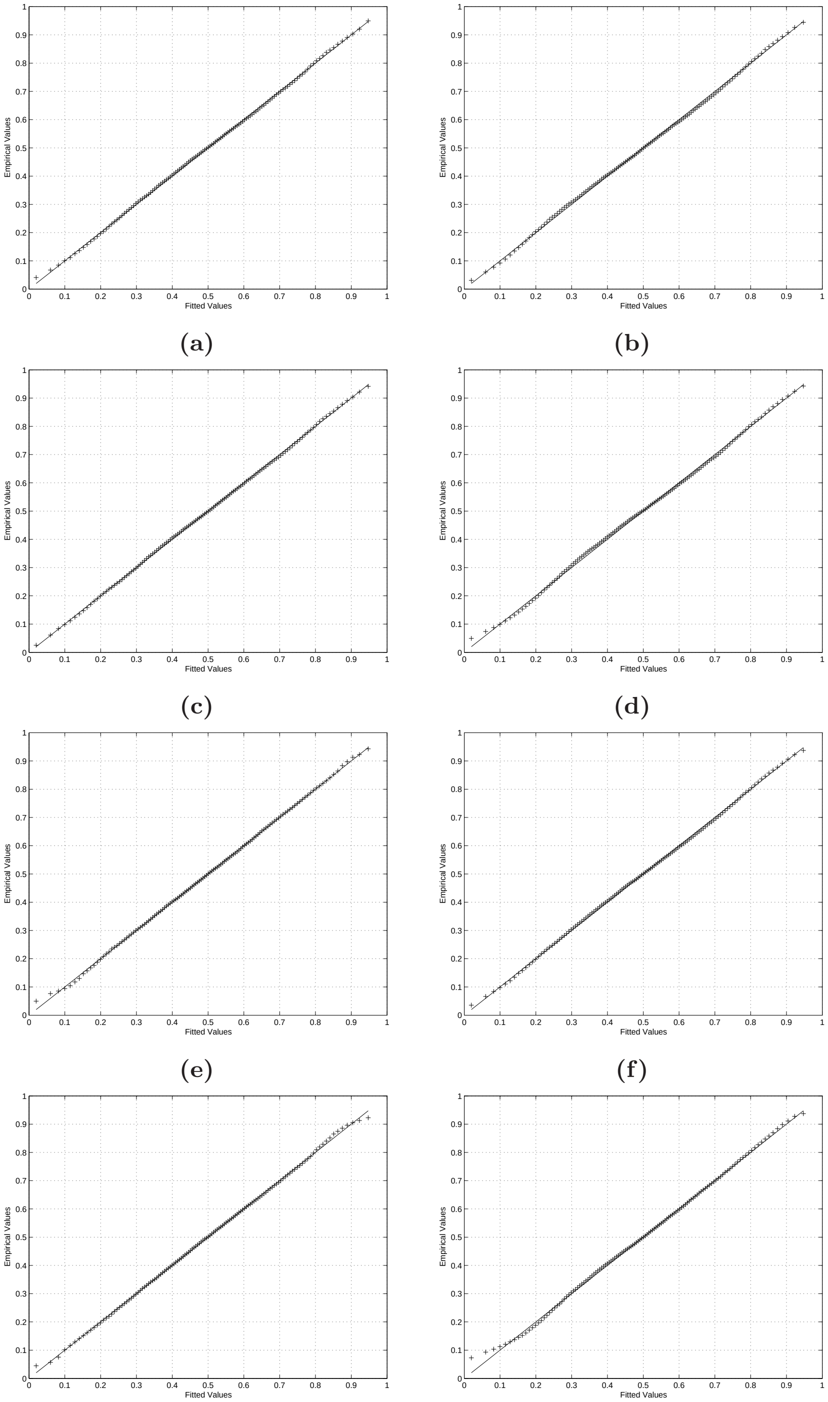

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.4: Stabilised probability plots of the imaginary-valued Gabor coefficients of image 5.11(a), plotted per orientation. Heavy lines are the observed distribution points, thin lines are those expected from the ML parameter estimates. (a) s=1, $\theta = 0$ rad, (b) s=1, $\theta = \frac{\pi}{4}$ rad, (c) s=1, $\theta = \frac{\pi}{2}$ rad, (d) s=1, $\theta = \frac{3\pi}{4}$ rad, (e) s=2, $\theta = 0$ rad, (f) s=2, $\theta = \frac{\pi}{4}$ rad, (g) s=2, $\theta = \frac{\pi}{2}$ rad, (h) s=2, $\theta = \frac{3\pi}{4}$ rad.

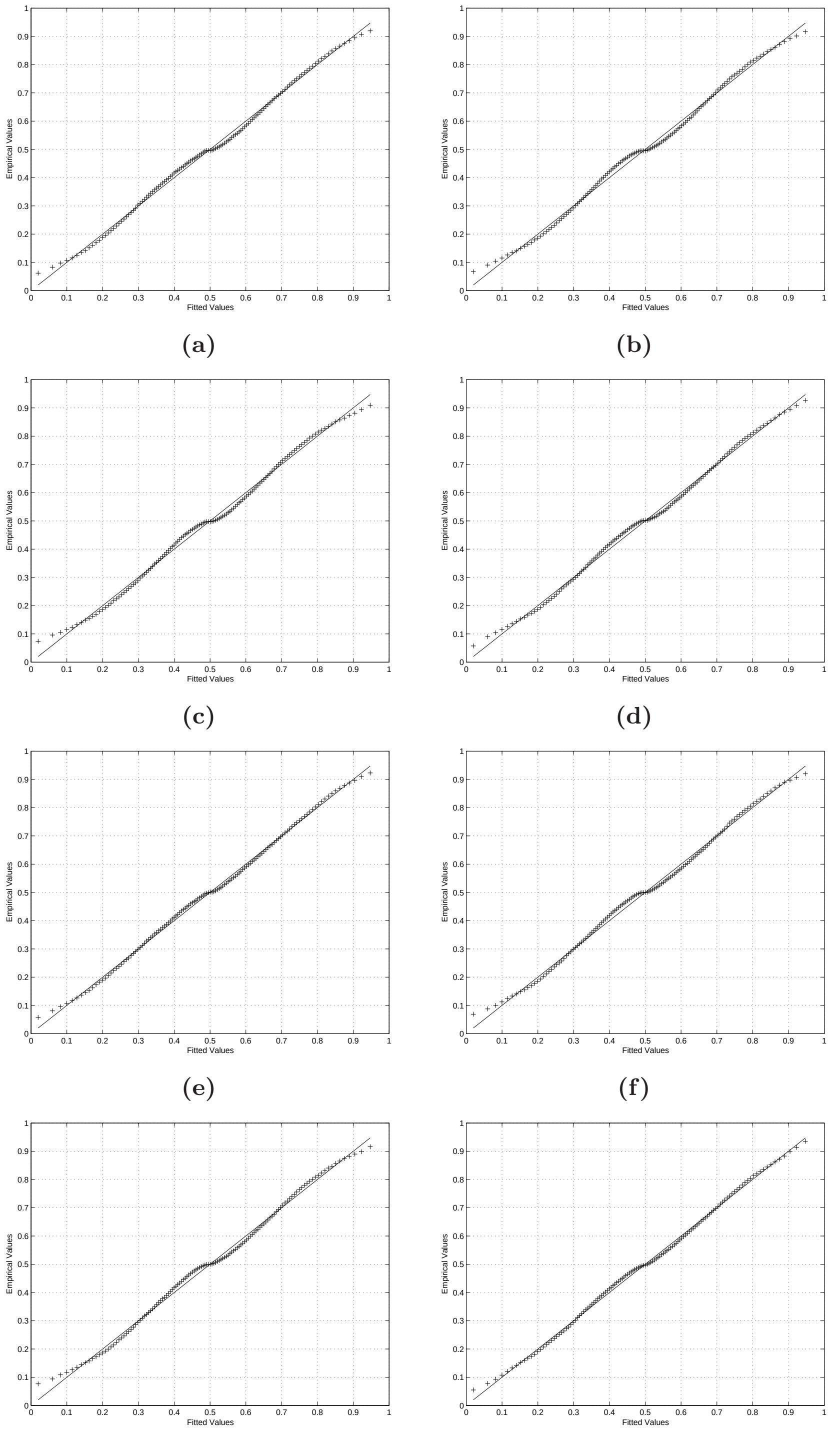

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.5: Stabilised probability plots of the real-valued Gabor coefficients of image 5.11(b), plotted per orientation. Heavy lines are the observed distribution points, thin lines are those expected from the ML parameter estimates. (a) s=1, $\theta = 0$ rad, (b) s=1, $\theta = \frac{\pi}{4}$ rad, (c) s=1, $\theta = \frac{\pi}{2}$ rad, (d) s=1, $\theta = \frac{3\pi}{4}$ rad, (e) s=1, $\theta = 0$ rad, (f) s=1, $\theta = \frac{\pi}{4}$ rad, (g) s=1, $\theta = \frac{\pi}{2}$ rad, (h) s=1, $\theta = \frac{3\pi}{4}$ rad.

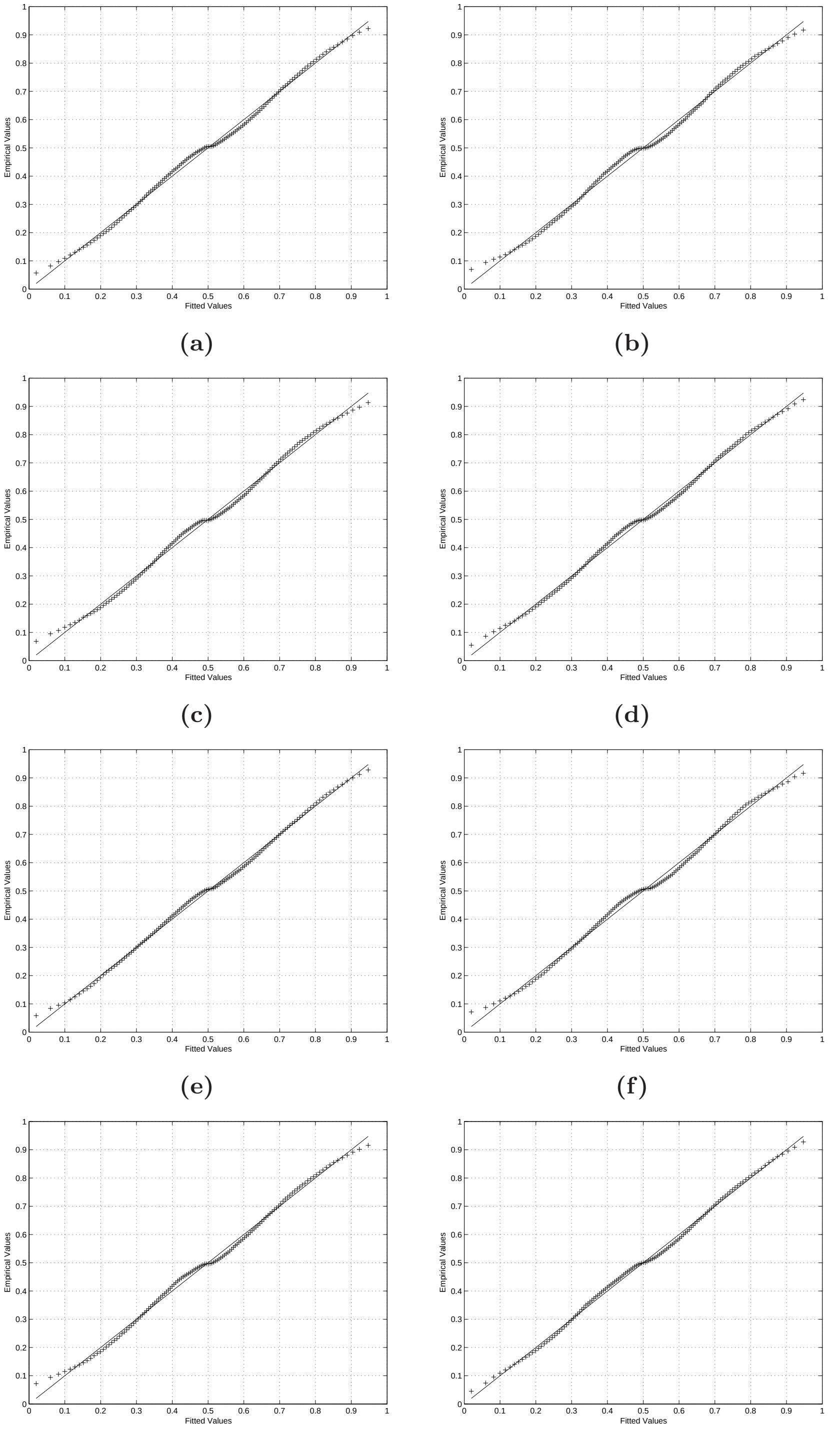

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.6: Stabilised probability plots of the imaginary-valued Gabor coefficients of image 5.11(b), plotted per orientation. Heavy lines are the observed distribution points, thin lines are those expected from the ML parameter estimates. (a) s=1, $\theta = 0$ rad, (b) s=1, $\theta = \frac{\pi}{4}$ rad, (c) s=1, $\theta = \frac{\pi}{2}$ rad, (d) s=1, $\theta = \frac{3\pi}{4}$ rad, (e) s=1, $\theta = 0$ rad, (f) s=1, $\theta = \frac{\pi}{4}$ rad, (g) s=1, $\theta = \frac{\pi}{2}$ rad, (h) s=1, $\theta = \frac{3\pi}{4}$ rad.

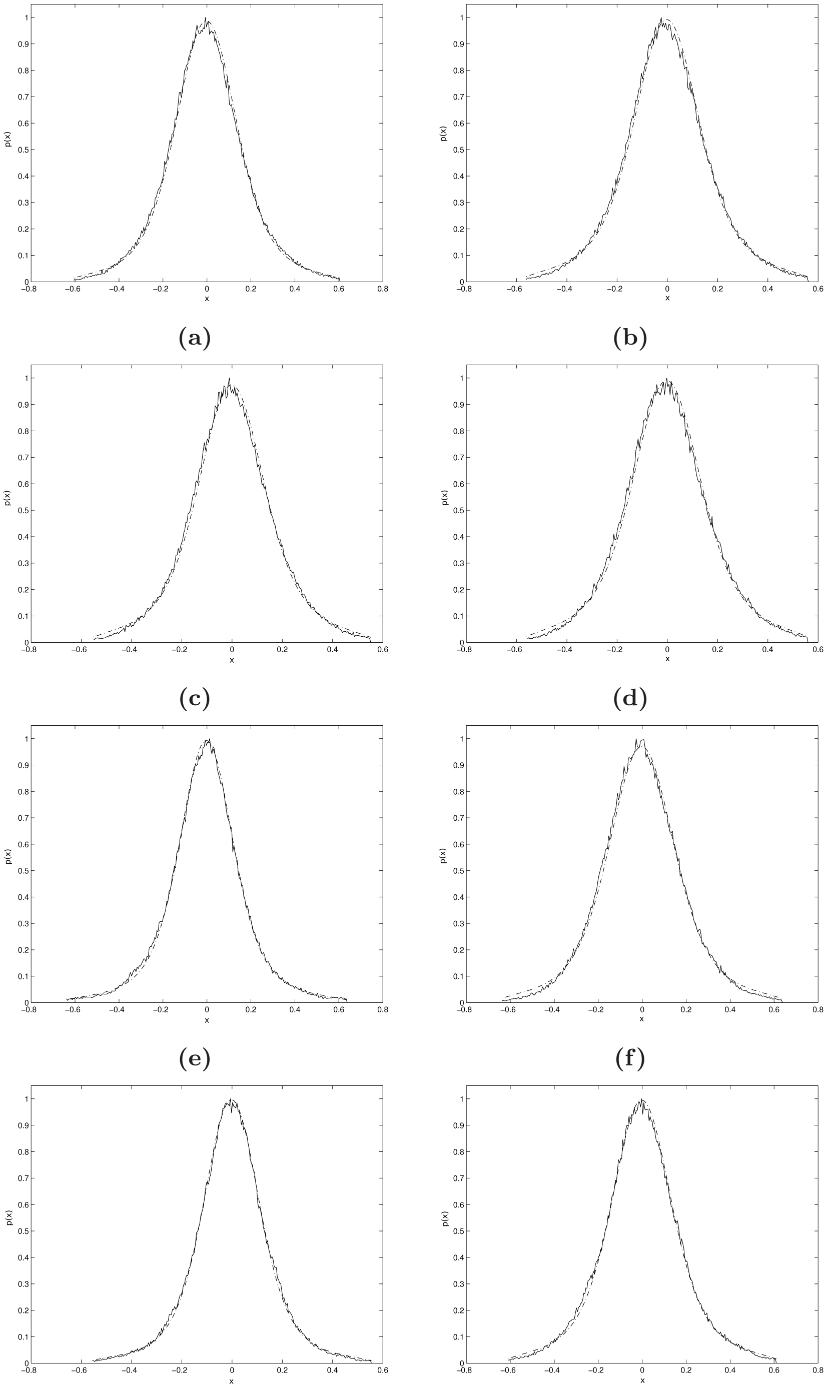

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.7: Observed densities of the real-valued Gabor coefficients of the speckled image of Fig. 5.16(a), and those fitted using the model of (5.21), for each scale and orientation. Solid lines are the observed densities, dash-dot lines are the fitted densities. (a) s=1, $\theta = 0$ rad, (b) s=1, $\theta = \frac{\pi}{4}$ rad, (c) s=1, $\theta = \frac{\pi}{2}$ rad, (d) s=1, $\theta = \frac{3\pi}{4}$ rad, (e) s=2, $\theta = 0$ rad, (f) s=2, $\theta = \frac{\pi}{4}$ rad, (g) s=2, $\theta = \frac{\pi}{2}$ rad, (h) s=2, $\theta = \frac{3\pi}{4}$ rad.

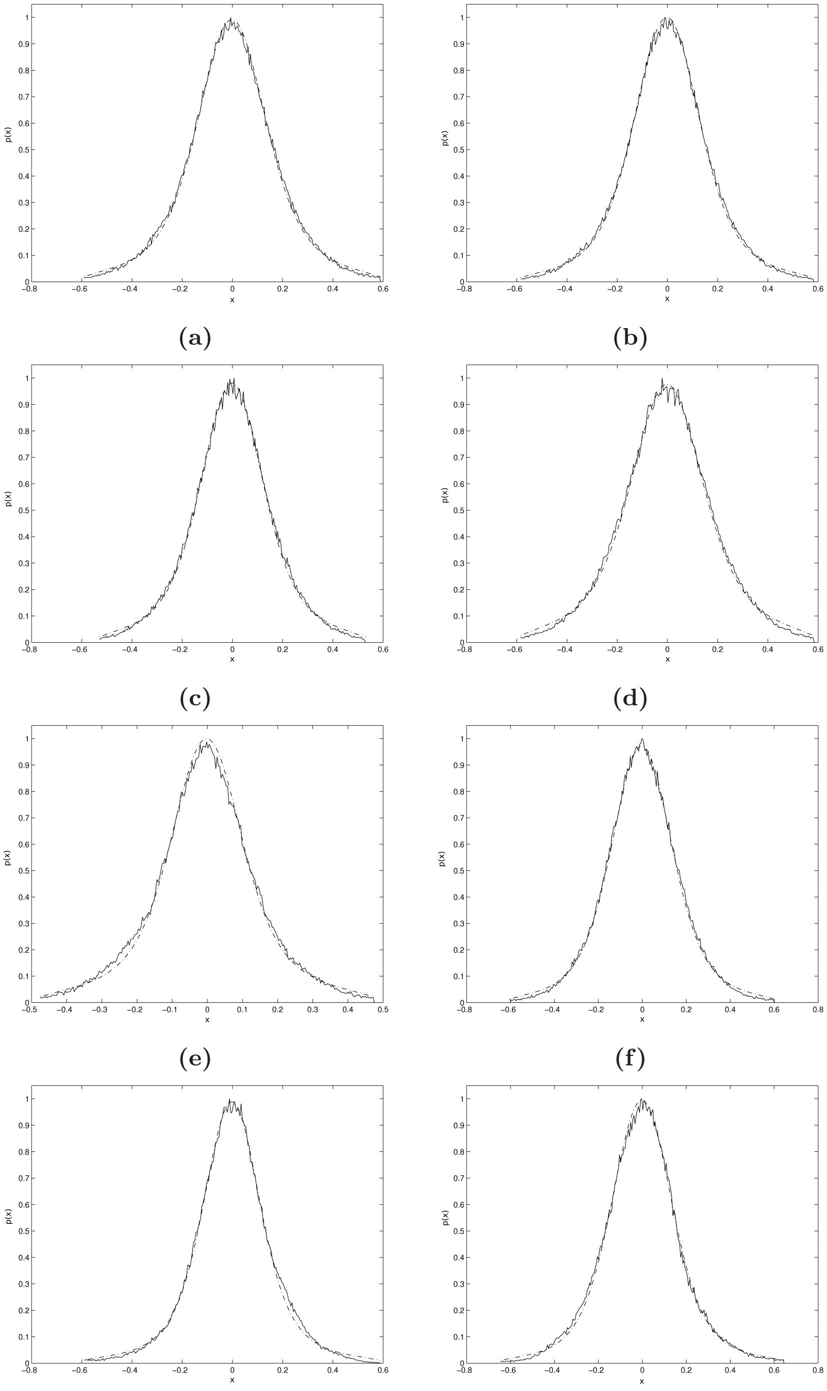

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.8: Observed densities of the imaginary-valued Gabor coefficients of the speckled image of Fig. 5.16(a), and those fitted using the model of (5.21), for each scale and orientation. Solid lines are the observed densities, dash-dot lines are the fitted densities. (a) s=1, $\theta = 0$ rad, (b) s=1, $\theta = \frac{\pi}{4}$ rad, (c) s=1, $\theta = \frac{\pi}{2}$ rad, (d) s=1, $\theta = \frac{3\pi}{4}$ rad, (e) s=2, $\theta = 0$ rad, (f) s=2, $\theta = \frac{\pi}{4}$ rad, (g) s=2, $\theta = \frac{\pi}{2}$ rad, (h) s=2, $\theta = \frac{3\pi}{4}$ rad.

