

Attribute-Driven Segmentation and Analysis of Mammograms

Sze Man Simon Kwok, B.E. (Hons.)

This thesis is presented for the degree of Doctor of Philosophy of
The University of Western Australia

Centre for Intelligent Information Processing Systems
School of Electrical, Electronic and Computer Engineering
The University of Western Australia
Crawley, WA 6009
Australia

September 2004

Abstract

In this thesis, we introduce a *mammogram analysis system* developed for the automatic segmentation and analysis of mammograms. This original system has been designed to aid radiologists to detect breast cancer on mammograms. The system embodies *attribute-driven segmentation* in which the attributes of an image are extracted progressively in a step-by-step, hierarchical fashion. Global, low-level attributes obtained in the early stages are used to derive local, high-level attributes in later stages, leading to increasing refinement and accuracy in image segmentation and analysis.

The proposed system can be characterized as:

- a bootstrap engine driven by the attributes of the images;
- a solid framework supporting the process of hierarchical segmentation;
- a universal platform for the development and integration of segmentation and analysis techniques; and
- an extensible database in which knowledge about the image is accumulated.

Central to this system are three major components:

1. a series of applications for *attribute acquisition*;
2. a standard format for *attribute normalization*; and
3. a database for *attribute storage* and data exchange between applications.

The first step of the automatic process is to segment the mammogram hierarchically into several distinctive regions that represent the anatomy of the breast. The adequacy and quality of the mammogram are then assessed using the anatomical features obtained from segmentation. Further image analysis, such as breast density classification and lesion detection, may then be carried out inside the breast region.

Several domain-specific algorithms have therefore been developed for the attribute acquisition component in the system. These include:

1. automatic pectoral muscle segmentation;
2. adequacy assessment of positioning and exposure; and

3. contrast enhancement of mass lesions.

An adaptive algorithm is described for automatic segmentation of the pectoral muscle on mammograms of mediolateral oblique (MLO) views. The pectoral margin is first approximated by a straight line, which is validated for correctness of location and orientation. This straight line is then refined into a curve, using an iterative edge detector, that delineates the pectoral margin more accurately.

After segmentation, and before analysis, it is necessary to ensure that the mammogram is of sufficient quality for further analysis by machine. This mirrors the requirement that a mammogram must be of sufficient quality for a radiologist to detect lesions or other abnormalities with high sensitivity and specificity. The adequacy of breast positioning is evaluated on MLO view mammograms using several quality criteria, including non-exclusion of breast tissue, visualization of the nipple in profile, inclusion of inframammary fold, and proper positioning of the pectoral muscle. The adequacy of film exposure is also evaluated.

Mass lesions are then highlighted using a novel contrast enhancement technique, developed for this purpose. This technique is a modification of the amplitude and phase (AMPM) demodulation transform introduced by Daugman and Downing. This new technique was compared with two other existing contrast enhancement techniques—morphological enhancement and histogram equalization—and was found to perform better than either.

The above applications need to interact with each other during the segmentation and analysis process. Attribute normalization and storage are necessary for this. The attributes of mammograms generated by the applications were normalized to a standard format, and stored in an Extensible Markup Language (XML) database. This database was specified by an open-ended and extensible schema, and was designed to be application- and platform-independent.

Contents

Abstract	iii
Acknowledgements	xiii
1 Introduction	1
1.1 Outline of Thesis	2
2 Literature Review	5
2.1 Mammogram Segmentation	5
2.2 Mammographic Image Quality	9
2.2.1 Quality Assurance and Quality Control	10
2.2.2 Mammographic positioning	10
2.2.3 Film Exposure	11
2.3 Image Enhancement of Mammograms	12
3 Mammogram Analysis System	15
3.1 System Overview	15
3.2 Mammogram Attributes	16
3.2.1 Image Data	16
3.2.2 Visual Features	17
3.2.3 Analytical Information	17
3.3 Acquisition of Mammogram Attributes	18
3.4 Normalization of Mammogram Attributes	20
3.4.1 Image Orientation	20
3.4.2 Co-ordinate System	21
3.4.3 Spatial Resolution	21
3.4.4 Measurement Units	22
3.4.5 Representations of Object Boundary	22
3.4.6 Resulting Images	22
3.4.7 Example: Normalization of the Breast Border	22
3.5 Storage of Mammogram Attributes	23
3.5.1 A Database of Mammogram Attributes	24

3.5.2	Example: XML Representation of the Breast Border	25
3.6	Conclusions	26
4	Automatic Pectoral Muscle Segmentation	29
4.1	Introduction	29
4.1.1	Literature Review	30
4.2	Overall Algorithm	31
4.3	Orientation and Notation	33
4.4	Straight Line Estimation and Validation	33
4.4.1	Straight Line Estimation	34
4.4.2	Straight Line Validation	37
4.5	Iterative Cliff Detection	38
4.5.1	Cliff Detection	38
4.5.2	Iterative Refinement	41
4.5.3	Region Enclosing	42
4.6	Experimental Setup	42
4.7	Results	44
4.7.1	Radiologists' Assessments	44
4.7.2	Algorithm Performance	46
4.8	Discussion	48
4.9	Suggestions for Further Work	50
4.10	Conclusions	51
5	Adequacy Assessment of Positioning and Exposure	53
5.1	The Need for Automatic Quality Assessment	53
5.2	Quality Attributes of Mammograms	53
5.3	Optimal Positioning in the MLO View	54
5.4	Adequate Exposure	56
5.5	Adequacy Assessment Algorithms	57
5.5.1	Image Orientation and Notation	57
5.5.2	Exclusion of Breast Tissue	58
5.5.3	Nipple in Profile	58
5.5.4	Locating the Inframammary Fold	60
5.5.5	Positioning of the Pectoral Muscle	60
5.5.6	Curvature of the Pectoral Margin	62
5.5.7	Optical Density Measures	63
5.6	Experimental Results	66
5.7	Discussion	68
5.8	Suggestions for Further Work	70
5.8.1	Improvements to Algorithms	70
5.8.2	Extension to Other Quality Attributes	71

5.8.3	Assessment of Other Mammogram Databases	72
5.9	Conclusions	73
6	Contrast Enhancement of Mass Lesions	75
6.1	Introduction	75
6.2	Amplitude and Phase Demodulation	76
6.2.1	Algorithm of AMPM Demodulation Transform	76
6.2.2	Results with Texture Images	78
6.2.3	Selection of Half-Plane and Carrier Wave	80
6.3	AMPM Demodulation with Multiple Carriers	81
6.3.1	Algorithm of Multi-carrier AMPM Demodulation	82
6.4	Modification for Contrast Enhancement of Mass Lesions	84
6.4.1	Carrier Wave of Zero Frequency	85
6.4.2	Partial Reconstruction	85
6.4.3	Image Subtraction	87
6.4.4	Contrast Stretching	87
6.5	Existing Techniques of Contrast Enhancement	88
6.5.1	Morphological Enhancement	88
6.5.2	Histogram Equalization	90
6.6	Quantitative Measures of Contrast Enhancement	91
6.6.1	Distribution Separation Measure (DSM)	92
6.6.2	Target to Background Contrast Ratio Using Standard Deviation (TBCs)	93
6.6.3	Target to Background Contrast Ratio Using Entropy (TBCe)	93
6.6.4	Combining the Quantitative Measures	94
6.7	Experimental Setup	94
6.8	Results	95
6.9	Discussion	98
6.10	Suggestions for Further Work	102
6.11	Conclusions	103
7	Conclusion	105
7.1	Summary of Research Outcomes	105
7.1.1	Mammogram Analysis System	105
7.1.2	Automatic Pectoral Muscle Segmentation	106
7.1.3	Adequacy Assessment of Positioning and Exposure	107
7.1.4	Contrast Enhancement of Mass Lesions	107
7.2	Suggestions for Further Work	107

A Abstract System Model	111
A.1 Introduction	111
A.2 Image Attributes	111
A.3 System Model	113
A.4 Three Major Components	114
A.5 Acquisition of Attributes	115
A.5.1 Attribute Dependence Graph	115
A.5.2 Sequential Attribute Acquisition	116
A.5.3 Parallel Attribute Acquisition	118
A.6 Normalization of Attributes	119
A.6.1 Image Orientation	119
A.6.2 Co-ordinate System	120
A.6.3 Spatial Resolution	121
A.6.4 Measurement Units	121
A.6.5 Representations of Object Boundary	122
A.7 Storage of Attributes	122
A.7.1 XML Database	123
A.7.2 XML Schema	126
A.8 Design Strategy for Actual Systems	126
A.8.1 Example: Bank Cheques	127
A.8.2 Example: Human Faces	128
A.8.3 Example: Mammograms	128
A.9 Conclusions	128
B XML Schema of Mammogram Attributes	131
B.1 Basic Elements	131
B.1.1 Point	131
B.1.2 Line	132
B.1.3 Array	132
B.1.4 Image	132
B.2 XML Elements of Mammogram Attributes	132
B.2.1 Mammogram	132
B.2.2 Breast Anatomy	133
B.2.3 Adequacy	134
B.2.4 Abnormality	134
B.3 XML Schema File	135
C Iterative Threshold Selection	143
Bibliography	145
List of Publications	155

List of Figures

1.1	Image Analysis System for Mammograms	2
3.1	Mammogram Analysis System	16
3.2	Dependence graph of mammogram attributes	18
3.3	Attribute priority assignments for adequacy measures and lesion description	20
3.4	Standard image orientations and co-ordinate system for mammogram at- tributes	21
3.5	Normalization of the breast border attribute	23
3.6	Element structure of mammogram attributes	24
4.1	Flowchart for automatic pectoral muscle segmentation	32
4.2	Notation used in pectoral muscle segmentation	33
4.3	Illustration of straight line estimation	34
4.4	Iterative threshold selection	35
4.5	Illustration of iterative cliff detection	39
4.6	Comparison of intensity profiles	40
4.7	Sigmoid model	41
4.8	Region enclosing	42
4.9	Pectoral muscle segmentation results	47
4.10	Comparison of straight line and curve segmentation results	49
5.1	Optimal positioning in the MLO view	55
5.2	Notation used in MLO view positioning	57
5.3	Breast tissue exclusions	59
5.4	Nipple in profile	59
5.5	Inframammary folds	61
5.6	Position of the pectoral muscle	61
5.7	Curvature of the pectoral muscle	63
5.8	Pixel value to optical density conversion for MIAS images	64
5.9	Optical density histograms	65
5.10	The MIAS image satisfied all quality criteria	68
5.11	Exclusions of nipple and inframammary fold	69

5.12	Optical density distribution for MIAS images	70
5.13	Poor segmentation due to overexposure	71
6.1	Demodulation of synthetic texture	79
6.2	Demodulation of natural texture	79
6.3	Half-plane orientations and centre-of-mass frequencies	80
6.4	Single-carrier demodulation of incoherent texture fields	81
6.5	Partitioning of Fourier spectrum in multi-carrier demodulation	82
6.6	Multi-carrier demodulation of incoherent texture fields	84
6.7	Contrast enhancement using texture demodulation	86
6.8	Contrast enhancement using morphological operation	91
6.9	Contrast enhancement using histogram equalization	92
6.10	Target and background regions in ROIs	95
6.11	Quantitative contrast measures for MIAS images	96
6.12	Average of normalized contrast measures for MIAS images	97
6.13	Results of enhanced masses from the MIAS images	99
6.14	Texture demodulation of normal tissue	102
A.1	Expanding set of image attributes	112
A.2	System model of attribute-driven segmentation and analysis	113
A.3	Three major components of the attribute-driven system	114
A.4	Attribute dependence graphs	115
A.5	Sequential attribute acquisition with single desired node	117
A.6	Sequential attribute acquisition with multiple desired node	118
A.7	Parallel attribute acquisition	119
A.8	Eight possible orientations of an image	120
A.9	Various image co-ordinate systems	121
B.1	Graphical representation of the <Point> element	131
B.2	Graphical representation of the element of <code>lineType</code>	132
B.3	Graphical representation of the element of <code>arrayType</code>	132
B.4	Graphical representation of the <Image> element	132
B.5	Graphical representation of the <Mammogram> element	133
B.6	Graphical representation of the <Breast_Anatomy> element	133
B.7	Graphical representation of the <Adequacy> element	134
B.8	Graphical representation of the <Abnormality> element	135

List of Tables

4.1	Five-point assessment scale with score descriptions	43
4.2	Numbers of segmentation accepted by the two radiologists	44
4.3	Contingency tables of straight line and curve segmentation	44
4.4	Segmentation accuracy ratings by the two radiologists	45
4.5	Contingency tables of radiologists' ratings	46
5.1	Number of MIAS images fulfilling the quality criteria	67
6.1	The number of highest scores achieved by enhancement techniques	98
6.2	Pearson correlation coefficients between contrast measures	98

Acknowledgements

I would like to thank my supervisors, Dr. Ramachandran Chandrasekhar, Prof. Yianni Attikiouzel and A/Prof. Thomas Bräunl, who have provided excellent guidance, encouragement and support during the course of my Ph.D. research. This thesis would never have been written without their help. I would also like to thank the Head of School, A/Prof. Gary Bundell, and other members in the research groups of which I have been a member: Centre for Intelligent Information Processing Systems (CIIPS) at The University of Western Australia (UWA), and Australian Research Centre for Medical Engineering (ARCME) at Murdoch University.

I specially thank Dr. M. Rickard and Dr. N. Borecky, of BreastScreen NSW, who assessed the segmentation results in Chapter 4. I would also like to thank Prof. I. James, of Murdoch University, and Dr. D. Milech, of UWA, for their expert guidance in statistics. Mark Sheriffs of UWA made the artwork shown in Fig. 1.1.

I am grateful to my fellow students for their friendship and assistance, particularly to Martin Masek, Pantazis Houlis, Thomas Hanselmann, James Young, Frederick Chee and Fangwei Zhao. They have provided valuable advice for my research, and interesting sporting activities to keep me healthy.

Finally, and most importantly, I am deeply thankful to my parents, my sister and Susanna for their understanding and patience. To them, any gratitude I may express will not be enough.

Chapter 1

Introduction

Mammograms are X-ray images of the breast. They are used in the early detection of breast cancer, which is the most common cause of death from cancer for women in Australia [1]. In the attempt to decrease breast cancer mortality rates, a *mammographic screening program* has been established nationally in Australia, where healthy women in the target age group of 50–69 years are encouraged to participate in mammographic examinations [1]. From 1993 onward, there has been a steady decline in mortality rates from breast cancer for Australian women in the target age group, indicating the success and effectiveness of the screening program [1].

Breast cancer screening produces a large number of mammograms everyday. Currently in Australia, each mammographic examination produces four X-ray images; two different views for each breast. Each of these mammograms must be viewed and diagnosed carefully by an expertly trained radiologist to ensure that breast cancer is correctly identified when the disease is present (high sensitivity), and no cancer is incorrectly detected when it is not really there (high specificity). This is a demanding task, and it would be beneficial if computers could be used to improve radiologists' ability to detect the few cancer cases in the large amount of normal-looking images that they view everyday [2].

Mammographic image analysis is a challenging task in the field of medical image processing. Due to the variable appearance of the compressed breast, superimposition of different types of tissue and subtle signs of cancer, the image processing problems to be solved are generally non-trivial [3]. It is even more different to develop a fully automatic image analysis system, that takes the digital image of a mammogram as its input, and presents an analytical report with annotated images as its output, without any human interaction during the entire process. An overview of such a system is shown in Fig. 1.1: it represents the overall goal of which this research forms a part.

In our approach, the development of this *mammogram analysis system* revolves around the concept of *attribute-driven image segmentation and analysis*, in which the knowledge about an image is described by a set of *image attributes*. These attributes are extracted progressively in a step-by-step, cyclic procedure where global, low-level attributes gathered in the early stages are used to derive local, high-level attributes in the later stages.

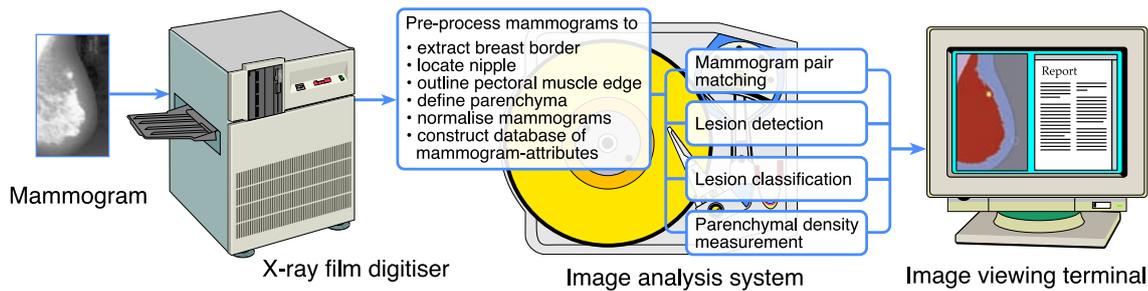


FIGURE 1.1: Overview of the Image Analysis System for Mammograms.

While the complexity of segmentation and analysis algorithms increases during the procedure, there is also a monotonic accumulation of knowledge inside the system, leading to increasing refinement and accuracy of the results.

The research described in this thesis forms part of the proposed automated image analysis system, shown in Fig. 1.1. The first part of the original work consists of a new attribute-driven system designed for mammogram segmentation and analysis. The second part of the original contribution pertains to the special-purpose techniques that were developed for the analysis of mammograms. They include:

1. automatic segmentation of the pectoral muscle;
2. adequacy assessment of breast positioning and film exposure; and
3. contrast enhancement of mass lesions.

In the process of developing the attribute-driven system, it was realized that the approach could be generalized into an *abstract system model*. We believe that this generic model may be customized for a wide range of image processing problems, following a design strategy. However, a full experimental evaluation of this discovery is beyond the scope of this thesis. Some fundamental concepts of the abstract system model are documented at the end of this thesis to aid further investigation. Nevertheless, it is worthwhile to note that the mammogram analysis system is actually an instantiation of the abstract system model, so the two are closely related.

1.1 Outline of Thesis

This thesis consists of seven chapters and three appendices. Following this introductory chapter, a review of the literature on mammogram segmentation and analysis is given in a single chapter. The next chapter is devoted to the design and implementation of the mammogram analysis system. The experiments that constitute the original work on mammographic image analysis are presented in the three succeeding chapters, which are followed by the concluding chapter. The various chapters are outlined below.

Chapter 2 gives a literature review of the existing segmentation and analysis techniques developed for mammograms. The purpose of this chapter is to provide an overview on the

previous work published in three main areas: mammogram segmentation, mammographic image quality, and image enhancement of mammograms. Each of these areas is directly related to the experimental work presented in Chapters 4 to 6.

Chapter 3 describes the actual system implemented for the attribute-driven segmentation and analysis of mammograms. An overview of this system is first given in this chapter. It is followed by the identification of three different types of image attributes that are extracted from the mammograms. The proposed system constitutes of three major components, namely, attribute acquisition, attribute normalization, and attribute storage. Attribute acquisition is controlled by a series of computer applications that extract image attributes progressively from an image, and derive high-level attributes from low-level ones. An attribute dependence graph is introduced here to represent the data dependence and priorities of the image attributes. Attribute normalization is used to standardize the data format of image attributes generated by different applications. Several aspects of the scale and representation of image attributes are considered. Attribute storage is accomplished by building an attribute database. An Extensible Markup Language (XML) database with a properly designed schema is proposed for the storage of image attributes.

The computer applications developed for extracting image attributes are not described in Chapter 3; instead they form the substance of the three succeeding chapters. The image processing techniques presented in Chapters 4 to 6 are self-contained with their own notational and experimental frameworks. They were developed as individual applications, but they were also fully integrated into the mammogram analysis system, exchanging image attributes with each other.

Chapter 4 is the first experimental chapter in which a fully automatic method of pectoral muscle segmentation is described. On the mediolateral oblique (MLO) view of mammograms, the pectoral muscle appears a high-intensity, triangular region across the upper posterior edge of the image. The boundary of the pectoral muscle is usually slightly curved. In the proposed method, the pectoral margin is first approximated by a straight line. The position and orientation of this straight line is validated using the endpoints of the breast border. Then the straight line is iteratively refined into a curve using an edge detector. The pectoral muscle is finally represented by an enclosed region.

Chapter 5 is the second experimental chapter in which the potential of computerized adequacy assessment of mammograms is addressed. A mammogram must be of sufficient quality for medical diagnosis with high sensitivity and specificity. In this experimental chapter, the positioning of the breast on a mammogram is assessed using a number of quality criteria, including non-exclusion of breast tissue, visualization of the nipple in profile, inclusion of inframammary fold, and proper positioning of the pectoral muscle. Algorithms are devised in this chapter to automatically evaluate these quality criteria of breast positioning on mammograms of MLO view. In addition, the quality of film exposure is also evaluated by examining the distribution of optical densities, which is generated from

the pixel values of the digitized mammogram.

Chapter 6 is the third experimental chapter in which the techniques of contrast enhancement of lesions is investigated. A novel contrast enhancement technique is presented in this chapter for highlighting different types of mass lesions. This technique was derived from the amplitude and phase (AMPM) demodulation transform introduced by Daugman and Downing [4] for texture analysis. The chapter therefore begins with the mathematical definition of AMPM demodulation transform. The original transform was extended from single-carrier to multi-carrier, but it was found ineffective when applied on mammograms. Hence the AMPM demodulation transform is further modified for the purpose of contrast enhancement of mass lesions. The modified demodulation transform was compared with two existing contrast enhancement techniques, which are morphological enhancement and histogram equalization. The experimental results were evaluated using three quantitative measures of mammographic contrast enhancement, developed by Bovis and Singh [5].

Chapter 7 concludes the thesis with a summary of the mammogram analysis system and the three problem-specific image processing techniques introduced in the experimental chapters. Some suggestions for further work are also given in this concluding chapter.

After the conclusion of this thesis, Appendix A introduces the abstract system model that may be applicable to the attribute-driven segmentation and analysis of a wide range of images. This appendix is closely related to Chapter 3, since the abstract system model is actually a generalization of the mammogram analysis system, but it is not in the scope of this thesis. The fundamental concept of knowledge accumulation using image attributes is first explained in this appendix, then the generic framework of the system model is presented. Like the actual system, this generic model is also composed of the three major components, i.e., attribute acquisition, normalization and storage. Each of these components is described at the abstract level with additional detail and examples. A design strategy is given to assist the instantiation of the abstract model into actual systems. The properties of XML and its advantages and disadvantages are also discussed in this appendix.

Because this appendix is written entirely from the abstract point of view, with the intention to separate all the general aspects from an actual system, the domain-specific issues of mammogram segmentation and analysis are not discussed in this appendix. The reader should keep in mind that the abstract system model is applicable to any class of images, not only to mammograms. As a result, the three major components of the system model and the design principles related to them are also presented from the generic perspective.

Chapter 2

Literature Review

A review of the literature on the segmentation and analysis of mammograms is given in this chapter. It focuses on three main topics: mammogram segmentation, mammographic image quality, and image enhancement of mammograms. These three topics correspond respectively to the work presented in Chapters 4 to 6 of this thesis.

A wide range of approaches for global segmentation of the mammogram are reviewed here. These approaches include the segmentation of the breast border, the nipple, the pectoral muscle and the fibroglandular region. For mammographic image quality, a list of quality factors published in the clinical and technical literature is presented, but the emphasis of this survey is on the adequacy of breast positioning and film exposure. For image enhancement of mammograms, several contrast enhancement techniques are compared, including global image enhancement, region-based enhancement, and window-based enhancement.

2.1 Mammogram Segmentation

Mammogram segmentation is usually considered the primary step in the computer processing of mammograms, and it is not a easy task to automate. In many cases, the accuracy of segmentation directly affects the outcome of the subsequent image analysis procedures. For this reason, a reliable segmentation scheme is certainly a necessary component in a mammogram analysis system. The segmentation techniques described in this section refer to the *global segmentation* of mammograms, in which the mammogram is partitioned into several disjoint regions. Segmentation of local objects on a mammogram, such as lesions and microcalcification clusters, is not included here.

Many attempts have been made in the literature [6–32] to segment a mammogram globally into distinct regions. These regions represent the anatomical features of the breast, such as the breast border, the nipple, the pectoral muscle and the fibroglandular region. The first three are often considered the *landmarks* on mammograms. Separating the breast object from the image background is almost always the first step in mammogram segmentation. This step produces a breast border, or skin-air interface, on the

mammogram.

Among the earlier papers, Miller and Astley [6] proposed the use of anatomical features for breast asymmetry detection. A semi-automated procedure was used to exclude the non-breast background and the pectoral muscle from the mammogram. Three types of feature images are generated for the breast region. A threshold is determined from each of these feature images and this threshold is used to segment the fat and non-fat regions. Specific detail on how to generate the feature images and how to select the threshold is not given in the paper. Only the segmentation accuracy of the non-fat region is computed in their work, and an average accuracy of 80.3% was reported for the best set of results.

Bick *et al.* [7] presented an automatic method for segmenting the breast region on digitized mammograms. In their approach, the original image is first subsampled to a lower resolution. The local grey-level range is then calculated for each image pixel as the difference between the local maximum and local minimum on an annulus centred on the pixel. A modified histogram is computed using all the pixels with small local grey-level range. The histogram is thresholded to label the image pixels as directly exposed, unexposed and potential breast regions. Region growing and morphological filtering are then used to remove artefacts on the image before the breast border is delineated. This method has been applied to 740 mammograms, including both mediolateral oblique (MLO) and craniocaudal (CC) views, scanned by three different digitizers. The results were rated by two mammographers and two physicists, and a success rate of 97% was reported.

Suckling *et al.* [8] used multiple linked self-organizing neural networks to segment the mammogram. They proposed that the mammogram can be segmented into four major components: background, pectoral muscle, fibroglandular region and adipose region. The algorithm developed by them is based on the classification of feature vectors constructed from statistical texture measures. Self-organizing neural networks are trained separately and linked together via a decision logic. They are used to classify the feature vectors extracted from the mammogram. The method has been tested on 30 digitized mammograms. Only the performance of parenchymal segmentation was measured, and the mean and standard deviation of the agreement measure of overlapping areas was 0.69 ± 0.12 .

Mendez *et al.* [9] described an fully automatic technique for breast border and nipple detection on mammograms and stated that mammogram segmentation is a pre-requisite for further image analysis. In their approach, the image is first thresholded to eliminate border misplacement due to the film digitization process. The image is also smoothed to filter out noise. Then five reference points are detected on the image in order to divide the breast into three regions, and in each region the breast border is searched for in different directions, using a tracking algorithm. After that, three algorithms are used to detect the nipple: (1) maximum height of the breast border; (2) maximum gradient of the grey-levels across the median-top section of the breast; and (3) maximum second derivative of the grey-levels across the median-top section of the breast. Algorithms (1) and (3) are then combined to derive the final nipple detection method. Their method has been tested on

156 mammograms obtained from both lateral oblique (AXI) and CC views. The breast segmentation results were evaluated by two expert radiologists and one physicist and 89% of them were reported to be accurate. The detected nipple positions were compared with the real positions indicated by radiologists, and a mean difference of 6 mm was reported.

Chandrasekhar and Attikiouzel [12,16,18] used background modelling and subtraction to extract the breast border automatically. In their approach [18], the background region is first estimated by examining the cumulative intensity histogram and calculating the *background monochromaticity*. The estimated background is then modelled by a two-dimensional polynomial with degree 0 to 2, and the modelled image is subtracted from the original to yield a binary image of the breast region. This method has been applied to 322 mammograms, and 32 of them were randomly chosen and evaluated by a radiologist. A success rate of 94% was reported. The segmented breast border was accurate enough for the nipple contour to be delineated.

Besides the segmentation of the breast border, Chandrasekhar and Attikiouzel [13,16] presented a simple and accurate method for automatically locating the nipple on the skin-air interface of a mammogram. In their approach [13], normal lines are drawn on the breast border and the average intensity gradient along the normal direction is computed. If the nipple is in profile, it is located by detecting the sudden change of normal direction. If the nipple is not in profile, it is located by detecting the distinctive change of average intensity gradient. This method has been tested on 24 mammograms from two databases, including both MLO and CC views, and a success rate of 96% was reported. Their method needs to be performed on a fairly smooth and accurate breast border. In other cases, some researchers, such as Georgsson [20] and Yam *et al.* [21], simply assumed that the point on the breast border furthest away from the pectoral muscle in the MLO view, and furthest away from the posterior image edge in the CC view, was the nipple location, although this was only an approximation.

Chandrasekhar and Attikiouzel [22] also investigated mask-based edge detection in order to delineate the pectoral muscle. In their approach, a families of tunable parametric edge detectors were tested on mammograms and the pectoral edge were highlighted significantly with one of the detectors. The edge enhanced images were generated on six pairs of mammograms, and they were visually inspected. Although promising results were shown, no methods were given to trace out the continuous pectoral margin on the edge images, and to yield a segmentation mask of the pectoral muscle.

Chandrasekhar and Attikiouzel [19] extended the Russ operator and presented a new range-based neighbourhood operator for texture analysis. This operator was initially developed to discriminate between the adipose and fibroglandular tissues on mammograms [11]. By applying their extended Russ operator to mammograms, three types of texture enhanced images were generated, namely, *c*-image, *m*-image and η^2 -image. They found that the *c*-image was useful in detecting weak edges with good noise immunity, and may be used to outline the skin-air interface on mammograms. They also suggested that

the η^2 -image may be used to enhance the pectoral edge, the boundary of parenchyma, and possibly some circumscribed lesions.