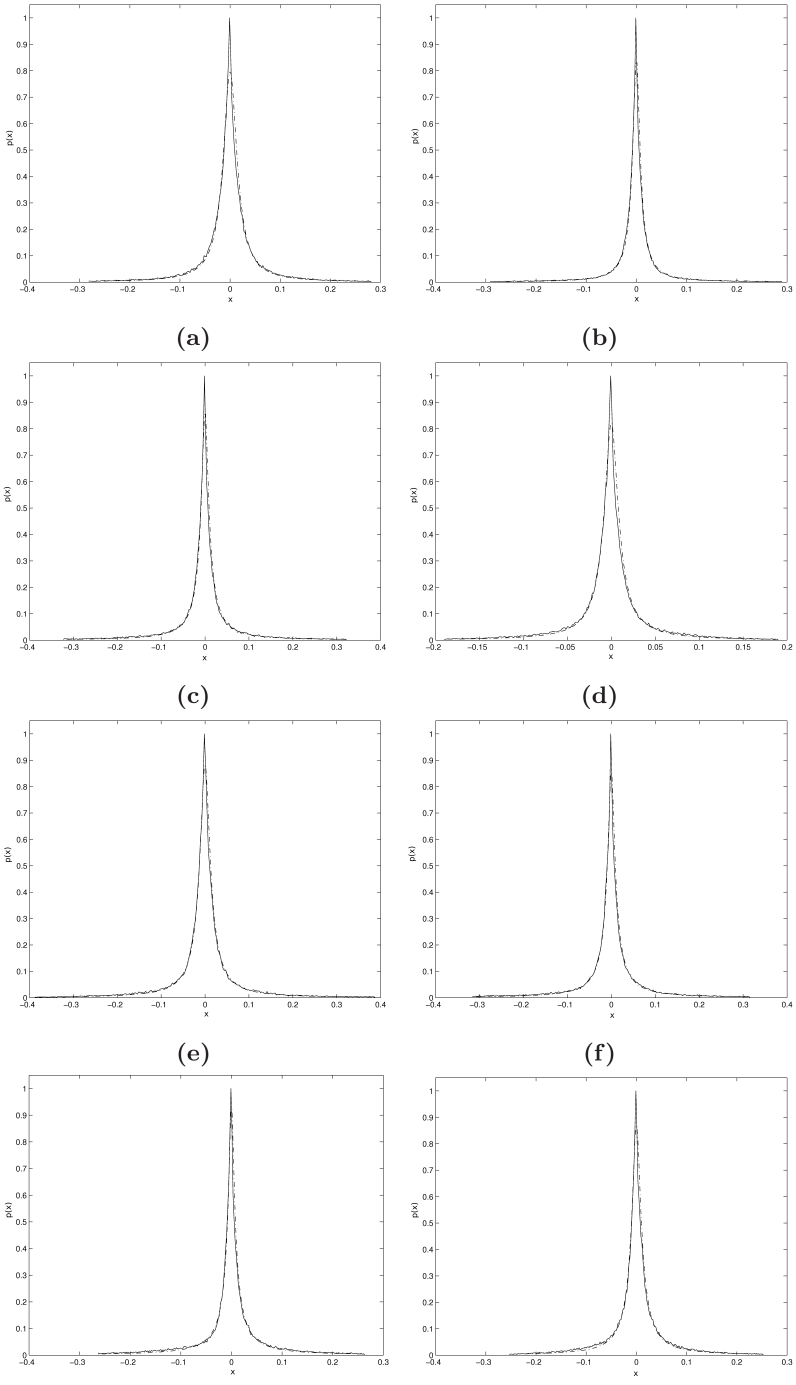

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.9: Observed densities of the real-valued Gabor coefficients of the speckled image of Fig. 5.16(b), and those fitted using the model of (5.21), for the real and imaginary components of each scale and orientation. Solid lines are the observed densities, dash-dot lines are the fitted densities. (a) s=1, $\theta = 0$ rad, (b) s=1, $\theta = \frac{\pi}{4}$ rad, (c) s=1, $\theta = \frac{\pi}{2}$ rad, (d) s=1, $\theta = \frac{3\pi}{4}$ rad, (e) s=2, $\theta = 0$ rad, (f) s=2, $\theta = \frac{\pi}{4}$ rad, (g) s=2, $\theta = \frac{\pi}{2}$ rad, (h) s=2, $\theta = \frac{3\pi}{4}$ rad.

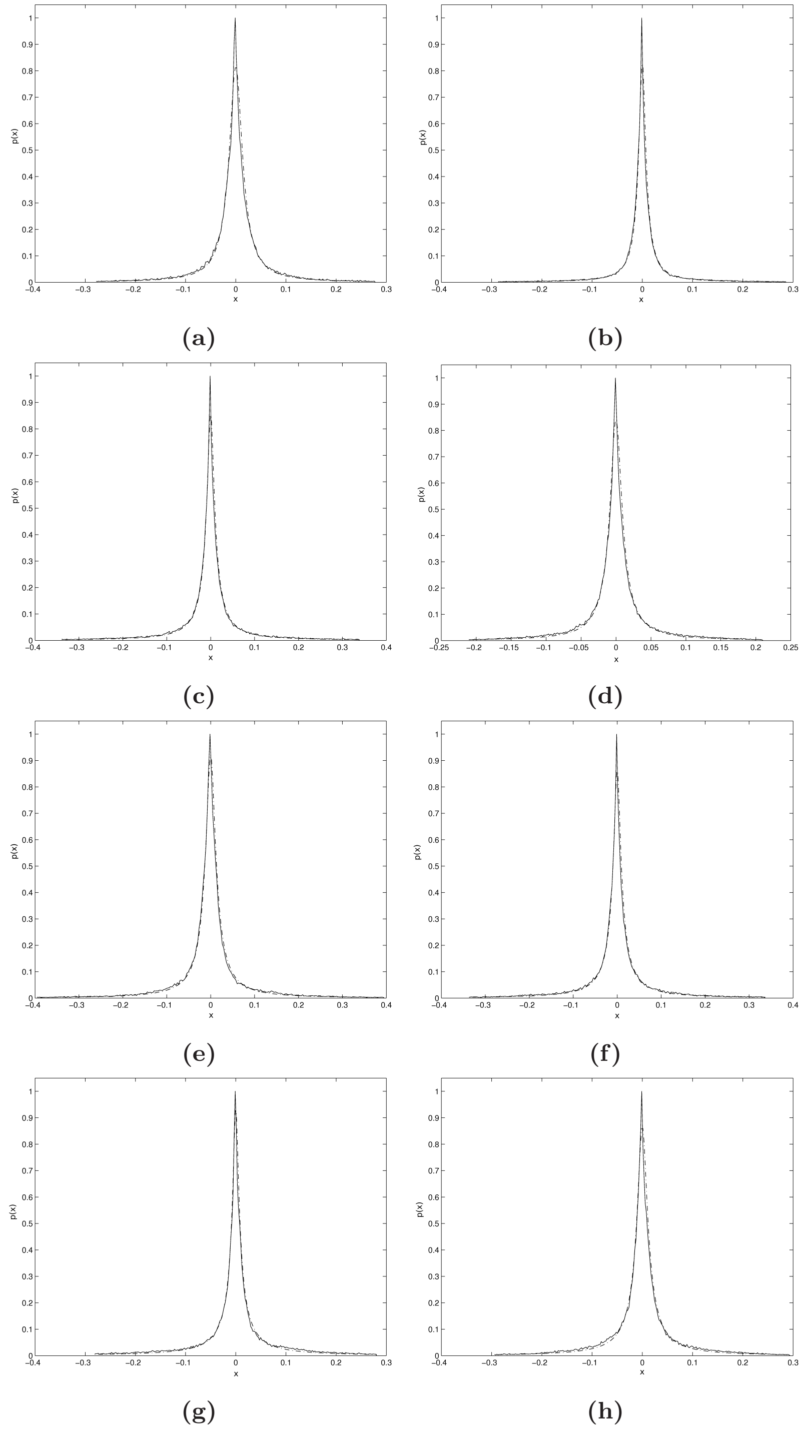

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure A.10: Observed densities of the imaginary-valued Gabor coefficients of the speckled image of Fig. 5.16(b), and those fitted using the model of (5.21), for the real and imaginary components of each scale and orientation. Solid lines are the observed densities, dash-dot lines are the fitted densities. (a) s=1, $\theta = 0$ rad, (b) s=1, $\theta = \frac{\pi}{4}$ rad, (c) s=1, $\theta = \frac{\pi}{2}$ rad, (d) s=1, $\theta = \frac{3\pi}{4}$ rad, (e) s=2, $\theta = 0$ rad, (f) s=2, $\theta = \frac{\pi}{4}$ rad, (g) s=2, $\theta = \frac{\pi}{2}$ rad, (h) s=2, $\theta = \frac{3\pi}{4}$ rad.

A.2 Parameter Estimation: Computed t_{max} Values

Selection of an appropriate value for t_{max} is crucial for accurate parameter estimation of stable processes using the empirical characteristic function. In Section 5.3 above, a method of parameter estimation was presented for the combined SαS and Gaussian model of the Gabor coefficients of log transformed speckled images. A procedure for computing optimal t_{max} values was also presented. Table A.1 below presents these t_{max} values, computed for various values of α, γ and σ.

Table A.1: Calculated optimal t_{max} values for parameter estimation by the proposed method of (5.25).

t_{max}, for $\sigma = 0.10$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	6.666	5.968	3.248	3.532	1.655	1.203	1.949	2.712
0.71	7.817	6.927	3.243	1.686	2.026	0.987	1.338	0.981
0.93	4.333	2.538	2.463	2.139	0.963	1.068	0.730	0.770
1.14	3.396	2.586	1.440	1.426	1.481	0.936	0.634	1.535
1.36	5.324	2.327	1.712	1.116	1.169	1.241	0.769	0.873
1.57	3.048	1.313	1.578	0.966	0.648	0.585	0.620	1.226
1.79	5.460	1.636	1.012	0.758	0.775	1.135	0.577	1.077
2.00	3.270	1.758	0.685	0.876	1.066	1.064	0.768	0.652

t_{max}, for $\sigma = 0.37$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	3.792	2.982	2.114	2.042	2.158	1.217	1.219	0.838
0.71	2.376	2.621	2.015	1.676	1.659	1.167	1.196	0.716
0.93	3.645	2.515	1.890	1.142	1.130	0.898	0.704	0.729
1.14	2.338	2.400	1.657	1.523	1.151	0.996	1.107	0.756
1.36	2.712	1.748	1.018	0.905	1.095	0.819	0.569	0.531
1.57	2.091	1.525	0.892	1.133	0.907	0.848	0.596	0.736
1.79	1.857	1.738	1.252	1.000	0.907	0.933	0.639	0.533
2.00	2.033	1.548	1.423	1.021	0.520	0.491	0.752	0.716

Continued overleaf

Table A.1: Calculated optimal t_{max} values for parameter estimation by the proposed method of (5.25), continued from previous page.

	t_{max}, for $\sigma = 0.64$							
α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	1.922	1.315	1.098	1.211	1.431	1.295	1.055	0.749
0.71	2.299	1.560	1.574	1.611	1.229	1.138	0.820	0.777
0.93	2.009	1.692	1.582	1.537	0.998	0.899	0.851	0.601
1.14	1.566	1.281	1.317	1.204	1.019	0.830	0.858	0.788
1.36	1.752	1.721	0.985	1.097	1.074	0.936	0.703	0.915
1.57	1.746	1.422	1.078	1.199	0.759	0.694	0.920	0.844
1.79	1.741	1.311	0.898	0.939	0.864	0.660	0.633	0.871
2.00	1.273	0.919	0.829	0.684	0.732	0.768	0.655	0.896

	t_{max}, for $\sigma = 0.91$							
α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	1.006	1.366	1.188	1.058	1.133	0.961	1.015	0.854
0.71	1.585	1.135	0.998	0.946	1.214	0.759	0.907	0.777
0.93	1.384	1.096	1.047	1.352	0.847	0.674	0.754	0.712
1.14	1.297	1.229	1.158	0.812	0.876	0.733	0.866	0.813
1.36	1.230	1.254	1.001	0.842	1.063	0.605	0.811	0.927
1.57	1.530	0.973	0.934	0.897	0.898	0.485	0.753	0.949
1.79	1.317	0.973	0.777	0.675	0.926	0.809	0.441	0.682
2.00	1.030	0.914	0.883	0.858	0.459	0.641	0.492	0.911

	t_{max}, for $\sigma = 1.19$							
α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	1.004	1.086	1.080	1.053	0.829	0.919	0.595	0.936
0.71	1.043	1.206	0.793	0.876	0.672	0.652	0.635	0.783
0.93	1.002	0.872	0.947	0.842	0.806	0.783	0.755	0.729
1.14	1.147	0.905	0.969	0.864	0.751	0.573	0.642	0.742
1.36	1.022	1.029	0.685	0.757	0.623	0.690	0.595	0.859
1.57	0.782	0.823	0.795	0.863	0.592	0.586	0.501	0.547
1.79	1.048	0.860	0.893	0.872	0.651	0.573	0.599	0.654
2.00	0.863	1.012	0.716	0.695	0.600	0.403	0.452	0.629

Continued overleaf

Table A.1: Calculated optimal t_{max} values for parameter estimation by the proposed method of (5.25), continued from previous page.

	t_{max}, for $\sigma = 1.46$							
α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	0.737	0.896	0.713	0.884	0.865	0.762	0.599	0.577
0.71	0.789	0.843	0.844	0.525	0.673	0.569	0.556	0.848
0.93	0.853	0.885	0.815	0.576	0.635	0.709	0.597	0.717
1.14	0.886	0.850	0.770	0.535	0.455	0.438	0.430	0.757
1.36	0.993	0.668	0.742	0.592	0.630	0.762	0.600	0.672
1.57	0.966	0.783	0.692	0.564	0.742	0.469	0.521	0.772
1.79	0.765	0.677	0.800	0.598	0.646	0.448	0.442	0.651
2.00	0.954	0.934	0.506	0.569	0.550	0.552	0.484	0.533

	t_{max}, for $\sigma = 1.73$							
α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	0.776	0.765	0.546	0.618	0.605	0.681	0.530	0.513
0.71	0.748	0.720	0.587	0.512	0.807	0.436	0.639	0.699
0.93	0.784	0.641	0.696	0.490	0.565	0.614	0.664	0.795
1.14	0.890	0.740	0.664	0.662	0.593	0.526	0.517	0.570
1.36	0.675	0.772	0.711	0.700	0.587	0.587	0.573	0.672
1.57	0.645	0.692	0.733	0.615	0.499	0.482	0.421	0.589
1.79	0.802	0.541	0.645	0.589	0.470	0.635	0.465	0.625
2.00	0.769	0.575	0.735	0.570	0.631	0.455	0.515	0.711

	t_{max}, for $\sigma = 2.00$							
α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	0.622	0.738	0.666	0.598	0.591	0.538	0.470	0.397
0.71	0.523	0.580	0.701	0.517	0.631	0.522	0.449	0.636
0.93	0.703	0.557	0.437	0.432	0.503	0.566	0.554	0.777
1.14	0.664	0.650	0.475	0.659	0.476	0.525	0.524	0.685
1.36	0.767	0.517	0.744	0.516	0.356	0.456	0.398	0.670
1.57	0.780	0.611	0.437	0.499	0.767	0.449	0.393	0.714
1.79	0.604	0.643	0.435	0.523	0.541	0.334	0.376	0.728
2.00	0.587	0.667	0.473	0.537	0.420	0.471	0.616	0.640

Appendix B

Supplementary Results from Chapter 6

B.1 Parameter Estimation

The method proposed in Section 5.3.2 for the estimation of the parameters of a combined SαS and Gaussian distribution was tested by Monte Carlo simulation as described in Section 6.2.1. The results of this test were reported and discussed in Section 6.3.1 above. Tables B.1-B.3 presented below contain the full set of results from this test. Each table contains the error measured in the estimation of one of the three parameters of the combined distribution (α, γ and σ).

Table B.1: α parameter estimation errors, for various values of α, γ and σ.

$\xi(\hat{\alpha})$, for $\sigma = 0.10$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	7.10×10^{-3}	8.80×10^{-3}	9.11×10^{-3}	6.25×10^{-3}	5.48×10^{-3}	4.93×10^{-3}	2.55×10^{-3}	1.50×10^{-2}
0.71	7.74×10^{-3}	9.23×10^{-3}	5.30×10^{-3}	4.65×10^{-3}	4.40×10^{-3}	3.49×10^{-3}	3.08×10^{-3}	1.68×10^{-2}
0.93	8.36×10^{-3}	6.16×10^{-3}	8.03×10^{-3}	6.00×10^{-3}	4.29×10^{-3}	4.96×10^{-3}	3.07×10^{-3}	1.78×10^{-2}
1.14	8.87×10^{-3}	9.31×10^{-3}	7.63×10^{-3}	6.43×10^{-3}	3.25×10^{-3}	4.93×10^{-3}	2.05×10^{-3}	1.66×10^{-2}
1.36	9.11×10^{-3}	7.89×10^{-3}	9.03×10^{-3}	9.35×10^{-3}	4.70×10^{-3}	3.81×10^{-3}	2.72×10^{-3}	1.81×10^{-2}
1.57	1.10×10^{-2}	8.84×10^{-3}	6.22×10^{-3}	8.09×10^{-3}	3.45×10^{-3}	3.71×10^{-3}	2.90×10^{-3}	1.66×10^{-2}
1.79	1.10×10^{-2}	9.17×10^{-3}	7.20×10^{-3}	5.78×10^{-3}	5.59×10^{-3}	4.12×10^{-3}	4.05×10^{-3}	1.62×10^{-2}
2.00	1.15×10^{-2}	7.48×10^{-3}	6.84×10^{-3}	5.72×10^{-3}	5.72×10^{-3}	3.85×10^{-3}	3.53×10^{-3}	1.44×10^{-2}

$\xi(\hat{\alpha})$, for $\sigma = 0.37$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	5.68×10^{-3}	6.36×10^{-3}	9.99×10^{-3}	1.02×10^{-2}	1.52×10^{-2}	1.49×10^{-2}	6.22×10^{-3}	1.52×10^{-2}
0.71	5.74×10^{-3}	1.11×10^{-2}	7.08×10^{-3}	1.17×10^{-2}	1.08×10^{-2}	1.67×10^{-2}	8.23×10^{-3}	1.39×10^{-2}
0.93	6.97×10^{-3}	7.91×10^{-3}	8.43×10^{-3}	9.91×10^{-3}	1.07×10^{-2}	8.62×10^{-3}	8.48×10^{-3}	1.20×10^{-2}
1.14	8.23×10^{-3}	1.05×10^{-2}	1.25×10^{-2}	1.48×10^{-2}	1.35×10^{-2}	1.70×10^{-2}	4.84×10^{-3}	1.71×10^{-2}
1.36	1.08×10^{-2}	9.90×10^{-3}	1.13×10^{-2}	1.25×10^{-2}	1.15×10^{-2}	7.73×10^{-3}	6.05×10^{-3}	1.70×10^{-2}
1.57	1.06×10^{-2}	7.67×10^{-3}	1.27×10^{-2}	1.68×10^{-2}	1.44×10^{-2}	9.01×10^{-3}	3.93×10^{-3}	1.53×10^{-2}
1.79	1.26×10^{-2}	1.61×10^{-2}	1.61×10^{-2}	8.02×10^{-3}	8.26×10^{-3}	4.00×10^{-2}	5.41×10^{-3}	1.30×10^{-2}
2.00	1.21×10^{-2}	1.35×10^{-2}	1.55×10^{-2}	1.57×10^{-2}	1.75×10^{-2}	9.67×10^{-3}	4.74×10^{-3}	1.33×10^{-2}

$\xi(\hat{\alpha})$, for $\sigma = 0.64$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	1.22×10^{-2}	1.03×10^{-2}	1.17×10^{-2}	1.39×10^{-2}	1.30×10^{-2}	2.44×10^{-2}	1.54×10^{-2}	1.48×10^{-2}
0.71	9.09×10^{-3}	1.03×10^{-2}	7.26×10^{-3}	1.26×10^{-2}	1.24×10^{-2}	2.10×10^{-2}	2.07×10^{-2}	1.59×10^{-2}
0.93	6.95×10^{-3}	9.43×10^{-3}	1.04×10^{-2}	1.09×10^{-2}	2.63×10^{-2}	1.52×10^{-2}	2.46×10^{-2}	1.68×10^{-2}
1.14	1.01×10^{-2}	1.14×10^{-2}	1.38×10^{-2}	1.30×10^{-2}	2.39×10^{-2}	1.21×10^{-2}	1.93×10^{-2}	1.39×10^{-2}
1.36	9.13×10^{-3}	1.29×10^{-2}	1.03×10^{-2}	1.20×10^{-2}	1.29×10^{-2}	1.76×10^{-2}	5.98×10^{-3}	1.38×10^{-2}
1.57	1.16×10^{-2}	1.01×10^{-2}	1.34×10^{-2}	1.50×10^{-2}	1.52×10^{-2}	1.50×10^{-2}	8.27×10^{-3}	1.26×10^{-2}
1.79	1.04×10^{-2}	1.10×10^{-2}	1.43×10^{-2}	1.41×10^{-2}	1.39×10^{-2}	1.47×10^{-2}	6.33×10^{-3}	1.15×10^{-2}
2.00	1.02×10^{-2}	1.37×10^{-2}	1.25×10^{-2}	9.78×10^{-3}	1.26×10^{-2}	1.67×10^{-2}	1.05×10^{-2}	1.14×10^{-2}

Continued overleaf

Table B.1: α parameter estimation errors, for various values of α, γ and σ, continued from previous page