Karssemeijer [14] segmented the mammogram for parenchymal pattern classification. He pointed out that accurate and robust segmentation is a pre-requisite of quantitative feature extraction on mammograms for classification. His method automatically segments a mammogram into three regions: breast tissue, pectoral muscle and background. In other words, only two anatomical features (breast border and pectoral muscle) are detected. The method for breast region segmentation is very simple; it is just a global thresholding technique using histogram peak detection. His major contribution pertains to the successful application of Hough transform to approximate the boundary of pectoral muscle using a straight line. This simple method has been adopted by several other researchers for pectoral muscle segmentation, and the details are discussed in the literature review of Chapter 4.

Aylward *et al.* [15] combined geometric and statistical modelling techniques and segmented the mammogram into five regions: background, uncompressed-fat, fat, dense and pectoral muscle. The breast border and pectoral muscle edge are first extracted using a gradient magnitude ridge traversal algorithm for the geometric models. After that, statistical models are used to define different types of breast tissue. These statistical models are formed using pixel intensities, the concept of distribution sampling, and Gaussian mixture models. They also devised three applications of the segmentation method: (1) general intensity windowing for breast cancer screening; (2) component-specific intensity windowing for breast lesion characterization; and (3) breast density estimation for breast cancer risk assessment. These applications were applied to 70 mammograms from three digital mammography units. The results were judged by expert mammographers to be “clinically interesting” [15]. They also claimed that the geometry models were generated successfully on all images, and on 95% of the images for the statistical models. However, quantitative evaluation of their results was not given.

Ferrari *et al.* [17] proposed a method to identify the skin-air interface, pectoral muscle and fibroglandular disc. In their approach, the mammogram is first enhanced in contrast and morphologically filtered. An approximate boundary is then extracted using a chain-code method and smoothed by a cubic B-spline interpolation technique. After that, the true boundary of the breast is detected by drawing normal lines on the approximate boundary, and by examining the grey-level histogram of the pixels on each normal line. Finally, the skin-air boundary is smoothed again. This method has been tested on 66 mammograms and the skin-air interface was detected accurately on 61 of them (a success rate of 92%). However, the nipple contour was eliminated in some images due to the smoothing of cubic B-splines. In addition, neither in this paper [17] nor in [27] was their method of breast border segmentation described in full detail, making it difficult for other researchers to implement their method. Furthermore, they modified Karssemeijer’s Hough transform method [14] to approximate the pectoral muscle using a straight line.

The pectoral muscle was accurately detected, within a difference of 1 to 3 mm, on 58/66 (88%) images. The segmentation of the fibroglandular disc was also described in their work, but the results were only visually compared with the histogram equalized original image. Nevertheless, a success rate of 55/66 (84%) was reported.

Recently, Ferrari *et al.* [27] presented another automatic method for pectoral muscle identification using Gabor wavelets. In their approach, the pectoral muscle edge is first enhanced by a bank of Gabor filters specially designed in the frequency domain. The magnitude and phase images are then computed using a vector-summation procedure, and the magnitude value of each pixel is propagated in the direction of the phase. Finally, the pectoral boundary is formed by detecting relevant edges, connecting disjoint segments and eliminating false edges. This method was applied to 84 MLO view mammograms, and the results were compared with the pectoral edges drawn by one of the authors under the supervision of a radiologist. The average false-positive and false-negative rates reported were $0.58 \pm 4.11\%$ and $5.77 \pm 4.83\%$, respectively. Their pectoral muscle segmentation method using Gabor wavelets was compared with that using Hough transform [17], for which the average false-positive and false-negative rates were $1.98 \pm 6.09\%$ and $25.19 \pm 19.14\%$, respectively. A summarized literature review of pectoral muscle segmentation is also given in Chapter 4.

Masek *et al.* [30,31] proposed a method based on expanding window to select a threshold for separating the pectoral muscle from other breast tissue on the image. In this approach, an image thresholding algorithm is applied to an expanding window, anchored at the image corner near the pectoral muscle, to determine a series of thresholds as a function of window size. The point of inflection in this function yields the threshold for segmenting the pectoral muscle. The edge of the pectoral muscle is then approximated by a straight line, but it is *not* refined into a curve. Their method was tested using two different thresholding algorithms: (1) a modified minimum cross-entropy thresholding algorithm [30], and (2) local median intensity [31]. Finally, the results of these two algorithms were combined with those obtained from multi-level global thresholding [32] to increase the segmentation accuracy.

2.2 Mammographic Image Quality

The image quality of a mammogram is influenced by a number of clinical and technical factors. Standardized *quality assurance programmes* have been established in clinical centres in the US and elsewhere to ensure that high mammographic image quality is achieved for every image. Specific testing routines have to be carried out by the radiologists, the technologists and the medical physicists in the facility [33]. It would be of great benefit if computerized image quality assessment could be integrated into the quality assurance programme of mammography.

2.2.1 Quality Assurance and Quality Control

The definitions of quality assurance and quality control in mammography is given in [33, p. 63] as: “quality assurance entails procedures that guarantee the quality of all facets of mammography practice. This includes equipment, radiation exposure, and film interpretation. A subset of this is quality control, which involves the technical procedures that guarantee a quality mammogram.” Quality control includes acceptance testing, establishment of baseline performance of equipment, assessing the reasons for performance changes in equipment, and documentation [33, p. 63].

Eklund *et al.* [34] emphasized the importance of constant attention to quality control in mammography. They suggested that a consistent *case-by-case, film-by-film, quality assessment*, with immediate corrective measures and repeat imaging, is essential for maintaining high image quality. The significant factors of mammographic image quality assessment include [34]:

1. clinical history, including the nature and specific location of clinical concerns;
2. availability of comparison images;
3. adequate visualization of areas of clinical or radiographic concern;
4. optimal amount of tissue included, particularly for MLO and CC views;
5. adequate exposure;
6. high contrast;
7. high resolution;
8. proper compression;
9. freedom from artefacts;
10. optimal viewing conditions, including freedom from distractions and high intensity light from view boxes; and
11. proper labelling of images.

Two of the eleven factors, (4) and (5) in the list above, are considered in Chapter 5 for automatic adequacy assessment using computer. The first factor is related to the *positioning* of the breast when the mammogram is acquired; the second factor is related to the *exposure* and *optical density* of the film.

2.2.2 Mammographic positioning

For breast positioning on mammograms, Eklund *et al.* [34] described in detail a number of quality criteria for optimum positioning in both MLO and CC views. These quality

criteria are also outlined in other literature [33,35–38], and they are summarized in Chapter 5. Breast positioning assessment relies on extracting anatomic landmarks accurately on mammograms. Thus automatic mammogram segmentation is particularly important to the automation of breast positioning assessment using computers. Moreover, the significance of breast positioning is such that it has been claimed in [38] that “invasive breast cancer detection by mammography may be improved through attention to correct positioning.”

From our literature survey, very few publication have tackled the problem of computerized breast positioning assessment. Among these is the work of Olsén, on *automatic determination of mammogram adequacy* [39]. She proposed a mathematical model, on the basis of fuzzy logic and pattern recognition, for combining quality criteria used in adequacy assessment. However, she implemented and tested only a small portion of the proposed system, and used only one anatomical feature, the pectoral muscle, for adequacy assessment on MLO view mammograms. So only the position and curvature of the pectoral boundary were automatically evaluated for adequacy in her work.

More anatomical features are used in our study of automatic breast positioning assessment. The breast border, the nipple location and the pectoral margin are all included in the algorithms described in Chapter 5.

2.2.3 Film Exposure

The exposure of mammograms is quantified by the optical density of the film. The importance of maintaining the mean optical density within a certain range is widely recognized in the literature [40]. An adequately exposed mammogram can be used to improve the image contrast of dense breast tissue [41]. It has also been reported in the literature that the mean film optical density is correlated to the success rate of small breast cancer detection [42]. Most importantly, “maintaining a constant and appropriate mean film optical density in mammography is an essential part of quality assurance in a breast screening programme” [40].

For the adequacy assessment of exposure, Eklund *et al.* [34] gave a rather subjective guideline, which requires that “the dense parenchymal areas of a mammogram should be adequately penetrated by radiation so that fibrous strands or vascular structures can be seen through the parenchyma,” and the film optical density of areas of highest glandular density should be no less than 0.7.

A more up-to-date guideline is given by Heywang-Köbrunner *et al.* [33] that requires the film to be exposed “in such a manner that all details relevant to the diagnosis are visualized in the optimum density range.” They have also stated that “The useful range of every film is limited to optical densities between 0.6 and 2.2–2.8” [33]. The mean optical density should be maintained approximately in the middle of the useful optical density range, i.e., between 1.4 and 1.8 [33]. These numerical constraints are useful for the computerized assessment of film exposure on digitized mammograms (Chapter 5).

2.3 Image Enhancement of Mammograms

Different image enhancement techniques have been used to improve the contrast of mammograms, so that more accurate abnormality detection and clinical diagnosis can be performed, by either human observers or computerized algorithms [5, 43–53]. Some of these enhancement techniques are reviewed in this section.

A commonly-used contrast enhancement technique for mammograms is *histogram equalization*. It is a simple, general and fast image transformation that is defined in many image processing textbooks [54, 55]. Gupta and Undrill [48] used histogram equalization to improve the visualization of lesions in their texture analysis. In some papers, such as [13, 27], histogram equalization is not part of the proposed algorithms, but it is used to enhance the visualization of the skin-air interface, the nipple contour and the pectoral margin on mammograms, simply for the readers to see them clearly. However, histogram equalization does not always give good results for mammograms. The study from Bovis and Singh [5], and also the work in Chapter 6 of this thesis, has shown that histogram equalization is often outperformed by other contrast enhancement techniques that are specially designed for mammograms. In addition, Chandrasekhar [11] described a zero-anchored histogram equalization, where the lowest intensity of the equalized image is mapped back to zero intensity, giving a better contrast enhanced image.

Morrow *et al.* [49] proposed a *region-based contrast enhancement* technique for mammograms. It is an adaptive neighbourhood processing technique that grows regions in a given image and enhances the regions with respect to their local backgrounds. In their approach, seed-fill region growing is used to define overlapping regions on the image. Two parameters of region growing, namely, *region growth tolerance* and *background width*, have to be selected manually using an interactive program. Each pixel on the image is taken as a seed, and region growing is carried out iteratively until a specific background width is achieved. The contrast of each region is then computed with respect to its background. Contrast enhancement is performed by applying an empirical transformation to the seed pixels. In addition, a quantitative measure of contrast enhancement was also defined based on the histogram of region contrast, and used for result comparison. A number of case studies were given for performance evaluation.

Rangayyan *et al.* [51] aptly named the region-based contrast enhancement technique described above [49] as *adaptive neighbourhood contrast enhancement*, and evaluated its effectiveness in increasing the sensitivity of breast cancer diagnosis. The enhancement method was applied to two sets of digitized mammograms: 78 images of 21 difficult cases; and 222 images of 28 interval cancer patients and 6 benign control cases. Three sets of images: the original films, digitized mammograms and enhanced images were presented to 6 radiologists for the difficult case set, and 3 radiologist for the interval cancer set. *Receiver operating characteristic* (ROC) analysis was used to compare the radiologists' performance in abnormality detection. The area under each binormal ROC curve was computed to represent the overall performance. The larger the area parameter, the more

accurate the detection. The area parameters for the original, digitized, and enhanced images were computed to be 0.6735, 0.6259, and 0.6745, respectively, for the difficult case set. The same figures for the interval cancer set were 0.3906, 0.4682, and 0.5407, respectively. The results show that the radiologists' performance with the enhanced image is the best among the three sets of images, and that the diagnosis sensitivity is improved by image enhancement.

Petrick *et al.* [52] proposed an *adaptive density-weighted contrast enhancement* filter for mass detection on mammograms. In their approach, the original mammogram is first passed through a map rescaler, in which the image is linearly scaled to $[0, 1]$ and globally thresholded using a fixed value to separate breast tissue. The largest object in the image is then extracted and rescaled to produce a normalized image. The normalized image is split into a density image, by applying a low-pass filter, and a contrast image, by applying a band-pass or high-pass filter. The density image is used to determine a multiplication factor, which is multiplied with the contrast image. The weighted contrast image is non-linearly rescaled to produce the final enhanced image. Their method was applied to 25 digitized mammograms with biopsy proven masses. The method detected 24 of the 25 true mass objects, but the number of false-positives needs to be further reduced.

Li *et al.* [53] used morphological filters to enhance patterns of suspected masses on mammograms, by cleaning up unrelated background texture. Using their method, the high intensity background, caused by dense breast tissue overlapped on the mass, is removed while preserving the patterns related to the mass. Hence it is a background correction procedure. In their approach, two basic morphological operations, erosion and dilation, are used to form a *dual morphological top-hat operation*. This top-hat operation is applied twice on the mammogram using structuring elements of two different sizes: one smaller than the minimum size of possible masses; the other larger than the maximum size of possible masses. As a result, most of the background texture and noise inside the mass is suppressed in the enhanced image. Note that their method is pattern-dependent since it relies on measuring the effective size of targeted mass patterns. In Chapter 6 of this thesis, their morphological enhancement technique is implemented and compared with our contrast enhancement technique.

Bovis and Singh [5] proposed three quantitative measures to evaluate the performance of contrast enhancement techniques applied to mammographic masses. These measures are: (1) distribution separation measure; (2) target to background contrast ratio using variance; and (3) target to background contrast ratio using entropy. They can be applied to region of interest (ROI) containing the mass, with the assumption that the target and background areas are accurately defined. Bovis and Singh tested these measures using 41 mass images, and compared the effectiveness of histogram equalization and fuzzy enhancement technique. The results have shown that fuzzy enhancement performed better. The detail of these quantitative measures are described in Chapter 6 of this thesis, and they are used to compare our contrast enhancement technique with two existing ones.

This concludes the review of the literature. In the next chapter, the design of an image analysis system is presented for the segmentation and analysis of mammograms.

Chapter 3

Mammogram Analysis System

A *mammogram analysis system* was developed in this research for the segmentation and analysis of mammograms. The design of such system is presented in this chapter. We consider that mammographic segmentation and analysis to be a hierarchical and cyclic process, which is driven by the image attributes of mammograms. Three types of image attributes are identified for mammograms in this chapter. These mammogram attributes are derived from the images progressively from gross to subtle, from low-level to high-level, from global to local, and from approximate to precise. The system consists of three major components, namely, attribute acquisition, attribute normalization, and attribute storage. The first component, attribute acquisition, corresponds to the mammogram segmentation and analysis algorithms presented later in Chapter 4, 5 and 6. The attributes generated by these algorithms are organized into a graph, to show the sequential order in which they should be derived in the system. For the other two components, the specification of a standard attribute format and the design of a attribute database are discussed in detail in this chapter.

3.1 System Overview

An overview of the mammogram analysis system is shown in Fig. 3.1. The system is attribute-driven and consists of three major components: (1) a collection of computer applications for *attribute acquisition*; (2) a standard attribute format for *attribute normalization*; and (3) a database for *attribute storage*. The input to the system is a digitized mammogram. The image is processed by a series of applications, which are executed in a pre-defined, sequential order. These applications include image segmentation, adequacy assessment, lesion detection, and other mammographic image analysis techniques. Each of the applications generate one or more attributes of the image, which are called *mammogram attributes*. Mammogram attributes are normalized to a standard format and stored into a database, so data exchange is possible between different applications. At any stage, the mammogram attributes stored in the database can be retrieved and used by an application to derive more attributes for further analysis. Consequently, a cyclic procedure is

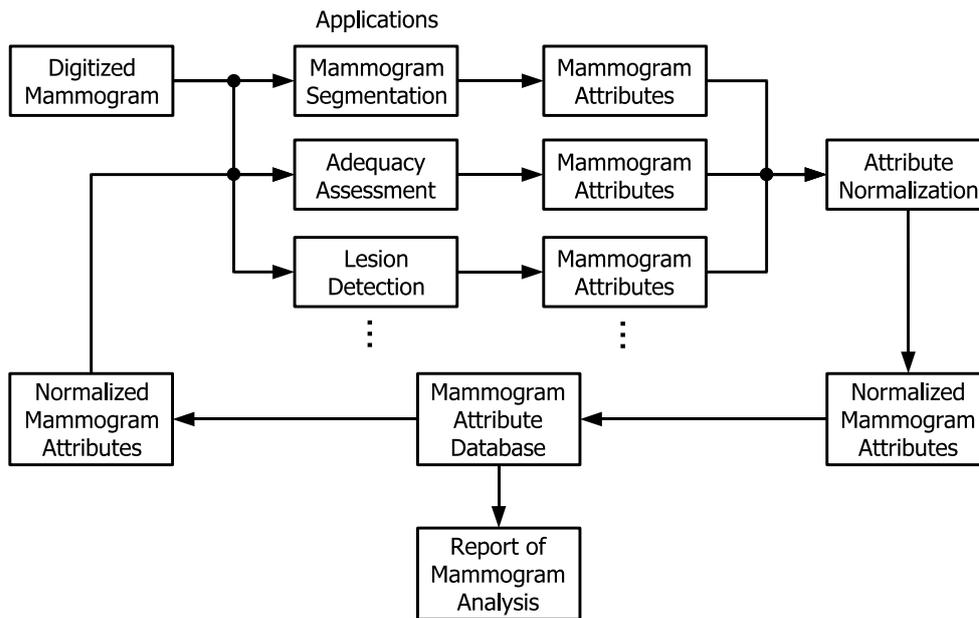


FIGURE 3.1: Mammogram analysis system. The design of this system follows the generic scheme presented in Fig. A.2.

formed in the system, where mammogram attributes are derived progressively from low-level to high-level and knowledge about the image is built up in the database. The output of the system is a report containing the results of image analysis.¹

3.2 Mammogram Attributes

Mammogram attributes are categorized into three types: image data, visual features and analytical information. These three types of attributes are also derived in the segmentation and analysis of mammograms. All the possible mammogram attributes that can be exchanged between applications through the database have to be carefully identified during the development of the system. The current set of mammogram attributes that we consider useful for our applications is described below. More mammogram attributes may be added to this set when new applications are developed.

3.2.1 Image Data

Image data includes the properties of the mammogram and its digital image. This type of attribute can be obtained when the mammogram is acquired and digitized, before any segmentation or analysis application is performed on the image. The following image data is used in our mammogram analysis system:

- image ID
- mammographic view (MLO/CC)

¹A generic form of the mammogram analysis system is documented in Appendix A.

- left/right breast
- film density
- peak kilovoltage (kVp)
- exposure (mAs)
- image file index
- image dimensions
- image resolution ($\mu\text{m}/\text{pixel}$)
- grey-level bit-depth
- response curve of digitizer (optical density to pixel value conversion)

3.2.2 Visual Features

Visual features consist of the description of breast anatomy and the texture of breast tissue. They are extracted directly from the image, using hierarchical segmentation and image texture analysis, at the earlier stages of mammographic analysis. The following are attributes of visual features:

- breast border
- nipple location
- pectoral margin (straight line and curve)
- outline of parenchyma
- density of breast tissue
- textural features

3.2.3 Analytical Information

Analytical information includes patient history, clinical findings and diagnostic results of the mammogram. This type of attributes are normally generated at the later stages of mammographic analysis, after the mammogram has been segmented into distinctive regions. For example, these attributes are considered analytical information:

- patient history
- adequacy of image quality
- adequacy of breast positioning

and, if applicable,

- locations of microcalcifications
- locations and sizes of masses
- types and severity of masses
- locations of breast asymmetry

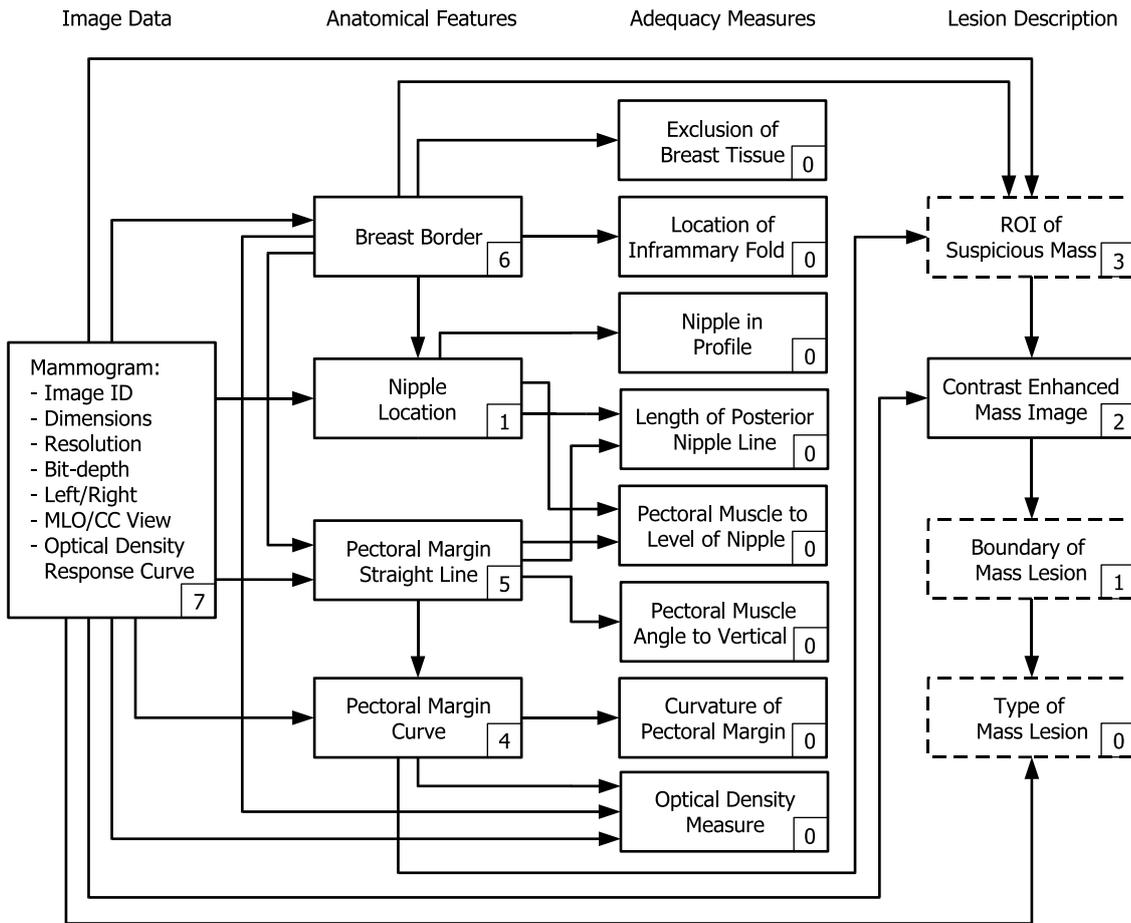


FIGURE 3.2: The dependence graph of mammogram attributes derived from hierarchical segmentation, adequacy measures and lesion detection. The number at the corner of each box indicates the attribute priority. The acquisition of the attributes in dashed line boxes has not been automated yet.

3.3 Acquisition of Mammogram Attributes

After the mammogram attributes have been categorized into three types, they are generated using a series of computer applications in the framework illustrated in Fig. 3.1. The mammogram attributes are derived progressively from low-level to high-level, i.e., from image data to visual features, and from visual features to analytical information. Low-level attributes extracted from the image at the earlier stages are used to generate high-level attributes at the later stages. Therefore, the output attributes of one application may be the input attributes of one or more subsequent applications. As a result, a very important property of mammogram attributes is revealed; this property is *dependence*.

The dependence of mammogram attributes, derived from hierarchical segmentation, adequacy assessment and mass lesion detection, is depicted in a *dependence graph* in Fig. 3.2. In this dependence graph, mammogram attributes are shown as *nodes* and the dependence relations are represented by *directed edges*. For example, the optical density measure depends on the image data, breast border and pectoral curve. This means that

the latter three attributes must be acquired first before the optical density measure is calculated.

In the mammogram analysis system, attributes are acquired sequentially. This means that at each cycle only one of the applications is executed and only the output attributes of that particular application are normalized and stored into the database. The *sequential order* of attribute acquisition can be determined by examining the dependence graph of mammogram attributes and assigning *priorities* to them (see Fig. 3.2). Attributes with higher priorities are acquired first, they are used to derive attributes with lower priorities. The algorithm that we use to assign priorities to attributes is called *attribute priority assignment*, which is described in Section A.5.2.

In Fig. 3.2, the attributes of image data have the highest priority in the graph, so they are acquired in the first step and stored into the database. Then hierarchical segmentation can be carried out on the mammogram using the attributes of image data. The order of segmentation is pre-defined as {breast border \rightarrow nipple location \rightarrow pectoral margin (straight line) \rightarrow pectoral margin (curve)}. However, this is not the only order allowed. The order can also be {breast border \rightarrow pectoral margin (straight line) \rightarrow pectoral margin (curve) \rightarrow nipple location}.

After the mammogram is segmented and the anatomical features are stored into the database, the adequacy and quality of the mammogram can be assessed. As shown in the dependence graph, the priorities of the attributes of adequacy measures are all zero (the lowest priority). This indicates that all these measures are independent of each other, and they can be acquired in any order. In fact, the adequacy measures constitute a good example of parallel attribute acquisition (Section A.5.3), although the actual system that we implemented only supported sequential attribute acquisition.

If the mammogram is assessed to be adequate, mass lesion detection may be performed. In our design, mass lesion detection is divided into four steps: suspicious mass detection, contrast enhancement, mass lesion segmentation and mass lesion classification. Only the second step, contrast enhancement, has been developed in our work (Chapter 6). We have assumed that the attributes of the other three steps can also be acquired, and have shown their relationship with other low-level attributes in the dependence graph. The first two steps of mass lesion detection require the input of image data and some anatomical features. But once the contrast enhanced ROI image of a mass is obtained, mass lesion segmentation and classification could be performed on the ROI image, independently of the original mammogram. In addition, when classifying the mass lesion, the features of other masses detected on other mammograms may be desired. In this case, the attributes of other mammograms can be accessed from the attribute database in the system. Hence the data exchange required by the analysis of multiple images maybe accomplished.

It is important to note that the attributes of adequacy assessment and those of mass lesion detection are independent of each other. This can be demonstrated by the two attribute dependence graphs shown in Fig. 3.3. These two dependence graphs are the

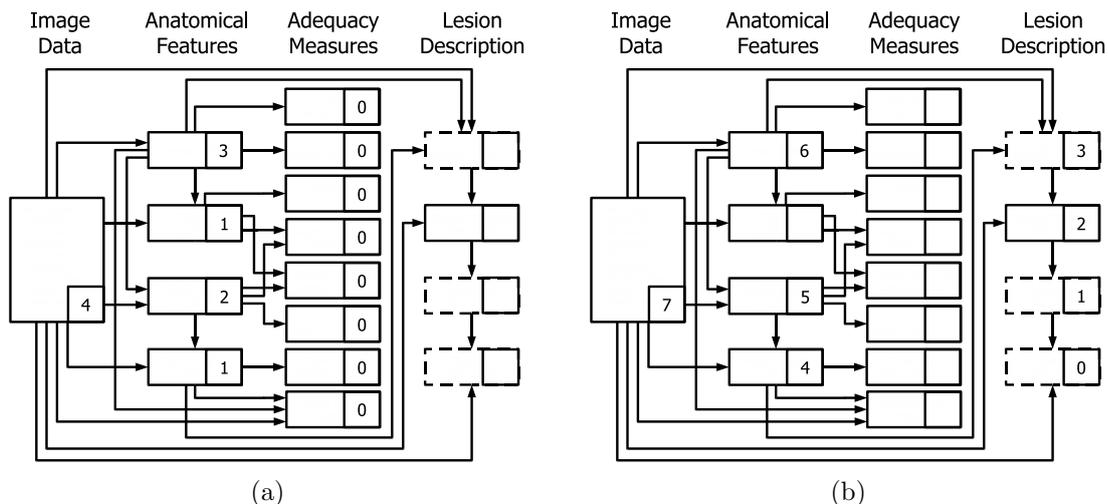


FIGURE 3.3: Referring to the dependence graph shown in Fig. 3.2. The number shown in each box is the attribute priority. (a) Attribute priority assignment for the acquisition of adequacy measures. (b) Attribute priority assignment for the acquisition of lesion description.

same as that shown in Fig. 3.2, except for the attribute priorities. In these two cases, we only want to derive a subset of the mammogram attributes. In Fig. 3.3(a), only the attributes of adequacy measures are desirable, and the graph shows that the attributes of lesion description are not required since they have no priorities assigned. On the other hand, in Fig. 3.3(b), only the attributes of lesion description are desirable, and the graph shows that the attributes of adequacy measures are not required. From these two graphs, we can conclude that adequacy assessment can be carried out independently of mass lesion detection, and vice versa.

3.4 Normalization of Mammogram Attributes

In order to exchange information between different applications in the system, the mammogram attributes have to be normalized to a standard format before they are stored into the database. The normalized attributes are expressed in XML, which is later explained in Section 3.5. When designing the specification of the standard format, all the image properties described in Section A.6 were considered. Each of them is discussed in the sections below.

3.4.1 Image Orientation

Digitized mammograms may be in different orientations due to two typical reasons: (1) a film digitization procedure may produce a misoriented image when the film is scanned in the wrong position; (2) the cassette used in computed radiography can be placed at various orientations to accommodate the examination condition, and as a result, misoriented image is produced [56, p. 789]. Algorithms to correct image misorientations may be used to rotate the original image back to standard orientation [56, pp. 789–793], [57].

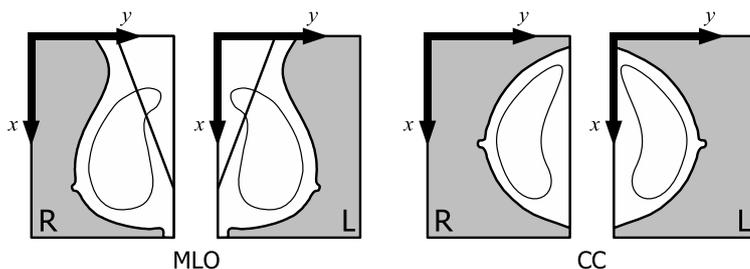


FIGURE 3.4: Standard image orientations and co-ordinate system for left and right mammograms of MLO and CC views.

For mammograms, the position of the film viewed by radiologists can be considered as the standard image orientation. Some radiologists place the mammograms of left and right breasts back-to-back on the viewbox to form mirror images [58], as shown in Fig. 3.4. So the standard image orientations of left and right mammograms are not the same, but symmetric. We should note that only one image orientation is used in the algorithms of mammogram segmentation and analysis presented in previous chapters. In these algorithms, it is easier for computation when all the right mammograms are mirrored at the vertical edge, so that the nipples face the same direction as all the left mammograms. Thus it is necessary to normalize all the attributes according to the standard image orientations shown in Fig. 3.4.

3.4.2 Co-ordinate System

The standard x - y co-ordinate system defined for the mammogram attributes has its origin at the top left corner of the image; x -axis is vertical pointing downwards and y -axis is horizontal pointing to the right (see Fig. 3.4). This setting was chosen because it complies with the right-hand rule of the Cartesian co-ordinate system. Note that a different co-ordinate system was used in the development of the algorithms presented in previous chapters. This co-ordinate system also has its origin at the top left corner, but its x -axis is horizontal and y -axis is vertical. The attributes produced by the applications of these algorithms must be converted to the standard co-ordinate system.

3.4.3 Spatial Resolution

The mammogram attributes derived from a subsampled image must be normalized according to the original spatial resolution of the mammogram. Since the mammograms from different databases may be digitized in different resolutions, no image resolution can be chosen as a single standard value for all digitized mammograms. So each mammogram has its own spatial resolution for all its attributes. Unlike the standard image orientation and co-ordinate system of mammogram attributes, spatial resolution is image-dependent and not universal.

3.4.4 Measurement Units

All mammogram attributes that involve measurements are described by standard measurement units. Any distance between two points on the mammogram is measured in millimetres (mm); and any angle between two lines is measured in degrees ($^{\circ}$). Technical data is also recorded in standard units. Peak kilovoltage set on the X-ray tube is recorded in kilovolts (kV), and exposure in milliampere-second (mAs) products [33].

3.4.5 Representations of Object Boundary

Some visual features of a mammogram, such as anatomical landmarks of the breast or mass lesions, are represented as attributes by their segmented boundaries. The requirements of object boundary normalization are addressed in Section A.6.5. An object boundary can be described using an array of points in standard image orientation, co-ordinate system and spatial resolution. If the boundary is closed (mass lesion), the boundary points are ordered in a clockwise direction. If the boundary is open (breast border, pectoral margin), the boundary points are ordered in a single direction from top to bottom. An example of object boundary normalization is given below.

3.4.6 Resulting Images

The images produced by an application, such as texture analysis, can also be treated as attributes in the system. Unlike numerical data, these *resulting images* are not normalized in any way. They are stored as image files in their specific format, possibly at a lower resolution than that of the original mammogram. The locations of these image files are recorded in the attribute database as Uniform Resource Locators (URLs), so that they can be retrieved by other applications. One example of a resulting image is the contrast enhanced ROI image of mass lesion.

3.4.7 Example: Normalization of the Breast Border

An example of object boundary normalization is shown in Fig. 3.5. When the breast border is extracted from the mammogram using background modelling and subtraction [18], the original digitized mammogram is subsampled and reduced in resolution from $50 \mu\text{m}/\text{pixel}$ to $400 \mu\text{m}/\text{pixel}$. If the original image dimensions are 2048×4320 , then after subsampling, the image dimensions are reduced to 256×540 . Also, the co-ordinate system defined in the algorithm of breast border extraction is different from the standard co-ordinate system defined in Section 3.4.2. The x and y axes are swapped.

In the application format, the breast border is defined as an array of points, (x_i, y_i) , on the image where i starts from 0. In this example, the total number of points is 586, so i is in the range $[0, 585]$. The order of the array entries runs from top to bottom. In the standard format, the breast border attribute is also described by an array of points, and only the x - y co-ordinates are different. Thus the normalization of the breast border is

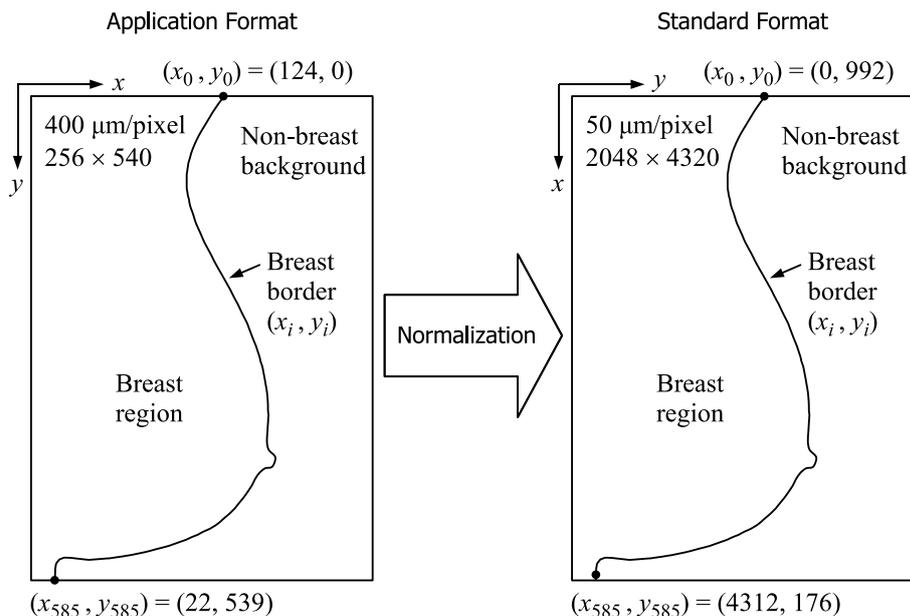


FIGURE 3.5: Normalization of the breast border attribute. The breast border is defined by the application as a series of points, (x_i, y_i) . There are 586 points on the border for this particular image. The same image orientation and the same order of array entries are used in both the standard and application formats. However different image resolutions ($\mu\text{m}/\text{pixel}$) and different co-ordinate systems are used in the two formats, hence normalization is required.

quite straight-forward. The two variables, x and y , are both multiplied by 8; and they are swapped in the co-ordinate system. Thus the top end-point $(124, 0)$ is mapped to $(0, 992)$, and the bottom end-point $(22, 539)$ is mapped to $(4312, 176)$, as shown in Fig. 3.5.

Due to the scale-up in image dimensions, the breast border in standard format is not a continuous boundary, and the bottom end-point does not coincide with the image edge. The 586 points derived by the application are instead evenly spread on the breast border in the standard format. Linear interpolation and extrapolation can be used to generate the intermediate points if it is necessary. However, these intermediate points are not stored in the database. The total number of array entries in the standard format is still 586.

3.5 Storage of Mammogram Attributes

The database of mammogram attributes basically consists of a collection of XML documents for *data storage*, and an XML schema for *data definition*. All the attributes of a mammogram are captured in an XML document and the structure of this XML document is constrained by an XML schema. A brief overview of XML database and XML Schema is presented in Section A.7. The implementation of the XML database and the design of the XML schema are discussed in the following section.

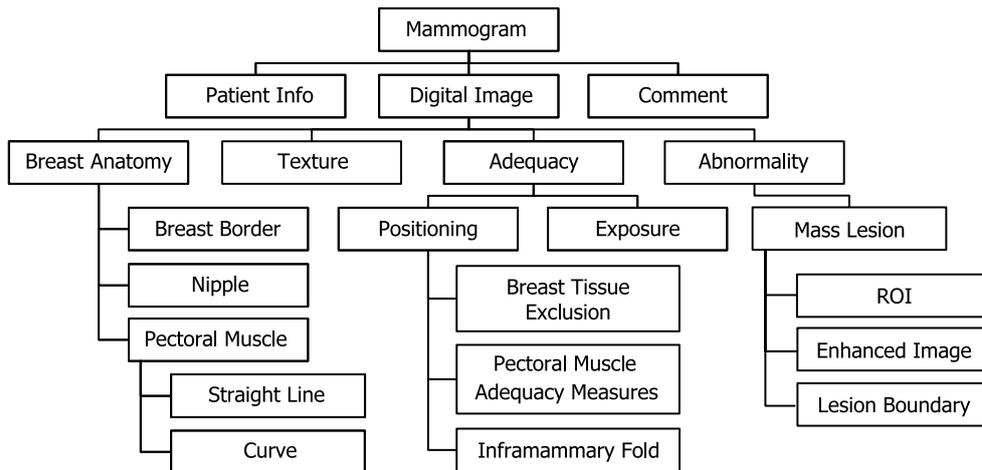


FIGURE 3.6: The element structure of some mammogram attributes defined by the XML schema.

3.5.1 A Database of Mammogram Attributes

The simplest way to implement the XML database of mammogram attributes is to use a *flat file database architecture*. In this architecture, the database provides hierarchical access to data first through the directory structure of a file system, and then through the element structure of the XML documents [59, p. 18]. Each XML document is simply a text file encoded in Unicode UTF-8 that can be easily accessed with an editor or programming interface.

All attributes derived from a single mammogram are captured in a single XML document. This data-processing-oriented document was primarily designed for data exchange and data manipulation. The XML schema of the mammogram attribute database was designed according to the standard data format specified for the mammogram attributes. Whenever an XML document is modified by the insertion of mammogram attributes, it is validated by the XML schema to ensure that the element structure and datatype are correct.

The *element structure* of mammogram attributes, defined by the XML schema, is shown in Fig. 3.6. It is a *tree* in the terminology of graph theory. At the root of the tree is the XML element `<Mammogram>`. Under `<Mammogram>` there are three subelements: `<Patient_Info>`, `<Digital_Image>` and `<Comment>`. `<Patient_Info>` and `<Comment>` are just for recording the patient's history and textual comments from radiologist. The mammogram attributes derived from image segmentation and analysis are all captured in `<Digital_Image>`. The four major categories of mammogram attributes are `<Breast_Anatomy>`, `<Texture>`, `<Adequacy>` and `<Abnormality>`. Some mammogram attributes of image segmentation, adequacy assessment, and mass lesion detection are included in three of the categories as shown in Fig. 3.6. Other mammogram attributes, such as image data, are described using XML attributes which are not shown in the figure. The full version of the XML schema is included in Appendix B.

Because mammogram attributes are progressively accumulated in the document, a very flexible and extensible element structure has to be used. The element structure is *not* dependent on the order of attribute acquisition or any application in the system. So it is not necessary to modify the element structure even if the applications have been updated in the system, as long as the attributes are normalized to the standard format. Furthermore, new mammogram attributes could be added to the database easily by attaching new XML elements at the end of the element tree. For example, an XML element `<Microcalcification>` can be attached under `<Abnormality>` to include the mammogram attributes produced by microcalcification detection. In this case, the XML schema has to be modified to include the new mammogram attributes.

The dependence graph (Fig. 3.2) and the element structure (Fig. 3.6) of mammogram attributes may look similar, but they are not the same. In graph theory, the dependence graph is an acyclic directed graph [60, p. 25], but the element structure of the XML document is a rooted tree [60, p. 13], which is acyclic and undirected. Also, some of the nodes in the dependence graph are tightly coupled, so there is only one sequential order to acquire those mammogram attributes. On the other hand, the elements in the XML schema were designed to be loosely coupled (flexible structure), hence mammogram attributes can be inserted into the XML document in any order, provided the top-level elements near the root have already been created. In other words, the attribute acquisition order is much more restrictive than the attribute insertion order to the database. Thus, altering the dependence graph does not necessarily mean that the database schema also has to be changed, unless new attributes are added.

3.5.2 Example: XML Representation of the Breast Border

After the attribute of the breast border is normalized to the standard format, it is represented using XML in the database. The breast border is an array of points as shown in the example in Section 3.4.7. It is therefore described in XML as:

```
<Breast_border array_length="586">
  <Point i="0" x="0" y="992"/>
  <Point i="1" x="8" y="984"/>
  <Point i="2" x="16" y="976"/>
  ...
  <Point i="585" x="4312" y="176"/>
</Breast_border>
```

The list of elements above has been cut short by omitting most of the intermediate points. The definition of this XML representation is expressed in the form of XML Schema:

```
<xs:element name="Breast_border" type="arrayType" minOccurs="0"/>

<xs:complexType name="arrayType">
  <xs:sequence>
    <xs:element ref="Point" maxOccurs="unbounded"/>
  </xs:sequence>
</xs:complexType>
```

```
</xs:sequence>
<xs:attribute name="array_length" type="xs:positiveInteger" use="required"/>
</xs:complexType>

<xs:element name="Point">
  <xs:complexType>
    <xs:attribute name="x" type="xs:nonNegativeInteger" use="required"/>
    <xs:attribute name="y" type="xs:nonNegativeInteger" use="required"/>
    <xs:attribute name="z" type="xs:double" use="optional"/>
    <xs:attribute name="i" type="xs:nonNegativeInteger" use="optional"/>
  </xs:complexType>
</xs:element>
```

The first line of this schema simply defines an XML element called `<Breast_border>`, which is of element type `arrayType`, and this element may occur zero or one time (`maxOccurs="1"` by default) in the entire XML document. That is, only one copy of the breast border is stored in the database for each mammogram.

The next section of the schema above defines an *XML element type*, which is a feature of XML Schema that allows users to re-use a particular element type in the definitions of different elements. This element type is called `arrayType` and it contains a sequence of `<Point>` elements. The maximum number of points allowed in the array is unbounded, but an XML attribute, named `array_length`, is required to indicate the actual number of points in the array. Note that the XML Schema definition only specifies that the value of `array_length` must be a `positiveInteger`, but it does not specify that the value of `array_length` must match the number of points in the array. In fact, there is no way to define the latter constraint in XML Schema due to its limitation. Those constraints that cannot be defined in XML Schema should be implemented as validation code in the applications.

The last section of the schema defines an XML element called `<Point>`, which has four XML attributes, two required and two optional. A `<Point>` element must contain the `x` and `y` attributes, which are both `nonNegativeInteger`. This element may also contain the `z` attribute of `double` type for data value, and the `i` attribute of `nonNegativeInteger` type for indexing.

3.6 Conclusions

A mammogram analysis system has been presented in this chapter. The system provides a means to acquire, normalize, store and re-use the mammogram attributes that are produced progressively, from low-level to high-level, by several segmentation and analysis applications. By identifying the useful mammogram attributes and working out the attribute dependence, the sequential order of attribute acquisition can be determined. We have shown that adequacy assessment and lesion detection are independent of each other,

although they both depend on the input of image data and anatomical features.

Although different applications may operate with data formats that are incompatible with each other, data exchange among these applications is still possible by standardizing the format of mammogram attributes, such as image orientation, co-ordinate system, spatial resolution, measurement units and object boundary representations.

XML is used to describe the mammogram attributes in the database, with some structure and datatype constraints defined by XML Schema. An open-ended, application-independent element structure is specified in the schema so that it can be extended easily in the future.

This concludes the description of the mammogram attribute system, and its specification of attribute normalization and attribute storage. The next three chapters present three applications of mammogram segmentation and analysis, designed for attribute acquisition. The first one presented is pectoral muscle segmentation.

Chapter 4

Automatic Pectoral Muscle Segmentation

4.1 Introduction

Segmentation is an unavoidable step in mammographic image analysis. We contend that before the digitized mammogram is analyzed by computer, it must be segmented into its representative anatomical regions. In the hierarchical approach of segmentation, three anatomical *landmarks* are first extracted automatically on a mammogram [61]: they are the breast border [18], the nipple [13] and the pectoral muscle (as previously described in Section 3.3). In this chapter, a method is proposed for automatically segmenting the pectoral muscle on mediolateral oblique (MLO) view mammograms. The work presented in this chapter is also published in [62].

When the MLO view is properly imaged, the pectoral muscle should always appear as a high-intensity, triangular region across the upper posterior margin of the image. The craniocaudal (CC) view is not considered here because the pectoral muscle is only seen in about 30%–40% of CC images [34]. Several factors complicate the segmentation of the pectoral muscle. Depending on anatomy and patient positioning during image acquisition, the pectoral muscle could occupy as much as half of the breast region, or as little as a few percent of it. The curvature of the muscle edge is usually convex, but it can also be concave, or a mixture of both. Although the pectoral muscle boundary is perceived to be visually continuous by humans, there are large variations in edge strength and texture. In many cases the upper part of the boundary is a sharp intensity edge while the lower part is more likely to be a texture edge, due to the fact that it is overlapped by fibroglandular tissue. In addition, the muscle edge may be obscured by artifacts on the digitized mammogram, such as sticky tapes. Because of all these factors, automatic segmentation of the pectoral muscle by computer is a demanding task.

Automatic pectoral muscle segmentation is useful in many areas of mammographic analysis. The work of Gupta and Undrill [48] indicates that mammographic parenchyma

and the pectoral region may have similar texture characteristics, causing a high number of false positives when detecting suspicious masses. In other words, the pectoral muscle could interfere with automated detection of cancers. Also the area overlying the pectoral muscle is a common area for cancers to develop and is particularly checked by radiologists to reduce false negatives. It is therefore necessary to segment out the pectoral muscle before lesion detection, as stated in [63]. Similarly, exclusion of the pectoral muscle is required for automatic breast tissue density quantification [14, 64]. The pectoral edge is also used as one of the axes in 3-dimensional reconstructions from multiple mammographic views [20, 21]; and it is one of the pivotal landmarks in mammogram-pair registration and comparison [6]. Furthermore, it is recommended by radiologists that, for a high quality MLO mammogram, the pectoral muscle should be seen to the level of the nipple or below [34, 35, 37]. Some authorities have recommended, in addition, that the pectoral margin should be convex [34, 35]; that the angle made by the margin with the image edge should not be less than 20° [33]; and that the quality of the pectoral margin may be assessed on a numerical scale [37]. Hence extracting the pectoral edge, as both straight line and curve, is particularly important in automatic evaluation of mammographic adequacy [65].

4.1.1 Literature Review

The Hough transform, used by Karssemeijer [14] to detect the pectoral edge, is a popular technique. This method assumes that the pectoral edge is approximately a straight line oriented in a certain direction. To ensure that the correct peak is selected in the Hough space, gradient magnitude and orientation, length of projected line, and corresponding pectoral area were taken into account. Because of this careful selection scheme, the results were claimed to be very robust and reliable [14]. Other recent studies based on the Hough transform include: Ferrari *et al.* [17] who segmented mammograms into skin-air boundary, fibroglandular tissue, and pectoral muscle; Yam *et al.* [21] who refined the Hough transform linear approximation into a curved pectoral boundary using a dynamic programming method; and Georgsson [20] who extracted the pectoral muscle by region growing, but later found the Hough transform to be more reliable.

There are several other approaches to segment the pectoral muscle. Suckling *et al.* [8] segmented mammograms into four major components: background, pectoral muscle, fibroglandular region and adipose region, using multiple, linked self-organizing neural networks. Aylward *et al.* [15] used a gradient magnitude ridge traversal algorithm at a small scale to extract multiple initial points and then resolved the resulting multiple edge definitions via a voting scheme. Their method parallels that of Karssemeijer [14]. Saha *et al.* [64, 66] reported a semi-automatic method that requires input from an operator to locate the pectoral muscle; delineation is then performed automatically. In addition, Chandrasekhar and Attikiouzel proposed two techniques to enhance the pectoral muscle region on mammograms: the extended Russ operator [19], and tunable parametric edge detection [22], although final segmentation on the enhanced images had not been carried out.

In summary, most of the previous work approximates the pectoral edge as a straight line and then refines that straight line into a more accurate curved boundary, if needed. The accuracy of straight line approximation is usually satisfactory. With the Hough transform, the pectoral muscle was segmented in 92.8% of the mammograms tested in [14]; and the pectoral muscle was *accurately* detected in 87.9% of the images within a difference of 1 to 3 mm in [17]. However the reliability of curve refinement had not been specifically tested and reported in the literature, until the publication of a recent study [27] in which Ferrari *et al.* compared their pectoral muscle segmentation method using Gabor wavelets with that using Hough transform on a subset of the Mammographic Image Analysis Society (MIAS) database [67]. They found that the percentage of false-positive and false-negative pixels given by the former method is less than the latter one, indicating that curve delineation is more accurate than straight line approximation. Their study comes close to the subject of this chapter.

In our method, the two-step approach of straight line approximation and curve refinement is taken as the basis, and extended by iteration to confer robustness. But the algorithms that we developed to derive the straight line and the curve are fundamentally different from the existing methods. Several techniques are also introduced to overcome failure due to poor breast positioning and artifacts, and to validate the straight line approximation before it is refined to a curve. The final segmentation results have been thoroughly assessed by two mammographic experts.

4.2 Overall Algorithm

From the observations made in section 4.1 regarding the appearance of the pectoral region, we frame the following four hypotheses that form the foundation for the algorithm:

1. On mediolateral oblique mammograms, the pectoral muscle is a roughly *triangular region* occupying a corner of the mammogram;
2. The pectoral muscle is defined as a region of *higher* intensity than the surrounding tissue;
3. The pectoral margin is characterized by a fairly sharp change in intensity, i.e., it is an *intensity edge*;
4. The intensity edge can be located by fitting step functions to a sequence of the intensity profiles which lie perpendicular to the pectoral boundary. The step-transition point represents the position of the edge.

Because the pectoral margin may be and is usually curved, it is first estimated as a straight line, which is later refined to a curve. The algorithm therefore consists of two main steps: (a) straight line estimation; and (b) iterative cliff detection. The flowchart of this algorithm is shown in Fig. 4.1.

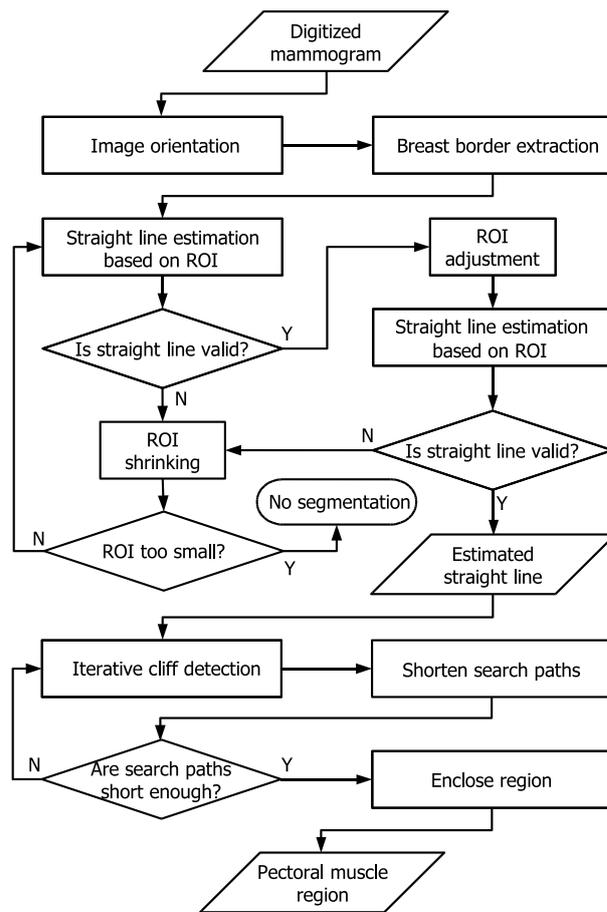


FIGURE 4.1: Flowchart for automatic pectoral muscle segmentation on MLO mammograms. The image orientation and breast border extraction are pre-processing steps that are not part of this algorithm.

There are two pre-processing steps: image orientation and breast border extraction. The image is first oriented in portrait mode to face the same direction for consistency, as shown in Fig. 4.2. In the MIAS database, the image dimensions and left/right labels are already given, so only image rotation and mirroring are required. The breast border is then extracted automatically by background modelling and subtraction [18]. The accuracy of the extracted border does not significantly affect the performance of pectoral muscle segmentation, since the breast border is only used for validation and the endpoints of the border are chosen cautiously (as explained in section 4.4.2).

The first part of the segmentation algorithm generates a straight line approximating the pectoral edge. The initial straight line estimation is carried out within a region of interest (ROI). The straight line is then tested for validity. If valid, the ROI is adjusted accordingly, and a second straight line estimation is performed in the new ROI. If the second straight line is also valid, it is used as the input to iterative cliff detection.

If the straight line is found to be not valid at any stage, the ROI is shrunk to a smaller size and the estimation cycle repeated. When the ROI is smaller than a certain size, the algorithm terminates with no segmentation of the pectoral muscle.

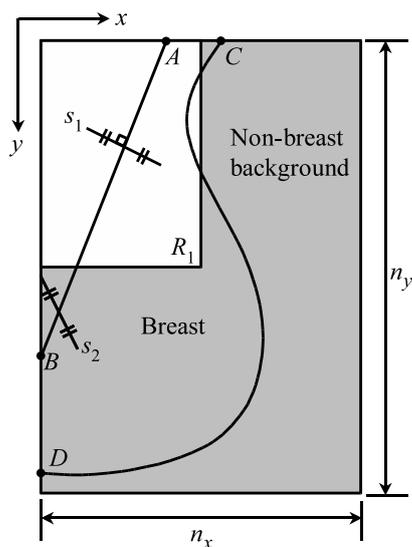


FIGURE 4.2: The mammogram is oriented so that the pectoral muscle is located at the top left corner. The co-ordinate axes are directed as shown with the origin also at the top left corner. The width and height of the whole image are denoted by n_x and n_y , respectively. R_1 is the initial region of interest, equivalent to one quarter of the image. The straight line \overline{AB} is an approximation to the pectoral edge. The end-points of the breast border are C and D . s_1 represents a search path perpendicular to \overline{AB} whereas s_2 represents a rotated search path.

The second part of the algorithm is iterative cliff detection in which the straight line is refined to a curve that more accurately depicts the pectoral margin. This is an iterative process whereby the detected curve is fed back to the cliff detection module for increased accuracy. Finally, the “triangle” with the detected pectoral margin as hypotenuse, is closed to give a segmentation mask for the pectoral muscle.

4.3 Orientation and Notation

In our segmentation algorithm, all digitized mammograms are oriented so that the nipple faces the right, i.e. all the right breast images are mirrored vertically. Therefore all input images are always upright with the pectoral muscle at the top left corner. The intensity of the image is denoted by $I(x, y)$ in the range $[0, I_{\max}]$. The origin of the co-ordinate system is at the top left corner of the image, where x is defined to be the horizontal axis and y to be the vertical one (see Fig. 4.2). The number of pixels of the image in the x - and y -directions are denoted by n_x and n_y , respectively. Parameters used in the algorithm are given in millimetres; they may be converted to pixels if divided by the image resolution in mm/pixel.

4.4 Straight Line Estimation and Validation

Straight line estimation is used to approximate the pectoral muscle with a straight line. This algorithm is based on iterative threshold selection and straight line fitting with a

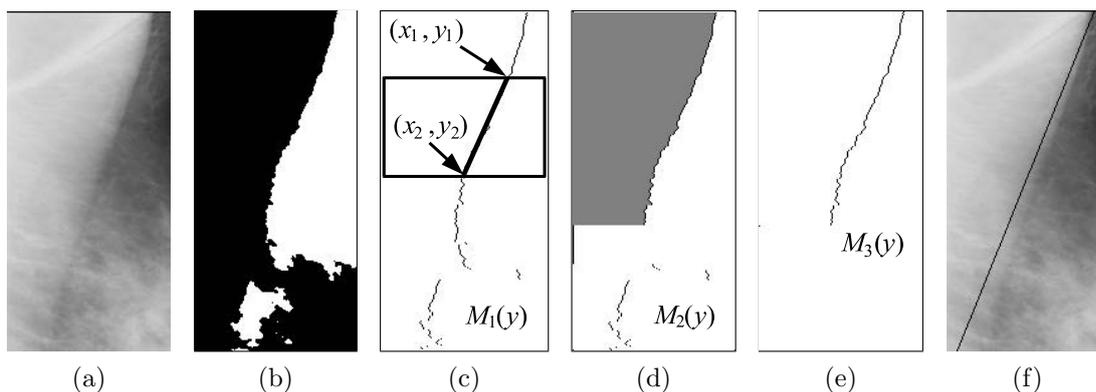


FIGURE 4.3: Illustration of straight line estimation. (a) Initial ROI of MIAS image mdb227lm. (b) Median filtered binary image produced by iterative threshold selection. (c) $M_1(y)$, obtained by tracing the border of black region. Its gradient is computed in the sliding window. (d) $M_2(y)$, result of removing positive gradient segments, with the largest area under the curve shaded. (e) $M_3(y)$, selected for straight line fitting. (f) Straight line approximation to the pectoral edge.

gradient test. The result is then validated by a simple criterion, independently of the straight line fit.

4.4.1 Straight Line Estimation

Defining the Region of Interest (ROI)

Since the pectoral muscle is located at the top left corner of the image, the *top left quarter* of the image is taken to be the initial region of interest (ROI), as shown in Fig. 4.2. It is assumed that the pectoral edge appears in this ROI (partially, if not fully) and that it intersects the top and left image edges. The first straight line estimation of pectoral edge is performed in this ROI, which is represented by R_1 :

$$R_1 = \{(x, y) : 0 \leq x < n_x/2 \quad \text{and} \quad 0 \leq y < n_y/2\} \quad (4.1)$$

Fig. 4.3(a) shows the initial ROI of mammogram mdb227lm taken from the MIAS database.

Iterative Threshold Selection

After setting the initial ROI, the pectoral muscle (pectoral region) should be separated from other tissues (non-pectoral region) as it has been done *manually* in Fig. 4.4(a). However, determining a global threshold automatically is not straightforward. In many MLO mammograms, the lower portion of the pectoral muscle is spatially superimposed on some glandular tissue known as the “tail of Spence” [34]. Also, the image intensity of the glandular tissue can be very near or identical to that of the pectoral muscle, causing intensity overlap of the pectoral and non-pectoral regions in the histogram (see Fig. 4.4(b)).