$\xi(\hat{\alpha})$, for $\sigma = 0.91$

α	$\gamma=0.50$	$\gamma=0.71$	$\gamma=0.93$	$\gamma=1.14$	$\gamma=1.36$	$\gamma=1.57$	$\gamma=1.79$	$\gamma=2.00$
0.50	1.29×10^{-2}	1.06×10^{-2}	1.45×10^{-2}	1.59×10^{-2}	2.06×10^{-2}	2.10×10^{-2}	2.99×10^{-2}	1.67×10^{-2}
0.71	1.07×10^{-2}	1.35×10^{-2}	7.92×10^{-3}	1.17×10^{-2}	1.57×10^{-2}	1.87×10^{-2}	1.38×10^{-2}	1.67×10^{-2}
0.93	1.04×10^{-2}	1.42×10^{-2}	1.37×10^{-2}	1.69×10^{-2}	1.57×10^{-2}	1.88×10^{-2}	1.51×10^{-2}	1.52×10^{-2}
1.14	1.07×10^{-2}	9.72×10^{-3}	9.90×10^{-3}	1.22×10^{-2}	1.43×10^{-2}	1.67×10^{-2}	2.74×10^{-2}	1.70×10^{-2}
1.36	9.75×10^{-3}	1.18×10^{-2}	1.45×10^{-2}	9.84×10^{-3}	1.21×10^{-2}	1.70×10^{-2}	1.36×10^{-2}	1.45×10^{-2}
1.57	1.35×10^{-2}	1.18×10^{-2}	1.33×10^{-2}	1.39×10^{-2}	1.40×10^{-2}	1.57×10^{-2}	1.46×10^{-2}	1.40×10^{-2}
1.79	1.24×10^{-2}	1.31×10^{-2}	1.34×10^{-2}	1.64×10^{-2}	1.49×10^{-2}	4.48×10^{-2}	1.10×10^{-2}	1.37×10^{-2}
2.00	1.36×10^{-2}	1.03×10^{-2}	1.32×10^{-2}	1.55×10^{-2}	1.44×10^{-2}	4.32×10^{-2}	6.70×10^{-3}	1.39×10^{-2}

$\xi(\hat{\alpha})$, for $\sigma = 1.19$

α	$\gamma=0.50$	$\gamma=0.71$	$\gamma=0.93$	$\gamma=1.14$	$\gamma=1.36$	$\gamma=1.57$	$\gamma=1.79$	$\gamma=2.00$
0.50	1.72×10^{-2}	1.30×10^{-2}	1.99×10^{-2}	2.20×10^{-2}	2.71×10^{-2}	2.76×10^{-2}	2.42×10^{-2}	1.44×10^{-2}
0.71	1.41×10^{-2}	1.08×10^{-2}	1.22×10^{-2}	1.82×10^{-2}	1.38×10^{-2}	2.33×10^{-2}	2.37×10^{-2}	1.62×10^{-2}
0.93	1.53×10^{-2}	9.08×10^{-3}	1.12×10^{-2}	1.48×10^{-2}	1.37×10^{-2}	3.03×10^{-2}	2.28×10^{-2}	1.57×10^{-2}
1.14	1.23×10^{-2}	1.24×10^{-2}	1.12×10^{-2}	1.40×10^{-2}	1.64×10^{-2}	1.88×10^{-2}	3.26×10^{-2}	1.44×10^{-2}
1.36	1.19×10^{-2}	1.10×10^{-2}	1.27×10^{-2}	1.67×10^{-2}	1.49×10^{-2}	2.09×10^{-2}	2.57×10^{-2}	1.35×10^{-2}
1.57	1.50×10^{-2}	1.27×10^{-2}	9.32×10^{-3}	1.24×10^{-2}	1.55×10^{-2}	3.35×10^{-2}	2.91×10^{-2}	1.50×10^{-2}
1.79	1.01×10^{-2}	1.01×10^{-2}	1.09×10^{-2}	1.35×10^{-2}	1.89×10^{-2}	2.13×10^{-2}	3.32×10^{-2}	1.30×10^{-2}
2.00	1.40×10^{-2}	1.13×10^{-2}	1.28×10^{-2}	1.40×10^{-2}	1.16×10^{-2}	1.74×10^{-2}	1.35×10^{-2}	1.46×10^{-2}

$\xi(\hat{\alpha})$, for $\sigma = 1.46$

α	$\gamma=0.50$	$\gamma=0.71$	$\gamma=0.93$	$\gamma=1.14$	$\gamma=1.36$	$\gamma=1.57$	$\gamma=1.79$	$\gamma=2.00$
0.50	1.75×10^{-2}	1.33×10^{-2}	2.56×10^{-2}	2.21×10^{-2}	2.33×10^{-2}	3.72×10^{-2}	2.30×10^{-2}	1.33×10^{-2}
0.71	1.31×10^{-2}	1.98×10^{-2}	1.61×10^{-2}	1.97×10^{-2}	2.42×10^{-2}	3.57×10^{-2}	2.74×10^{-2}	1.37×10^{-2}
0.93	1.32×10^{-2}	1.39×10^{-2}	1.26×10^{-2}	1.58×10^{-2}	2.06×10^{-2}	3.98×10^{-2}	1.12×10^{-2}	1.48×10^{-2}
1.14	1.13×10^{-2}	1.18×10^{-2}	1.29×10^{-2}	1.75×10^{-2}	2.87×10^{-2}	2.63×10^{-2}	3.44×10^{-2}	1.41×10^{-2}
1.36	1.05×10^{-2}	1.49×10^{-2}	1.80×10^{-2}	2.41×10^{-2}	2.02×10^{-2}	1.54×10^{-2}	2.52×10^{-2}	1.15×10^{-2}
1.57	1.26×10^{-2}	1.21×10^{-2}	9.47×10^{-3}	1.56×10^{-2}	1.65×10^{-2}	2.26×10^{-2}	3.63×10^{-2}	1.35×10^{-2}
1.79	1.34×10^{-2}	1.45×10^{-2}	1.18×10^{-2}	1.09×10^{-2}	1.63×10^{-2}	1.88×10^{-2}	4.13×10^{-2}	1.38×10^{-2}
2.00	1.32×10^{-2}	9.61×10^{-3}	1.42×10^{-2}	1.38×10^{-2}	1.47×10^{-2}	2.42×10^{-2}	3.77×10^{-2}	1.32×10^{-2}

Continued overleaf

Table B.1: α parameter estimation errors, for various values of α, γ and σ, continued from previous page

$\xi(\hat{\alpha})$, for $\sigma = 1.73$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	2.52×10^{-2}	1.82×10^{-2}	2.24×10^{-2}	2.17×10^{-2}	3.01×10^{-2}	4.62×10^{-2}	2.45×10^{-2}	1.10×10^{-2}
0.71	1.60×10^{-2}	1.70×10^{-2}	1.00×10^{-2}	2.05×10^{-2}	2.74×10^{-2}	4.61×10^{-2}	2.68×10^{-2}	1.39×10^{-2}
0.93	1.83×10^{-2}	1.56×10^{-2}	1.66×10^{-2}	1.73×10^{-2}	2.04×10^{-2}	2.36×10^{-2}	2.88×10^{-2}	1.06×10^{-2}
1.14	1.46×10^{-2}	1.34×10^{-2}	1.74×10^{-2}	1.80×10^{-2}	2.01×10^{-2}	3.29×10^{-2}	2.07×10^{-2}	1.34×10^{-2}
1.36	1.44×10^{-2}	1.40×10^{-2}	1.21×10^{-2}	1.48×10^{-2}	2.27×10^{-2}	3.02×10^{-2}	3.93×10^{-2}	1.31×10^{-2}
1.57	1.29×10^{-2}	1.63×10^{-2}	1.11×10^{-2}	1.06×10^{-2}	2.29×10^{-2}	2.06×10^{-2}	3.84×10^{-2}	1.40×10^{-2}
1.79	1.34×10^{-2}	1.37×10^{-2}	1.32×10^{-2}	1.16×10^{-2}	1.54×10^{-2}	2.53×10^{-2}	2.19×10^{-2}	1.40×10^{-2}
2.00	1.44×10^{-2}	1.77×10^{-2}	1.24×10^{-2}	1.14×10^{-2}	1.55×10^{-2}	2.40×10^{-2}	4.64×10^{-2}	1.48×10^{-2}

$\xi(\hat{\alpha})$, for $\sigma = 2.00$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	2.44×10^{-2}	3.18×10^{-2}	2.00×10^{-2}	3.38×10^{-2}	3.29×10^{-2}	6.09×10^{-2}	2.08×10^{-2}	1.29×10^{-2}
0.71	1.47×10^{-2}	1.53×10^{-2}	2.32×10^{-2}	2.22×10^{-2}	2.73×10^{-2}	4.22×10^{-2}	5.38×10^{-2}	1.34×10^{-2}
0.93	2.17×10^{-2}	1.91×10^{-2}	1.19×10^{-2}	1.91×10^{-2}	2.29×10^{-2}	4.71×10^{-2}	3.46×10^{-2}	1.16×10^{-2}
1.14	1.21×10^{-2}	1.42×10^{-2}	1.41×10^{-2}	1.73×10^{-2}	3.35×10^{-2}	3.67×10^{-2}	3.05×10^{-2}	1.18×10^{-2}
1.36	1.53×10^{-2}	1.19×10^{-2}	1.77×10^{-2}	1.45×10^{-2}	1.96×10^{-2}	2.29×10^{-2}	4.10×10^{-2}	1.23×10^{-2}
1.57	1.20×10^{-2}	1.29×10^{-2}	1.23×10^{-2}	2.83×10^{-2}	1.98×10^{-2}	6.09×10^{-2}	4.58×10^{-2}	1.37×10^{-2}
1.79	1.15×10^{-2}	1.41×10^{-2}	1.21×10^{-2}	1.41×10^{-2}	1.98×10^{-2}	4.53×10^{-2}	3.05×10^{-2}	1.20×10^{-2}
2.00	1.14×10^{-2}	1.43×10^{-2}	9.57×10^{-3}	1.59×10^{-2}	2.29×10^{-2}	2.66×10^{-2}	2.27×10^{-2}	1.51×10^{-2}

Table B.2: γ parameter estimation errors, for various values of α, γ and σ.

$\xi(\hat{\gamma})$, for $\sigma = 0.10$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	5.69×10^{-3}	7.51×10^{-3}	9.63×10^{-3}	6.79×10^{-3}	6.28×10^{-3}	6.10×10^{-3}	8.10×10^{-3}	2.16×10^{-2}
0.71	4.19×10^{-3}	6.50×10^{-3}	5.12×10^{-3}	5.79×10^{-3}	6.45×10^{-3}	5.54×10^{-3}	5.89×10^{-3}	1.26×10^{-2}
0.93	4.80×10^{-3}	4.94×10^{-3}	5.45×10^{-3}	5.97×10^{-3}	6.59×10^{-3}	6.33×10^{-3}	6.47×10^{-3}	1.46×10^{-2}
1.14	2.95×10^{-3}	5.10×10^{-3}	4.90×10^{-3}	5.25×10^{-3}	5.65×10^{-3}	7.11×10^{-3}	4.96×10^{-3}	1.05×10^{-2}
1.36	2.84×10^{-3}	3.89×10^{-3}	6.08×10^{-3}	9.76×10^{-3}	5.19×10^{-3}	5.87×10^{-3}	5.03×10^{-3}	7.60×10^{-3}
1.57	4.18×10^{-3}	4.48×10^{-3}	5.11×10^{-3}	5.92×10^{-3}	6.16×10^{-3}	5.93×10^{-3}	6.03×10^{-3}	8.43×10^{-3}
1.79	3.97×10^{-3}	4.97×10^{-3}	5.15×10^{-3}	5.32×10^{-3}	6.44×10^{-3}	6.44×10^{-3}	7.06×10^{-3}	8.17×10^{-3}
2.00	4.82×10^{-3}	6.01×10^{-3}	6.32×10^{-3}	5.50×10^{-3}	6.59×10^{-3}	6.09×10^{-3}	7.28×10^{-3}	7.32×10^{-3}

$\xi(\hat{\gamma})$, for $\sigma = 0.37$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	5.00×10^{-3}	5.75×10^{-3}	9.31×10^{-3}	1.44×10^{-2}	2.95×10^{-2}	3.10×10^{-2}	6.38×10^{-3}	1.14×10^{-2}
0.71	5.19×10^{-3}	8.26×10^{-3}	8.88×10^{-3}	1.58×10^{-2}	2.04×10^{-2}	5.26×10^{-2}	6.87×10^{-3}	5.29×10^{-3}
0.93	4.96×10^{-3}	7.28×10^{-3}	9.12×10^{-3}	1.16×10^{-2}	2.53×10^{-2}	2.50×10^{-2}	7.91×10^{-3}	3.73×10^{-3}
1.14	4.60×10^{-3}	6.14×10^{-3}	7.01×10^{-3}	1.75×10^{-2}	1.83×10^{-2}	2.64×10^{-2}	6.99×10^{-3}	5.56×10^{-3}
1.36	4.20×10^{-3}	6.56×10^{-3}	6.12×10^{-3}	1.29×10^{-2}	1.40×10^{-2}	1.27×10^{-2}	6.30×10^{-3}	4.48×10^{-3}
1.57	4.43×10^{-3}	5.65×10^{-3}	6.05×10^{-3}	1.29×10^{-2}	2.01×10^{-2}	1.15×10^{-2}	5.82×10^{-3}	4.12×10^{-3}
1.79	4.96×10^{-3}	7.83×10^{-3}	9.35×10^{-3}	1.08×10^{-2}	1.37×10^{-2}	1.73×10^{-2}	7.19×10^{-3}	5.13×10^{-3}
2.00	4.61×10^{-3}	6.61×10^{-3}	1.14×10^{-2}	1.28×10^{-2}	3.84×10^{-2}	1.27×10^{-2}	6.14×10^{-3}	5.08×10^{-3}

$\xi(\hat{\gamma})$, for $\sigma = 0.64$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	7.70×10^{-3}	1.04×10^{-2}	1.60×10^{-2}	1.84×10^{-2}	3.27×10^{-2}	9.78×10^{-2}	1.64×10^{-1}	1.59×10^{-2}
0.71	6.18×10^{-3}	7.98×10^{-3}	8.64×10^{-3}	1.21×10^{-2}	2.86×10^{-2}	9.35×10^{-2}	4.77×10^{-2}	8.44×10^{-3}
0.93	3.87×10^{-3}	7.49×10^{-3}	9.83×10^{-3}	1.45×10^{-2}	3.80×10^{-2}	6.96×10^{-2}	5.76×10^{-2}	7.27×10^{-3}
1.14	5.37×10^{-3}	1.00×10^{-2}	1.45×10^{-2}	1.55×10^{-2}	3.68×10^{-2}	5.50×10^{-2}	2.44×10^{-2}	6.93×10^{-3}
1.36	4.27×10^{-3}	8.94×10^{-3}	1.04×10^{-2}	1.60×10^{-2}	2.88×10^{-2}	5.77×10^{-2}	2.24×10^{-2}	6.41×10^{-3}
1.57	4.29×10^{-3}	6.14×10^{-3}	8.73×10^{-3}	1.83×10^{-2}	2.92×10^{-2}	5.09×10^{-2}	2.06×10^{-2}	4.69×10^{-3}
1.79	5.46×10^{-3}	7.03×10^{-3}	1.07×10^{-2}	2.61×10^{-2}	3.17×10^{-2}	5.40×10^{-2}	1.64×10^{-2}	4.88×10^{-3}
2.00	5.16×10^{-3}	9.73×10^{-3}	1.05×10^{-2}	1.26×10^{-2}	3.46×10^{-2}	4.27×10^{-2}	1.42×10^{-2}	5.75×10^{-3}

Continued overleaf

Table B.2: γ parameter estimation errors, for various values of α, γ and σ, continued from previous page.

$\xi(\hat{\gamma})$, for $\sigma = 0.91$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	1.04×10^{-2}	1.20×10^{-2}	2.07×10^{-2}	3.04×10^{-2}	5.25×10^{-2}	8.47×10^{-2}	2.75×10^{-1}	1.55×10^{-2}
0.71	7.45×10^{-3}	1.49×10^{-2}	1.21×10^{-2}	1.72×10^{-2}	4.01×10^{-2}	8.76×10^{-2}	1.86×10^{-1}	1.30×10^{-2}
0.93	7.30×10^{-3}	1.13×10^{-2}	2.00×10^{-2}	2.26×10^{-2}	3.35×10^{-2}	6.74×10^{-2}	1.58×10^{-1}	1.54×10^{-2}
1.14	6.31×10^{-3}	9.25×10^{-3}	1.33×10^{-2}	1.77×10^{-2}	3.03×10^{-2}	7.97×10^{-2}	8.83×10^{-2}	9.79×10^{-3}
1.36	4.39×10^{-3}	7.23×10^{-3}	9.85×10^{-3}	1.49×10^{-2}	2.66×10^{-2}	7.61×10^{-2}	6.64×10^{-2}	9.18×10^{-3}
1.57	5.95×10^{-3}	1.02×10^{-2}	1.77×10^{-2}	2.36×10^{-2}	2.68×10^{-2}	6.99×10^{-2}	7.84×10^{-2}	7.42×10^{-3}
1.79	7.20×10^{-3}	8.09×10^{-3}	1.09×10^{-2}	2.32×10^{-2}	4.01×10^{-2}	5.88×10^{-2}	4.33×10^{-2}	7.82×10^{-3}
2.00	5.53×10^{-3}	7.65×10^{-3}	1.20×10^{-2}	1.90×10^{-2}	3.83×10^{-2}	6.59×10^{-2}	4.41×10^{-2}	7.52×10^{-3}

$\xi(\hat{\gamma})$, for $\sigma = 1.19$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	1.32×10^{-2}	1.30×10^{-2}	3.14×10^{-2}	4.25×10^{-2}	9.09×10^{-2}	1.25×10^{-1}	3.83×10^{-1}	1.43×10^{-2}
0.71	1.06×10^{-2}	1.07×10^{-2}	1.83×10^{-2}	3.26×10^{-2}	4.53×10^{-2}	1.38×10^{-1}	2.75×10^{-1}	1.40×10^{-2}
0.93	1.00×10^{-2}	1.06×10^{-2}	1.72×10^{-2}	3.10×10^{-2}	4.15×10^{-2}	1.17×10^{-1}	2.61×10^{-1}	1.05×10^{-2}
1.14	7.67×10^{-3}	1.22×10^{-2}	1.33×10^{-2}	1.89×10^{-2}	5.13×10^{-2}	8.95×10^{-2}	2.13×10^{-1}	1.07×10^{-2}
1.36	8.34×10^{-3}	1.10×10^{-2}	1.77×10^{-2}	2.65×10^{-2}	4.42×10^{-2}	1.07×10^{-1}	1.72×10^{-1}	1.09×10^{-2}
1.57	7.66×10^{-3}	1.13×10^{-2}	1.18×10^{-2}	2.48×10^{-2}	3.70×10^{-2}	9.02×10^{-2}	1.14×10^{-1}	1.27×10^{-2}
1.79	6.56×10^{-3}	8.84×10^{-3}	1.48×10^{-2}	2.20×10^{-2}	6.09×10^{-2}	8.62×10^{-2}	1.21×10^{-1}	1.07×10^{-2}
2.00	8.83×10^{-3}	9.77×10^{-3}	1.74×10^{-2}	2.84×10^{-2}	3.85×10^{-2}	8.50×10^{-2}	9.46×10^{-2}	1.15×10^{-2}

$\xi(\hat{\gamma})$, for $\sigma = 1.46$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	1.43×10^{-2}	1.73×10^{-2}	4.21×10^{-2}	5.55×10^{-2}	7.23×10^{-2}	2.03×10^{-1}	4.17×10^{-1}	2.05×10^{-2}
0.71	9.55×10^{-3}	1.90×10^{-2}	2.33×10^{-2}	4.26×10^{-2}	7.59×10^{-2}	1.87×10^{-1}	2.09×10^{-1}	1.35×10^{-2}
0.93	9.55×10^{-3}	1.41×10^{-2}	1.77×10^{-2}	3.36×10^{-2}	6.26×10^{-2}	1.23×10^{-1}	4.43×10^{-2}	1.20×10^{-2}
1.14	6.49×10^{-3}	1.26×10^{-2}	1.55×10^{-2}	3.11×10^{-2}	8.59×10^{-2}	1.35×10^{-1}	2.24×10^{-1}	1.18×10^{-2}
1.36	6.64×10^{-3}	9.33×10^{-3}	2.27×10^{-2}	4.52×10^{-2}	6.67×10^{-2}	1.25×10^{-1}	2.04×10^{-1}	8.44×10^{-3}
1.57	7.99×10^{-3}	1.14×10^{-2}	1.22×10^{-2}	3.09×10^{-2}	5.83×10^{-2}	1.19×10^{-1}	1.92×10^{-1}	9.96×10^{-3}
1.79	6.21×10^{-3}	1.50×10^{-2}	1.77×10^{-2}	1.98×10^{-2}	5.55×10^{-2}	8.29×10^{-2}	2.36×10^{-1}	7.79×10^{-3}
2.00	9.73×10^{-3}	9.03×10^{-3}	1.36×10^{-2}	3.24×10^{-2}	5.29×10^{-2}	1.49×10^{-1}	1.71×10^{-1}	6.63×10^{-3}

Continued overleaf

Table B.2: γ parameter estimation errors, for various values of α, γ and σ, continued from previous page.