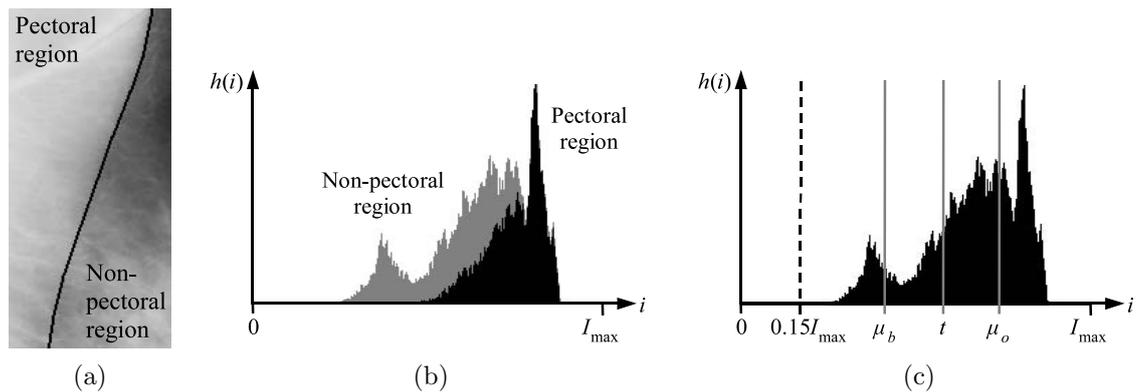


FIGURE 4.4: (a) Pectoral and non-pectoral regions in the initial ROI, R_1 . The spatial overlap of pectoral muscle and glandular tissue can be seen in the lower portion. The region boundary was traced out by hand. (b) Intensity histogram, $h(i)$, of R_1 showing the intensity overlap of the pectoral (black) and non-pectoral (grey) regions corresponding to (a). (c) Iterative threshold selection, where grey-levels below 15% of I_{\max} are excluded (left of dotted line); μ_b and μ_o are the mean values of the background and object grey-levels, respectively; and t is the final threshold.

Due to both *spatial and intensity overlaps* of the two regions, it is not always possible to find a single threshold that completely separates the pectoral muscle from other tissues. However, *iterative threshold selection* can be used to optimize the conversion of the grey-scale image to a binary image in the sense that the image average luminance is preserved; this is explained and mathematically proven in [68] (see appendix C). This method usually works well even if the histogram is not bi-modal [69, p. 129]. For the images affected by intensity overlap, iterative threshold selection is less likely to over- or under-estimate the threshold when compared with other techniques that place the threshold at the valley between two peaks in the histogram.

The algorithm given below has been slightly modified from that given in [69, pp. 129–130] and assumes that the image has two regions of dominant grey-levels:

1. All grey-levels below 15% of I_{\max} are removed from the histogram, $h(i)$, of the region R_1 . It is assumed that the non-breast background and the majority of the breast-edge tissue have been excluded to ensure that the segmentation result is more reliable.
2. A threshold t is determined as the mean of all remaining pixel values in R_1 :

$$t = \frac{\sum_{i \geq 0.15I_{\max}} i \cdot h(i)}{\sum_{i \geq 0.15I_{\max}} h(i)} \quad (4.2)$$

3. The region R_1 is segmented into background and object by thresholding at t .
4. The mean values of the background and object grey-levels, denoted by μ_b and μ_o

respectively, are calculated by the following equations:

$$\mu_b = \frac{\sum_{0.15I_{\max} \leq i < t} i \cdot h(i)}{\sum_{0.15I_{\max} \leq i < t} h(i)} \quad ; \quad \mu_o = \frac{\sum_{i \geq t} i \cdot h(i)}{\sum_{i \geq t} h(i)} \quad (4.3)$$

5. t is then updated as the mid-point of μ_b and μ_o :

$$t = \frac{\mu_o + \mu_b}{2} \quad (4.4)$$

6. If the new t remains unchanged, it is the final threshold; otherwise steps 3 to 6 are repeated.

Pixel Selection

After thresholding, the edge of the pectoral muscle has to be traced out on the binary image (Fig. 4.3(b)) by a *pixel selection* operation. First, impulse noise on the binary image is removed by applying a 5×5 median filter. Then each horizontal line of the binary image is scanned from left to right, and the *first* background pixel on each scan line is selected. The positions of all the selected pixels define the function $M_1(y)$, that roughly represents the pectoral edge.

Gradient Test

If the selected pixels $M_1(y)$ represent the actual pectoral edge accurately, straight line fitting can be applied to it directly. However, in some cases, the curve $M_1(y)$ deviates toward the right and forms a concave segment, whenever the glandular tissue overlaps the pectoral edge. The deviation from the actual edge may lead to an inaccurate straight line estimation.

A *gradient test* was therefore designed to eliminate the concave segments on the function $M_1(y)$. A sliding window of height 20 mm and width equal to the ROI is used in the test. As the window slides from top to bottom, a straight line is fitted to the portion of $M_1(y)$ that lies within the window, and the gradient of the fitted line is computed (see Fig. 4.3(c)). The gradient function, $g(y)$, is given by:

$$g(y) = \frac{x_2 - x_1}{y_2 - y_1} \quad \text{for} \quad \frac{y_2 - y_1}{2} < y < \frac{n_y}{2} - \frac{y_2 - y_1}{2} \quad (4.5)$$

where (x_1, y_1) and (x_2, y_2) are the end-points of the fitted line, and $n_y/2$ is the height of R_1 .

Normally, $g(y)$ is *negative* when $M_1(y)$ is a decreasing function which represents the actual pectoral edge. If there is a deviation from the pectoral edge, $g(y)$ becomes *positive*. Hence in order to eliminate the concave deviations, $M_1(y)$ is set to zero whenever $g(y)$ is non-negative. Consequently the remaining pixels form a new function $M_2(y)$, which may

consist of discontinuous segments. Note that $g(y)$ is undefined at both ends of the ROI and $M_1(y)$ would not be set to zero there.

Straight Line Fitting

Although the concave deviations have been removed, some small, discontinuous segments left in $M_2(y)$ may also affect the accuracy of the straight line estimation. Therefore only the *continuous segment with the largest area under the curve* (shown shaded in Fig. 4.3(d)) is used for straight line fitting because it is most likely to be the actual pectoral edge. This segment is represented by a third function $M_3(y)$ in Fig. 4.3(e). Straight line fitting with least squared error is then applied to $M_3(y)$ and results in the *first* straight line approximation to the pectoral edge, as shown in Fig. 4.3(f). This line is shown as \overline{AB} in Fig. 4.2.

4.4.2 Straight Line Validation

Validation Criterion

A simple criterion is used to validate the straight line estimation. Line \overline{AB} must intersect the top and left image edges inside the breast region, but the intersections may not be inside the ROI. To define the breast region, the breast border $(x_{\text{border}}, y_{\text{border}})$ was extracted automatically by polynomial modelling [18]. Since this method is not immune to artifacts from tapes and misplaced labels, the end-points of the breast border (denoted as C and D in Fig. 4.2) are determined by the maxima within the top and left margins of the image:

$$x_C = \max\{x_{\text{border}} : 0 \leq y_{\text{border}} \leq 0.15n_y\} \quad (4.6)$$

$$y_D = \max\{y_{\text{border}} : 0 \leq x_{\text{border}} \leq 0.15n_x\} \quad (4.7)$$

If for any reason the breast border is not available, x_C and y_D can be replaced by n_x and n_y respectively. The validation criterion can then be described by the following expressions:

$$0 < x_A < x_C \quad \text{and} \quad 0 < y_B < y_D \quad (4.8)$$

where $(x_A, 0)$, $(0, y_B)$, $(x_C, 0)$ and $(0, y_D)$ are the co-ordinates of points A , B , C and D respectively. If the line is valid, *ROI adjustment* is invoked; otherwise *ROI shrinking* is performed. Details of these two methods are given in the following sections.

ROI Adjustment

The first ROI, R_1 , is only an initial estimate of the location of the pectoral edge. The ROI has to be *adjusted* so that the entire pectoral muscle is included, resulting in a more accurate straight line approximation. Therefore a new ROI, \hat{R}_1 , is defined so that \overline{AB}

runs diagonally from the top right corner to the left bottom corner in \hat{R}_1 , i.e.,

$$\hat{R}_1 = \{(x, y) : 0 \leq x < x_A \quad \text{and} \quad 0 \leq y < y_B\} \quad (4.9)$$

Then a *second* straight line estimation is performed on \hat{R}_1 , following the same procedure as described in section 4.4.1. The result is used to update \overline{AB} . If the new straight line is also valid, it represents the best approximation to the pectoral edge from this stage of the algorithm; it is then ready to be refined into a curve by the iterative cliff detection, which is described in later sections.

ROI Shrinking

ROI shrinking is used when the straight line estimation is not valid. The result of invalid estimation could be due to internal texture or large artifacts on the pectoral muscle, but in most cases, the main cause is the breakdown of the assumption that the pectoral muscle occupies approximately half of the ROI. This smaller than expected pectoral muscle leads to an underestimated threshold. Shrinking the ROI so that the assumption is upheld is the basis for this step. If R_m is the current ROI, then the new ROI, R_{m+1} , is defined as the *top left quarter* of R_m , i.e.,

$$R_m = \{(x, y) : 0 \leq x < n_x/2^m \quad \text{and} \quad 0 \leq y < n_y/2^m\} \quad (4.10)$$

The same straight line estimation (described in section 4.4.1) is performed on the new ROI in the hope that the result would be valid. The smallest possible ROI in this algorithm is R_4 , which is 1/256 of the original image size. If no valid straight line is found after R_4 is used, it is concluded that the pectoral edge cannot be detected, perhaps because it is absent altogether from the mammogram.

4.5 Iterative Cliff Detection

Cliff detection is designed to refine the straight line approximation into a curve that delineates the pectoral edge more accurately. The resulting curve is further refined by applying cliff detection multiple times; hence it is an iterative process. Finally, a roughly triangular region is enclosed to represent the segmented pectoral muscle.

4.5.1 Cliff Detection

Cliff detection is performed on the whole mammogram and not on any ROI. It consists of the following steps:

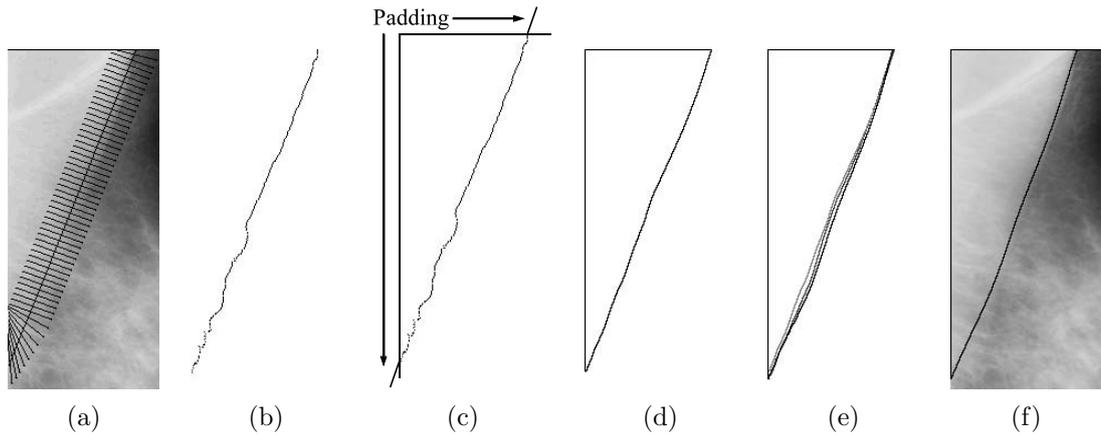


FIGURE 4.5: Illustration of iterative cliff detection. (a) Search paths placed across the straight line \overline{AB} (only every fifth path is shown). (b) All the detected cliff locations. (c) Two pieces of padding, parallel to the straight line, are added at both ends of the detected curve for smoothing. (d) Detected curve after first iteration. (e) Detected curves after 1st, 5th and 9th iterations, respectively from left to right. (f) Final segmentation of the pectoral muscle.

Defining Search Paths

If the straight line estimation is accurate enough, the actual pectoral edge should be in close proximity to \overline{AB} . Therefore, *search paths* are used to find the exact position of the muscle edge in the vicinity of the straight line. The whole image is first smoothed by an average filter of size $m \times m$. On every pixel of \overline{AB} , a search path of length $2d$ is defined so that it is perpendicular to \overline{AB} and the distances on both sides of \overline{AB} are equal.

However, there is an exception at the image edges. Part of a search path could lie outside of the image if it is placed perpendicularly to \overline{AB} . In such cases, the search path is rotated in the appropriate direction so that its outlying end-point is moved back into the range of the image, just touching the image edge (see Fig. 4.5(a)).

An example of perpendicular and rotated search paths is shown as s_1 and s_2 respectively in Fig. 4.2. The rotated search paths may overlap each other, but this is acceptable. Although other orientations for the search paths were tried, the perpendicular orientation gave the best results, and was best suited to detect the intensity edge representing the pectoral margin.

Keeping the search path in full length is important because search path length and intensity surface smoothness are related. The smoother the intensity surface, the longer the search path that should be used to detect the intensity cliff. That is why search path length, d , and smoothing filter size, m , are *reduced together* in each iteration (see also Section 4.5.2). Experiment has shown that shortening the search path at the image edge results in less accurate cliff detection; while rotating the full length search path has a higher chance to include the cliff in the intensity profile.

Extracting Intensity Profiles

The next step is to extract *intensity profiles* along the search paths. The intensity profile, $P(k)$, is a function that represents the cross-section of the smoothed image surface on each search path. If there is a pectoral edge, a significant intensity drop, or a *cliff*, is expected on the intensity profile. Cliff detection is actually designed to locate the intensity cliff of the pectoral margin from a collection of intensity profiles across the straight line \overline{AB} . Fig. 4.6 illustrates the difference between intensity profiles of the original and smoothed images.

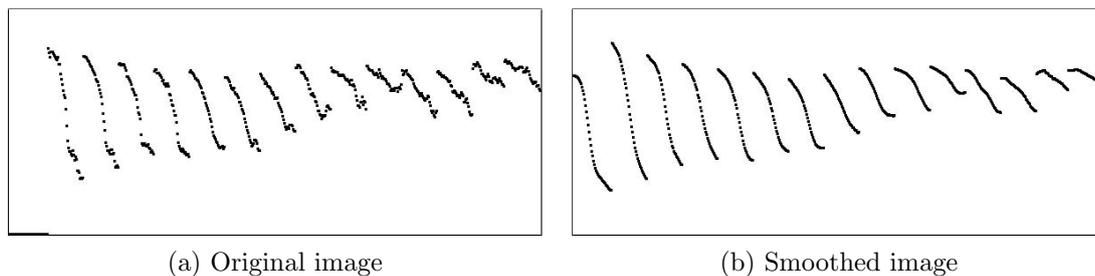


FIGURE 4.6: Comparison of intensity profiles of original and smoothed images. Only every 20th profile is shown here.

Determining Cliff Locations

The determination of the spatial co-ordinates of the intensity cliff is the next step in the algorithm. It was observed that most of the smoothed profiles exhibit a typical sigmoidal shape with various slopes and intensity changes. However, that typical shape did not necessarily appear in all profiles, as may be seen in the two rightmost profiles in Fig. 4.6(b). To confer robustness on the algorithm, it was decided to locate the cliff by fitting the smoothed profile to a *sigmoid function*.

This sigmoid function, $S(k; \alpha, \beta, p_{\max}, p_{\min})$, is a hyperbolic tangent function:

$$S(k; \alpha, \beta, p_{\max}, p_{\min}) = \frac{p_{\max} - p_{\min}}{2} \tanh(\alpha - \beta k) + \frac{p_{\max} + p_{\min}}{2} \quad \text{for } k \in [0, 2d]; \beta > 0 \quad (4.11)$$

where p_{\max} and p_{\min} are the maximum and minimum intensities in $P(k)$, and α and β are fitting parameters solved by least squared error, i.e., $\varepsilon^2 = [S(k; \alpha, p_{\max}, p_{\min}) - P(k)]^2$ is minimized. Fig. 4.7 shows a plot of the sigmoid. The cliff location, k_c , is determined at the point of inflection of $S(k)$, i.e.,

$$k_c = \begin{cases} 0 & \text{if } \frac{\alpha}{\beta} < 0 \\ \frac{\alpha}{\beta} & \text{if } 0 \leq \frac{\alpha}{\beta} \leq 2d \\ 2d & \text{if } \frac{\alpha}{\beta} > 2d \end{cases} \quad (4.12)$$

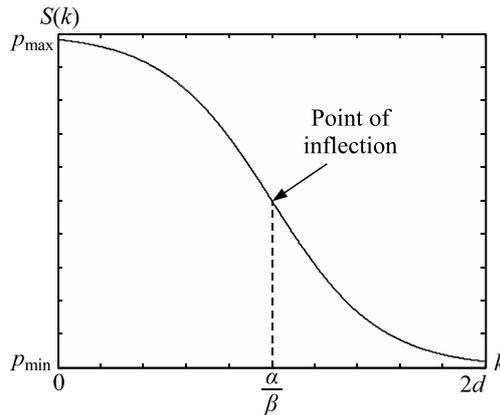


FIGURE 4.7: The sigmoid model, $S(k)$, where α and β are fitting parameters; p_{\max} and p_{\min} are the maximum and minimum values of the smoothed profile, $P(k)$, respectively; $2d$ is the length of the corresponding search path. The point of inflection at $k_c = \alpha/\beta$ is assumed to be the cliff location on the intensity profile.

The cliff locations of all the extracted profiles are marked on a blank image, which has the same size as the original (see Fig. 4.5(b)). This set of cliff locations is denoted by \mathcal{C} .

Smoothing the Detected Curve

The next step is to express the set of cliff locations as a function of y . This is done by eliminating certain cliff locations from \mathcal{C} , and adding intermediate values where necessary. If there is more than one cliff location for a given y value on a horizontal scan line, only that with the largest x value is retained; the others are eliminated. On the other hand, if there is no cliff location for a given y value on a horizontal scan line, two-point linear interpolation is used to generate an intermediate point. The resulting function is smoothed by a 24 mm wide moving-window average filter. Also, two strips of padding, 12 mm each, and parallel to line \overline{AB} , are added at both ends of the function to reduce distortion (see Fig. 4.5(c)). This smoothed function is called the *detected curve* and is the output of cliff detection.

4.5.2 Iterative Refinement

Cliff detection can be applied iteratively to refine the detected curve. Since the image is smoothed before determining cliff locations, the sharpness of the pectoral edge is reduced, and hence the detected curve from the *first* cliff detection may be slightly inaccurate. Two previously defined variables of the algorithm, m and d , are set to 4 and 8 mm respectively in the first cycle of cliff detection. In the next iteration, m is decremented by 1 pixel and d by 2 pixels. Search paths are defined in the same way, but this time on the detected curve rather than the straight line. The same procedure is carried out as described in section 4.5.1 to produce a new curve. The iterative process stops when either m or d is reduced to 1 pixel. During the curve refinement, the image surface is gradually *sharpened*

together with a reduction in the maximum search distance. Consequently the final detected curve delineates the pectoral edge to a higher degree of accuracy. The final segmentation result for mammogram mdb227lm is depicted in Fig. 4.5(f).

4.5.3 Region Enclosing

Since the detected curve is a function of y , its top end always coincides with the top edge of the image. However, in some cases, its bottom end may not be aligned with the left edge of the image. In order to form a closed region, the bottom end is extended by a straight line parallel to \overline{AB} (Fig. 4.2), after the final iteration, if necessary. The extended curve is then smoothed again by the 24 mm average filter. Finally, this enclosed region can be used to segment the pectoral muscle on the mammogram. It is pointed out that this extrapolation has been deliberately chosen to guarantee region closure using a line segment that is consistent with the overall direction of the pectoral margin. An example is illustrated in Fig. 4.8.

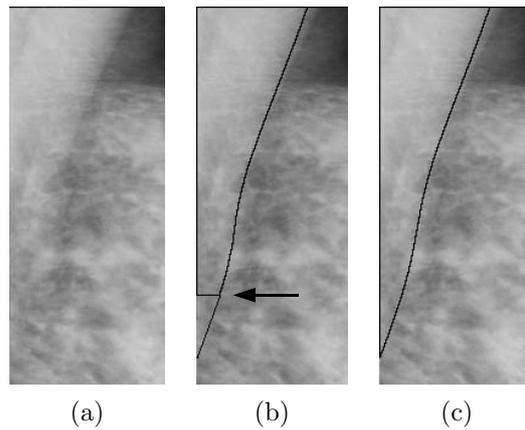


FIGURE 4.8: (a) Pectoral muscle of MIAS image mdb215ll. (b) A straight line parallel to \overline{AB} is extended from the end of detected curve (marked by arrow). (c) Enclosed region for final segmentation.

4.6 Experimental Setup

This new method for pectoral muscle segmentation was tested on 322 digitized mammograms from the MIAS database [67]. The original $50 \mu\text{m}/\text{pixel}$ images were low-pass filtered and reduced in resolution to $400 \mu\text{m}/\text{pixel}$. The original bit-depth of 8 bits [0–255] was retained.

Two expert mammographic radiologists¹ were invited to assess the goodness of the segmentation. In order to better objectify the results and ensure consistency, at the beginning of the assessment, the radiologists were acquainted with the purpose of segmentation:

¹Dr. Mary Rickard and Dr. Natacha Borecky

The purpose of the segmentation is to localize the pectoral margin with sufficient accuracy so as to segment out the “pectoral triangle” on mediolateral oblique mammograms. This will facilitate exclusion of muscular tissue with non-exclusion of parenchymal tissue, mammogram density analysis, parenchymal characterization, mammogram adequacy determination via nipple to pectoral margin perpendicular lines, etc.

The radiologists were then presented with the original mammograms and two types of segmented images: the straight line and curve segmentation of the pectoral muscles generated by straight line estimation and cliff detection respectively. The segmented images were displayed on a computer screen in a random order. The segmentation boundary was coloured in red and superimposed on the original mammogram, without any contrast enhancement. The radiologists were asked to assess all the images of straight line segmentation first. After that the images of curve segmentation were assessed in the same manner. For each image, the radiologists determined whether the segmentation was acceptable, which is a binary decision, and then, during a second pass, they rated the goodness of segmentation using a five-point scale, as explained in Table 4.1. A score of 3 or less indicates an adequate segmentation.

TABLE 4.1: Five-Point Assessment Scale with Score Descriptions

Score	Meaning	Description
1	Exact	The segmented straight line or curve delineates the pectoral margin exactly. Any deviations from the visually perceived margin are imperceptible or insignificant.
2	Optimal	The segmented straight line or curve delineates the pectoral margin exactly for at least half its length and adequately for the other half.
3	Adequate	The segmented straight line or curve delineates the pectoral margin inexactly but with sufficient accuracy for localizing the pectoral margin for the intended purpose.
4	Sub-optimal	The segmented straight line or curve delineates the pectoral margin inadequately for at least half its length.
5	Inadequate	The segmented straight line or curve is either absent or is so inaccurate as to be inadequate for localizing the pectoral margin for the intended purpose.

4.7 Results

4.7.1 Radiologists' Assessments

The numbers of straight line and curve segmentation images accepted by the two radiologists are listed in Table 4.2. It shows that radiologist 1 rated as acceptable 243 (75.5%) images for straight line segmentation and 280 (87.0%) images for curve segmentation; while radiologist 2 rated as acceptable 170 (52.8%) images for straight line segmentation and 216 (67.1%) images for curve segmentation. Although the segmentation images accepted by radiologist 2 are fewer than those accepted by radiologist 1, both radiologists agreed that after applying curve refinement the number of acceptable images increased.

TABLE 4.2: Numbers of Segmentation Accepted by the Two Radiologists

Radiologist 1	Straight Line Segmentation		Curve Segmentation	
	Decision	Number	Percentage	Number
Acceptable	243	75.5%	280	87.0%
Not Acceptable	79	24.5%	42	13.0%
Radiologist 2	Straight Line Segmentation		Curve Segmentation	
	Decision	Number	Percentage	Number
Acceptable	170	52.8%	216	67.1%
Not Acceptable	152	47.2%	106	32.9%

TABLE 4.3: Contingency Tables of Straight Line and Curve Segmentation

Radiologist 1	Curve	
	Straight Line	Acceptable
Acceptable	242	1
Not Acceptable	38	41
Radiologist 2	Curve	
	Straight Line	Acceptable
Acceptable	151	19
Not Acceptable	65	87

The significance of the improvement can be determined statistically by the McNemar test [70, pp. 75–80]. This test is useful for detecting changes in initial and final responses (acceptable images) due to experimental intervention (curve segmentation). The test was performed using the contingency tables shown in Table 4.3. The McNemar chi-squared test statistics and approximated P-values are 33.23 and 8.18×10^{-9} for radiologist 1, and 24.11 and 9.11×10^{-7} for radiologist 2. Since the P-values are very small in both cases, we conclude that *curve segmentation has improved the accuracy of pectoral muscle segmentation*. The contingency tables also show that after curve segmentation, 38 images were upgraded from not acceptable to acceptable for radiologist 1, and 65 images for

TABLE 4.4: Segmentation Accuracy Ratings by the Two Radiologists

Radiologist 1 Rating (Score)	Straight Line Segmentation		Curve Segmentation	
	Number	Percentage	Number	Percentage
Exact (1)	84	26.1%	169	52.5%
Optimal (2)	106	32.9%	85	26.4%
Adequate (3)	53	16.5%	32	9.9%
Sub-optimal (4)	33	10.2%	10	3.1%
Inadequate (5)	46	14.3%	26	8.1%
Radiologist 2 Rating (Score)	Straight Line Segmentation		Curve Segmentation	
	Number	Percentage	Number	Percentage
Exact (1)	34	10.6%	87	27.0%
Optimal (2)	50	15.5%	54	16.8%
Adequate (3)	113	35.1%	117	36.3%
Sub-optimal (4)	59	18.3%	30	9.3%
Inadequate (5)	66	20.5%	34	10.6%

radiologist 2. On the other hand, only 1 image was downgraded from acceptable to not acceptable for radiologist 1, and 19 images for radiologist 2.

The radiologists’ five-point ratings for the 322 straight line segmentation images and 322 curve segmentation images are shown in Table 4.4. The results show that radiologist 1 rated the straight line segmentation *adequate or better* on 243 (75.5%) images and the curve segmentation adequate or better on 286 (88.8%) images. The same figures for radiologist 2 are 197 (61.2%) and 258 (80.1%) respectively. Again, the ratings of radiologist 2 are lower than that of radiologist 1, but in both cases the ratings were improved (by 13.3% and 18.9%) after applying curve segmentation. By taking the average score of each image, 67.4% of the straight line segmentation and 83.9% of the curve segmentation were rated adequate or better (≤ 3).

Furthermore, after curve refinement, the ratings given by radiologist 1 increased on 164 (50.9%) images, remained unchanged on 141 (43.8%) images and decreased on 17 (5.3%) images. The same figures for radiologist 2 are 147 (45.7%), 152 (47.2%) and 23 (7.1%) respectively. Therefore by refining the straight lines into curves, almost half of the image ratings increased by one level or more. These results once again demonstrated that pectoral muscle segmentation was more accurate after applying iterative cliff detection.

As mentioned previously in Section 4.6, the radiologists assessed the images in two passes using two different scales. Therefore, “acceptable” in Table II is *not necessarily* equivalent to “adequate or better” in Table IV. The difference between the two sets of results illustrates the uncertainty encountered in human observer studies, especially when two different rating scales are given.

The 5×5 contingency tables of the segmentation ratings are shown in Table 4.5 to give the readers an insight into the *correlation* or *agreement* between both radiologists.

TABLE 4.5: Contingency Tables of Radiologists' Ratings

Straight Line	Radiologist 2				
	Exact	Optimal	Adequate	Sub-optimal	Inadequate
Radiologist 1					
Exact	30	27	25	1	1
Optimal	2	21	53	29	1
Adequate	1	1	25	19	7
Sub-optimal	0	0	9	7	17
Inadequate	1	1	1	3	40

Curve	Radiologist 2				
	Exact	Optimal	Adequate	Sub-optimal	Inadequate
Radiologist 1					
Exact	85	34	50	0	0
Optimal	2	18	53	12	0
Adequate	0	2	14	13	3
Sub-optimal	0	0	0	5	5
Inadequate	0	0	0	0	26

Since our 5-point scale is an ordinal measurement, the Spearman rank-order correlation coefficient [70, pp. 235–244], r_s , was used to evaluate the association between the radiologists' ratings. For the straight line segmentation ratings, r_s is 0.740 and the P-value is 4.54×10^{-57} . For the curve segmentation ratings, r_s is 0.748 and the P-value is 6.27×10^{-59} . Since the P-values are very small in both cases, we conclude that the positive correlation between the ratings given by both radiologists is statistically significant at the 1% level.

4.7.2 Algorithm Performance

The average processing time for a single mammogram was around 5.9 seconds, in which 0.4 second was spent on straight line estimation and 5.5 seconds on iterative cliff detection. The computational time for curve refinement is much longer than that for straight line estimation because the former is an iterative process, in which the whole image has to be smoothed multiple times with filters of different sizes.

In straight line estimation, ROI shrinking was performed on 40 images, i.e., the straight line did not pass the first validation in those cases. In the end, a valid straight line could be found on 34 of those images: so shrinking the ROI is a necessary step when the pectoral muscle is smaller than normal. The straight line estimation was not successful on the other 6 images, hence there was no curve segmentation for these images as well. In addition, the curve extension used to enclose the pectoral region was performed on 195 images. The average length of the extended line was approximately 5.5 mm, which is relatively small on the mammogram.

Some examples of the final segmentation are shown in Fig. 4.9.

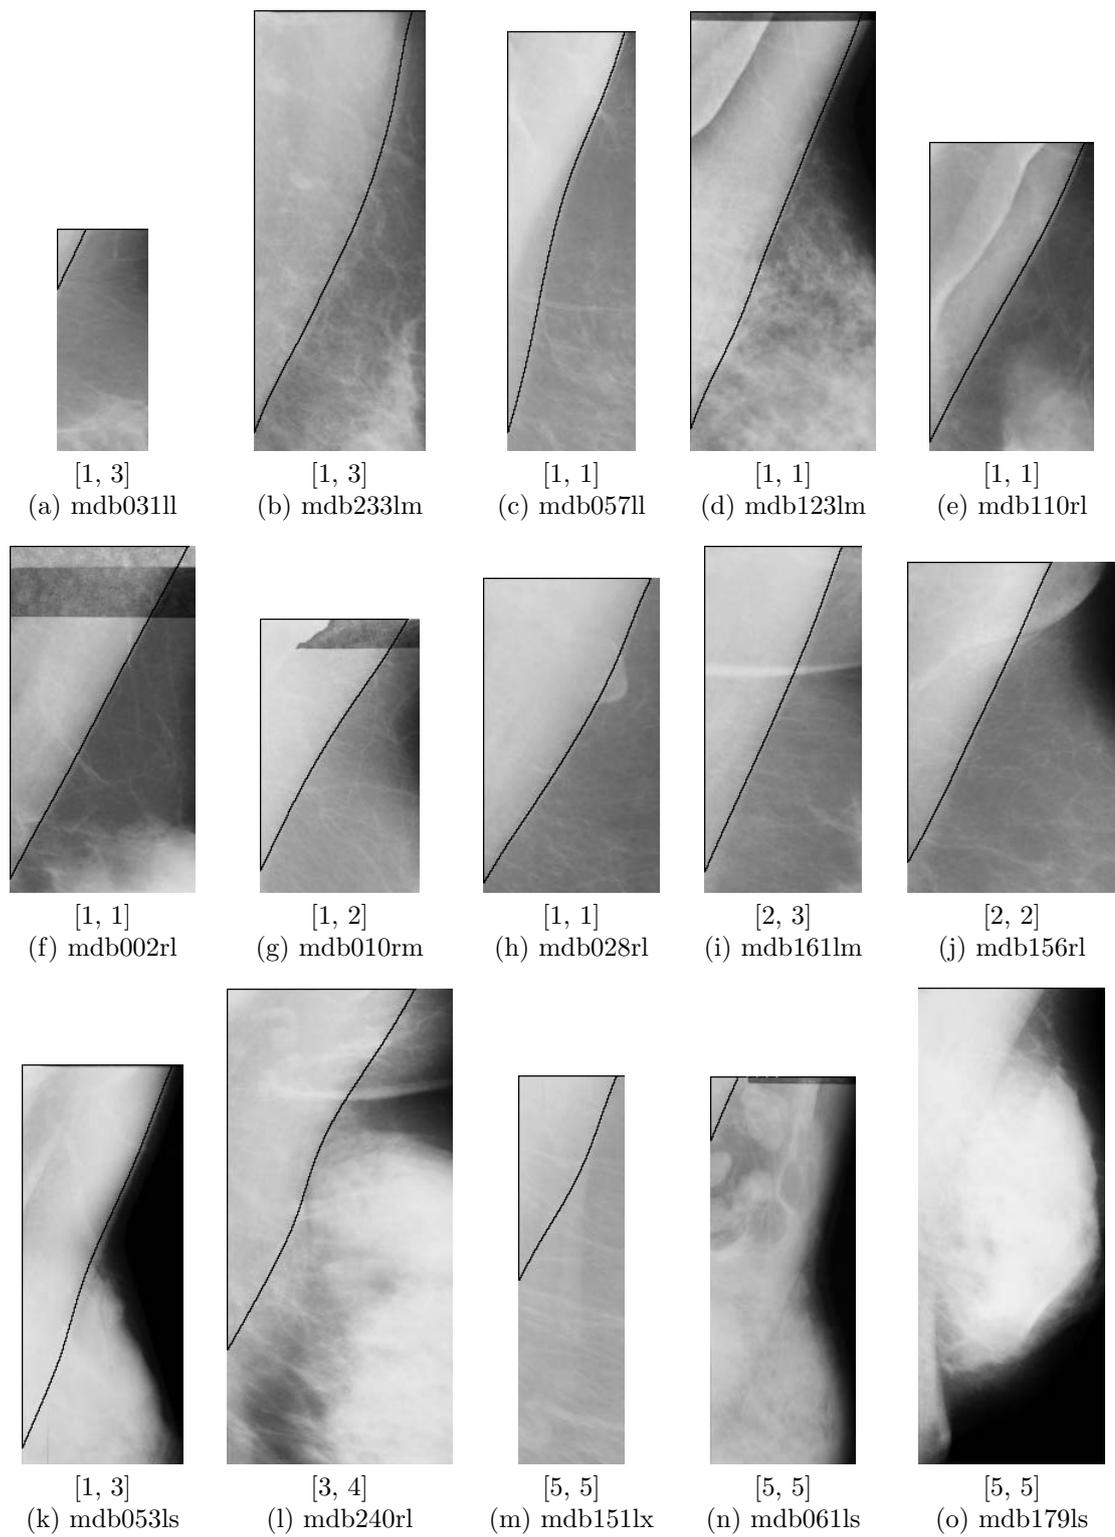


FIGURE 4.9: Segmentation results on MLO mammograms from the MIAS database. The scores of each image rated by two radiologists are shown in $[r_1, r_2]$. All images are shown in the same scale. There was no segmentation for (o).

4.8 Discussion

The algorithm is adaptive to variations in pectoral muscle size, density and curvature. In an adequate mammogram, the pectoral muscle should be visible down to the level of the nipple [34]. In some extreme cases, however, the muscle may appear on the mammogram as very small or very large. Two examples of tiny and large pectoral muscles are shown in Fig. 4.9(a) and (b). They were both segmented accurately because multiple ROI's, from large to small, were used in the straight line estimation. Also, there is a variation on the image intensities of different pectoral muscles. We found that the method of iterative threshold selection can adaptively compute an appropriate threshold level for each image, provided that the pectoral region occupies roughly 50% of the ROI. Furthermore the segmentation method was accurate for tracing pectoral boundaries of different curvatures, whether convex, concave or both.

Sometimes there may be more than one layer of tissue in the pectoral region and more than one edge may be depicted. In such cases the straight line approximation is more likely to be placed at the outer edge rather than at the inner edge. The reasons for this are that (i) the straight line estimation method uses histogram thresholding instead of edge detection to find the edge; and (ii) the straight line is fitted only to that segment of the selected pixels which encloses the largest area. Therefore the outer edge is more likely to be chosen. Fig. 4.9(c)–(e) show pectoral muscles with internal edges of different edge strengths and they are all delineated at the outer edges.

The results also demonstrate that the segmentation method is robust against artifacts such as sticky tapes. Fig. 4.9(f) and (g) show that although the top parts of the pectoral muscles are obscured by tapes of different thickness, both were segmented successfully. These images show one of the advantages of cliff detection: the refinement of the straight line is restricted in certain directions, since the search paths are set perpendicularly, and only intensity *decreases* in the profiles are searched for. Also, the detected curve in each iteration is smoothed so that no sharp corners on the curve are allowed. More examples of different kinds of small edges crossing the pectoral boundary are shown in Fig. 4.9(h)–(j).

In Fig. 4.9(k), a very dense breast is shown and the lower half of the pectoral muscle is almost completely obscured by other tissues. In this case, the extrapolative power of the segmentation method is demonstrated. The pectoral boundary was extended following the direction of the estimated straight line. And the refinement process did not alter the extended line significantly since there were no other strong intensity changes within the reach of the search paths.

One disadvantage of the proposed, intensity-based method is its weakness in detecting texture edges. In some cases the pectoral boundary is not a clear intensity edge but a fuzzy texture edge. Cliff detection is not a suitable method to detect edges between different textures. However it can still be used to approximate the pectoral boundary based on the intensity changes alone. An example is illustrated in Fig. 4.9(l).

Another shortcoming relates to the detection of vertical pectoral edges, which are

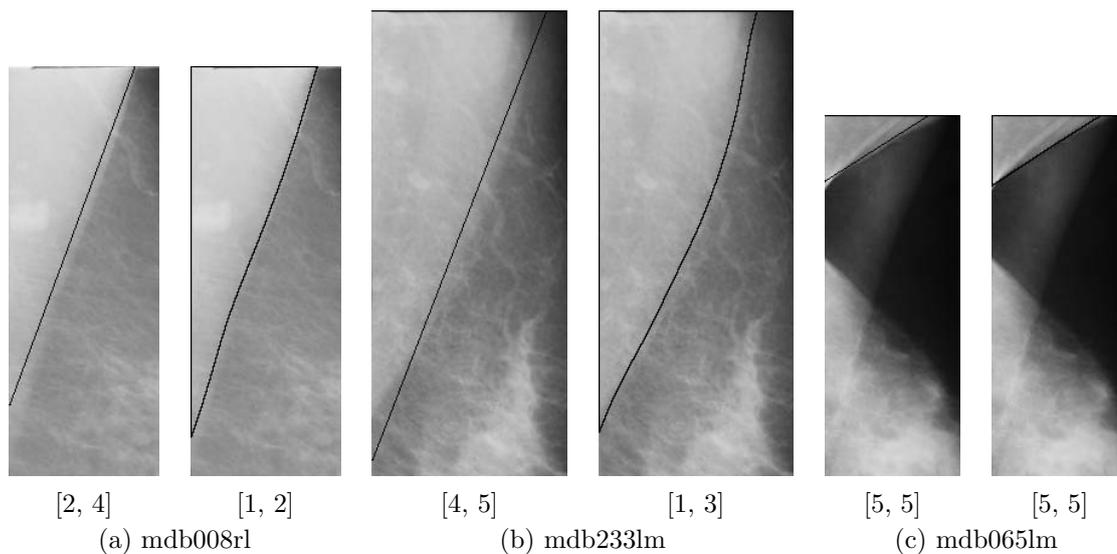


FIGURE 4.10: Comparison of straight line and curve segmentation results on mammograms from the MIAS database. The scores of each image rated by two radiologists are shown in $[r_1, r_2]$. (a) The straight line (left) is roughly parallel to the pectoral margin but not exactly on it. This is corrected in curve segmentation (right). (b) The pectoral muscle is poorly segmented by the straight line (left) but accurately segmented by the curve (right). (c) The axillary fold is mistakenly segmented by the straight line (left) and hence the curve (right) is also inadequate. Moreover, hypothesis (2) stated in page 31 is not valid for this image.

excluded by assumption. Because a valid straight line approximation is assumed to intersect the top and left image edges, a vertical line cannot satisfy the validation criterion. As shown in Fig. 4.9(m) the pectoral edge is almost vertical but it was not segmented correctly.

A very poor segmentation is shown in Fig. 4.9(n), which is a particularly difficult image due to the coarse texture in the muscle region. Finally, no segmentation was produced for Fig. 4.9(o). The pectoral edge of this image is barely visible. The primary reason for no segmentation is that the pectoral muscle does not exist on the image, or that a large part of the pectoral margin is obscured by dense tissue.

The radiologists also made the following comments for poorly segmented images. For some images, the straight lines were placed *parallel* to the pectoral margins but not exactly on them. Poor straight line estimation also resulted when the pectoral margins were very curved. These kinds of inaccuracy, especially in the straight line estimation, were normally corrected by the iterative refinement in curve segmentation (see Fig. 4.10(a), (b)). However, in a small number of cases, some straight lines completely missed the pectoral margins due to the influence of the *axillary fold*, which appears as a high-intensity triangular region, similar to the pectoral muscle (see Fig. 4.10(c)). The radiologists also noted that the axillary fold should not normally be seen on the MLO view and that its presence is often the result of poor positioning.

4.9 Suggestions for Further Work

The proposed algorithm could be improved in future by incorporating a texture edge detector together with cliff detection. This is particularly targeted at extracting the texture edge at the lower portion of the pectoral muscle. Such a texture edge detector should be sensitive to texture direction since the texture direction of the pectoral margin is usually perpendicular to that of the overlying glandular tissue. Also texture detection should be carried out at higher image resolutions than the one used in cliff detection.

In straight line estimation, iterative threshold selection was proposed for separating the pectoral muscle from other tissue, which is a typical binary-tree splitting problem [71]. We have tried other techniques that use histogram shape analysis and place the threshold at the valley between two peaks. But due to the extensive intensity overlap of pectoral and non-pectoral regions, the performance of these techniques are not satisfactory. In the future work, other thresholding methods, such as binary Gaussian mixture model and least mean square error optimization, should be investigated and compared with the performance of iterative threshold selection.

The accuracy of pectoral muscle segmentation may be further improved if we combine iterative cliff detection with other straight line approximation methods. Masek [32] critically compared the performance of his straight line approximation method with our straight line and curve segmentation algorithms presented in [72]. He devised an automatic voting scheme to compare the segmentation results, image by image, obtained from 322 mammograms from the MIAS database. The results were interesting: (1) our straight line segmentation algorithm performed better than his straight line segmentation algorithm on 39.6% of the images; (2) our curve segmentation algorithm performed better than his straight line segmentation algorithm on 59.8% of the images. These findings have shown that our curve segmentation was more accurate than his straight line approximation, which was also more accurate than our straight line approximation. It is therefore worthwhile to combine his straight line approximation method with our curve refinement algorithm, so that segmentation accuracy can be further improved.

Most of the other difficulties in segmentation result from a poor positioning of the patient during image acquisition or to poor image contrast. Thus, one way to overcome these problems would be to ensure acquisition of good quality images in the first place. We have suggested elsewhere [65] that automatic quality assurance during image acquisition is a feasible solution to this problem. When the breast is correctly positioned, a sufficient amount of pectoral muscle should be seen to the level of the nipple at an angle larger than 20° to the vertical edge [33]. When the image is properly exposed at optimal contrast, the pectoral muscle should be differentiated from other dense tissue more easily.

4.10 Conclusions

A new method has been developed for automatic segmentation of the pectoral muscle on the MLO views of mammograms. The method consists of two major components: straight line estimation and iterative cliff detection. By approximating the pectoral boundary with a straight line and refining it iteratively to a curve, it is possible to accurately delineate an enclosed pectoral region. The method was adaptive to large variations in appearance of the pectoral muscle and margin. The method remained effective when parts of the pectoral edge were obscured by superimposed glandular tissue or artifacts. The method was tested out on the 322 digitized mammograms of the MIAS database and two mammographic radiologists assessed the segmentation results. Their findings show that segmentation accuracy was improved after refining the straight line into a curve using iterative cliff detection and that 83.9% of the curve segmentations were adequate or better.

This concludes the first experimental chapter on mammogram segmentation and analysis. In the next chapter, some algorithms for mammogram adequacy assessment are presented and they are also tested on all the mammograms from the MIAS database.

Chapter 5

Adequacy Assessment of Positioning and Exposure

The adequacy of mammographic image quality may be assessed with the aid of computers. The term adequacy generally covers many aspects of quality in mammography, but in this chapter we focus on the automatic evaluation of breast positioning and film exposure. These two quality aspects were measured on digitized mammograms without the use of any phantom. The algorithms for this automatic assessment are described in this chapter. Anatomical features of the mammogram are used in these algorithms to generate the adequacy measures, as shown in the attribute dependence graph in Fig. 3.2.

5.1 The Need for Automatic Quality Assessment

There are several reasons for wanting to assess mammographic adequacy using computers. If a mammogram is assessed for adequacy immediately after it has been acquired, a repeat examination can be conducted, if necessary, before the subject leaves the screening centre. This would save the patient's anxiety and worry when told to return for a repeat examination at a later date, without knowing the cause.

Adequacy assessment is also an important component in the computerized analysis of mammograms. Low-quality mammograms could greatly influence the performance of automatic segmentation and the accuracy of cancer detection. Inadequate images should be identified as early as possible so that they are not subjected to needless further analysis. In addition, automatic quality assessment may be integrated into an ongoing quality assurance program at the facility, or into a radiographer training scheme.

5.2 Quality Attributes of Mammograms

Achieving high image quality is essential to the radiologist who wants to interpret the mammogram with high sensitivity and specificity [34]. Although the assessment of mammographic image quality usually requires some subjective considerations, there is a con-

sensus in the literature [33–35, 37, 38] on the *quality attributes* that an ideal mammogram should incorporate. A high quality mammogram should exhibit correct positioning, optimal compression, good contrast, adequate exposure, low noise, high sharpness, and absence of artefacts [38]. Chief among these attributes is positioning, which crucially determines the amount of tissue inclusion and correlates with the overall quality of the mammogram [38].

Images of standardized test object, or phantoms, are routinely acquired and analyzed as part of the quality assurance program on a film-screen or full-field digital mammography system. Phantom images are widely used in the research of automated quality control of mammographic imaging. Blot *et al.* [73] used grey-level co-occurrence matrices to analysis the textural structures of digitized phantom films, for quantitative objective assessment of different imaging systems. Thomas *et al.* [74] reported an automated method for measuring contrast-to-noise ratio on phantom images with superior precision. Bijkerk *et al.* [75] used the computer to interpret phantom images and to determine the corresponding contrast-detail curves. They found that computer detection was more efficient than human perception, since the latter is influenced by the image presentation and viewing conditions. An automated scoring algorithm was developed by Muller *et al.* [76] for contrast-detail phantoms used in digital mammography. Images acquired at different dose levels or with different detector calibration status could be differentiated by their algorithm.

There are also some papers related to noise detection and removal on mammograms. An iterative noise detection scheme was described in [77] that locates and discriminates shot noise from real calcifications on film-screen mammograms. Masek *et al.* [78] targeted the removal of high-intensity background noise and rectangular labels, and as a result, improved the performance of their skin-air interface segmentation.

The approach taken in this chapter is different from the previous work above, in that it is more relevant to the clinical rather than technical aspects of mammography. Several algorithms have been developed to automatically evaluate two unexplored areas: *adequacy of mammographic positioning* and *exposure*. Measurable quality criteria related to these two quality aspects are described in the following sections.

5.3 Optimal Positioning in the MLO View

In mammography, usually two standard views are taken for each breast; they are the mediolateral oblique (MLO) and craniocaudal (CC) views. However, only MLO view mammograms are provided in the MIAS database [67]. For this reason, the adequacy assessment in this research is restricted to MLO view mammograms.

The MLO view is considered by the radiologists the most important view, in which all of the breast tissue is most likely included on the film [34, 35]. To ensure that the entire glandular body is imaged with best possible compression, *correct positioning* of the breast

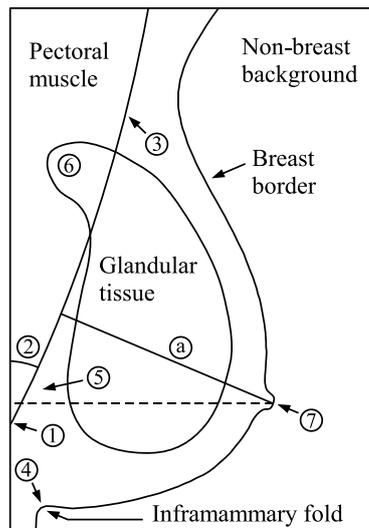


FIGURE 5.1: Optimal positioning in the MLO view. (1) The pectoral muscle is visible at least to the level of the nipple (dotted line). (2) The angle between the pectoral margin and the posterior image edge is at least 20° . (3) The anterior pectoral margin is convex. (4) The inframammary fold is included on the image. (5) The posterior glandular tissue is separated from the image edge by retroglandular fat. (6) Upper glandular tissue superimposed over the upper pectoral muscle is not excluded in the view. (7) The nipple is in profile. (a) The *posterior nipple line* is drawn along the nipple axis from the nipple-skin junction to the pectoral muscle or to the back of the image, whichever comes first [34]. This line is not needed in MLO view assessment.

is an absolute necessity [33]. Exclusion of glandular tissues from the field of view increases the risk of lesions being missed because they do not appear on the mammograms. In fact, a recent study on clinical image quality and the risk of interval cancer has shown that “invasive breast cancer detection by mammography may be improved through attention to correct positioning” [38].

Optimal positioning can be achieved in the MLO view when a number of *quality criteria* are fulfilled. These criteria are described in most detail in [34] and also outlined in other literature [33, 35–38]. They are demonstrated in Fig. 5.1 and summarized below:

1. The pectoral muscle should be visible at least to the level of the nipple. A lesser amount of the pectoral muscle seen on the image indicates that some posterior glandular tissue may have been excluded. It is reported in [36] that this criterion could not be fulfilled by all patients. In exceptional cases, the technologist should determine whether any posterior and lateral tissue has been omitted.
2. The pectoral muscle should course diagonally along the posterior image edge at an angle of at least 20° .¹
3. The pectoral margin usually exhibits anterior convexity when the pectoral muscle was fully mobilized and maintained medially during compression.²

¹This criterion has been stated only in [33].

²This criterion has been stated only in [34].

4. The inframammary fold should be included on the image. This ensures that the lower posterior tissue has not been excluded.
5. Retroglandular fat should be present between the posterior image edge and the glandular tissue. This ensures that posterior glandular tissue has not been excluded at the image edge.
6. Upper glandular tissue superimposed over the upper pectoral muscle, known as the *tail of Spence*, should be projected properly into the field of view. This extension of tissue toward the axilla is only prominent in some women [34]. When the tail is seen being cutoff at the image edge, it indicates that the pectoral muscle might not be fully mobilized medially.
7. The nipple should be seen in profile. This has been emphasized for proper positioning but never at the expense of not imaging posterior breast tissue [34]. The nipple should preferably be seen in profile in at least one of the standard views.

Quality criteria (1)–(3) are for assessing the positioning of the pectoral muscle; while (4)–(6) are for preventing the exclusion of glandular tissues at the posterior image edge. It is worthwhile to note that glandular tissues can also be excluded at the anterior and inferior image edges, usually when the breast is larger than the film. These exclusions are even more obvious than the posterior one since part of the skin-air interface is cutoff from the image (see Fig. 5.3(a)).

All of the criteria mentioned above, together with the anterior and inferior breast tissue exclusions, were evaluated by the computer in this assessment, except for (5) and (6) because obtaining an accurate outline of glandular tissues is a non-trivial problem.

The posterior nipple line on the MLO view is shown in Fig. 5.1. It is compared with the one measured on the CC view when assessing the amount of tissue included on the CC view. The posterior nipple line on the MLO view should not exceed that on the CC view by more than 1 cm [34, 37, 38]; otherwise optimal tissue inclusion has not been achieved on the CC view. This comparative assessment is only for evaluating the adequacy of the CC view and is not required for the MLO view.

5.4 Adequate Exposure

For an adequately exposed mammogram, the glandular tissues should be well penetrated so that the observer can see detailed structures through dense parenchyma, and the fatty areas and skin line can be visualized and evaluated in high-intensity light [34]. Proper exposure is a prerequisite of good contrast; uniformly high contrast would be required in every density range [33].

The requirements above are quite subjective. In terms of technical measures, the optical density of the mammogram is examined. The *useful optical density range*, also

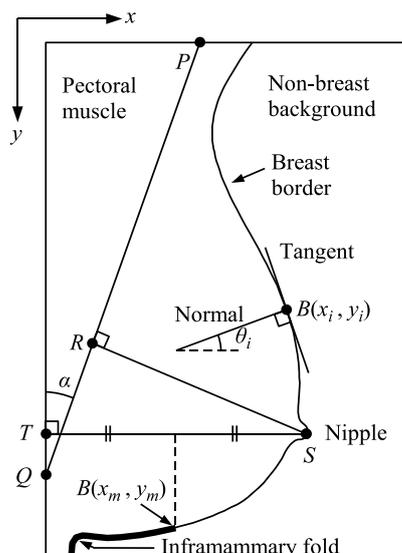


FIGURE 5.2: The mammogram is oriented so that the image is upright and the nipple is faced to the right. The co-ordinate axes are directed as shown with the origin placed at the top left corner. The breast border is a discrete function denoted by $B(x, y)$. The normal to the tangent of $B(x_i, y_i)$ is directed inwards to the breast. The angle between the normal and the positive x -axis is shown as θ_i . The straight line \overline{PQ} is an approximation to the pectoral margin and its angle to the y -axis is shown as α . The posterior nipple line \overline{RS} is perpendicular to \overline{PQ} . The level of the nipple is represented by a horizontal line \overline{ST} .

known as *latitude*, of every film is limited between 0.6 and 2.2 (or 2.8 to 3.0 in bright light) [33]. Therefore breast tissue in the density range below 0.6 or above 2.2 is not visualized optimally to human observers. To ensure that most of the image details relevant to the diagnosis are captured within the useful density range, the *mean optical density* should lie approximately in the middle of the useful density range, i.e., between 1.4 and 1.8 [33]. If the mean value is below 1.4, the mammogram is probably *underexposed*. On the other hand, if the mean value is above 1.8, the mammogram is possibly *overexposed*. Underexposure results in the risk of masking abnormalities, such as low-opacity lesions and microcalcifications, in dense parenchyma; while overexposure limits the visualization of skin line and fatty areas on the mammogram [34].

5.5 Adequacy Assessment Algorithms

5.5.1 Image Orientation and Notation

The image orientation and notation used in the adequacy assessment algorithms are illustrated in Fig. 5.2. For consistency, the image orientation and co-ordinate system are the same as those used in pectoral muscle segmentation (see Chapter 4). Every image is oriented so that it is upright and the nipple faces the right. The origin of the co-ordinate system is at the top left corner of the image, where x is defined to be the horizontal axis and y to be the vertical one. The breast border, that consists of n pixels, is represented

by a discrete function $B(x_i, y_i)$ for $i = 0, \dots, n - 1$, running from top to bottom. All the normals to the border are directed inwards to the breast. The angle of each normal, θ_i , in the range $[180^\circ, -180^\circ)$, is positive when measured anti-clockwise from positive x -axis and negative when measured clockwise.

Parameters used in the algorithm are given in millimetres; they may be converted to pixels if divided by the image resolution in mm/pixel.

5.5.2 Exclusion of Breast Tissue

Although the exclusion of breast tissue at the posterior edge of the image cannot be detected unless an accurate outline of dense glandular tissue is available, the breast tissue exclusion at the anterior and inferior edges of the image can be identified by examining the breast border. The breast border was extracted from the image by background modelling and subtraction [18]. If part of the breast is excluded at the image edge, a segment of the extracted breast border would not be a smooth curve but in the form of horizontal or vertical straight lines. These abnormal segments can be identified by examining the normal directions of all pixels on the breast border, and by locating consecutive normals that point to the same horizontal or vertical direction. Also these segments should be found very close to the image edge (see Fig. 5.3).

The algorithm for detecting breast tissue exclusion is given as follows:

1. For every point on the breast border, the tangent at (x_i, y_i) is estimated by fitting a straight line (using least squared error) to $B(x, y)$ within a neighbourhood of 10 mm centred on (x_i, y_i) .
2. The normal to $B(x_i, y_i)$ is perpendicular to the tangent, pointing inwards to the breast region.
3. The angle measured from the positive x -axis to the normal is denoted by θ_i in the range $[180^\circ, -180^\circ)$.
4. Exclusion of breast tissue is detected when there exists a segment on $B(x, y)$ that:
 - (a) all the normals in that segment are either horizontal or vertical, i.e., θ_i is either 0° or -90° ; and
 - (b) the segment is at least 10 mm long; and
 - (c) the segment is within a 5 mm margin of the anterior and inferior image edges.

5.5.3 Nipple in Profile

The algorithm used to locate the nipple and to determine whether it is in profile is described in [13]. A brief explanation of the basis of that method is given here. Whether the nipple is in profile can be determined by examining the change of normal direction

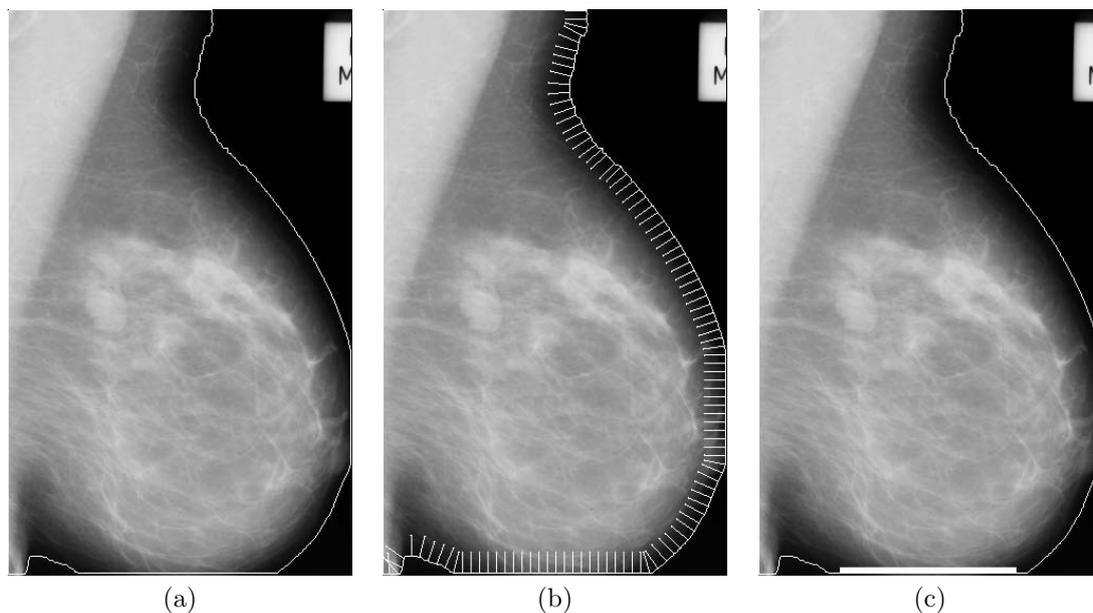


FIGURE 5.3: (a) Breast border extracted from MIA image mdb019ll. (b) Normals to the breast border (only every 8th normal is shown). (c) Breast tissue exclusions detected on the breast border are shown as thickened white lines.