$\xi(\hat{\gamma})$, for $\sigma = 1.73$

α	$\gamma=0.50$	$\gamma=0.71$	$\gamma=0.93$	$\gamma=1.14$	$\gamma=1.36$	$\gamma=1.57$	$\gamma=1.79$	$\gamma=2.00$
0.50	2.32×10^{-2}	2.32×10^{-2}	3.80×10^{-2}	5.64×10^{-2}	9.13×10^{-2}	3.37×10^{-1}	9.38×10^{-1}	2.01×10^{-2}
0.71	1.32×10^{-2}	2.20×10^{-2}	1.82×10^{-2}	5.12×10^{-2}	9.16×10^{-2}	3.06×10^{-1}	6.42×10^{-1}	1.46×10^{-2}
0.93	1.41×10^{-2}	1.85×10^{-2}	2.79×10^{-2}	4.38×10^{-2}	6.79×10^{-2}	9.53×10^{-2}	2.26×10^{-1}	1.17×10^{-2}
1.14	1.18×10^{-2}	1.45×10^{-2}	2.20×10^{-2}	4.16×10^{-2}	5.38×10^{-2}	1.99×10^{-1}	3.86×10^{-1}	1.20×10^{-2}
1.36	1.09×10^{-2}	1.23×10^{-2}	1.83×10^{-2}	3.48×10^{-2}	7.69×10^{-2}	1.56×10^{-1}	2.84×10^{-1}	9.64×10^{-3}
1.57	9.44×10^{-3}	1.66×10^{-2}	2.20×10^{-2}	2.07×10^{-2}	8.34×10^{-2}	1.11×10^{-1}	2.24×10^{-1}	8.91×10^{-3}
1.79	1.11×10^{-2}	1.51×10^{-2}	1.86×10^{-2}	2.50×10^{-2}	5.15×10^{-2}	1.42×10^{-1}	2.17×10^{-1}	8.61×10^{-3}
2.00	8.25×10^{-3}	1.87×10^{-2}	2.10×10^{-2}	2.50×10^{-2}	4.92×10^{-2}	1.43×10^{-1}	1.96×10^{-1}	6.87×10^{-3}

$\xi(\hat{\gamma})$, for $\sigma = 2.00$

α	$\gamma=0.50$	$\gamma=0.71$	$\gamma=0.93$	$\gamma=1.14$	$\gamma=1.36$	$\gamma=1.57$	$\gamma=1.79$	$\gamma=2.00$
0.50	2.25×10^{-2}	4.29×10^{-2}	4.29×10^{-2}	8.87×10^{-2}	1.40×10^{-1}	6.52×10^{-1}	9.36×10^{-1}	2.35×10^{-2}
0.71	1.40×10^{-2}	1.83×10^{-2}	4.36×10^{-2}	5.22×10^{-2}	1.04×10^{-1}	2.64×10^{-1}	9.95×10^{-1}	1.63×10^{-2}
0.93	1.95×10^{-2}	2.49×10^{-2}	2.33×10^{-2}	5.04×10^{-2}	7.19×10^{-2}	2.66×10^{-1}	3.19×10^{-1}	1.42×10^{-2}
1.14	1.07×10^{-2}	1.79×10^{-2}	2.33×10^{-2}	4.26×10^{-2}	9.82×10^{-2}	2.03×10^{-1}	4.80×10^{-1}	1.32×10^{-2}
1.36	1.19×10^{-2}	1.23×10^{-2}	3.15×10^{-2}	3.92×10^{-2}	6.64×10^{-2}	1.50×10^{-1}	3.21×10^{-1}	1.00×10^{-2}
1.57	1.02×10^{-2}	1.16×10^{-2}	1.77×10^{-2}	4.84×10^{-2}	6.78×10^{-2}	2.13×10^{-1}	2.23×10^{-1}	8.36×10^{-3}
1.79	6.66×10^{-3}	1.60×10^{-2}	2.20×10^{-2}	3.53×10^{-2}	7.46×10^{-2}	1.55×10^{-1}	1.70×10^{-1}	7.27×10^{-3}
2.00	7.76×10^{-3}	1.60×10^{-2}	1.41×10^{-2}	3.61×10^{-2}	9.24×10^{-2}	1.91×10^{-1}	1.93×10^{-1}	7.60×10^{-3}

Table B.3: σ parameter estimation errors, for various values of α, γ and σ.

$\xi(\hat{\sigma})$, for $\sigma = 0.10$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	7.58×10^{-3}	1.19×10^{-2}	6.91×10^{-2}	1.11×10^{-2}	1.01×10^{-1}	2.35×10^{-2}	2.36×10^{-2}	3.24×10^{-2}
0.71	1.59×10^{-2}	1.12×10^{-2}	2.18×10^{-2}	4.03×10^{-2}	2.09×10^{-2}	2.19×10^{-2}	2.62×10^{-2}	3.53×10^{-2}
0.93	1.65×10^{-2}	3.61×10^{-2}	1.20×10^{-2}	1.07×10^{-1}	1.99×10^{-2}	2.54×10^{-2}	2.48×10^{-2}	3.29×10^{-2}
1.14	1.46×10^{-2}	9.37×10^{-3}	1.26×10^{-1}	1.97×10^{-2}	1.99×10^{-2}	2.45×10^{-2}	2.51×10^{-2}	4.53×10^{-2}
1.36	1.38×10^{-2}	1.03×10^{-2}	1.60×10^{-2}	3.94×10^{-1}	2.21×10^{-2}	2.52×10^{-2}	2.69×10^{-2}	4.03×10^{-2}
1.57	3.65×10^{-2}	3.97×10^{-2}	1.67×10^{-2}	2.06×10^{-2}	2.17×10^{-2}	2.55×10^{-2}	2.82×10^{-2}	2.89×10^{-2}
1.79	1.80×10^{-2}	1.65×10^{-2}	2.33×10^{-2}	2.58×10^{-2}	2.58×10^{-2}	2.67×10^{-2}	2.78×10^{-2}	3.07×10^{-2}
2.00	8.78×10^{-2}	2.07×10^{-1}	2.06×10^{-2}	2.35×10^{-2}	2.79×10^{-2}	2.76×10^{-2}	2.75×10^{-2}	4.06×10^{-2}

$\xi(\hat{\sigma})$, for $\sigma = 0.37$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	6.61×10^{-3}	1.08×10^{-2}	2.30×10^{-2}	3.27×10^{-2}	1.08×10^{-1}	1.27×10^{-1}	1.67×10^{-2}	2.66×10^{-2}
0.71	8.33×10^{-3}	1.80×10^{-2}	2.28×10^{-2}	6.09×10^{-2}	1.26×10^{-1}	3.15×10^{-1}	1.95×10^{-2}	2.54×10^{-2}
0.93	1.15×10^{-2}	1.85×10^{-2}	3.56×10^{-2}	6.53×10^{-2}	1.88×10^{-1}	1.97×10^{-1}	2.15×10^{-2}	3.07×10^{-2}
1.14	1.08×10^{-2}	2.50×10^{-2}	2.75×10^{-2}	1.19×10^{-1}	1.79×10^{-1}	1.74×10^{-1}	2.13×10^{-2}	2.58×10^{-2}
1.36	1.45×10^{-2}	2.90×10^{-2}	4.62×10^{-2}	9.29×10^{-2}	1.87×10^{-1}	1.29×10^{-1}	2.46×10^{-2}	2.62×10^{-2}
1.57	1.42×10^{-2}	2.80×10^{-2}	4.82×10^{-2}	1.01×10^{-1}	1.84×10^{-1}	1.31×10^{-1}	1.94×10^{-2}	2.72×10^{-2}
1.79	2.56×10^{-2}	4.56×10^{-2}	7.79×10^{-2}	1.41×10^{-1}	1.58×10^{-1}	2.24×10^{-1}	2.02×10^{-2}	2.66×10^{-2}
2.00	3.05×10^{-2}	5.60×10^{-2}	1.36×10^{-1}	9.07×10^{-2}	3.89×10^{-1}	2.08×10^{-1}	1.92×10^{-2}	3.00×10^{-2}

$\xi(\hat{\sigma})$, for $\sigma = 0.64$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	6.38×10^{-3}	9.10×10^{-3}	1.40×10^{-2}	1.92×10^{-2}	3.78×10^{-2}	1.31×10^{-1}	3.22×10^{-1}	1.37×10^{-2}
0.71	6.87×10^{-3}	1.09×10^{-2}	1.37×10^{-2}	1.94×10^{-2}	4.95×10^{-2}	2.14×10^{-1}	1.26×10^{-1}	1.46×10^{-2}
0.93	6.96×10^{-3}	1.25×10^{-2}	2.02×10^{-2}	3.07×10^{-2}	8.42×10^{-2}	2.13×10^{-1}	2.12×10^{-1}	1.64×10^{-2}
1.14	1.08×10^{-2}	2.17×10^{-2}	3.77×10^{-2}	4.47×10^{-2}	1.20×10^{-1}	2.50×10^{-1}	1.21×10^{-1}	1.96×10^{-2}
1.36	9.89×10^{-3}	2.63×10^{-2}	3.44×10^{-2}	5.68×10^{-2}	9.81×10^{-2}	2.55×10^{-1}	9.53×10^{-2}	2.05×10^{-2}
1.57	1.45×10^{-2}	1.91×10^{-2}	3.07×10^{-2}	7.34×10^{-2}	1.15×10^{-1}	2.17×10^{-1}	1.21×10^{-1}	1.92×10^{-2}
1.79	1.44×10^{-2}	2.28×10^{-2}	5.20×10^{-2}	1.42×10^{-1}	1.73×10^{-1}	2.75×10^{-1}	1.21×10^{-1}	1.77×10^{-2}
2.00	1.66×10^{-2}	4.10×10^{-2}	5.14×10^{-2}	6.28×10^{-2}	2.42×10^{-1}	2.97×10^{-1}	1.21×10^{-1}	2.01×10^{-2}

Continued overleaf

Table B.3: σ parameter estimation errors, for various values of α, γ and σ, continued from previous page.

$\xi(\hat{\sigma})$, for $\sigma = 0.91$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	6.97×10^{-3}	8.35×10^{-3}	1.21×10^{-2}	1.84×10^{-2}	3.08×10^{-2}	5.20×10^{-2}	2.85×10^{-1}	9.97×10^{-3}
0.71	6.79×10^{-3}	1.20×10^{-2}	1.12×10^{-2}	1.61×10^{-2}	3.60×10^{-2}	8.68×10^{-2}	2.79×10^{-1}	1.25×10^{-2}
0.93	8.30×10^{-3}	1.40×10^{-2}	2.19×10^{-2}	2.64×10^{-2}	4.08×10^{-2}	6.99×10^{-2}	2.91×10^{-1}	1.23×10^{-2}
1.14	1.09×10^{-2}	1.28×10^{-2}	2.10×10^{-2}	2.57×10^{-2}	4.34×10^{-2}	1.24×10^{-1}	2.20×10^{-1}	1.11×10^{-2}
1.36	7.19×10^{-3}	1.49×10^{-2}	2.02×10^{-2}	2.35×10^{-2}	4.75×10^{-2}	1.62×10^{-1}	1.05×10^{-1}	1.07×10^{-2}
1.57	1.45×10^{-2}	2.12×10^{-2}	3.15×10^{-2}	5.12×10^{-2}	5.85×10^{-2}	2.00×10^{-1}	2.82×10^{-1}	1.20×10^{-2}
1.79	1.61×10^{-2}	1.66×10^{-2}	2.60×10^{-2}	5.53×10^{-2}	9.26×10^{-2}	1.32×10^{-1}	1.79×10^{-1}	1.14×10^{-2}
2.00	1.27×10^{-2}	1.97×10^{-2}	2.71×10^{-2}	5.53×10^{-2}	9.82×10^{-2}	2.16×10^{-1}	1.79×10^{-1}	1.36×10^{-2}

$\xi(\hat{\sigma})$, for $\sigma = 1.19$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	6.14×10^{-3}	6.68×10^{-3}	1.37×10^{-2}	1.65×10^{-2}	3.48×10^{-2}	4.80×10^{-2}	2.21×10^{-1}	5.12×10^{-3}
0.71	7.65×10^{-3}	7.66×10^{-3}	1.19×10^{-2}	1.82×10^{-2}	2.38×10^{-2}	8.46×10^{-2}	2.59×10^{-1}	6.14×10^{-3}
0.93	8.28×10^{-3}	1.07×10^{-2}	1.51×10^{-2}	2.28×10^{-2}	3.00×10^{-2}	8.18×10^{-2}	3.19×10^{-1}	8.76×10^{-3}
1.14	8.48×10^{-3}	1.33×10^{-2}	1.25×10^{-2}	1.72×10^{-2}	4.90×10^{-2}	8.02×10^{-2}	3.00×10^{-1}	1.02×10^{-2}
1.36	1.08×10^{-2}	1.54×10^{-2}	1.98×10^{-2}	2.92×10^{-2}	5.14×10^{-2}	1.70×10^{-1}	2.48×10^{-1}	1.61×10^{-2}
1.57	1.18×10^{-2}	1.67×10^{-2}	1.66×10^{-2}	3.58×10^{-2}	4.55×10^{-2}	1.10×10^{-1}	1.99×10^{-1}	9.11×10^{-3}
1.79	1.18×10^{-2}	1.80×10^{-2}	2.65×10^{-2}	3.21×10^{-2}	1.14×10^{-1}	1.06×10^{-1}	3.01×10^{-1}	8.03×10^{-3}
2.00	1.25×10^{-2}	1.76×10^{-2}	3.06×10^{-2}	5.24×10^{-2}	6.45×10^{-2}	1.73×10^{-1}	2.01×10^{-1}	9.39×10^{-3}

$\xi(\hat{\sigma})$, for $\sigma = 1.46$

α	$\gamma = 0.50$	$\gamma = 0.71$	$\gamma = 0.93$	$\gamma = 1.14$	$\gamma = 1.36$	$\gamma = 1.57$	$\gamma = 1.79$	$\gamma = 2.00$
0.50	5.67×10^{-3}	6.67×10^{-3}	1.32×10^{-2}	1.49×10^{-2}	1.95×10^{-2}	5.43×10^{-2}	1.86×10^{-1}	3.19×10^{-3}
0.71	5.37×10^{-3}	9.18×10^{-3}	1.01×10^{-2}	1.81×10^{-2}	3.06×10^{-2}	7.30×10^{-2}	1.32×10^{-1}	4.05×10^{-3}
0.93	8.89×10^{-3}	9.70×10^{-3}	1.10×10^{-2}	1.67×10^{-2}	3.01×10^{-2}	6.43×10^{-2}	2.13×10^{-2}	6.15×10^{-3}
1.14	5.60×10^{-3}	9.58×10^{-3}	1.29×10^{-2}	2.04×10^{-2}	5.64×10^{-2}	9.71×10^{-2}	2.31×10^{-1}	6.33×10^{-3}
1.36	7.25×10^{-3}	8.61×10^{-3}	1.87×10^{-2}	3.38×10^{-2}	5.03×10^{-2}	1.42×10^{-1}	2.37×10^{-1}	1.01×10^{-2}
1.57	9.71×10^{-3}	1.36×10^{-2}	1.34×10^{-2}	3.03×10^{-2}	5.12×10^{-2}	1.44×10^{-1}	2.73×10^{-1}	9.85×10^{-3}
1.79	1.06×10^{-2}	1.90×10^{-2}	1.99×10^{-2}	2.26×10^{-2}	5.51×10^{-2}	7.51×10^{-2}	3.13×10^{-1}	1.09×10^{-2}
2.00	1.46×10^{-2}	1.33×10^{-2}	1.75×10^{-2}	4.18×10^{-2}	6.63×10^{-2}	2.19×10^{-1}	2.38×10^{-1}	1.38×10^{-2}

Continued overleaf

Table B.3: σ parameter estimation errors, for various values of α, γ and σ, continued from previous page.

$\xi(\hat{\sigma})$, for $\sigma = 1.73$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	7.05×10^{-3}	6.31×10^{-3}	1.04×10^{-2}	1.21×10^{-2}	1.81×10^{-2}	6.60×10^{-2}	2.99×10^{-1}	3.15×10^{-3}
0.71	6.88×10^{-3}	9.05×10^{-3}	7.85×10^{-3}	1.48×10^{-2}	2.65×10^{-2}	8.78×10^{-2}	2.83×10^{-1}	4.38×10^{-3}
0.93	8.57×10^{-3}	1.08×10^{-2}	1.31×10^{-2}	1.71×10^{-2}	2.45×10^{-2}	3.05×10^{-2}	1.36×10^{-1}	4.09×10^{-3}
1.14	8.09×10^{-3}	1.03×10^{-2}	1.26×10^{-2}	2.12×10^{-2}	2.35×10^{-2}	9.04×10^{-2}	2.99×10^{-1}	6.10×10^{-3}
1.36	9.10×10^{-3}	1.07×10^{-2}	1.42×10^{-2}	2.06×10^{-2}	4.05×10^{-2}	8.14×10^{-2}	2.67×10^{-1}	7.35×10^{-3}
1.57	9.01×10^{-3}	1.61×10^{-2}	1.99×10^{-2}	1.70×10^{-2}	5.32×10^{-2}	6.47×10^{-2}	1.90×10^{-1}	7.24×10^{-3}
1.79	1.15×10^{-2}	1.40×10^{-2}	1.70×10^{-2}	2.06×10^{-2}	3.95×10^{-2}	1.03×10^{-1}	1.85×10^{-1}	6.70×10^{-3}
2.00	1.19×10^{-2}	2.01×10^{-2}	2.35×10^{-2}	2.54×10^{-2}	4.34×10^{-2}	1.21×10^{-1}	1.64×10^{-1}	7.80×10^{-3}

$\xi(\hat{\sigma})$, for $\sigma = 2.00$

α	γ =0.50	γ =0.71	γ =0.93	γ =1.14	γ =1.36	γ =1.57	γ =1.79	γ =2.00
0.50	5.42×10^{-3}	9.95×10^{-3}	1.03×10^{-2}	1.60×10^{-2}	2.05×10^{-2}	1.44×10^{-1}	2.33×10^{-1}	2.83×10^{-3}
0.71	5.64×10^{-3}	6.09×10^{-3}	1.16×10^{-2}	1.38×10^{-2}	2.23×10^{-2}	5.33×10^{-2}	3.38×10^{-1}	4.00×10^{-3}
0.93	9.24×10^{-3}	1.02×10^{-2}	9.11×10^{-3}	1.53×10^{-2}	2.12×10^{-2}	7.91×10^{-2}	1.38×10^{-1}	4.40×10^{-3}
1.14	7.23×10^{-3}	8.59×10^{-3}	8.63×10^{-3}	1.87×10^{-2}	3.22×10^{-2}	6.73×10^{-2}	2.10×10^{-1}	4.47×10^{-3}
1.36	8.98×10^{-3}	8.03×10^{-3}	1.80×10^{-2}	1.92×10^{-2}	2.85×10^{-2}	6.26×10^{-2}	1.46×10^{-1}	5.77×10^{-3}
1.57	1.07×10^{-2}	8.57×10^{-3}	1.07×10^{-2}	2.17×10^{-2}	3.14×10^{-2}	1.09×10^{-1}	1.16×10^{-1}	5.46×10^{-3}
1.79	8.50×10^{-3}	1.10×10^{-2}	1.63×10^{-2}	2.23×10^{-2}	4.26×10^{-2}	7.98×10^{-2}	9.32×10^{-2}	7.06×10^{-3}
2.00	9.90×10^{-3}	1.70×10^{-2}	1.44×10^{-2}	2.49×10^{-2}	5.69×10^{-2}	1.21×10^{-1}	1.11×10^{-1}	6.52×10^{-3}

B.2 Edge Detection

The performance of the Gabor zero crossing edge detection method proposed in Chapter 5 was evaluated in the tests discussed in Section 6.2.2. These tests included the application of the proposed edge detector, and four other methods from the literature, to sets of simulated, phantom and clinical images. The results of these tests are discussed in Section 6.3.2 above. The full set of edge maps are presented here. Figs. B.1–B.6 display the output of the edge detectors when applied to simulated images, Figs. B.7–B.10 display the edge maps produced for the phantom images, and Figs. B.11-B.20 show the edge maps for the clinical images.

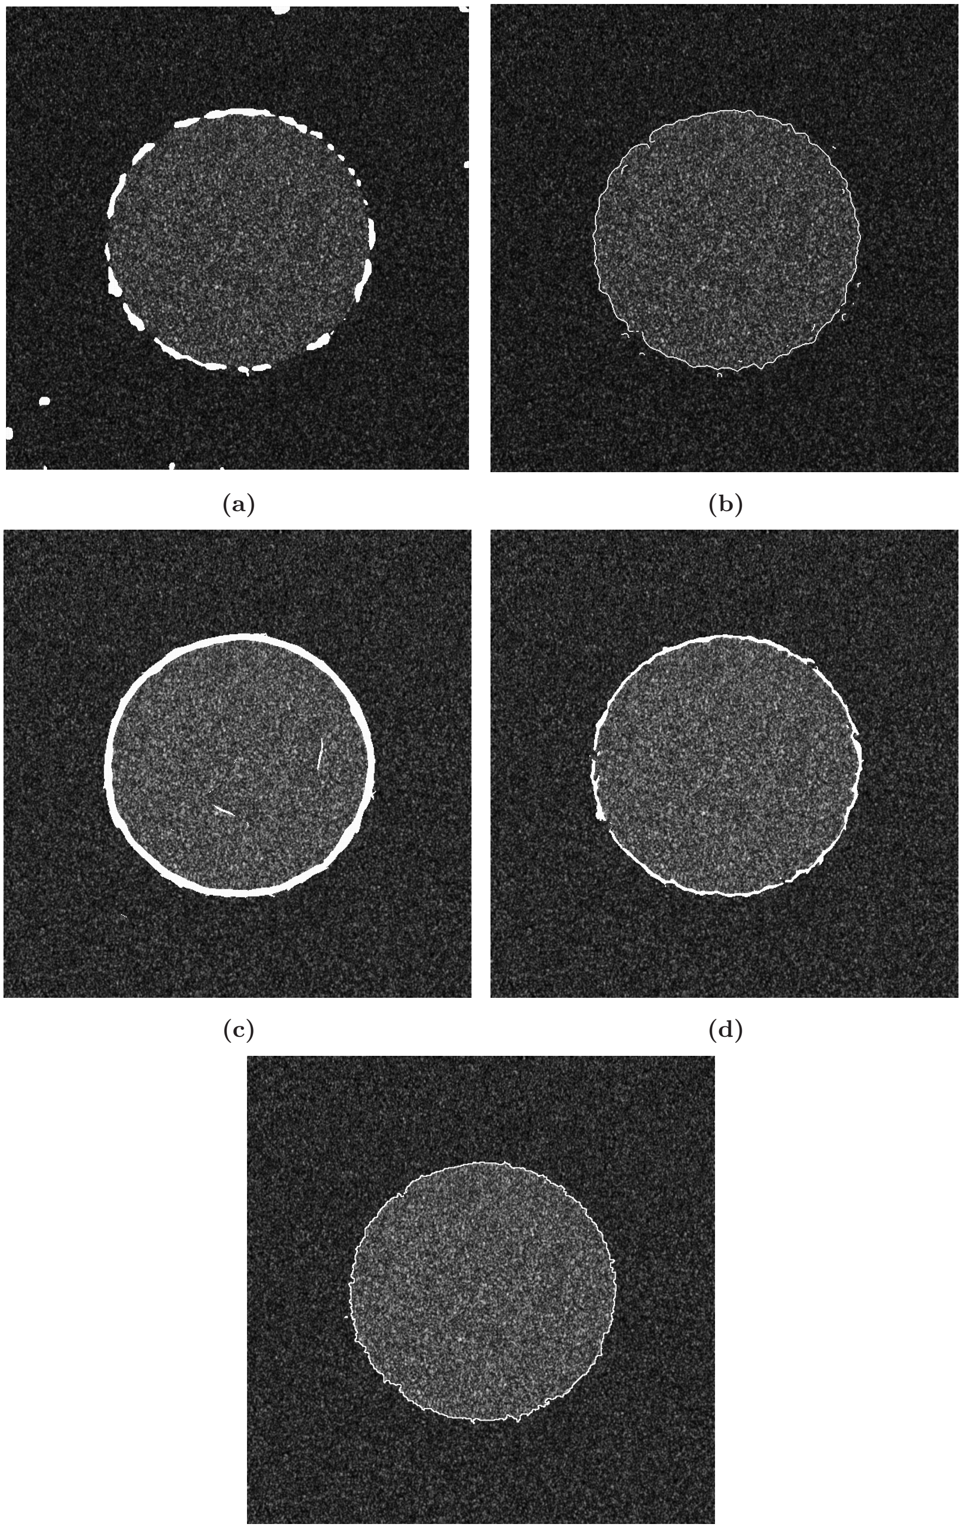

(a) **(b)**

(c) **(d)**

(e)