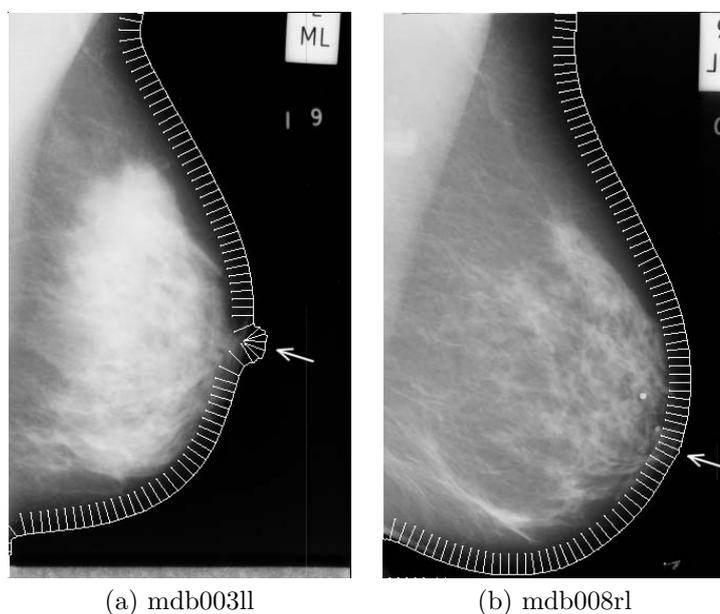


FIGURE 5.4: Breast borders and nipple locations (white arrows) are illustrated on both images. Only every 8th normal is shown. (a) Nipple is in profile, detected by the distinct changes in normal directions. (b) Nipple is not in profile; changes of normal directions are very small.

on the breast border. If the nipple is in profile, it is depicted as a small semi-circle on the breast border; otherwise it is not seen on the breast border. With this distinctive feature, the nipple is inferred to be in profile by the algorithm when the change of normal direction near the nipple is higher than usual, i.e., when the maximum change is over an absolute threshold, defined in [13]. Two mammograms, one with the nipple in profile and one without, are shown in Fig. 5.4. Note the difference on the border curvature near the nipples of both images.

5.5.4 Locating the Inframammary Fold

The method used to locate the inframammary fold is also based on examining the normal direction on the breast border. Generally, when the inframammary fold exists, a distinctive concave curve (i.e., the centre of curvature lies outside the breast region) is formed on the border of the lower posterior breast region; otherwise the breast border is only a slightly convex curve. Therefore the inframammary fold can be located by the following algorithm:

1. Because the inframammary fold is at the lower posterior breast region, only a portion of the breast border is included for examination (shown as thickened line in Fig. 5.2). This portion is defined on $B(x_i, y_i)$ for $i = m, \dots, n - 1$ where x_m is half of the maximum horizontal distance between the breast border and posterior image edge. The maximum horizontal distance does not necessarily coincide with the nipple location.
2. The change of the normal directions is represented by θ'_i which is given by:

$$\theta'_i = \begin{cases} \theta_i - \theta_{i-1} & \text{if } i > 0 \\ \theta_{i+1} - \theta_i & \text{if } i = 0 \end{cases} \quad (5.1)$$

3. All concave segments on $B(x, y)$ which have θ' greater than zero are labelled.
4. The inframammary fold is detected when there exists one or more labelled segments which are at least 10 mm long.
5. The location of the inframammary fold is determined at the midpoint of the labelled segment that is maximum in length.

Some examples of inframammary folds that were detected by the algorithm are shown in Fig. 5.5.

5.5.5 Positioning of the Pectoral Muscle

To determine whether the pectoral muscle is visible to the level of nipple or below, two straight lines are drawn on the image. The first straight line approximating the pectoral margin is given by the automatic segmentation algorithm described in Chapter 4. It is

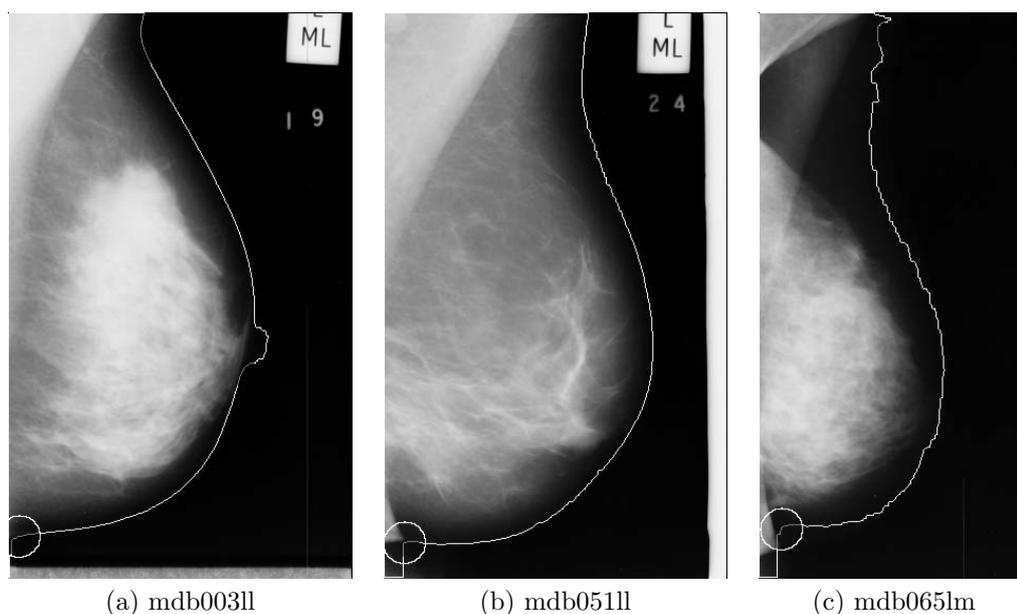


FIGURE 5.5: The inframammary fold detected on each of these three images is located at the centre of the white circle. The concave segments that are at least 10 mm long were labelled in grey colour. (a) Three concave segments were detected but only the bottom one was included for examination. (b) The inframammary fold was detected correctly although the turning point is very shape. (c) The inframammary fold was detected correctly although the border is not smooth.

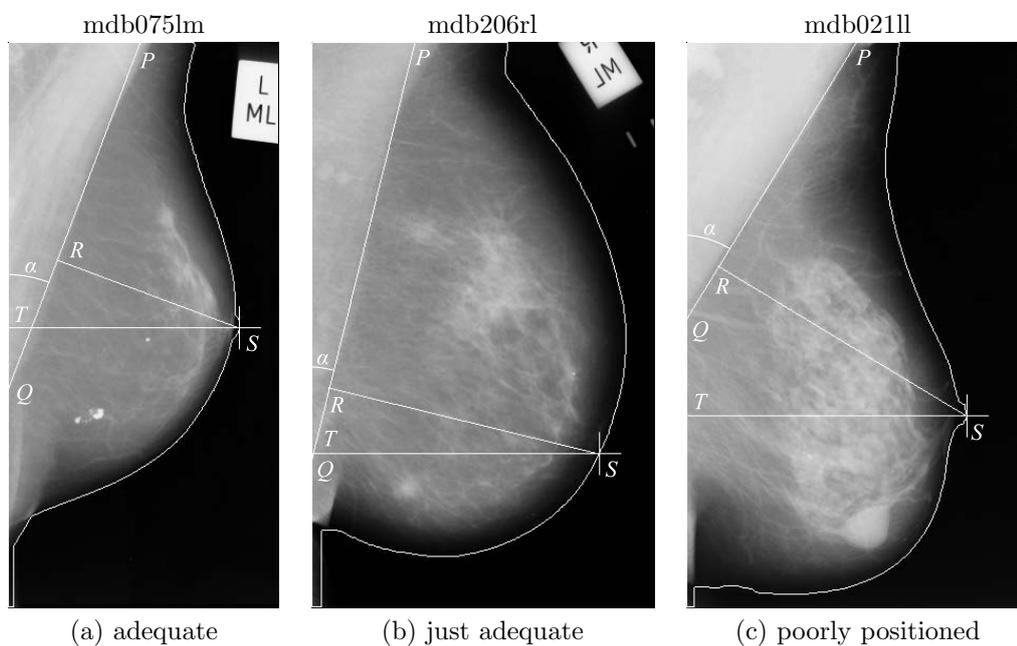


FIGURE 5.6: Straight line approximation of pectoral muscle, level of the nipple and posterior nipple line are shown as \overline{PQ} , \overline{ST} and \overline{RS} respectively on the images. (a) The pectoral muscle is visible to the level of nipple. $\alpha = 13.8^\circ$. (b) The pectoral muscle extends below the level of nipple by 22.8 mm. $\alpha = 20.7^\circ$. (c) The pectoral muscle is not adequately depicted; its inferior extent is 37.2 mm above the level of nipple. $\alpha = 31.3^\circ$.

denoted as \overline{PQ} in Fig. 5.2. The second straight line representing the level of nipple is drawn horizontally from the nipple to the posterior image edge, denoted as \overline{ST} . If the pectoral muscle is adequately imaged, it will be visible to the level of the nipple or below, and point Q should be at the same level or below point T .

The angle between the pectoral margin and the posterior image edge is denoted by α in Fig. 5.2. This angle is measured anti-clockwise from \overline{PQ} to the vertical edge. For adequate positioning, the angle should be 20° or greater [33].

The posterior nipple line for the MLO view is defined in [34] as “the distance from the nipple-skin junction to the pectoral muscle or to the back of the image, whichever comes first.” This line should be drawn along the nipple axis, but can be approximated by a straight line, shown as \overline{RS} in Fig 5.2, which is drawn perpendicularly from the nipple to \overline{PQ} . If \overline{PQ} and \overline{RS} intersect outside of the image domain, the posterior nipple line is only measured from the nipple to the posterior image edge.

Examples of adequately and poorly positioned pectoral muscles are shown in Fig. 5.6.

5.5.6 Curvature of the Pectoral Margin

When the pectoral muscle is fully mobilized and maintained medially during breast compression, the anterior margin of the pectoral muscle usually exhibits a convex curve [34] (i.e., the centre of curvature lies inside the pectoral region). The curvature of the pectoral margin is measured by examining the normal directions of all the pixels on the pectoral curve. The method to delineate the pectoral margin as a curve is described in Chapter 4.

The algorithm designed to distinguish between convexity and concavity of the pectoral margin includes the following steps (see Fig. 5.7):

1. For every point on the pectoral curve, $C(y)$, the tangent at y_i is estimated by fitting a straight line (using least squared error) to $C(y)$ within a neighbourhood of 40 mm centred on y_i .
2. The normal to each point on $C(y)$, directed inwards to the pectoral region, is computed by finding the line perpendicular to the corresponding tangent.
3. The angle of each normal is denoted by ϕ_i in the range $[180^\circ, -180^\circ]$. It is positive when measured anti-clockwise from the positive x -axis and negative when measured clockwise.
4. The change of the normal direction at y_i is represented by ϕ'_i which is given by:

$$\phi'_i = \begin{cases} \phi_i - \phi_{i-1} & \text{if } i > 0 \\ \phi_{i+1} - \phi_i & \text{if } i = 0 \end{cases} \quad (5.2)$$

5. If ϕ'_i is negative or zero, the pectoral curve is determined convex at y_i ; otherwise it is concave.

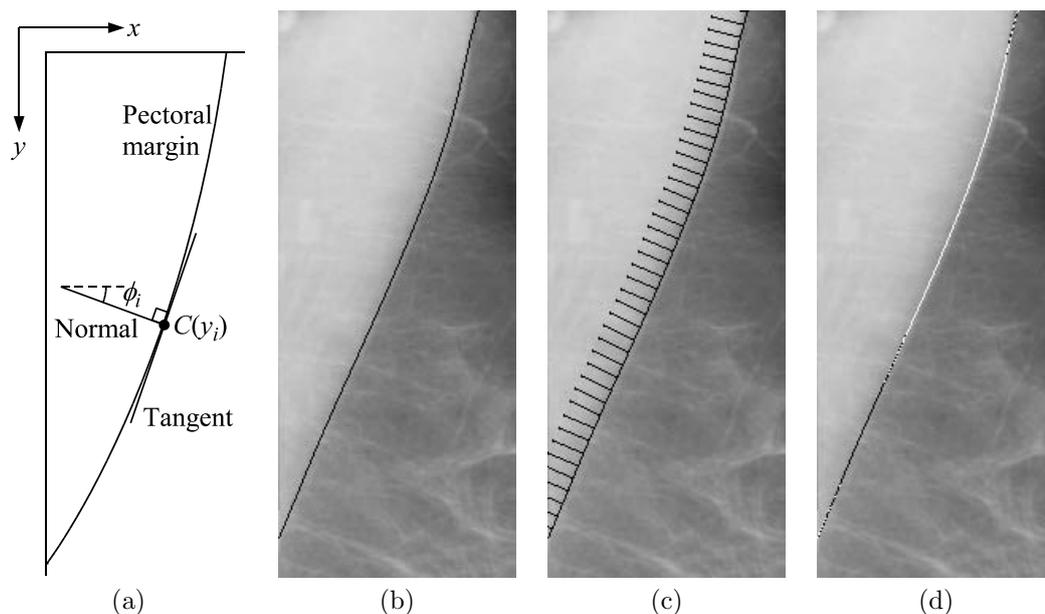


FIGURE 5.7: (a) The mammogram orientation, co-ordinate system and angle measurement is the same as Fig. 5.2. The pectoral margin is represented by a discrete function $C(y)$. The normal to the tangent of $C(y_i)$ is directed inwards to the pectoral region. The angle between the normal and the positive x -axis is shown as ϕ_i , which is negative when below the horizontal line. (b) Pectoral margin of MIAS image mdb026rl. (c) Normals to the pectoral margin (only every 8th normal is shown). (d) Convexity (white) and concavity (black) of the pectoral margin; 66% of the curve was determined convex.

6. The segmented pectoral margin usually exhibits a mixture of convexity and concavity. Therefore its overall curvature can be represented by the percentage of convex sections in the curve $C(y)$.

5.5.7 Optical Density Measures

Optical density is the amount of blackening produced on the film. It is determined by measuring the ratio of incident to transmitted light intensity, when the film is trans-illuminated, and expressing the value as a logarithm:

$$\text{Optical density} = \log(\text{intensity of incident light}/\text{intensity of transmitted light}) \quad (5.3)$$

When a film-screen mammogram is scanned by a digitizer, optical densities of the film are converted to grey-scale pixel values of a certain bit-depth. The relationship between the optical density and pixel value is represented by the response curve of the digitizer. If the response curve is linear, i.e., pixel values are related linearly to optical densities, its inverse could be expressed in the form:

$$\text{Optical density} = d_{\max} - \frac{d_{\max} p}{p_{\max}} \quad (5.4)$$

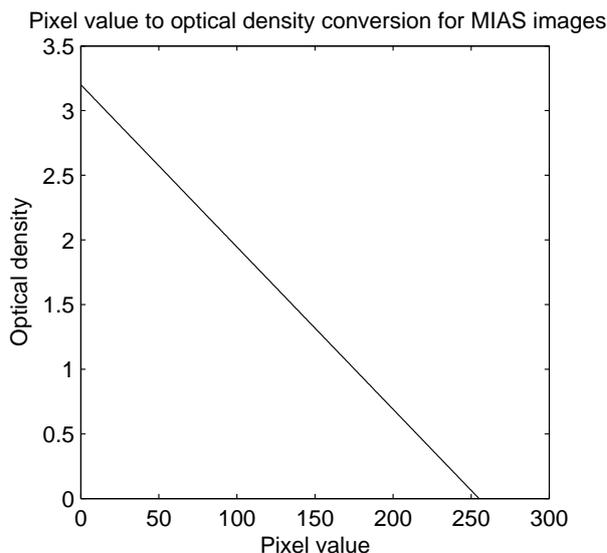


FIGURE 5.8: Inverse of the response curve of the digitizer used to produce images in MIAS database. It is used to convert pixel values back to optical density.

where d_{\max} is the maximum of the optical density range, p is the pixel value and p_{\max} is the maximum pixel value.

For example, the images in MIAS database were scanned with a Joyce-Loebl microdensitometer SCANDIG-3, which has a linear response in the optical density range $[0, 3.2]$ and 8-bit grey levels [67]. By substituting 3.2 to d_{\max} and 255 to p_{\max} into Equation 5.4, we have:

$$\text{Optical density} = 3.2 - \frac{3.2p}{255} \quad (5.5)$$

This equation is plotted in Fig. 5.8. It shows that the higher (brighter) the pixel value, the lower the optical density.

Using the inverse of the response curve, the optical density of the original film can be obtained from the pixel values of the image, and hence the density range of the mammogram can be determined. As a result, two optical density measures were used in the algorithm to assess the adequacy of film exposure: (i) mean optical density of breast tissue; and (ii) percentage of breast tissue lying within the useful optical density range. The algorithm includes the following steps:

1. Only the non-pectoral breast tissue is included in the assessment, so all the pixels lying outside the breast border and those lying inside the pectoral curve are excluded.
2. The optical density of each pixel within the breast region is calculated using (5.4).
3. The mean of the optical densities is then calculated. From the guideline given in [33], the mean optical density should lie between 1.4 and 1.8, i.e., in the middle of the useful optical density range. If the mean is outside of this range, the film is probably underexposed or overexposed.

- The percentage of breast tissue lying within the useful optical density range is computed. The useful optical density range defined in [33] is between 0.6 and 2.2. Within that range, density differences can be visualized with good contrast so that different types of breast tissue are distinguishable to human observers. Therefore the higher the percentage, the more adequate the film exposure.

Three mammograms and their optical density histograms are shown in Fig. 5.9. For all 3 cases, around 70% of the breast tissue is in the useful density range. However, mammogram (a) was well exposed with mean optical density at 1.54, (b) was overexposed with mean optical density at 2.08, and (c) was underexposed with mean optical density at 1.16.

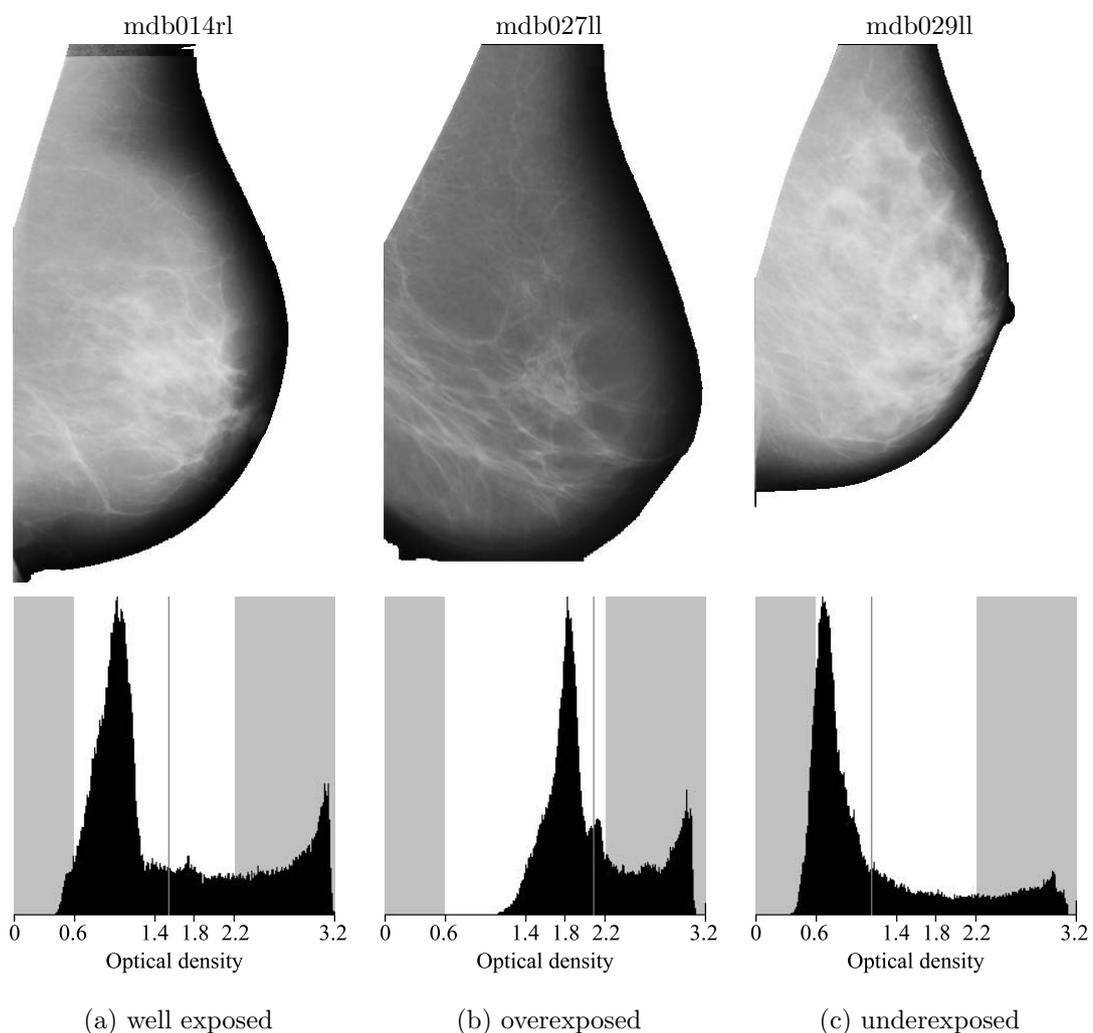


FIGURE 5.9: Optical density histograms for only the non-pectoral breast region of the mammograms. The mean optical density, marked by a grey vertical line on each histogram, should lie in the range $[1.4, 1.8]$. The useful optical density range $[0.6, 2.2]$ is represented by the white band between the shaded regions. (a) Adequate exposure; mean optical density 1.54; 73% of breast tissue in $[0.6, 2.2]$. (b) Overexposure; mean optical density 2.08; 69% of breast tissue in $[0.6, 2.2]$. (c) Underexposure; mean optical density 1.16; 74% of breast tissue in $[0.6, 2.2]$.

5.6 Experimental Results

The adequacy assessment algorithms was tested on all 322 images in the MIAS database. The setting was the same as in the previous chapter. The original $50\ \mu\text{m}/\text{pixel}$ images were reduced in resolution to $400\ \mu\text{m}/\text{pixel}$ by averaging 8×8 pixels to one. The original 8-bit grey levels [0–255] was retained.

The author reviewed all the images³ to identify some of the anatomic features, including the breast tissue exclusion, nipple in profile and inframammary fold. The author's observation was then compared with the results generated by the corresponding algorithms to determine their accuracy or success rates.

Breast tissue exclusion was detected on 66 images at the anterior and/or inferior edges, but with 7 false positives. Also the algorithm missed 5 cases with exclusion undetected. Therefore the success rate was 310/322 (96.3%).

The nipples were detected in profile by the algorithm on 93 of the images. After human assessment, 86 of the 93 cases were confirmed to be correct, so there were 7 false positives. The algorithm however left 25 true cases undetected. The success rate was therefore 290/322 (90.1%).

The inframammary folds located by the computer were indicated on the images by a white circle of radius 20 pixels (as shown in Fig. 5.5). The algorithm found the presence of inframammary folds on 121 of the 322 images. However in 12 of those 121 cases, the actual inframammary folds either were too far from the detected locations (outside of the circle) or did not exist at all. Also the algorithm failed to detect the inframammary folds on the breast borders of 25 images. In total, 285/322 (88.5%) cases were considered acceptable by the author.

Six images without the pectoral muscle segmentation results (see Chapter 4) were excluded from the measurements related to the pectoral margin. The distances from inferior extent of the pectoral margin to the level of the nipple were measured by the presented algorithm on the remaining 316 images. Only 6 of the them fulfilled the criterion that the pectoral muscle has to be visible to at least the level of the nipple. Therefore majority of the images were classified inadequate in this aspect. Even with the more relaxed criterion that the pectoral muscle should be depicted within 10 mm of the level of nipple or below, only 21 images passed it. On average, the inferior extent of the pectoral margin was 59 mm *above* the level of the nipple.

The angles of the pectoral margin were also measured on the 316 images as well. The angles were greater than 20° in 216 cases, which met the positioning criterion. On average, the angle was 22.6° to the vertical edge.

The lengths of the posterior nipple lines, measured perpendicularly from the nipple to the pectoral muscle or to the posterior image edge, were computed on all 316 images. The

³Due to difficulties in securing expert radiologists to assess the results, the author reviewed all the images independently of the results of the algorithms. Attempts were made to avoid bias by viewing the original mammograms without looking at the experimental results. Although this was not ideal, it was preferable to leaving the results unevaluated.

TABLE 5.1: Number of MIAS Images Fulfilling the Quality Criteria

Quality Criteria	Number/Total	Percentage
(a) No breast tissue excluded	256/322	79.5%
(b) Nipple in profile	93/322	28.9%
(c) Inframammary fold included	121/322	37.6%
(d) Pectoral margin to level of nipple	6/316	1.9%
(e) Angle of pectoral margin $\geq 20^\circ$	216/316	68.4%
(f) Convexity of pectoral margin $\geq 50\%$	207/316	65.5%
(g) Mean optical density in [1.4, 1.8]	193/322	59.9%
(h) Breast tissue in useful density range $\geq 60\%$	281/322	87.3%
All criteria satisfied	1/322	0.3%
Positioning criteria satisfied (including (a)–(f) except (d))	15/322	4.7%
Exposure criteria satisfied (including (g) and (h))	172/322	53.4%

average length for all the images was 104 mm. Unfortunately, there are no CC views in the MIAS database, so the posterior nipple lines of the MLO views could not be compared with those of the CC views.

The algorithm that measures the pectoral margin curvature was applied to the segmented pectoral curves of the 316 images. The percentage of convexity of each pectoral margin was calculated accordingly. On average, this convexity percentage was 57.8% for all the images. If we require at least half of the pectoral margin to be convex, i.e., 50% convexity, then 207/316 (65.5%) images would fulfill this requirement.

The adequacy of the film exposure of all 322 mammograms was determined by the mean optical density and the percentage of tissue visualized in the useful optical density range. After evaluating all the images by the computer, the mean optical densities were found in the range [1.4, 1.8] for 193 (59.9%) cases. On average, 69.8% of the breast tissue was visualized in the useful optical density range [0.6, 2.2] across all images.

The individual criteria used in the adequacy assessment of the MIAS images are summarized in Table 5.1. They are labelled from (a) to (h). Some of the criteria were satisfied by majority of the images; while others were fulfilled by relatively few images. Only *one* image in the database was found satisfying all the criteria, due to the fact that the pectoral muscle was rarely seen to the level of the nipple ((d) in Table 5.1). This outstanding mammogram is shown in Fig. 5.10. The cut-off points used in two criteria, (f) and (h), are not specified in the clinical references, but were rather subjectively set for this assessment. When only the positioning criteria were included in the assessment, i.e., (a) to (f) except (d), they could be fulfilled by 15 images. This is still relatively small for the whole set of images. On the other hand, when only the exposure criteria (g) and (h) were evaluated, 172 images were found with adequate exposure.

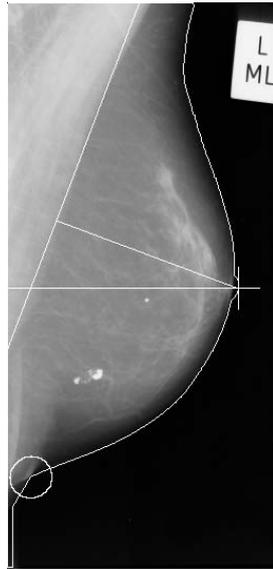


FIGURE 5.10: MIAS image mdb075lm. The only image that satisfied all quality criteria: no breast tissue exclusion; nipple in profile; inframammary fold included; pectoral margin extended below the level of nipple by 22.8 mm at an angle of 20.7° with 64% convexity; mean optical density at 1.74; and 82% of tissue in useful density range.

5.7 Discussion

The algorithm for breast tissue exclusion was found reliable when assessed visually as described in Section 5.6. Provided the breast border was accurately segmented, most of the exclusion could be detected. The detected exclusion also serves as an important cue to the accuracy of other extracted features. If the breast is excluded at the anterior image edge, it is possible that the nipple is excluded too. This could be the cause of failure in the automatic nipple detection. Likewise, if the breast tissue exclusion is at the inferior image edge, it is possible that the inframammary fold is also excluded from the image. Any inframammary fold located by the computer would then be a false positive. However, there is a chance that the nipple or inframammary fold is still included on the image when breast tissue exclusion is detected (see Fig. 5.11 for some examples).

The method described in [13] was also found reliable to determine whether the nipple is in profile. Its accuracy was affected by poor segmentation of the breast border. Sometimes the nipple was actually in profile but was incorrectly classified as not in profile, because the absolute threshold of θ' was set too high in the algorithm.

High success rate was achieved in the inframammary fold detection. Again the method was sometimes affected by poor segmentation of the breast border, since the basis is solely on the examination of normal directions. The inframammary folds could appear in different shapes: some are rounded and some are very sharp, as shown in Fig. 5.5. Most of them could be detected by the algorithm with few exceptions. The fold could not be detected when it was a very sharp and adjacent to a vertical *straight* line, on which θ' was zero (no curvature). In those cases, the concave segment was shorter than 10 mm

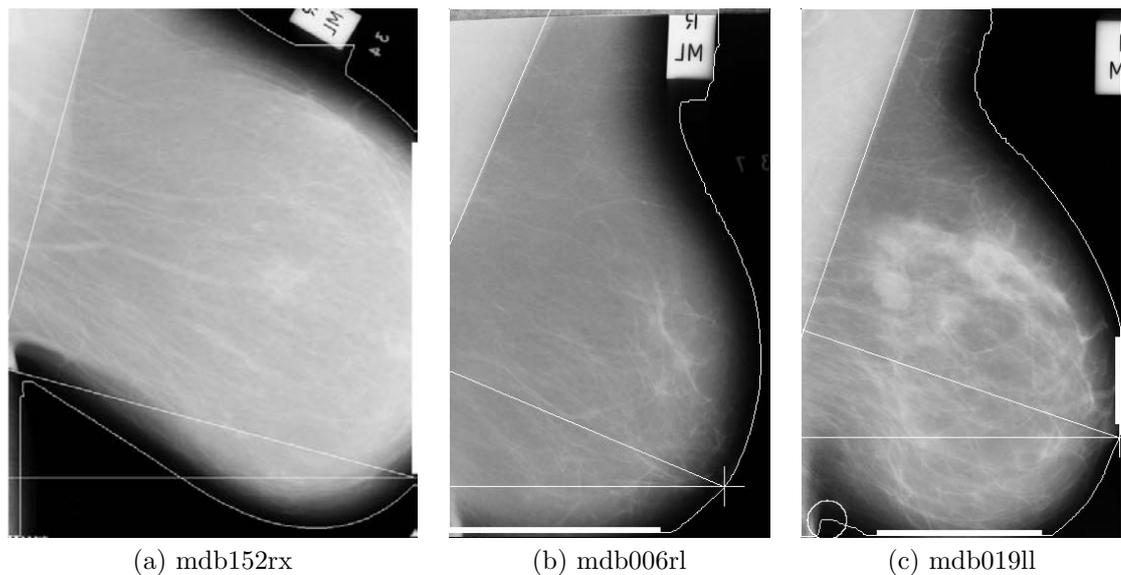


FIGURE 5.11: (a) Nipple is excluded at the anterior image edge and incorrectly located by the computer. (b) Inframammary fold is excluded at the inferior image edge. (c) Nipple and inframammary fold are both correctly located even though some breast tissue is excluded.

and left undetected. This could be overcome by reducing the required length for concave segments, but it may produce more false positives on poorly extracted breast borders. This is typical of the tradeoffs that need to be made to render algorithm robust.

The position of the pectoral muscle on the MLO view may be measured precisely by the proposed method. The amount of pectoral muscle included in the view, the angle of the pectoral margin and the posterior nipple line may be obtained by simply drawing *three straight lines* on the image.

However these measurements are subject to the accuracy of the pectoral muscle segmentation and the nipple location. The method used to measure the curvature of the pectoral margin is similar to that used on the breast border. Since the pectoral curvature is usually less obvious compared to the breast border, a larger neighbourhood is needed to estimate the tangents. Provided the pectoral curve was accurately segmented, the curvature measurement was found to be very accurate. The convexity and concavity could be labelled on the curve, as shown in Fig. 5.7(d), so that we know exactly which part of the curve is convex.

The scatter plot in Fig. 5.12 shows the relationship between the two optical density measures used in the algorithm. Mammograms with their mean optical density between 1.4 and 1.8 are more likely to include higher percentage of breast tissue in the useful optical density range. The maximum point of the trend line also lies within [1.4, 1.8]. However, the low r^2 value (r is correlation coefficient) of 0.11 indicates that the correlation is not strong, since the mean optical density is not the only factor that determines the amount of tissue visualized in the useful density range.

In fact, a high percentage of tissue in the useful density range does not necessarily

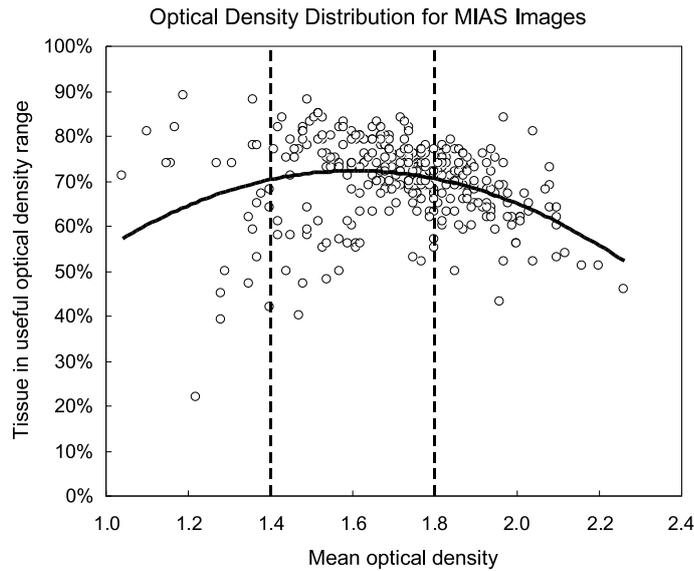


FIGURE 5.12: The scatter plot of the mean optical density versus the percentage of tissue visualized in the useful optical density range for all images in the MIAS database. The trend line is shown as a polynomial with degree 2; $r^2 = 0.11$ (r is correlation coefficient). In 193 cases, the mean optical density lies between 1.4 and 1.8 indicating adequate exposure.

ensure adequate exposure. As shown in Fig 5.9, although all three images have around 70% of breast tissue in the useful density range, only one of them with adequate mean optical density was properly exposed. Therefore both optical density measures have to be assessed when determining the adequacy of exposure.

Inadequate exposure does not only limit the visualization of breast tissues, it also affects the performance of the segmentation of the digitized mammograms. In Fig. 5.13, the mammogram was poorly segmented by the computer. The breast border is inaccurate; the nipple location is not on the spot and is actually in profile; the axillary fold is outlined instead of the pectoral margin, which is barely visible; the inframammary fold is circled but it does not really exist on the image. Any result produced by the positioning assessment are consequently not valid. The cause of poor segmentation is mainly due to the overexposure and incorrect positioning of the mammogram, in which the skin line is not properly visualized, and too little pectoral muscle can be seen. The measured mean optical density is too high at 2.10 and only 53% of breast tissue is in the useful density range.

5.8 Suggestions for Further Work

5.8.1 Improvements to Algorithms

Two criteria related to the posterior glandular tissue ((5) and (6) in Section 5.3) were not assessed in this work because an outline of glandular tissue was not available, due to the difficulties of segmenting dense tissue accurately in the breast region. There are several

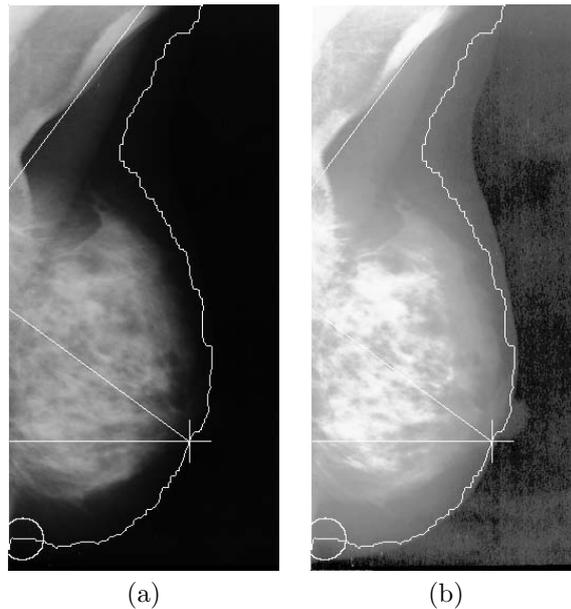


FIGURE 5.13: Result of MIAS image mdb066rm. (a) Overexposed and poorly segmented mammogram. Breast border, nipple location, inframammary fold and pectoral margin are all inaccurate. The level of nipple and posterior nipple line are consequently incorrect. (b) After histogram equalization, the skin-air interface is better visualized, showing that the nipple is in profile and the inframammary fold is actually excluded from the image.

methods in the literature for the segmentation of dense tissue [8, 15, 17]. In these methods, the pectoral muscle is usually excluded first before the dense tissue is segmented. However, the glandular tissue may sometimes appear superimposed on the pectoral muscle, i.e., the tail of Spence. This is a significant feature in adequacy assessment and should not be excluded. Because its density is only slightly different from that of the pectoral muscle, it is difficult to segment the tail of Spence accurately.

For the optical density measures, all the pixels on the image were converted to optical densities. Rather than doing that, it is equivalent if we convert the upper and lower limits of the useful optical density range to pixel values and use these limits to evaluate the image pixels. Only two optical density conversions are required in the latter case, hence computational time can be reduced. Similarly, if the conversion is linear, the mean optical density is equivalent to that converted from the mean pixel value.

5.8.2 Extension to Other Quality Attributes

Besides the adequacy of positioning and exposure, certainly more quality attributes mentioned in Section 5.2 should be included in the assessment. Not all quality attributes can be measured using the computer, since some of them depend on subjective considerations. For example, the measurements of contrast-detail curve and signal-to-noise ratio require the use of phantoms. However, some quantitative attributes, such as contrast, compression, sharpness and artefacts, are likely to be measured on individual mammograms using

the computer. These attributes are not considered in this chapter but should be included in future work.

High contrast should be achieved in every density range so that tissues of different types can be well differentiated. However, the contrast should not be too high otherwise the latitude will be too narrow [33]. Without phantom images, simple contrast measures such as the relative density difference between object and background is not really applicable to the mammogram since it is formed by the superposition of all types of tissues. One possible method to evaluate mammographic contrast may consist of two steps: (i) segment the image into several regions according to tissue density, such as uncompressed-fat, fat, dense and muscle, by mixture modelling [15] or other means; (ii) calculate the contrast in each region and determine whether uniformly high contrast is maintained in every density range.

The breast is adequately compressed when the glandular tissue is well spread out; simply knowing the thickness of the compressed breast is not enough. One of the possible methods to evaluate compression is by measuring the coherence of the ductal structure of glandular body. However, segmenting glandular tissues on the mammogram is already a demanding task; extracting its ductal structure is likely to prove even more difficult.

Quality of image sharpness is affected by motion blurring, geometric blurring and screen-film blurring [33]. Motion blurring may be detected by examining the sharpness of microcalcifications on the mammogram, although it is quite difficult to measure when the microcalcifications present in areas of dense tissue. Geometric blurring and screen-film blurring should be measured on phantom images. The former relates to the size of focal spot and focus-film distance; the latter depends on the resolution of the intensifying screen. The clarity of details can be expressed by the modulation transfer function (MTF).

Small artefacts can be mistaken by the observer as microcalcifications. One of the methods that detects tiny artefacts such as dust or hair is presented in [77]. Large artefacts could cover some of the breast tissue and prevent accurate diagnosis. Two algorithms for the removal of large, high-intensity artefacts are described in [78]. Further investigation could be carried out on the detection of other types of artefacts, such as fingerprints, powder, or defects of the screen.

If the CC view is available, its adequacy should be examined also. Most of the quality criteria can be applied to both MLO and CC views except for positioning. The criteria for optimal positioning on the CC view are clearly described in [34]. The posterior nipple line of the CC view should be compared with that of the MLO view as stated in Section 5.3. The posterior nipple line on the CC view should not be shorter than that on the MLO view by more than 1 cm.

5.8.3 Assessment of Other Mammogram Databases

Although only the images from the MIAS database were used for the results in this chapter, the principles of automatic adequacy assessment are applicable to digitized mammograms

in general. Because more recent image databases of mammograms were not available during this research, its scope extended only to the MIAS database. It is suggested that further work should be carried out on the adequacy assessment of mammograms from newer databases, which may contain a larger percentage of acceptable mammograms.

5.9 Conclusions

In this chapter, the adequacy of mammographic positioning and exposure was examined for the 322 images in the MIAS database. The purpose of this assessment is not only to investigate the quality of MIAS images, but also to demonstrate a proof-of-concept that mammographic image quality can be assessed by computer. The assessment algorithms developed were found to be reliable and robust, but they are more or less affected by the accuracy of the segmented anatomic features, i.e., breast border, nipple location and pectoral margin. In the assessment, only 4.7% of the images were found adequate in positioning and 53.4% exhibited proper exposure (Table 5.1). The results also showed that breast tissue exclusion, overexposure and poor positioning could be a source of errors for the automatic segmentation algorithms. Automatic assessment of the quality of CC view mammograms should be included in the future.

This concludes our description on mammogram adequacy assessment. In the next chapter, a novel technique is presented for highlighting mass lesions on mammograms.

Chapter 6

Contrast Enhancement of Mass Lesions

6.1 Introduction

After the mammogram has been segmented and assessed to be adequate in breast positioning and exposure, some diagnostic algorithms may be performed for the detection of abnormalities, such as mass lesions, in the breast region of the image. This chapter focuses on the contrast enhancement of mammographic masses, which is a necessary step for mass detection, segmentation and classification [51]. It is also one of the four steps that are performed in the mass lesion detection process, described in Section 3.3 and illustrated in Fig. 3.2.

Masses are often difficult to identify when they are obscured by dense breast tissue, or when the image quality is very poor. In these cases, the contrast between the object of interest and the background may be insufficient for high sensitivity mass detection. To improve the contrast of the image, some enhancement techniques specifically designed for mammograms are required.

A new contrast enhancement technique is presented in this chapter. This technique is a modification of the amplitude and phase (AMPM) demodulation transform developed by Daugman and Downing [4,79]. Based on their work, we modified the AMPM demodulation transform specifically for the purpose of highlighting mass lesions on mammograms. In addition, two existing contrast enhancement techniques, morphological enhancement and histogram equalization, are also described in this chapter. Our new contrast enhancement technique was compared with the existing ones by applying them to all mammographic masses from the MIAS database. The results were evaluated using three quantitative measures that were proposed by Bovis and Singh [5].

First of all, the original AMPM demodulation transform is described in the next section.

6.2 Amplitude and Phase Demodulation

The amplitude and phase demodulation transform was introduced by Daugman and Downing [4, 79], who suggested that any 2D pattern can be represented by a single carrier wave and a complex phasor that specifies over the space domain its amplitude modulation and phase modulation (AMPM). Therefore, an image of any pattern can be decomposed into three components:

1. carrier wave,
2. phasor's AM component, and
3. phasor's PM component.

These three components can also be used to completely reconstruct the original image through the AMPM modulation process. The algorithm of the demodulation transform [4] is described below.

6.2.1 Algorithm of AMPM Demodulation Transform

An arbitrary real-valued image, $I(x, y)$, can be represented as

$$I(x, y) = S(x, y) + s_0 \quad (6.1)$$

where s_0 is the mean value of $I(x, y)$, and $S(x, y)$ is a 2D pattern with zero mean value.

$S(x, y)$ is taken as the original pattern for demodulation. It is first expressed by the 2D Fourier series expansion using *paired conjugate spatial frequencies*

$$S(x, y) = \sum_{k=-N}^N \alpha_k \exp [j(\mu_k x + \nu_k y)] \quad (6.2)$$

In the expansion, there are $2N$ paired conjugate 2D frequency components $(\mu_k, \nu_k) = (-\mu_{-k}, -\nu_{-k})$ and their associated complex coefficients $\alpha_k = a_k + jb_k$ with $\alpha_0 = 0$. Moreover, because $S(x, y)$ is real rather than complex, the coefficients α_k are conjugate symmetric, i.e., $\alpha_k = \alpha_{-k}^*$. The polar form of α_k is $\|\alpha_k\| \exp(j\theta_k)$ where $\|\alpha_k\| = \sqrt{a_k^2 + b_k^2}$ and $\theta_k = \tan^{-1}(b_k/a_k)$.

Since the Fourier spectrum is conjugate symmetric, the Fourier half-plane may be defined in any orientation such that all values of the index k sharing the same sign (+/-) are in the same half-plane. Over the half-plane of all positive frequency components, the spectral centre of mass (μ_c, ν_c) is determined by:

$$\mu_c = \frac{\sum_{k=1}^N \|\alpha_k\| \mu_k}{\sum_{k=1}^N \|\alpha_k\|} \quad (6.3)$$

and

$$\nu_c = \frac{\sum_{k=1}^N \|\alpha_k\| \nu_k}{\sum_{k=1}^N \|\alpha_k\|} \quad (6.4)$$

The *predictive carrier wave*, $C(x, y)$, is then selected at the centre-of-mass frequency (μ_c, ν_c) :

$$C(x, y) = \exp [j(\mu_c x + \nu_c y)] \quad (6.5)$$

For simplicity, only a single centre-of-mass frequency is chosen here for the single carrier demodulation. However, more than one such gravitational centre can be chosen in multi-carrier demodulation, as explained later in Section 6.3.

The *complex modulation phasor*, $Z(x, y)$, is generated by *demodulating* the original signal $S(x, y)$ by the carrier $C(x, y)$. Thus all the positive frequency components in the half-plane are *shifted* to have their centre of mass (μ_c, ν_c) at the origin.

$$Z(x, y) = \sum_{k=1}^N \alpha_k \exp [j(\Delta\mu_k x + \Delta\nu_k y)] \quad (6.6)$$

where $\Delta\mu_k = \mu_k - \mu_c$ and $\Delta\nu_k = \nu_k - \nu_c$. It is worthwhile to note that

$$\begin{aligned} Z(x, y)C(x, y) &= \left\{ \sum_{k=1}^N \alpha_k \exp [j(\Delta\mu_k x + \Delta\nu_k y)] \right\} \exp [j(\mu_c x + \nu_c y)] \\ &= \sum_{k=1}^N \alpha_k \exp [j(\mu_k x + \nu_k y)] \end{aligned} \quad (6.7)$$

which is half of the series (6.2), and for the other half

$$\begin{aligned} Z^*(x, y)C^*(x, y) &= \left\{ \sum_{k=1}^N \alpha_k^* \exp [-j(\Delta\mu_k x + \Delta\nu_k y)] \right\} \exp [-j(\mu_c x + \nu_c y)] \\ &= \sum_{k=1}^N \alpha_k^* \exp [-j(\mu_k x + \nu_k y)] \\ &= \sum_{k=-N}^{-1} \alpha_k \exp [j(\mu_k x + \nu_k y)] \end{aligned} \quad (6.8)$$

The original pattern $S(x, y)$ can then be reconstructed without loss using (6.7) and (6.8) so

$$S(x, y) = Z(x, y)C(x, y) + Z^*(x, y)C^*(x, y) \quad (6.9)$$

Alternatively, taking only the real part

$$S(x, y) = 2 \Re[Z(x, y)C(x, y)] \quad (6.10)$$

Note that the mean value s_0 must be added back to $S(x, y)$ to obtain the initial image $I(x, y)$, in accordance with (6.1).

Furthermore, the complex modulation phasor (6.6) can be expressed in complex polar form to obtain as a product of amplitude modulation (AM) and phase modulation (PM) components, $A(x, y)$ and $\phi(x, y)$ respectively:

$$Z(x, y) = A(x, y) \exp[j\phi(x, y)] \quad (6.11)$$

where

$$A(x, y) = \left\{ \left[\sum_{k=1}^N \|\alpha_k\| \cos(\Delta\mu_k x + \Delta\nu_k y + \theta_k) \right]^2 + \left[\sum_{k=1}^N \|\alpha_k\| \sin(\Delta\mu_k x + \Delta\nu_k y + \theta_k) \right]^2 \right\}^{1/2} \quad (6.12)$$

and

$$\phi(x, y) = \tan^{-1} \left[\frac{\sum_{k=1}^N \|\alpha_k\| \sin(\Delta\mu_k x + \Delta\nu_k y + \theta_k)}{\sum_{k=1}^N \|\alpha_k\| \cos(\Delta\mu_k x + \Delta\nu_k y + \theta_k)} \right] \quad (6.13)$$

The AM component $A(x, y)$ may also be expressed just in terms of all the vector difference frequencies among all the components of $S(x, y)$ in the half-plane:

$$A(x, y) = \left\{ \sum_{m=1}^N \sum_{n=1}^N \|\alpha_m\| \|\alpha_n\| \cos[(\mu_m - \mu_n)x + (\nu_m - \nu_n)y + (\theta_m - \theta_n)] \right\}^{1/2} \quad (6.14)$$

This equation reveals the important property that *the phasor's AM component depends only on the original pattern $S(x, y)$ and is independent of the carrier frequency (μ_c, ν_c)* . Therefore the image demodulation process always produces the same *filter-independent* AM component that is particularly useful later in our approach to image enhancement (Section 6.4).

6.2.2 Results with Texture Images

The AMPM demodulation transform was applied to a synthetic texture composed by three frequency components and the results are shown in Fig. 6.1. The figure illustrates the derived components, $C(x, y)$, $A(x, y)$ and $\phi(x, y)$, together with the complete reconstruction of the image. The Fourier spectrum is very well-concentrated and therefore the phasor's PM component (Fig. 6.1(e)) contains a pattern with much less complexity. The AMPM demodulation transform was also tested on natural texture as shown in Fig. 6.2. For this texture, the Fourier spectrum is less concentrated and hence the resulting PM component is almost as complex as the original image. These two examples reveal that the complexity

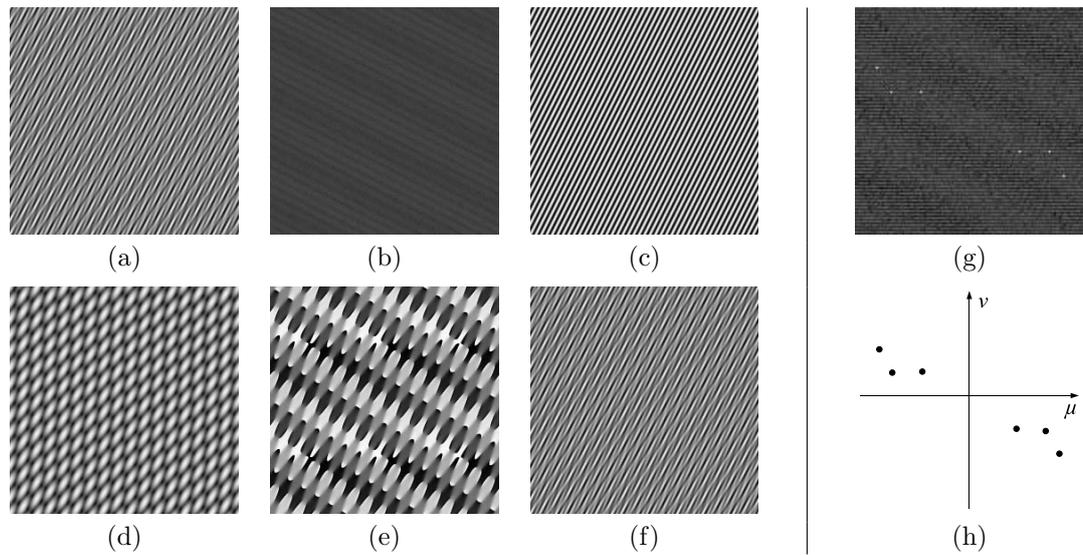


FIGURE 6.1: (a) Original synthetic three-component texture. (b) 2D Fourier spectrum of texture. (c) Derived carrier wave. (d) Phasor's AM component. (e) Phasor's PM component. (f) AMPM complete reconstruction. (g) Enlarged view of the origin of the Fourier spectrum (b) showing the spectral concentration. (h) A plot of the three frequency components on the Fourier plane; it is drawn to the same scale as (g). The bright dots in (g) correspond to the positions of frequency components in (h).

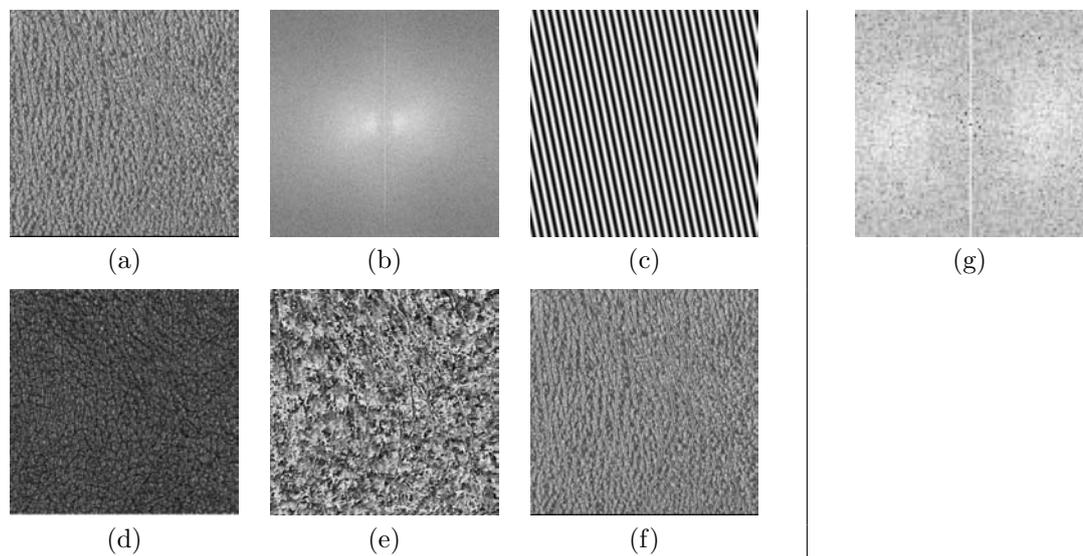


FIGURE 6.2: (a) Brodatz natural texture image illustrated in [80, p. D24]. (b) 2D Fourier spectrum of texture. (c) Derived carrier wave. (d) Phasor's AM component. (e) Phasor's PM component. (f) AMPM complete reconstruction. (g) Enlarged view of the origin of the Fourier spectrum (b) showing very little spectral concentration. The black spots in this image represent frequencies with very low magnitudes.

of the phasor's PM component $\phi(x, y)$ depends on the *spectral concentration* of pattern and the chosen carrier wave. The underlying image texture is predicted by the carrier wave much better if the Fourier spectrum of the image is well concentrated [4]. In other words, *if the predictive power of the carrier is large, the structure of the PM component will be much simpler than the original pattern* [4].

6.2.3 Selection of Half-Plane and Carrier Wave

The demodulation transform algorithm presented by Daugman and Downing does not specify the orientation of the Fourier half-plane on the spectrum [4]. Due to conjugate symmetry of the paired coefficients, the half-plane can be set in *any* orientation as long as all values of the index k sharing the same sign ($+/-$) are in the same half-plane. The carrier wave is selected as the centre-of-mass frequency in the positive half-plane using (6.3) and (6.4).

The centre-of-mass frequency should lie *inside* a spectral cluster in the Fourier spectrum, so that the textural pattern is effectively demodulated. This is usually the case when the boundary of the half-plane does not partition the spectral clusters, and there is only one spectrally concentrated cluster in the positive half-plane. An example is shown in Fig. 6.3(a). However, in a special case, the centre-of-mass frequency does lie *outside* of the spectral clusters, when the boundary of the half-plane cut across the spectral clusters, and two smaller clusters are depicted in the positive half-plane. This is illustrated in Fig. 6.3(b). So *the choice of half-plane orientation affects the computation of the centre-of-mass frequency*, and hence different orientations of half-planes may yield different carrier waves.

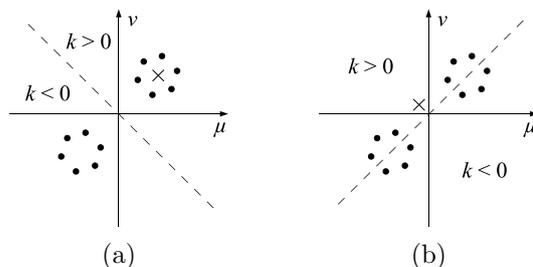


FIGURE 6.3: (a) The centre-of-mass frequency (cross) lies inside the spectral cluster (dots) in the positive Fourier half-plane ($k > 0$). (b) The centre-of-mass frequency lies outside of the spectral clusters in the positive Fourier half-plane.

In previous work [81], we have addressed this observation and introduced a method to select the half-plane. The idea is to set the boundary of the half-plane perpendicular to the dominant direction of the Fourier spectrum, which is also orthogonal to the dominant direction of the image texture [69, p. 603]. The selected carrier frequency is therefore closest to the concentrated portion of the spectrum and hence is expected to have more predictive power.

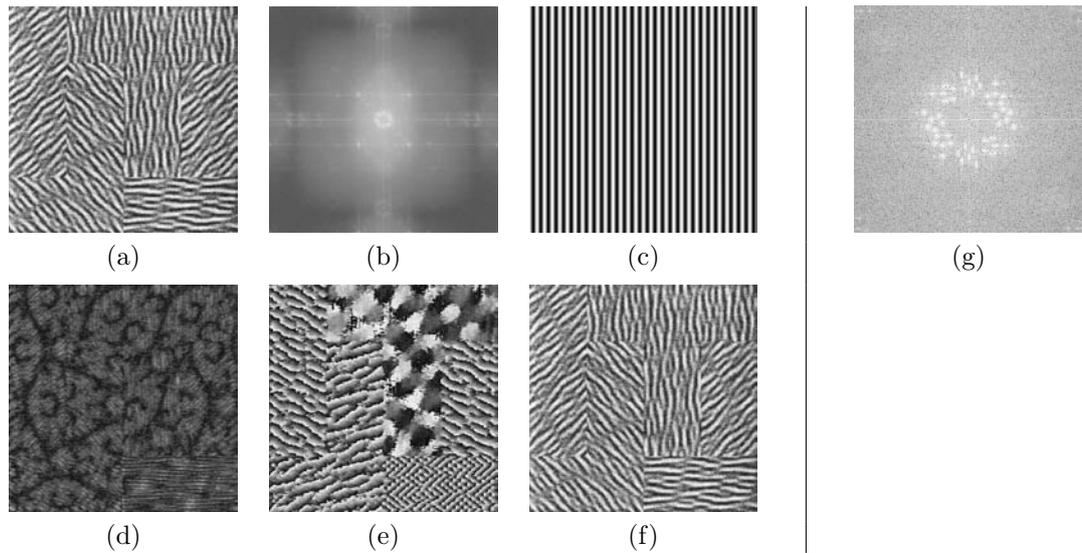


FIGURE 6.4: (a) Original pattern of incoherent texture fields, scanned from [4]. (b) 2D Fourier spectrum of texture. (c) Carrier wave was forced to be vertical. (d) Phasor's AM component. (e) Phasor's PM component. Its complexity is lowest in the texture region with vertical correlation (T-shaped); its complexity is highest in the texture region with horizontal correlation (bottom right). (f) AMPM complete reconstruction. (g) Enlarged view of the origin of the Fourier spectrum (b) showing spectral concentration in all directions.