Figure B.1: Edge detector output for the simulated edge images of Fig. 6.1(a). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

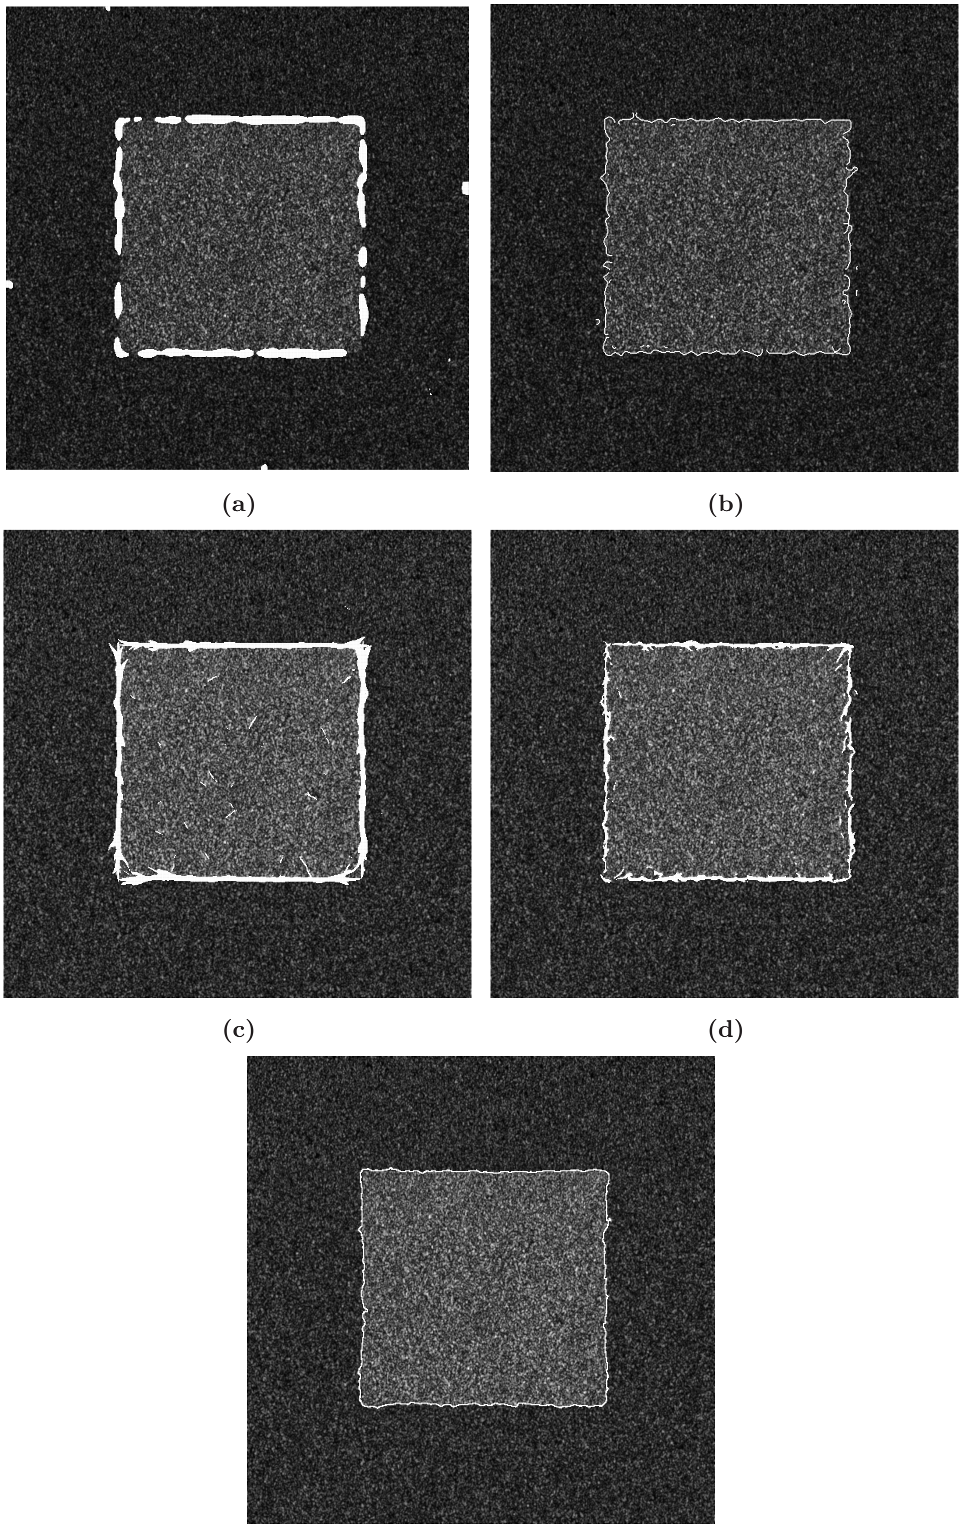

(a) **(b)**

(c) **(d)**

(e)

Figure B.2: Edge detector output for the simulated edge images of Fig. 6.1(b). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

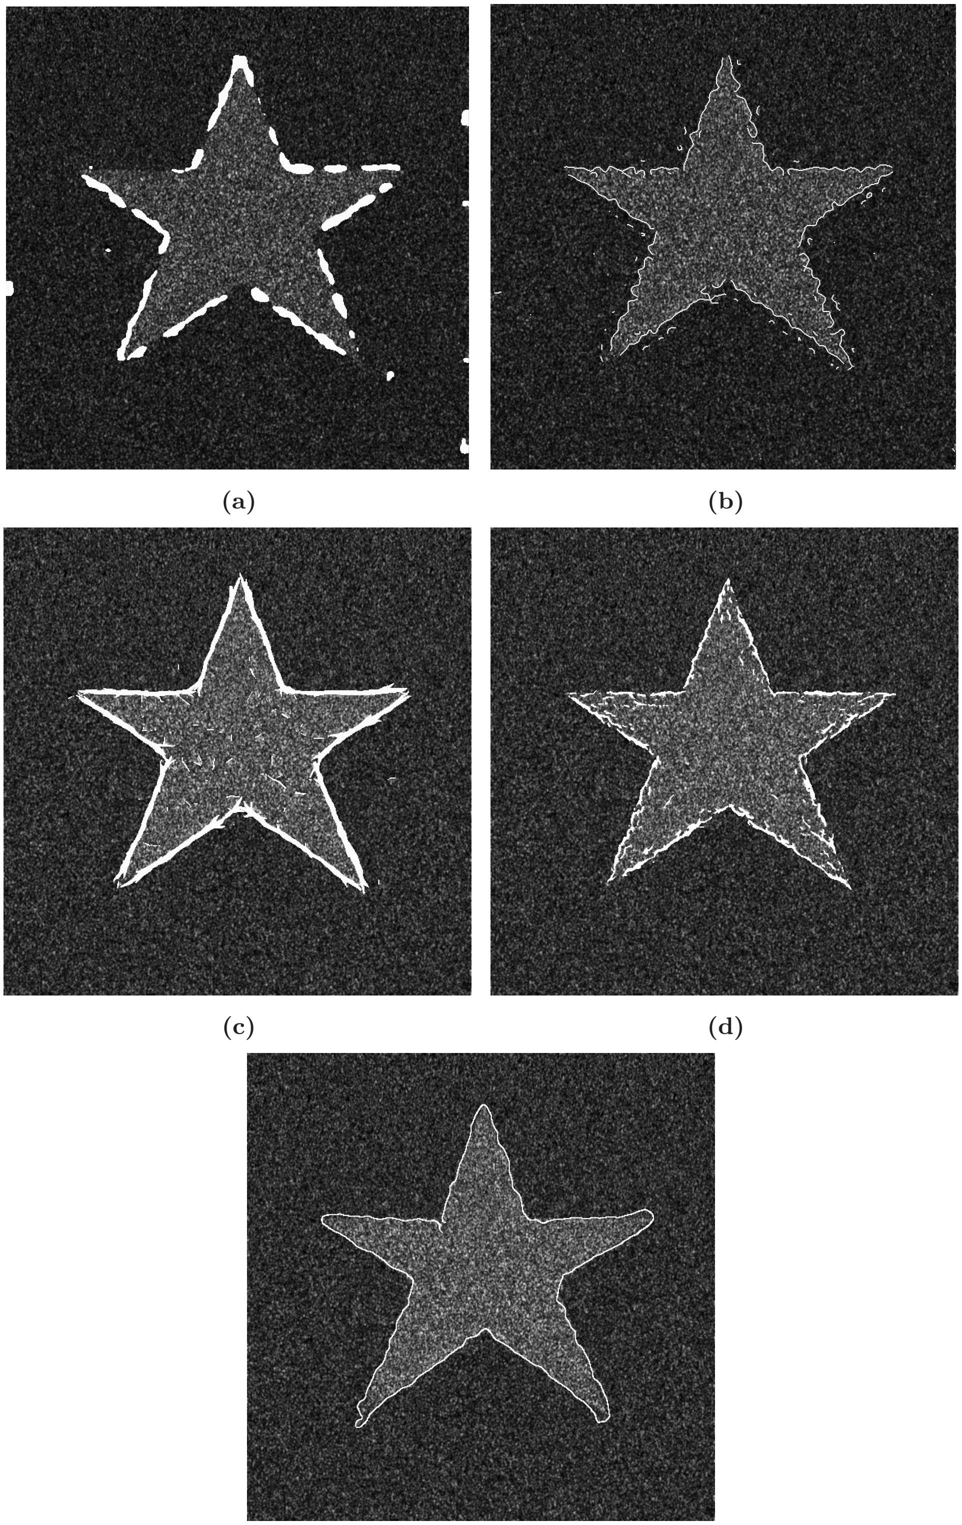

(a) **(b)** **(c)** **(d)** **(e)**

Figure B.3: Edge detector output for the simulated edge images of Fig. 6.1(d). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

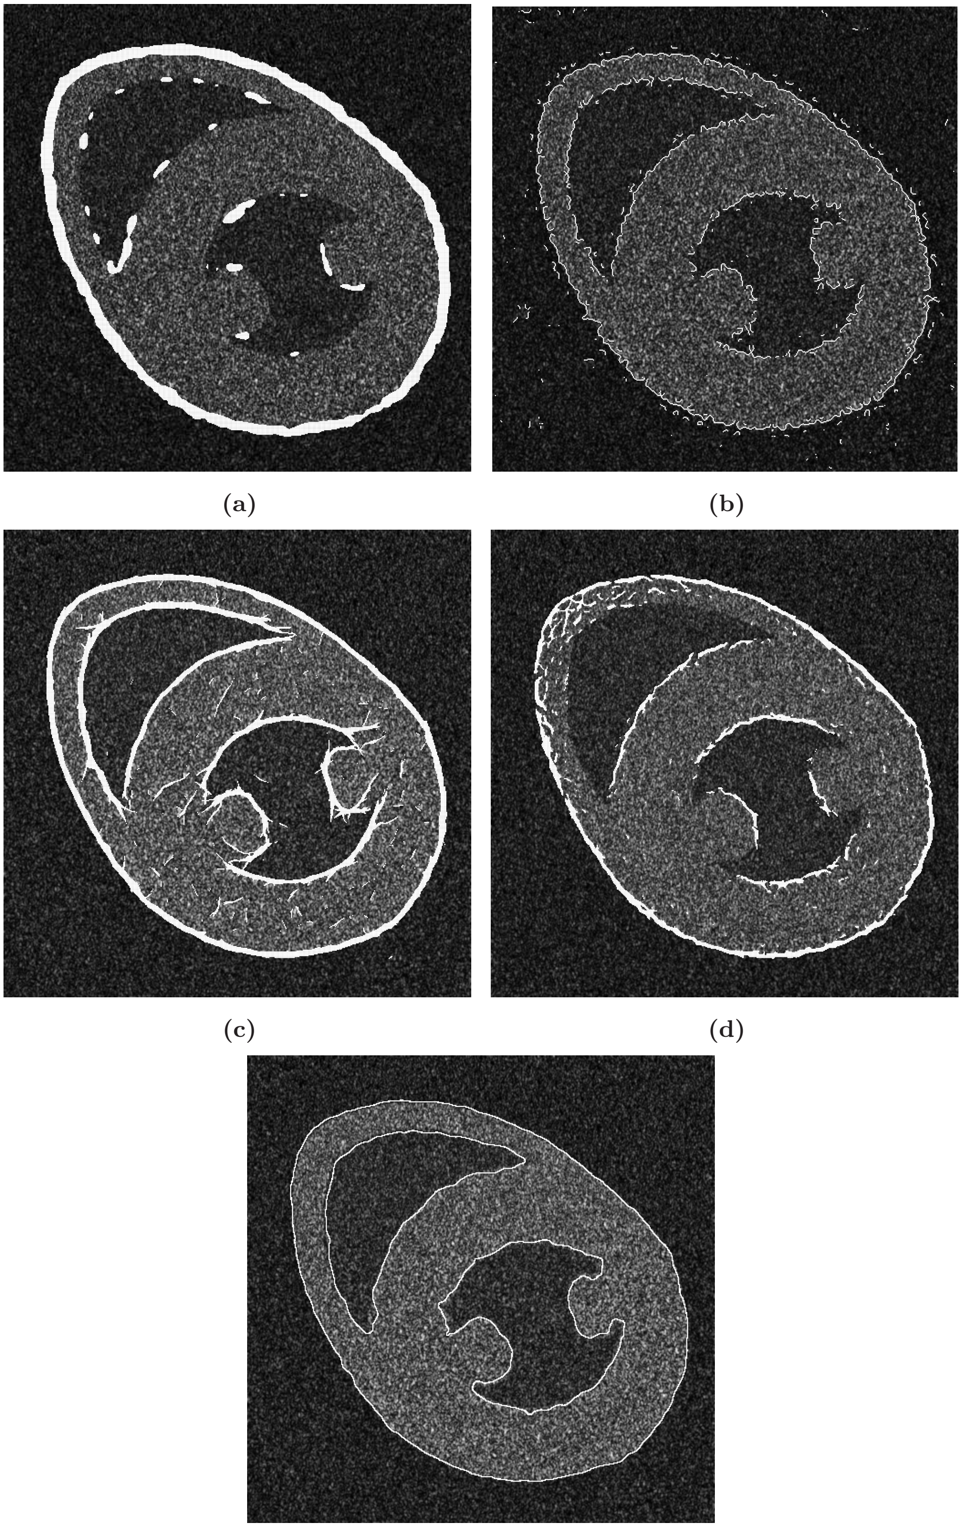

(a) **(b)**

(c) **(d)**

(e)

Figure B.4: Edge detector output for the simulated edge images of Fig. 6.3(b). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

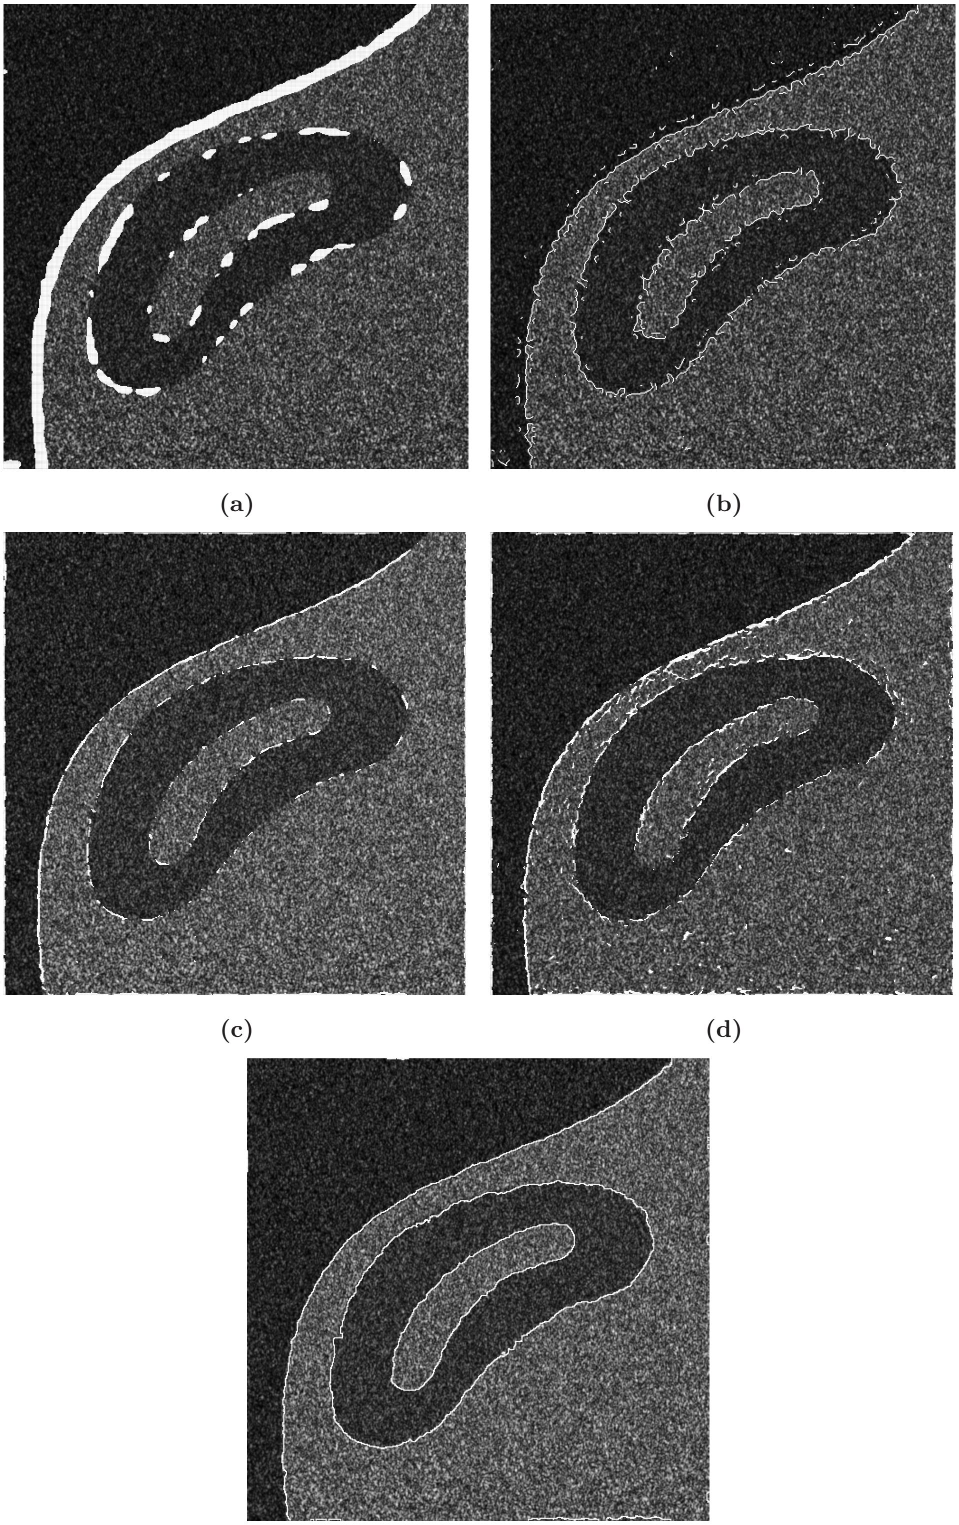

(a)

(b)

(c)

(d)

(e)

Figure B.5: Edge detector output for the simulated edge images of Fig. 6.3(c). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

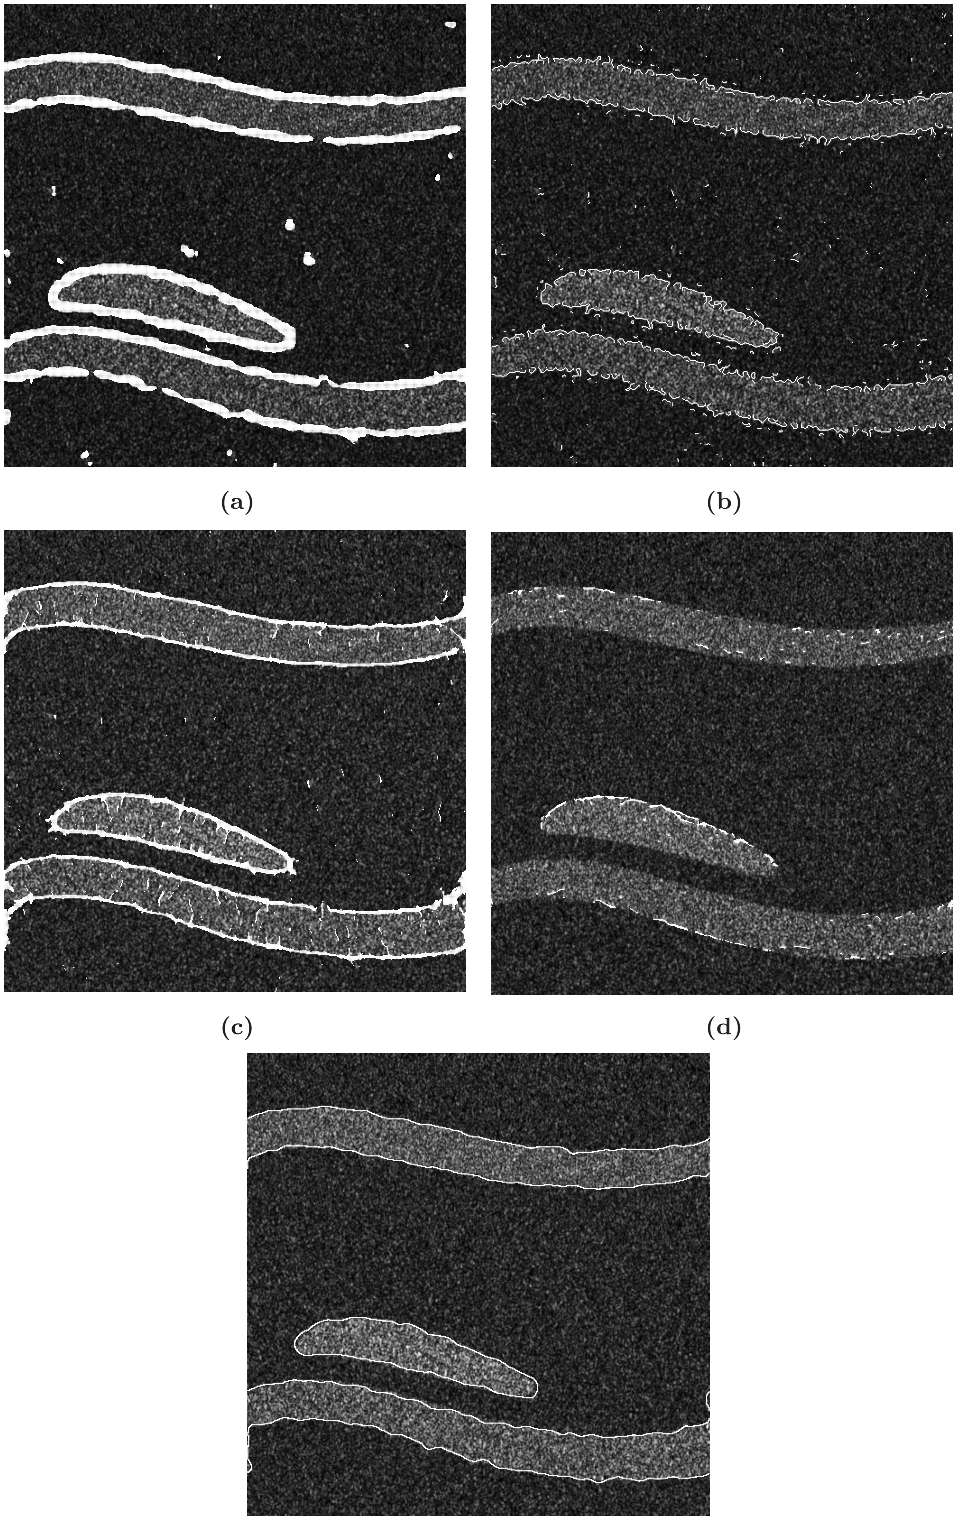

(a) (b) (c) (d) (e)

Figure B.6: Edge detector output for the simulated edge images of Fig. 6.3(d). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

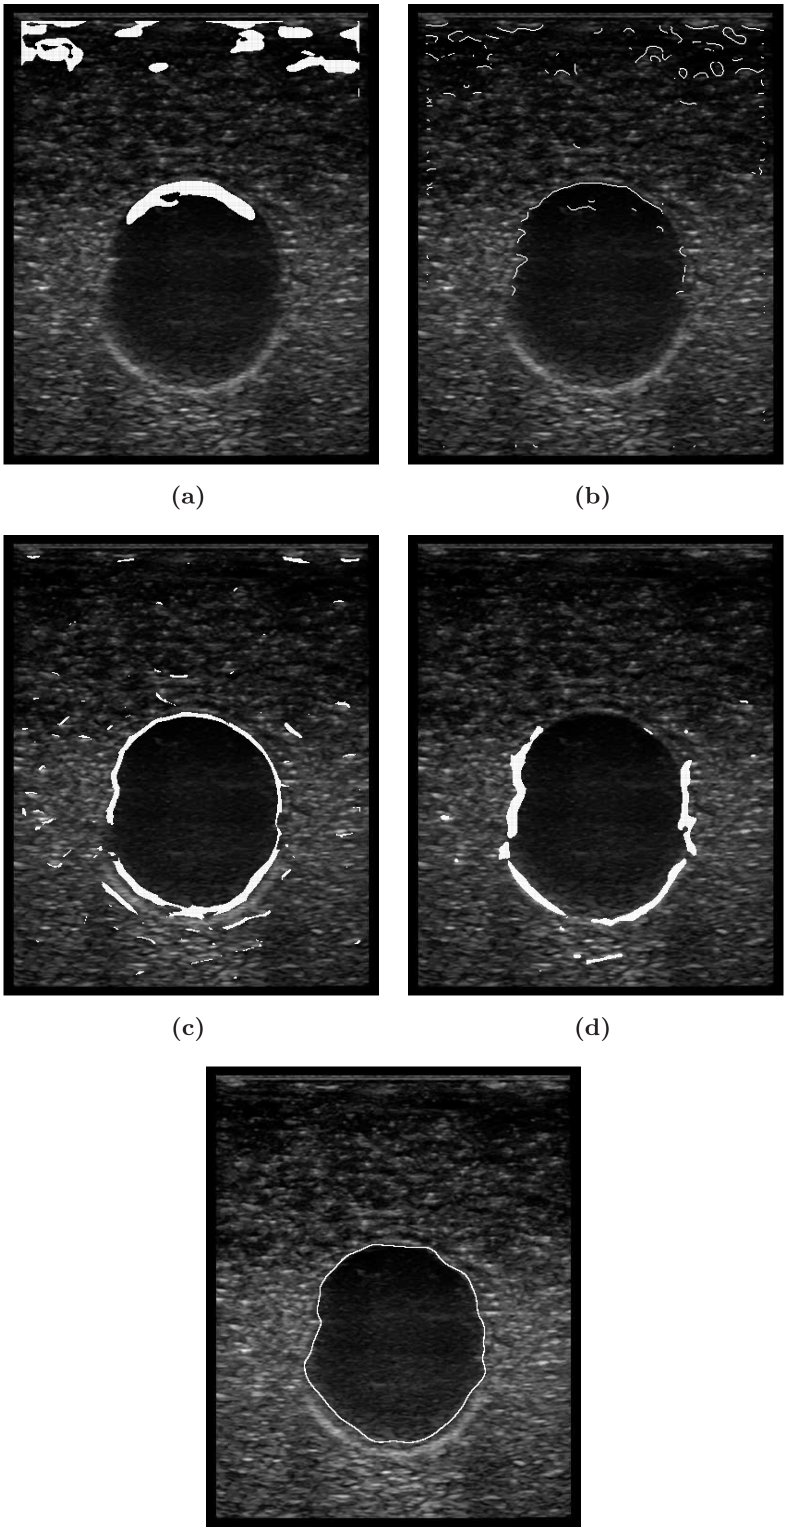

(a) **(b)**

(c) **(d)**

(e)

Figure B.7: Edge detector output for the phantom ultrasound images of Figs. 6.5(a). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

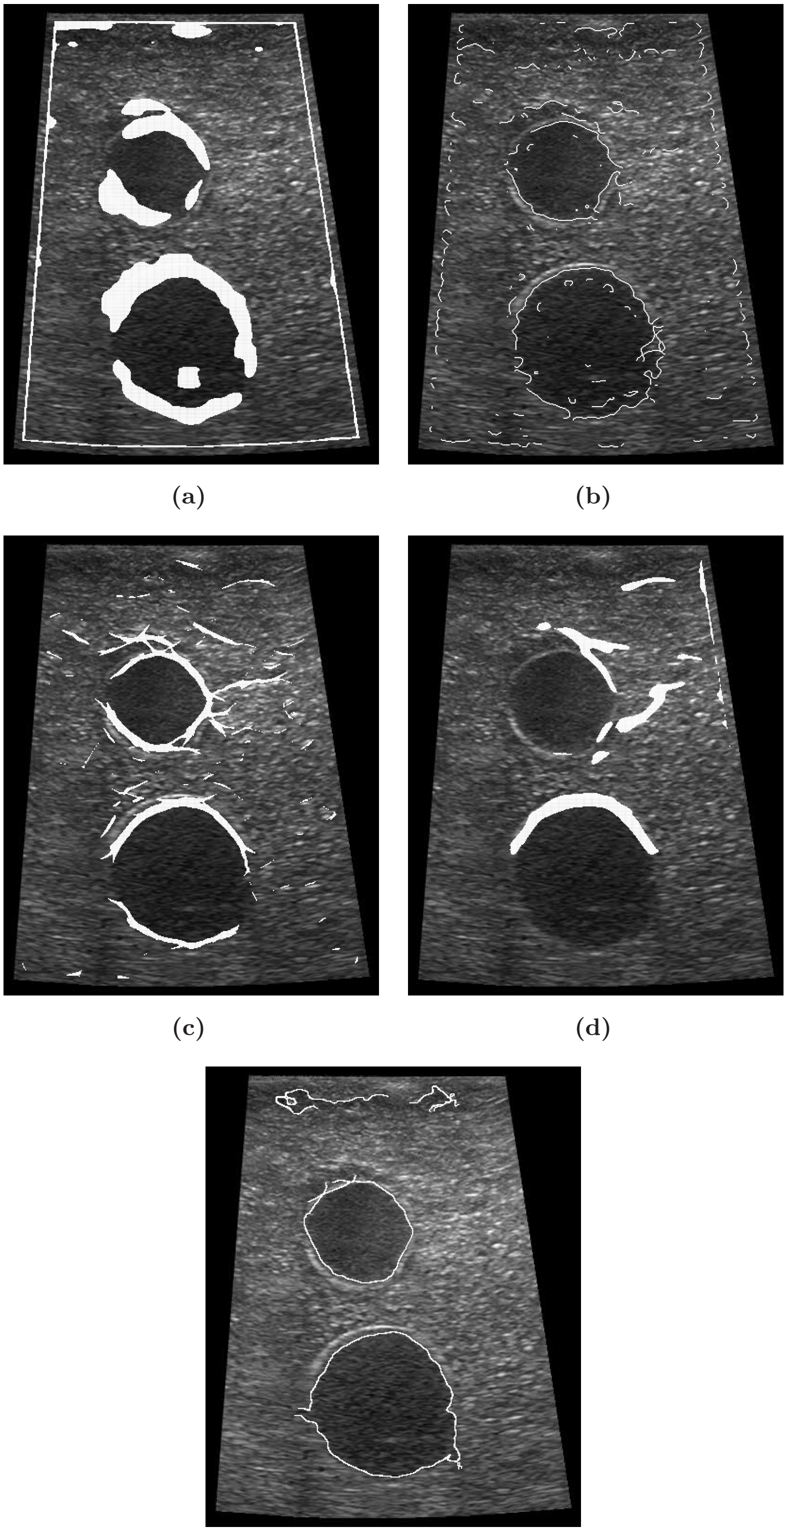

(a) **(b)**

(c) **(d)**

(e)

Figure B.8: Edge detector output for the phantom ultrasound images of Figs. 6.5(c). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

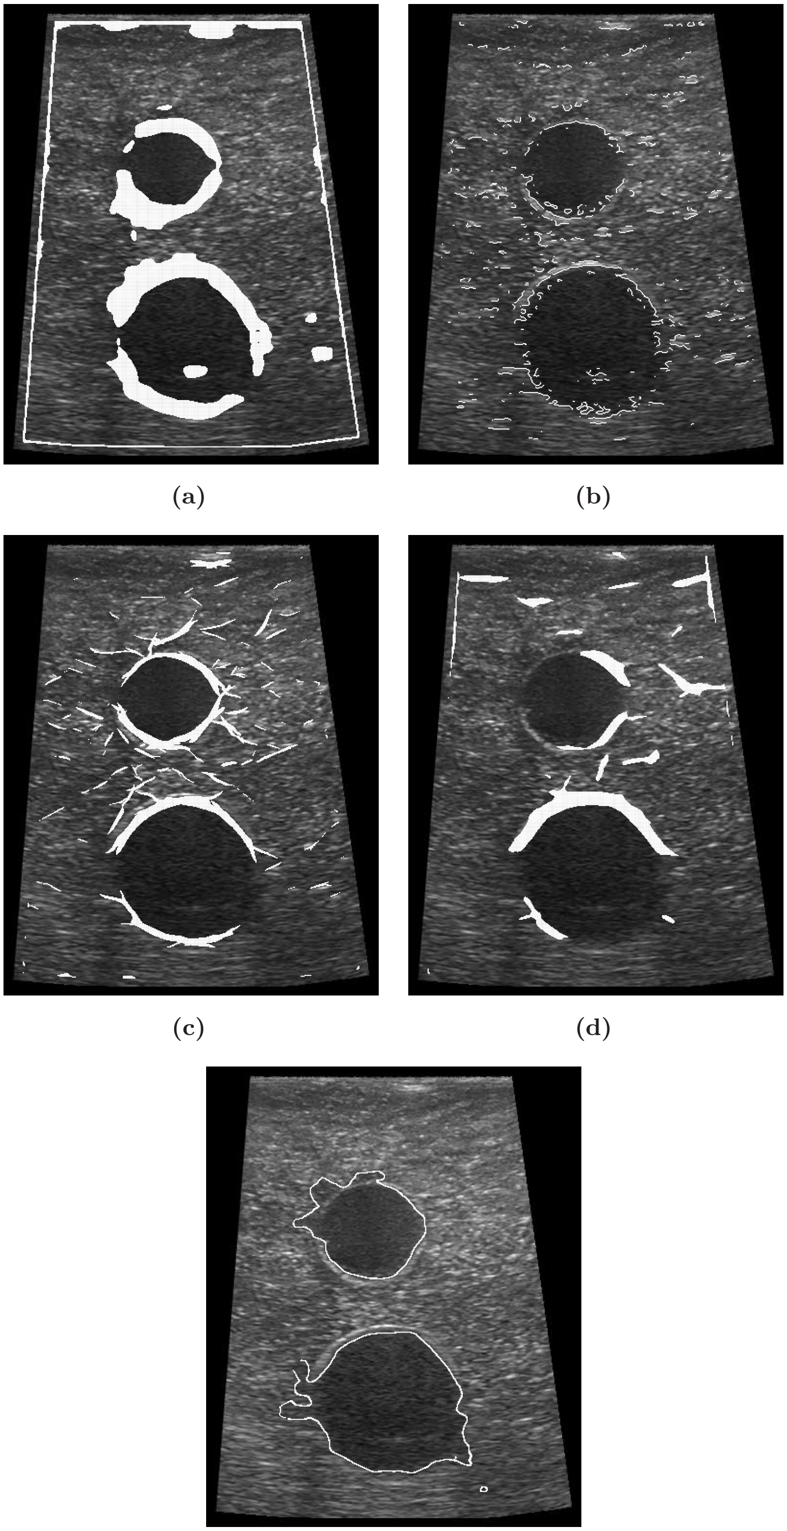

(a) **(b)** **(c)** **(d)** **(e)**

Figure B.9: Edge detector output for the phantom ultrasound images of Figs. 6.5(d). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

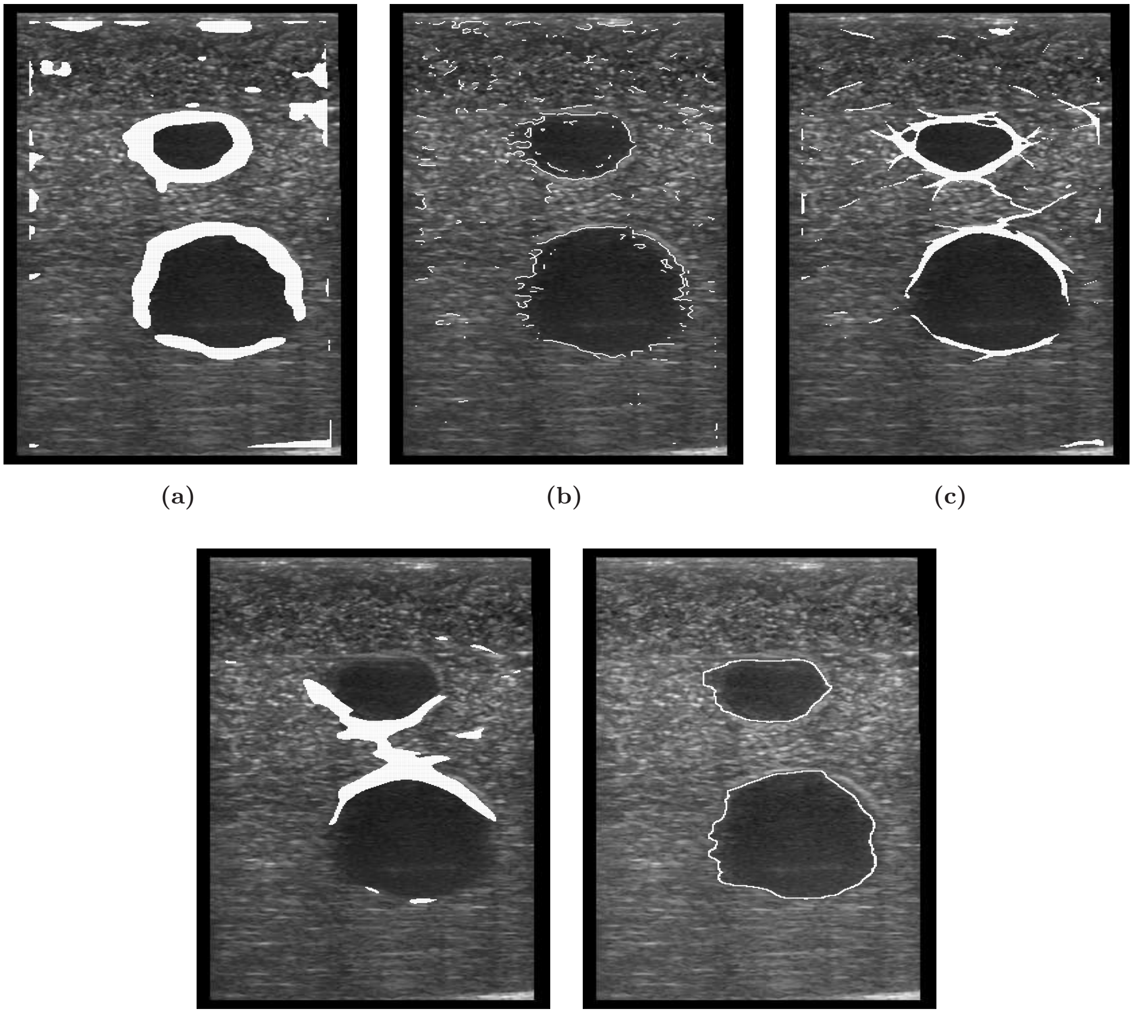

(a) **(b)** **(c)**

(d) **(e)**

Figure B.10: Edge detector output for the phantom ultrasound images of Figs. 6.5(e). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

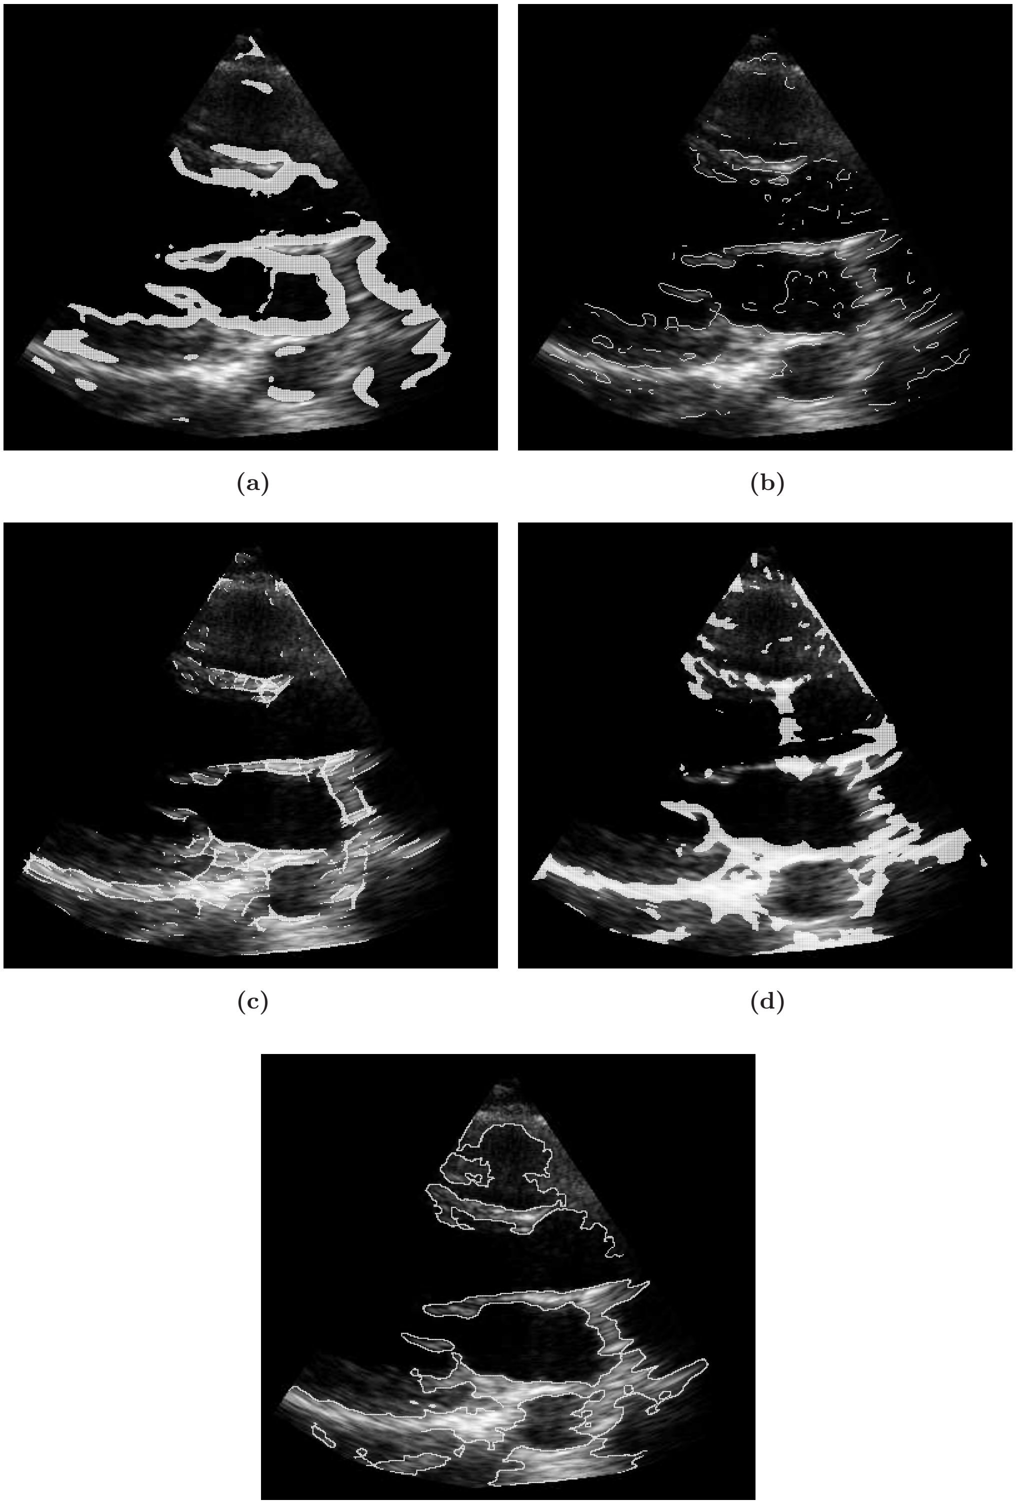

(a) **(b)**

(c) **(d)**

(e)

Figure B.11: Edge detector output for the clinical ultrasound images of Fig. 6.7(a). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

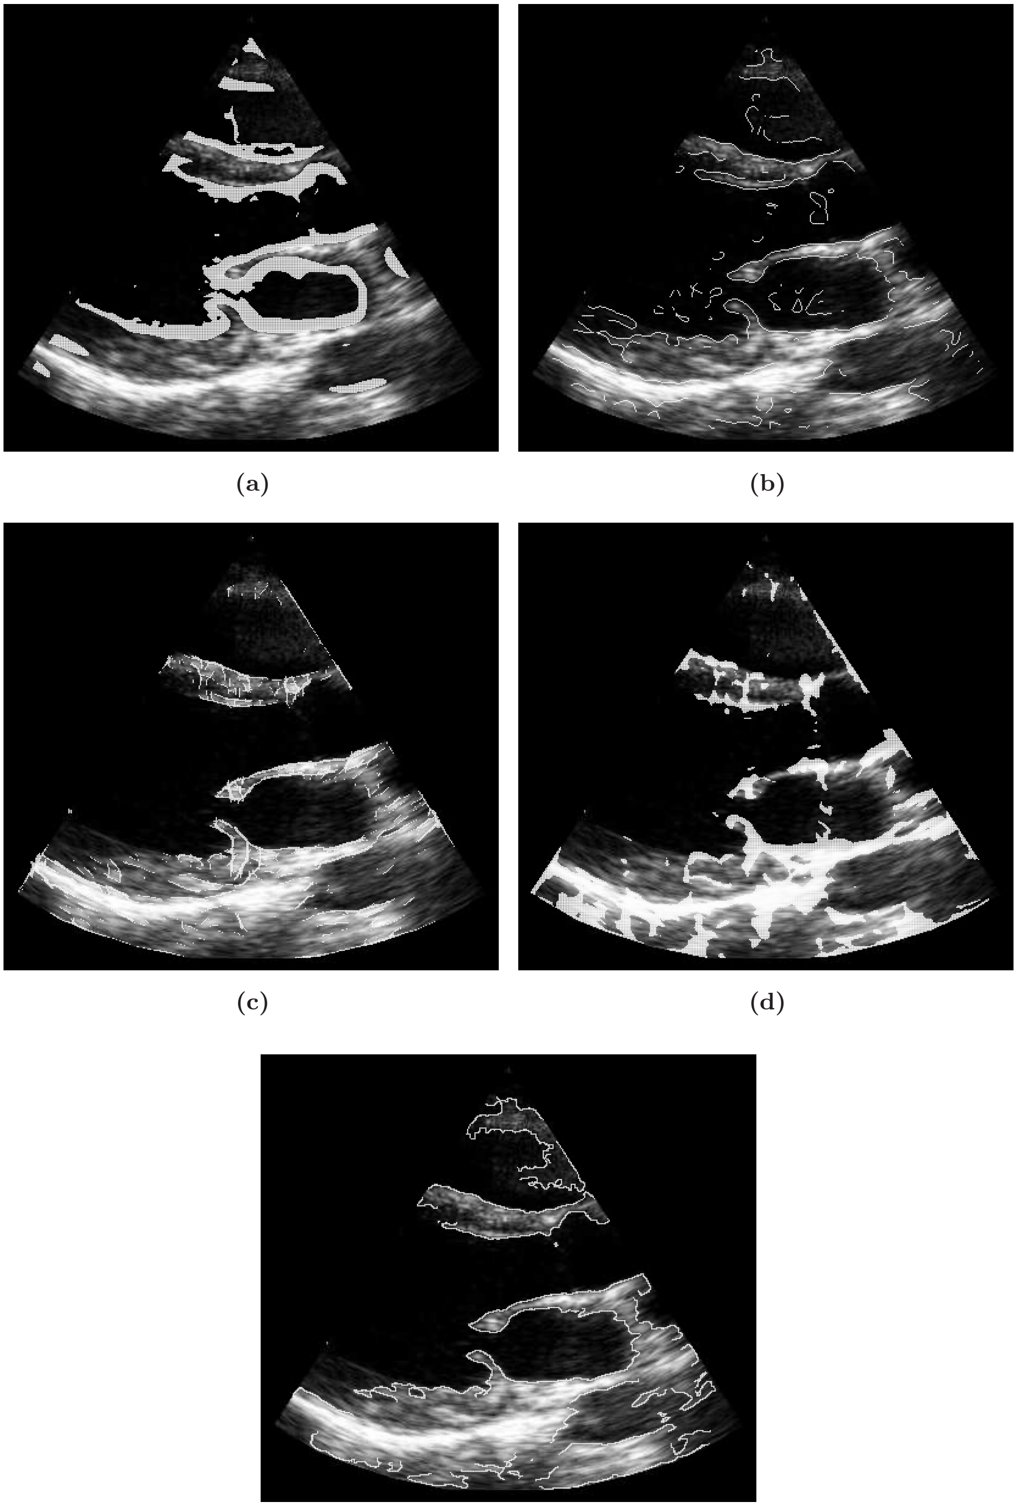

(a) **(b)** **(c)** **(d)** **(e)**

Figure B.12: Edge detector output for the clinical ultrasound images of Fig. 6.7(b). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

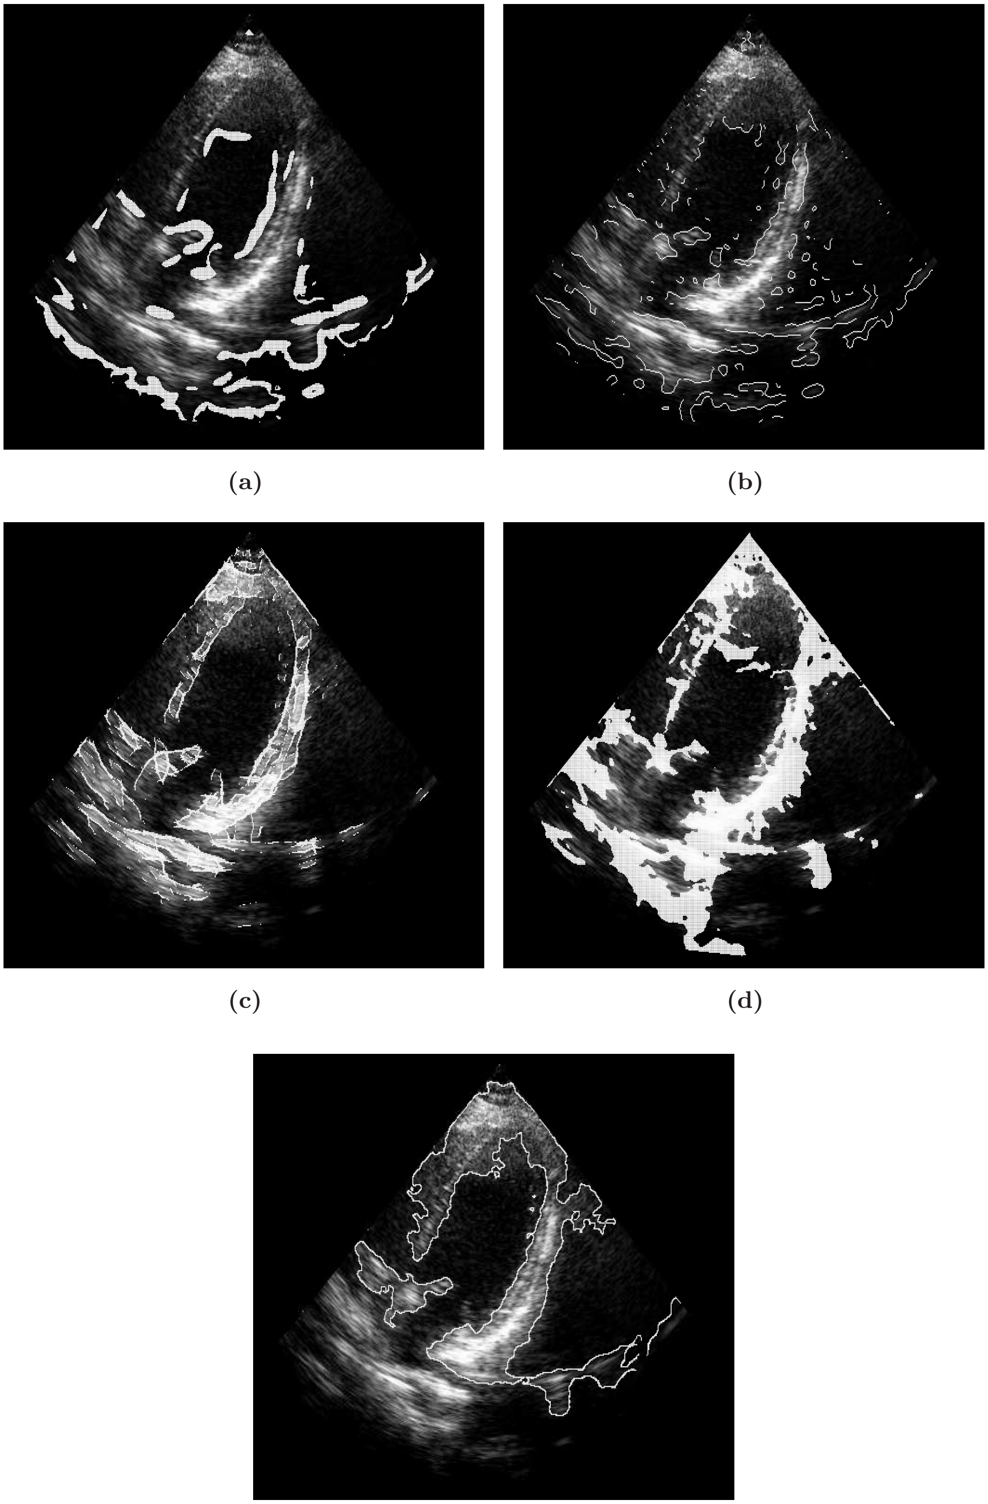

(a) **(b)** **(c)** **(d)** **(e)**

Figure B.13: Edge detector output for the clinical ultrasound images of Fig. 6.7(d). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

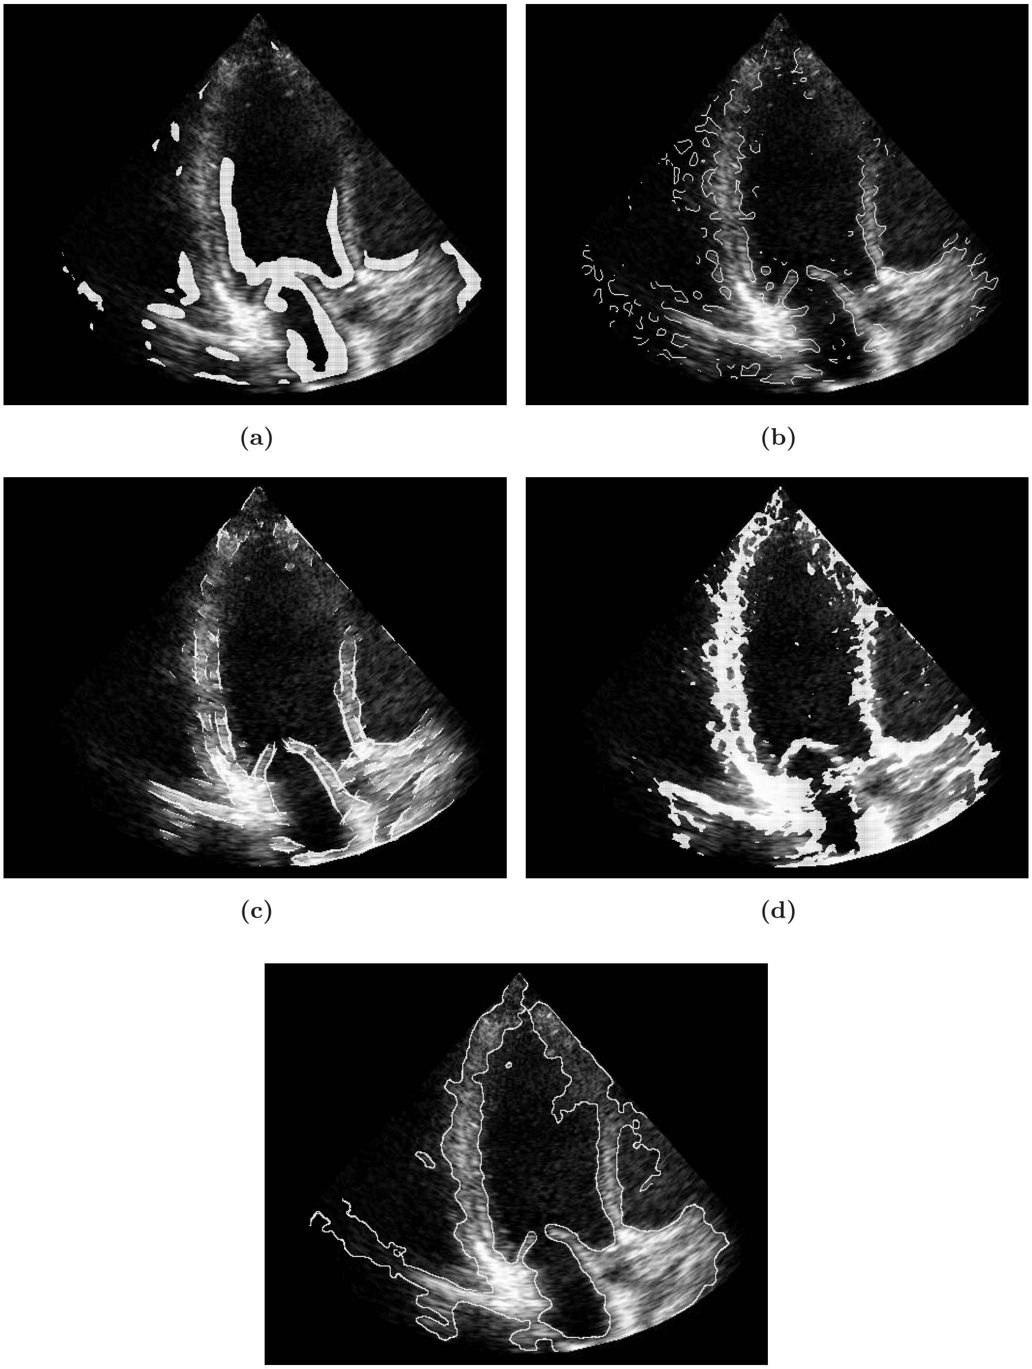

(a) **(b)** **(c)** **(d)** **(e)**

Figure B.14: Edge detector output for the clinical ultrasound images of Fig. 6.7(f). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

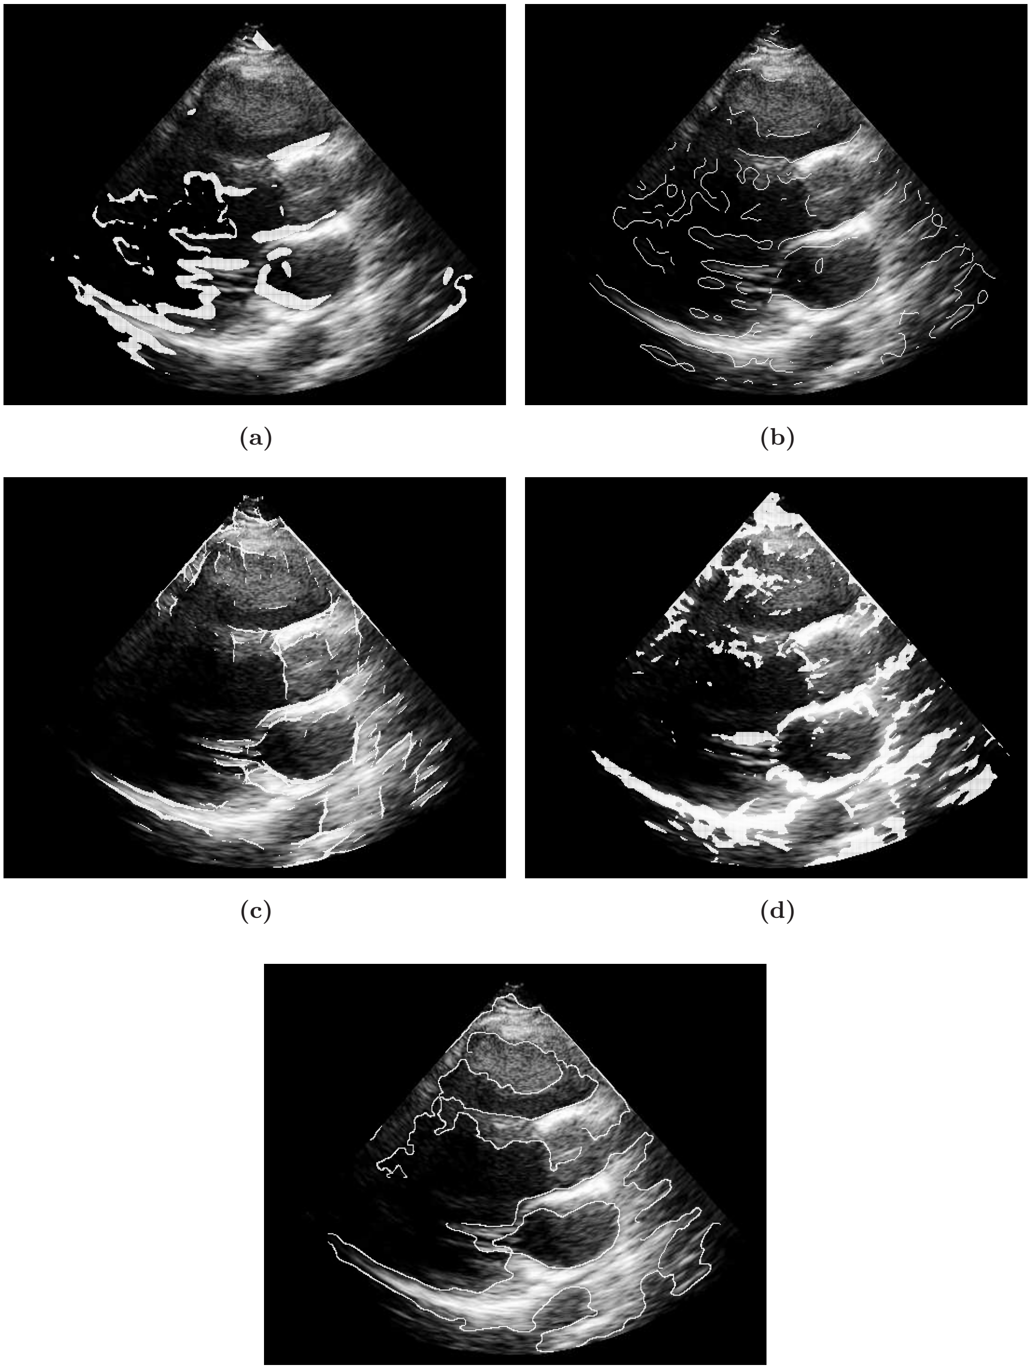

(a) (b) (c) (d) (e)

Figure B.15: Edge detector output for the clinical ultrasound images of Fig. 6.8(a). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

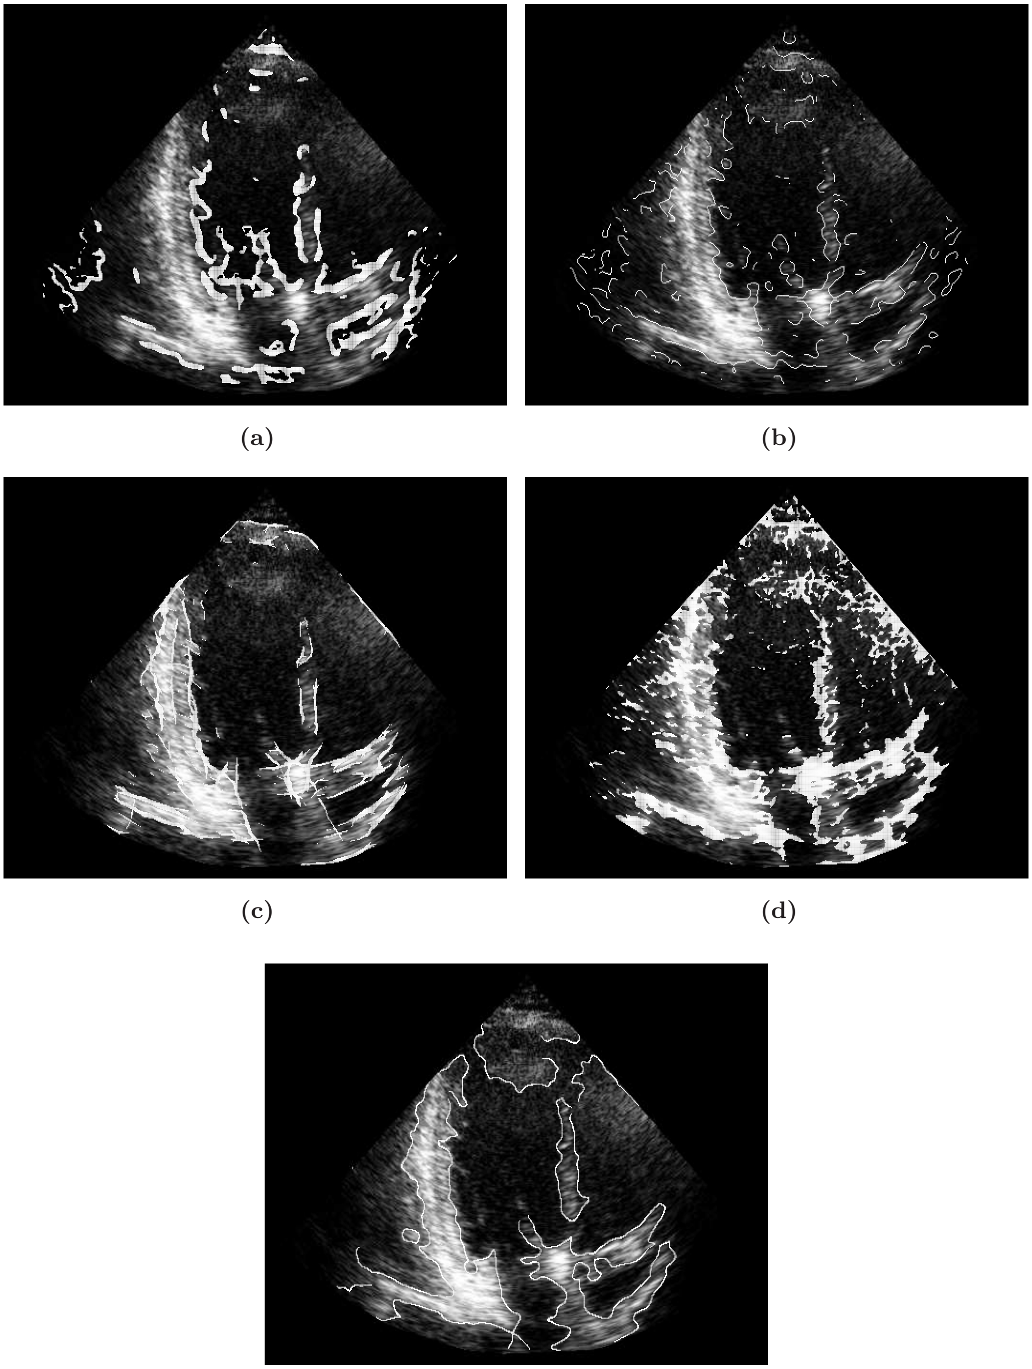

(a) **(b)**

(c) **(d)**

(e)

Figure B.16: Edge detector output for the clinical ultrasound images of Fig. 6.8(b). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

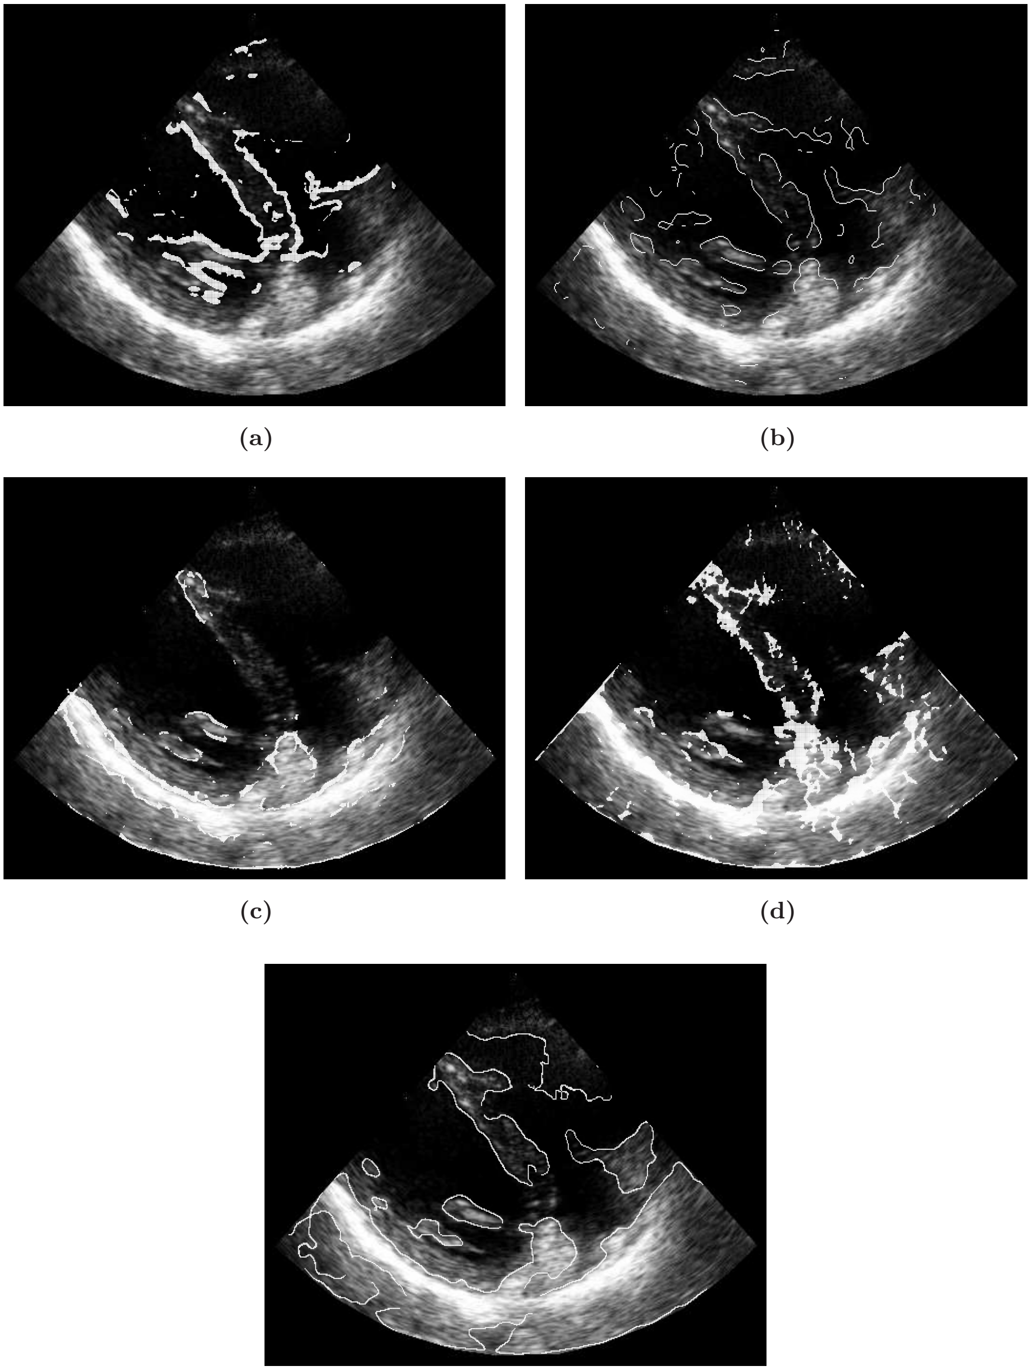

(a) (b)

(c) (d)

(e)

Figure B.17: Edge detector output for the clinical ultrasound images of Fig. 6.8(c). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

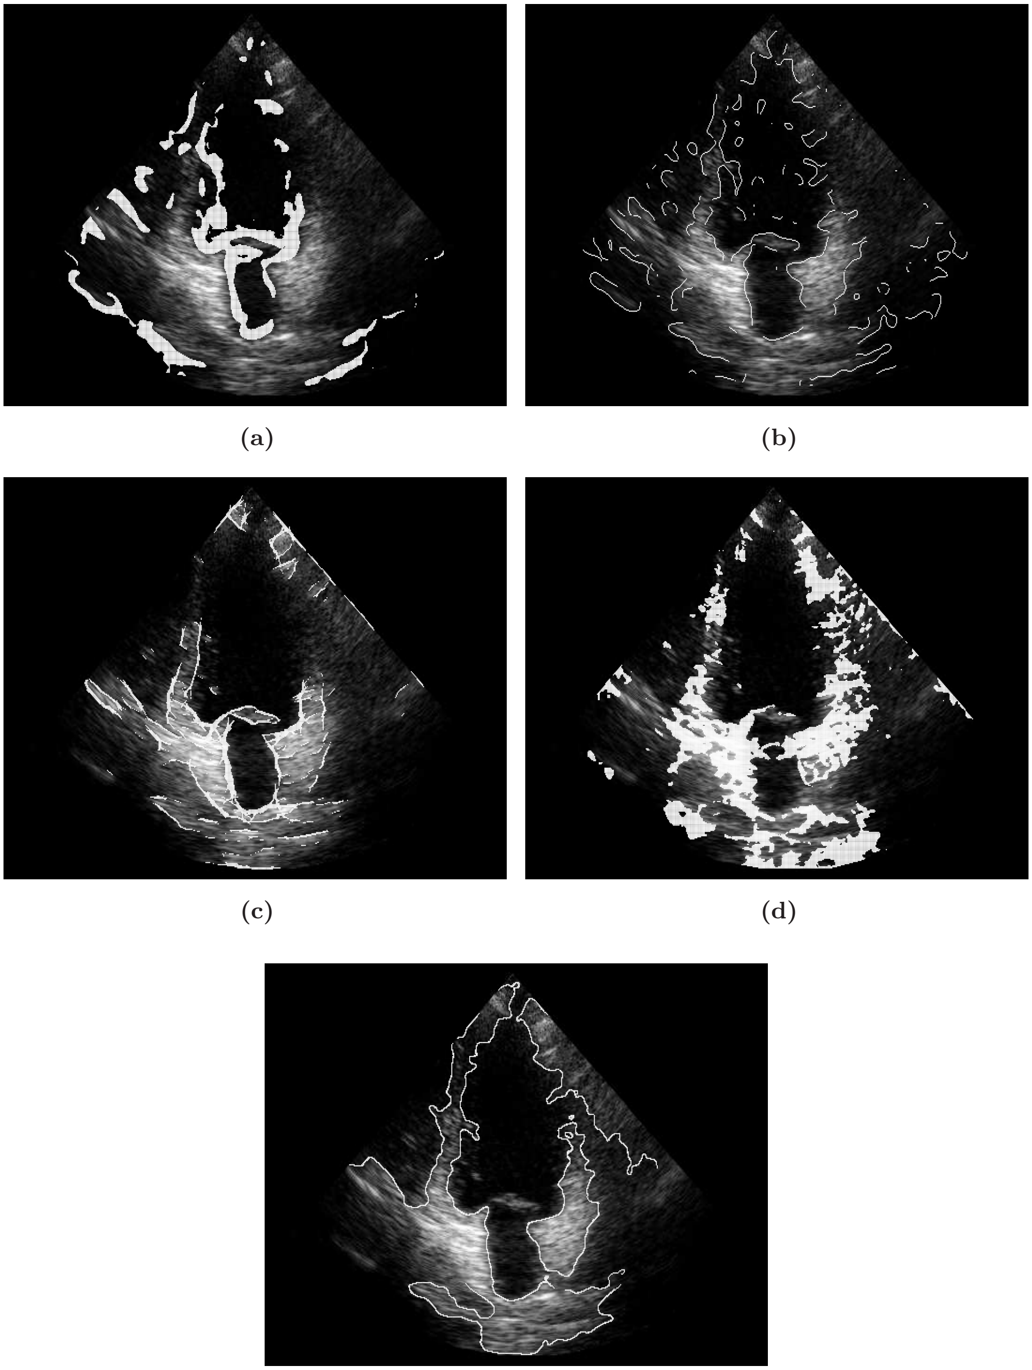

(a) **(b)** **(c)** **(d)** **(e)**

Figure B.18: Edge detector output for the clinical ultrasound images of Fig. 6.8(d). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

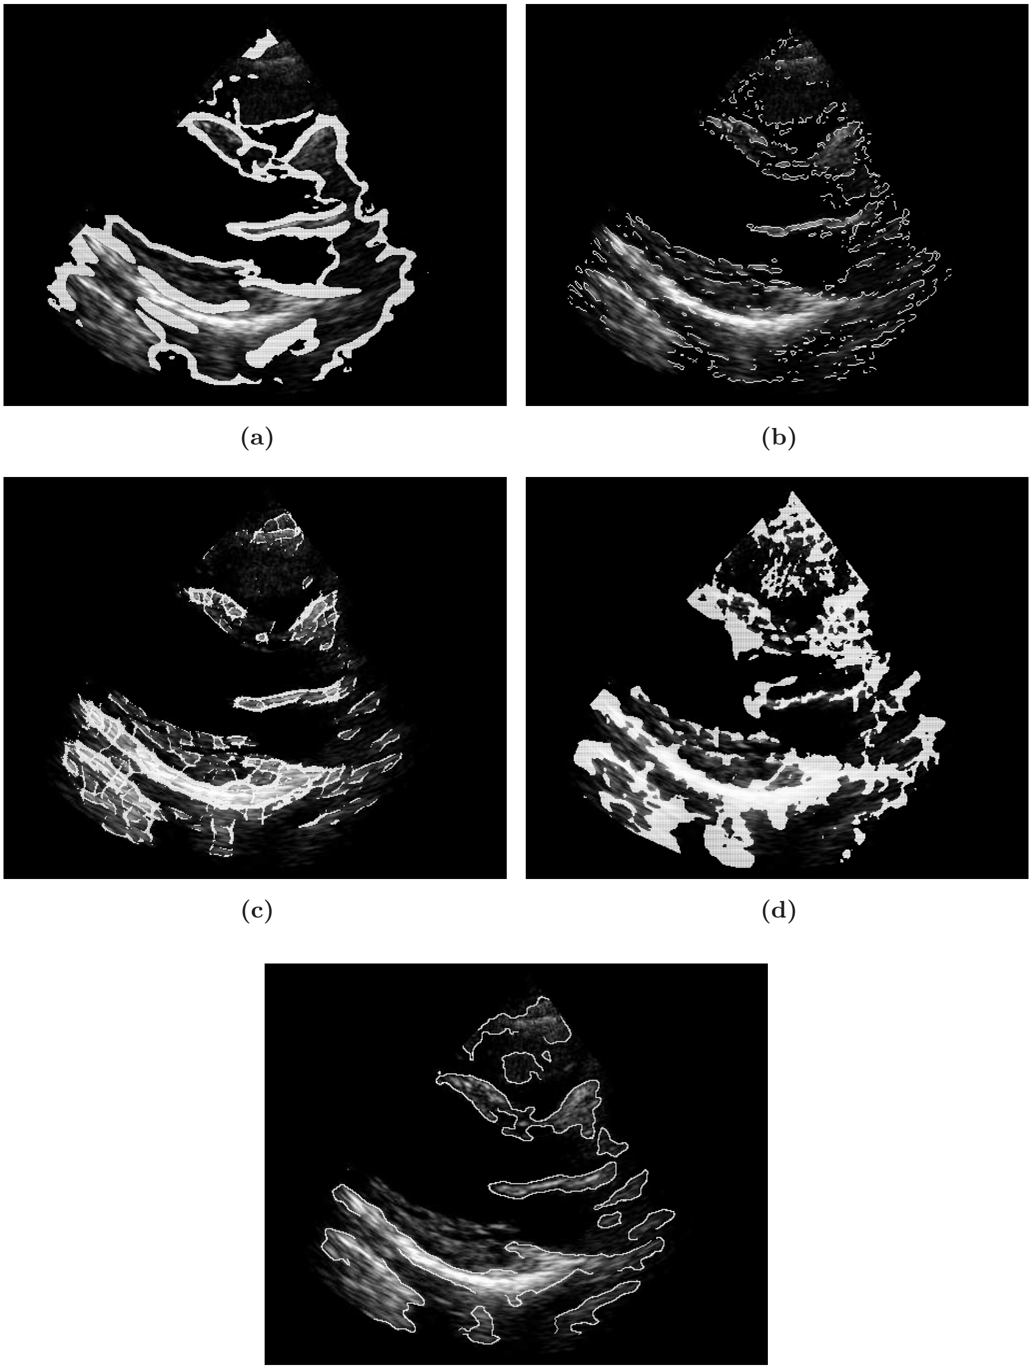

(a) (b) (c) (d) (e)

Figure B.19: Edge detector output for the clinical ultrasound images of Fig. 6.8(e). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

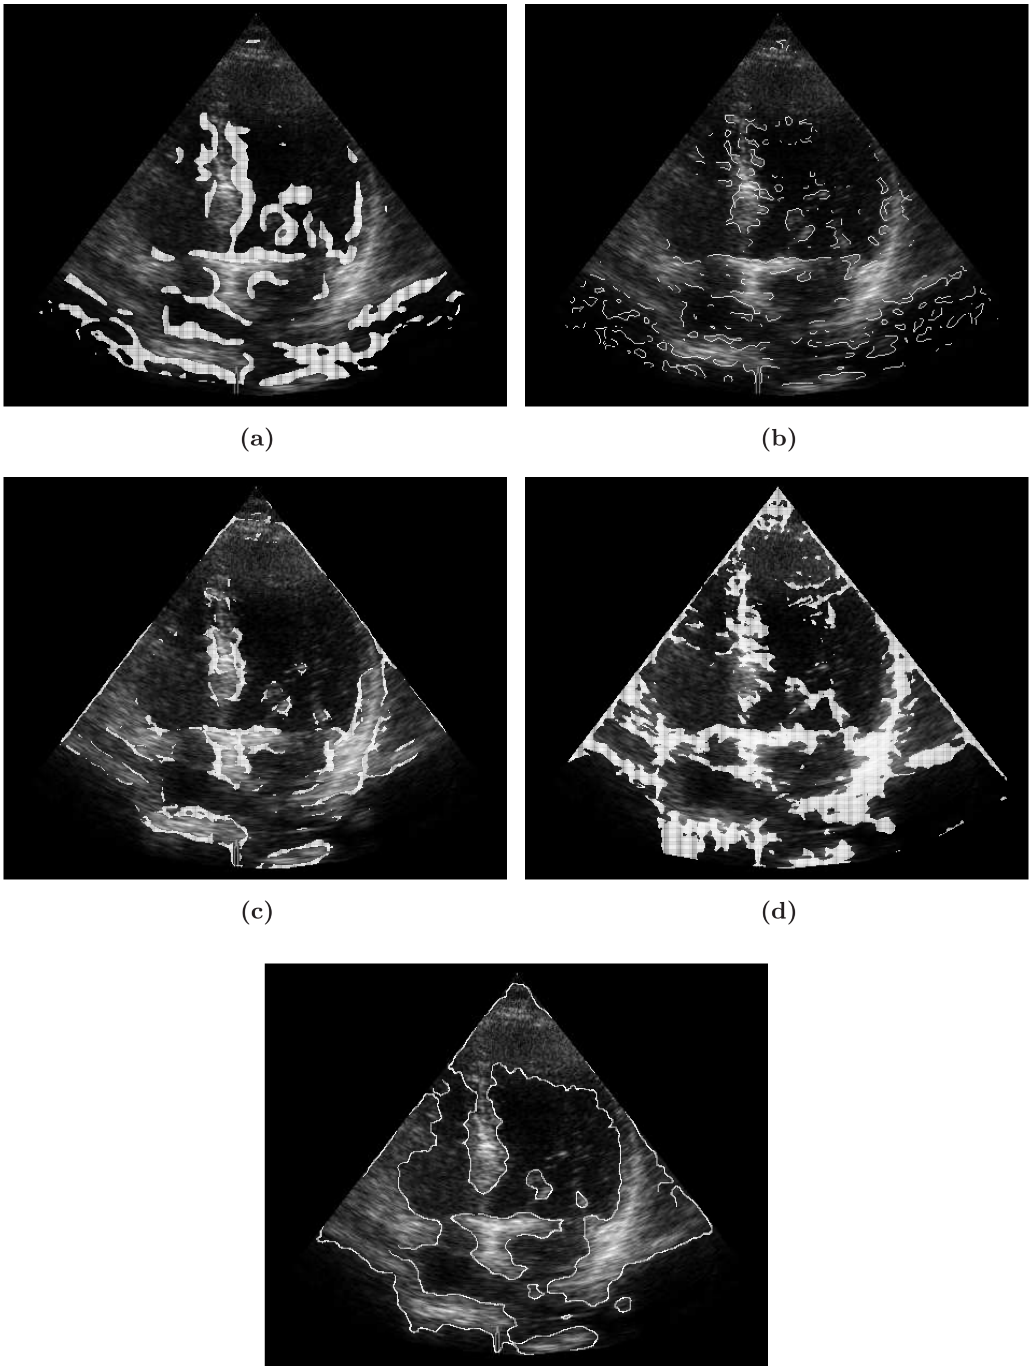

(a) **(b)**

(c) **(d)**

(e)

Figure B.20: Edge detector output for the clinical ultrasound images of Fig. 6.8(f). (a) CV, (b) RoA, (c) FA, (d) PC, (e) ZC.

Appendix C

Publications

The research presented in this thesis has resulted in the following publications:

Journal Publications:

- S. Finn, M. Glavin, and E. Jones, "Echocardiographic Speckle Reduction Comparison", *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 58, no. 1, pp. 82-2011, Jan. 2011; doi: 10.1109/TUFFC.2011.1776. Full text included.
- S. Finn, E. Jones, P, Nolan, K. Daly, and M. Glavin, "Ultrasound Edge Detection using Stable Modeling and Zero Crossings of Gabor Coefficients", *Med. Eng. Phys.*, 2011. Submitted;
- S. Finn, M. Glavin, and E. Jones, "Speckle Reducing Diffusion using a Static Edge Map", 2011. In preparation.

Conference Proceedings:

- S. Finn, E. Jones, and M. Glavin, "Objective and subjective evaluations of quality for speckle reduced echocardiography", in *Proc. IEEE Intl. Annu. Conf. Engineering Medicine and Biology Society*, Sep. 2009, pp. 503-506. doi: 10.1109/IEMBS.2009.5333572.