Nevertheless, with this half-plane selection method, the texture of mammographic mass images are still not simplified to a satisfactory level using the single-carrier demodulation. This is due to the fact that the predictive power of the carrier wave is insufficient, and the complexity of the decomposed components is not substantially reduced. Because the texture of mammographic masses is far from being spectrally concentrated, very little information of interest could be obtained from texture demodulation with a single carrier wave. We have therefore taken a different approach that is *independent of the half-plane orientation*, for the purpose of image enhancement. This is described in Section 6.4.

6.3 AMPM Demodulation with Multiple Carriers

The AMPM demodulation transform presented in [4] was restricted to the aspects of single carrier although the possibility of *multi-carrier demodulation* was left to future investigation. Nevertheless, Daugman and Downing have addressed the limitation of single-carrier demodulation using both synthetic texture and natural scenes in [4]. They have shown that if the original pattern consists of incoherent texture fields defined in different primary orientations of correlation, then a single carrier can only be chosen in one of the textural directions and consequently only a portion of the PM component is simplified while other portions remain complex (see Fig. 6.4). However, such a pattern may be demodulated using multiple carriers, each corresponding to one textural direction, so that the resulting PM component is completely simplified. Daugman and Downing have suggested the par-

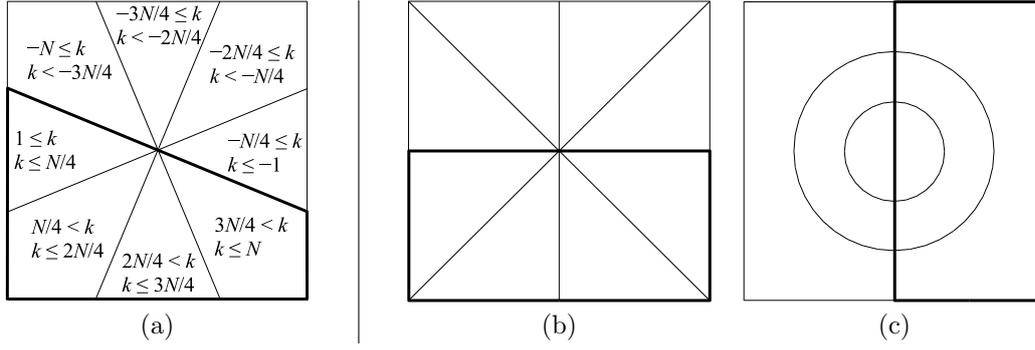


FIGURE 6.5: (a) Partitioning of the Fourier spectrum for the original pattern shown in Fig. 6.4. The k labelling in each sector corresponds to the Fourier series given in (6.15). (b)–(c) Other possible ways to partition the Fourier spectrum for other textural patterns.

partitioning of either the image domain or the Fourier domain, or both, into distinct regions that use different demodulation carriers [4]. The author therefore took this further step into the investigation of multi-carrier demodulation by partitioning the Fourier spectrum, as described below.

6.3.1 Algorithm of Multi-carrier AMPM Demodulation

In multi-carrier demodulation, the Fourier spectrum is divided into several regions. One carrier wave is selected in each region. These regions may have any shape but they must be symmetric on both sides of the half-planes. It is not necessary to have the regions in equal size, but for the pattern shown in Fig. 6.4(a), we divide the Fourier half-plane into four distinct sectors, which correspond to the four primary orientations of texture correlations, as shown in Fig. 6.5(a).

The Fourier series of the original pattern (6.2) is therefore expanded into *eight* summations:

$$\begin{aligned}
 S(x, y) &= \sum_{k=-N}^N \alpha_k \exp[j(\mu_k x + \nu_k y)] \\
 &= \sum_{k=1}^{N/4} \alpha_k \exp[j(\mu_k x + \nu_k y)] + \sum_{k=-N/4}^{-1} \alpha_k \exp[j(\mu_k x + \nu_k y)] + \\
 &\quad \sum_{k=N/4+1}^{2N/4} \alpha_k \exp[j(\mu_k x + \nu_k y)] + \sum_{k=-2N/4}^{-N/4-1} \alpha_k \exp[j(\mu_k x + \nu_k y)] + \\
 &\quad \sum_{k=2N/4+1}^{3N/4} \alpha_k \exp[j(\mu_k x + \nu_k y)] + \sum_{k=-3N/4}^{-2N/4-1} \alpha_k \exp[j(\mu_k x + \nu_k y)] + \\
 &\quad \sum_{k=3N/4+1}^N \alpha_k \exp[j(\mu_k x + \nu_k y)] + \sum_{k=-N}^{-3N/4-1} \alpha_k \exp[j(\mu_k x + \nu_k y)] \quad (6.15)
 \end{aligned}$$

with $\alpha_0 = 0$.

The centre-of-mass frequency, (μ_{c1}, ν_{c1}) , of the first sector for $1 \leq k \leq N/4$ is defined as

$$\mu_{c1} = \frac{\sum_{k=1}^{N/4} \|\alpha_k\| \mu_k}{\sum_{k=1}^{N/4} \|\alpha_k\|} \quad ; \quad \nu_{c1} = \frac{\sum_{k=1}^{N/4} \|\alpha_k\| \nu_k}{\sum_{k=1}^{N/4} \|\alpha_k\|} \quad (6.16)$$

The predictive carrier wave, $C_1(x, y)$, of the first sector for $1 \leq k \leq N/4$ is specified at (μ_{c1}, ν_{c1}) so

$$C_1(x, y) = \exp[j(\mu_{c1}x + \nu_{c1}y)] \quad (6.17)$$

The AMPM modulation phasor, $Z_1(x, y)$, of the first sector for $1 \leq k \leq N/4$ is computed as

$$Z_1(x, y) = \sum_{k=1}^{N/4} \alpha_k \exp[j(\Delta\mu_k x + \Delta\nu_k y)] \quad (6.18)$$

where $\Delta\mu_k = \mu_k - \mu_{c1}$ and $\Delta\nu_k = \nu_k - \nu_{c1}$. Similarly, the carriers waves and phasors for the other three sectors can be derived as in (6.16), (6.17) and (6.18).

As a result, the positive Fourier half-plane can be expressed as

$$\sum_{k=1}^N \alpha_k \exp[j(\mu_k x + \nu_k y)] = \sum_{q=1}^4 Z_q(x, y) C_q(x, y) \quad (6.19)$$

where q corresponds to the four sectors. Also, the negative Fourier half-plane can be expressed as

$$\sum_{k=-N}^{-1} \alpha_k \exp[j(\mu_k x + \nu_k y)] = \sum_{q=1}^4 Z_q^*(x, y) C_q^*(x, y) \quad (6.20)$$

The original pattern can be completely reconstructed using (6.19) and (6.20):

$$S(x, y) = \sum_{q=1}^4 Z_q(x, y) C_q(x, y) + Z_q^*(x, y) C_q^*(x, y) \quad (6.21)$$

Alternatively

$$S(x, y) = 2 \Re \left[\sum_{q=1}^4 Z_q(x, y) C_q(x, y) \right] \quad (6.22)$$

It is worthwhile to realize that the phasor of the whole pattern given in (6.6) is the sum of the phasors of the four sectors:

$$Z(x, y) = \sum_{k=1}^N \alpha_k \exp[j(\Delta\mu_k x + \Delta\nu_k y)]$$

$$= \sum_{q=1}^4 Z_q(x, y) \quad (6.23)$$

therefore $A(x, y)$ and $\phi(x, y)$ can be computed using (6.12) and (6.13) respectively.

The original pattern shown previously in Fig. 6.4 was demodulated by the four carriers chosen in the four spectral sectors, and the results are shown in Fig. 6.6. Note that the PM component is now simplified for texture fields of *all* orientations.

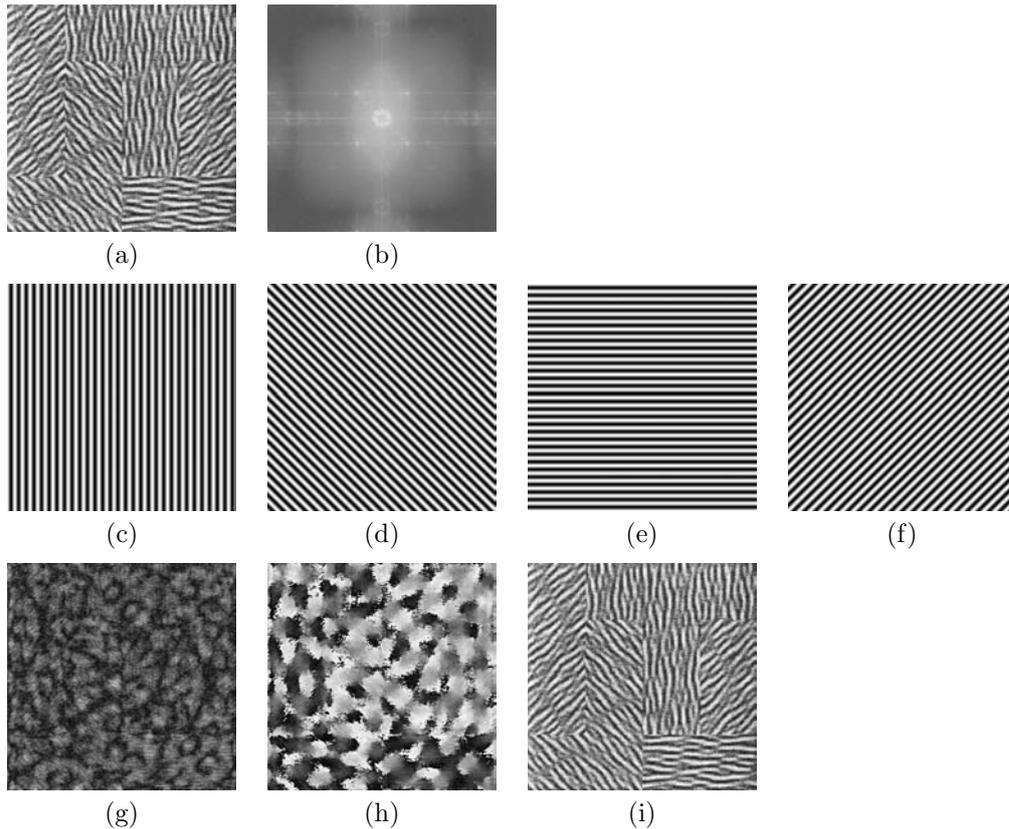


FIGURE 6.6: (a) Original pattern of incoherent texture fields illustrated in [4]. (b) 2D Fourier spectrum of texture. (c)–(f) Carrier waves chosen in the four sectors of Fourier half-plane. (g) Phasor’s AM component. (h) Phasor’s PM component. Its complexity is reduced for all texture fields. (i) AMPM complete reconstruction.

6.4 Modification for Contrast Enhancement of Mass Lesions

Although the single- or multi-carrier demodulation transform is able to decompose synthetic texture images into simplified AM and PM components, we have found that for spectrally less concentrated texture, such as natural or medical images, the complexity of the derived components still remains high. When the Fourier spectrum is very broad, no carrier waves can be chosen to have high predictive power. Even if the carrier is placed at the spectral centre of mass of the half-plane or one of the partitions, the resulting

PM component will contain little information of interest, achieving no economy through decomposition.

Therefore, rather than seeking simplification of image texture, we modified the AMPM demodulation algorithm [4] for the purpose of *contrast enhancement* of mass lesions on mammograms. The modified demodulation method was designed to enhance the contrast of ROI images of mass lesions, and it should be applied to the ROI image only, not to the whole mammogram.

The single-carrier demodulation algorithm, described previously in Section 6.2.1, was modified in four steps:

1. carrier wave of zero frequency;
2. phase-only and amplitude-only reconstructions;
3. subtraction of partially reconstructed component images; and
4. contrast stretching.

These modifications are explained in the next four sections and illustrated throughout in Fig. 6.7. Image (a) in the figure is the ROI of a mass lesion; images (b)–(f) were produced by the single-carrier demodulation transform with zero carrier frequency; and images (g)–(j) were derived by the modified algorithm for contrast enhancement.

6.4.1 Carrier Wave of Zero Frequency

As stated previously with the expression (6.14), the AM component $A(x, y)$ depends only on the original pattern $S(x, y)$ and is independent of the carrier frequency (μ_c, ν_c) . Thus one can obtain the same AM component even with *zero carrier frequency*, i.e.,

$$\begin{aligned} C(x, y) &= \exp[j(0x + 0y)] \\ &= 1 \end{aligned} \tag{6.24}$$

Therefore (6.10) becomes

$$S(x, y) = 2 \Re[Z(x, y)] \tag{6.25}$$

The phasor is simply projected into its AM and PM components using (6.11) without being demodulated. Although the AM and PM components are not simplified, they were found very useful in image enhancement (discussed later in Section 6.4.3).

6.4.2 Partial Reconstruction

The motivation for this step is the observation that a phase-only reconstruction of the discrete Fourier transform of an image results in an edge detected image of the original [82]. The intention here was to explore the range of images that resulted from amplitude-only

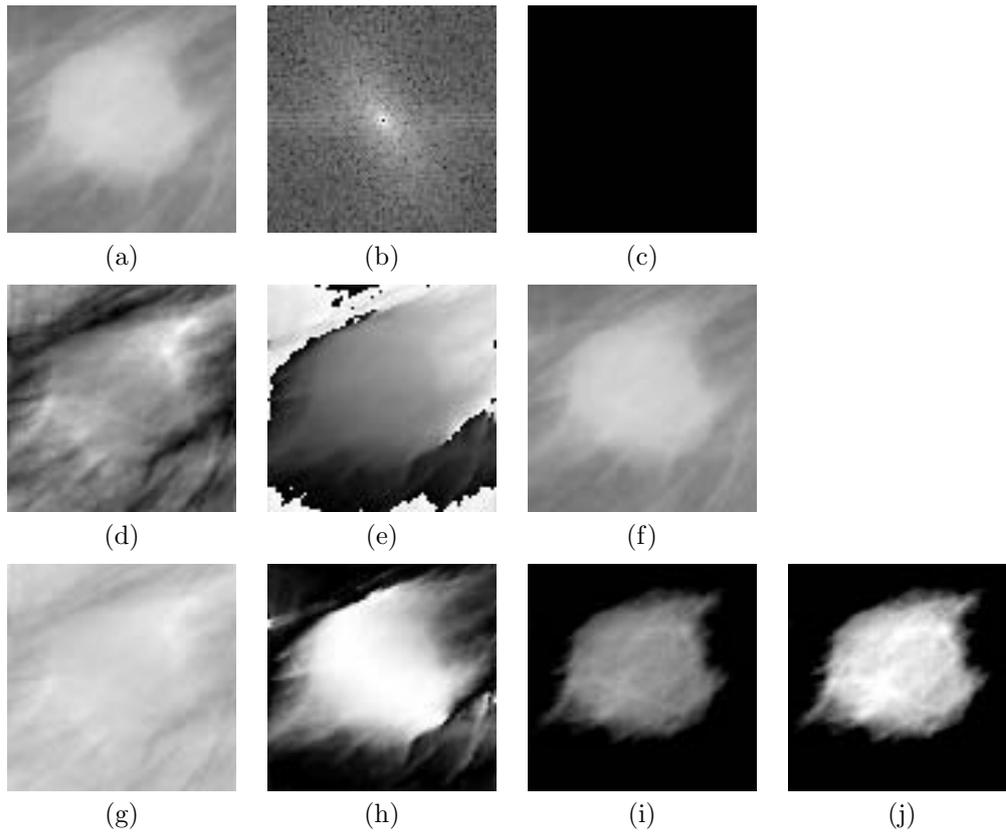


FIGURE 6.7: Contrast enhancement using texture demodulation. (a) Original ROI showing the circumscribed mass and texture of MIAS image mdb028rl. (b) 2D Fourier spectrum of texture. (c) Carrier wave of zero frequency. (d) Phasor's AM component. (e) Phasor's PM component. (f) AMPM complete reconstruction. (g) Amplitude-only reconstruction. (h) Phase-only reconstruction. (i) Result of image subtraction. (j) Result of contrast stretching.

and phase-only reconstruction and to determine if diagnostically useful information about the mammograms could be extracted from them.

When only the AM or PM component of the phasor is modulated with the carrier wave, the pattern is not completely but partially reconstructed. By omitting the AM or PM component in the reconstruction, the demodulation transform becomes *non-invertible*, and hence new components are derived.

As shown in (6.10) and (6.11), the original pattern $S(x, y)$ is completely reconstructed by the carrier $C(x, y)$ and the phasor $Z(x, y) = A(x, y) \exp[j\phi(x, y)]$. The pattern can also be partially reconstructed using either the AM component or PM component of the phasor. For *amplitude-only reconstruction*, all values of $\phi(x, y)$ are set to zero so that $Z(x, y) = A(x, y)$ and therefore

$$R_{AM}(x, y) = 2 \Re[A(x, y)C(x, y)] \quad (6.26)$$

where R_{AM} is the partially reconstructed pattern with AM component only.

On the other hand, for *phase-only reconstruction*, all values of $A(x, y)$ are set to one

so that $Z(x, y) = \exp[j\phi(x, y)]$ and hence

$$R_{\text{PM}}(x, y) = 2 \Re\{\exp[j\phi(x, y)]C(x, y)\} \quad (6.27)$$

where R_{PM} is the partially reconstructed pattern with PM component only.

In both (6.26) and (6.27), $C(x, y)$ is 1 if zero carrier frequency is used, so

$$R_{\text{AM}}(x, y) = 2 \Re[A(x, y)] \quad (6.28)$$

and

$$R_{\text{PM}}(x, y) = 2 \Re\{\exp[j\phi(x, y)]\} \quad (6.29)$$

6.4.3 Image Subtraction

The AM component of the phasor is specified in the space domain by and can be converted into a grey-scale image by linear transformation, i.e., the values of the AM component $A(x, y)$ in the range $[A_{\min}, A_{\max}]$ are mapped to $[0, I_{\max}]$ where I_{\max} is the maximum grey level. Similarly, the phase-only reconstructed component $R_{\text{PM}}(x, y)$ defined by can be converted into a grey-scale image by linearly mapping $[R_{\text{PM},\min}, R_{\text{PM},\max}]$ to $[0, I_{\max}]$.

The image representations of $A(x, y)$ and $R_{\text{PM}}(x, y)$ are shown in Fig. 6.7(d) and (h). In the AM component, the circumscribed mass is almost completely suppressed and the surrounding glandular tissue appears in dark colour. In the phase-only reconstructed image, the mass is very well highlighted with the background darkened, hence the image contrast is much higher than the original. These two meaningful images are therefore combined using image subtraction to produce the mass enhanced image with a pure black background.

In image subtraction, the *digital negative* of the AM component image is subtracted from the phase-only reconstructed image, i.e.,

$$H(x, y) = \max\{0, \hat{R}_{\text{PM}}(x, y) - [I_{\max} - \hat{A}(x, y)]\} \quad (6.30)$$

where $H(x, y)$ is the subtracted image; $\hat{R}_{\text{PM}}(x, y)$ and $\hat{A}(x, y)$ are the image representations of $R_{\text{PM}}(x, y)$ and $A(x, y)$ respectively. The image of $H(x, y)$ is shown in Fig. 6.7(i). We can see that the surrounding glandular tissue is removed, the shape of the circumscribed mass is revealed, and the background is in pure black colour. The idea of this step is based on intuition, and the motivation for using image subtraction is to remove some of the low intensity texture around the object, so that it is highlighted in the image.

6.4.4 Contrast Stretching

Although the background texture is removed and the mass is highlighted, the contrast of the subtracted image $H(x, y)$ may still be poor. This is because the intensity range of $H(x, y)$ obtained after image subtraction does not necessarily span the full range of

grey levels $[0, I_{\max}]$. To further enhance the contrast, *contrast stretching* [54, p. 235] (or intensity scaling) is carried out in the *final* step. Contrast stretching is simply a linear transformation that scale the image intensities to a desired range. If the desired range is $[0, I_{\max}]$, then the following equation can be used:

$$\hat{H}(x, y) = \frac{[H(x, y) - H_{\min}]I_{\max}}{H_{\max} - H_{\min}} \quad (6.31)$$

where H_{\min} and H_{\max} are the minimum and maximum values of $H(x, y)$ respectively.

The background of H_{\max} is already black, so H_{\min} is usually zero. But sometimes there are tiny bright spots in $H(x, y)$ that could lead to unsatisfactory scaling. To avoid the effect of outlying, high intensity pixels, a slight modification is introduced here. Rather than setting H_{\max} at the maximum value of $H(x, y)$, the 99th percentile of the histogram is used instead, i.e.,

$$\hat{H}(x, y) = \begin{cases} \frac{[H(x, y) - H_{\min}]I_{\max}}{H_{99\%} - H_{\min}} & \text{if } H(x, y) \leq H_{99\%} \\ I_{\max} & \text{if } H(x, y) > H_{99\%} \end{cases} \quad (6.32)$$

where $H_{99\%}$ is the maximum grey level in the 99th percentile of the histogram. As a result, 1% of the histogram, that most likely contains the bright spots in the image, is omitted from the computation of contrast stretching, and hence the enhancement process is more robust.

6.5 Existing Techniques of Contrast Enhancement

Two other existing techniques for contrast enhancement in digital mammography were implemented and compared with our modified texture demodulation method. The first technique is *morphological enhancement* which was developed by Li *et al.* [53], and the second one is the commonly-used *histogram equalization* [48].

6.5.1 Morphological Enhancement

The morphological enhancement technique described in [53] was designed to remove background texture from the mammogram and enhance the mass pattern by background correction. The algorithm was implemented by *dual morphological top-hat operations* and image subtraction [53]. It is applied to the entire image pixel by pixel with the same operations.

The morphological filters used in the algorithm [53] are represented as the combination of two basic operations: erosion and dilation.¹ Let f be a binary or grey-scale image. The

¹The symbols \ominus and \oplus were used to denote erosion and dilation in [53], but they are actually the symbols of *Minkowski subtraction* and *Minkowski addition*, respectively, as stated in [83, p. 76]. Minkowski subtraction and addition are equivalent to erosion and dilation, respectively, if and only if the structuring element is symmetric with respect to its origin.

erosion of f by a structuring element B is denoted by $\varepsilon_B(f)$ and is defined as the minimum of the translations of f by the vectors $-\mathbf{b}$ of B [83, pp. 65–66]:

$$\varepsilon_B(f) = \bigwedge_{\mathbf{b} \in B} f_{-\mathbf{b}} \quad (6.33)$$

Therefore, the eroded value at a given pixel \mathbf{x} is the minimum value of the image in the window defined by the structuring element when its origin is at \mathbf{x} [83, p. 66]:

$$[\varepsilon_B(f)](\mathbf{x}) = \min_{\mathbf{b} \in B} f(\mathbf{x} + \mathbf{b}) \quad (6.34)$$

The *dilation* of an image f by a structuring element B is denoted by $\delta_B(f)$ and is defined as the maximum of the translations of f by the vectors $-\mathbf{b}$ of B [83, p. 68]:

$$\delta_B(f) = \bigvee_{\mathbf{b} \in B} f_{-\mathbf{b}} \quad (6.35)$$

Hence, the dilated value at a given pixel \mathbf{x} is the maximum value of the image in the window defined by the structuring element when its origin is at \mathbf{x} [83, pp. 68–69]:

$$[\delta_B(f)](\mathbf{x}) = \max_{\mathbf{b} \in B} f(\mathbf{x} + \mathbf{b}) \quad (6.36)$$

Morphological opening and closing are defined in terms of erosion and dilation. The *opening* of an image f by a structuring element B is denoted by $\gamma_B(f)$ and is defined as the erosion of f by B followed by the dilation with the reflected structuring element $\check{B} = \{-\mathbf{b} \mid \mathbf{b} \in B\}$ [83, pp. 105–106]:

$$\gamma_B(f) = \delta_{\check{B}}[\varepsilon_B(f)] \quad (6.37)$$

The *closing* of an image f by a structuring element B is denoted by $\phi_B(f)$ and is defined as the dilation of f by B followed by the erosion with the reflected structuring element \check{B} [83, p. 108]:

$$\phi_B(f) = \varepsilon_{\check{B}}[\delta_B(f)] \quad (6.38)$$

After defining the filters, the morphological *white top-hat*, WTH, can be performed on an image f by subtracting the opening of f from its original [83, p. 121]:

$$\text{WTH}_B(f) = f - \gamma_B(f) \quad (6.39)$$

This operation is useful to extract the objects of sizes smaller than the structuring element from the original image. The dual top-hat operations developed for mammograms [53] consists of two WTH using two different disc shaped structuring elements that are *flat* and *symmetric* [83, pp. 64–65].

The first top-hat operation is used to extract the background texture and noise from

the grey-scale image f of a mammogram [53]:

$$r_1 = \text{WTH}_{B_1}(f) \quad (6.40)$$

where r_1 is the residue image containing the background texture; B_1 is a structuring element chosen to be *smaller* than the size of masses.

The second top-hat operation is used to enhance the mass pattern by background correction [53]:

$$r_2 = \text{WTH}_{B_2}(f) \quad (6.41)$$

where r_2 is the mass patten enhanced image; B_2 is a structuring element chosen to be *larger* than the size of masses. The actual sizes of B_1 and B_2 are specified later in this section.

Finally, the enhanced image f_1 can be derived as [53]:

$$f_1 = \max\{0, [r_2 - r_1]\} \quad (6.42)$$

The resulting image should have the mass pattern enhanced, and the background texture and noise inside the mass pattern removed. This enhanced image is post-processed by applying morphological opening and closing with a disk of five pixels in diameter so that other texture noise is eliminated [53]. An extra post-processing step was added by the author: the ROI of the highlighted mass is selected from the whole image and this ROI image is further enhanced by contrast stretching described in Section 6.4.4.

The sizes of mammographic masses were measured in [53] to derive the appropriate sizes for the structuring elements. When the image resolution is at $400 \mu\text{m}$, a disk with a diameter of 7 pixels was chosen to be the structuring element B_1 , because the smallest mass in the MIAS database also has 7 pixels in its diameter. On the other hand, a disk with a diameter of 75 pixels was chosen to be the structuring element B_2 , because the largest mass is of size 75 pixels in diameter also. These parameters were also used in our experiment later in this chapter.

The morphological enhancement technique described above was applied to the same mammographic mass used in Fig. 6.7. The results are shown in Fig. 6.8.

6.5.2 Histogram Equalization

Histogram equalization is a simple and commonly-used method for improving the contrast of an image. It was found to be useful in visualizing lesions on mammograms and increasing the sensitivity of subsequent analysis [48]. The basis of this method involves probability theory, where the image histogram is treated as a probability density function [54, pp. 241–242]. The goal is to transform the image so that the cumulative intensity histogram is linearized. The grey-level transformation function is derived from the cumulative his-

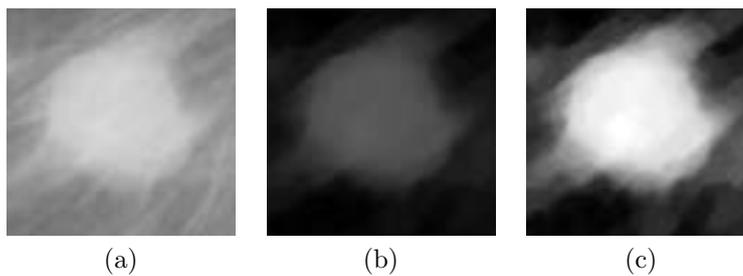


FIGURE 6.8: Contrast enhancement using morphological operation [53]. (a) Original ROI showing the circumscribed mass of MIAS image mdb028rl. (b) The morphologically enhanced image. (c) The result of contrast stretching.

togram, $c(i)$, which is defined as

$$c(i) = \begin{cases} h(i) & \text{if } i = 0 \\ c(i-1) + h(i) & \text{if } i = 1, \dots, L-1 \end{cases} \quad (6.43)$$

where h is the histogram of L grey levels.

Then the grey-level transformation function, $T(i)$, is derived as

$$T(i) = \frac{c(i) \cdot (L-1)}{c(L-1)} \quad (6.44)$$

Histogram equalization is performed by applying the transformation $T(i)$ to every pixel in the original image.

For our experiment on mammograms, the histogram equalization should not be applied to the whole image *but only to the breast region*. In other words, the non-breast background pixels were excluded from the formation of the grey-level transformation, so the enhanced image is not affected by the image background.

Histogram equalization was applied to the same mass image as previous examples and the result is shown in Fig. 6.9.

6.6 Quantitative Measures of Contrast Enhancement

The performance of the contrast enhancement techniques presented in previous sections should be evaluated by some quantitative measures. The three measuring methods developed by Bovis and Singh in [5] were used for this purpose. Their quantitative measures were designed to evaluate the contrast improvement of a target region against its surrounding background. In our case, the target regions are the highlighted masses on the mammograms. The computation of these quantitative measures are explained in the following sections.

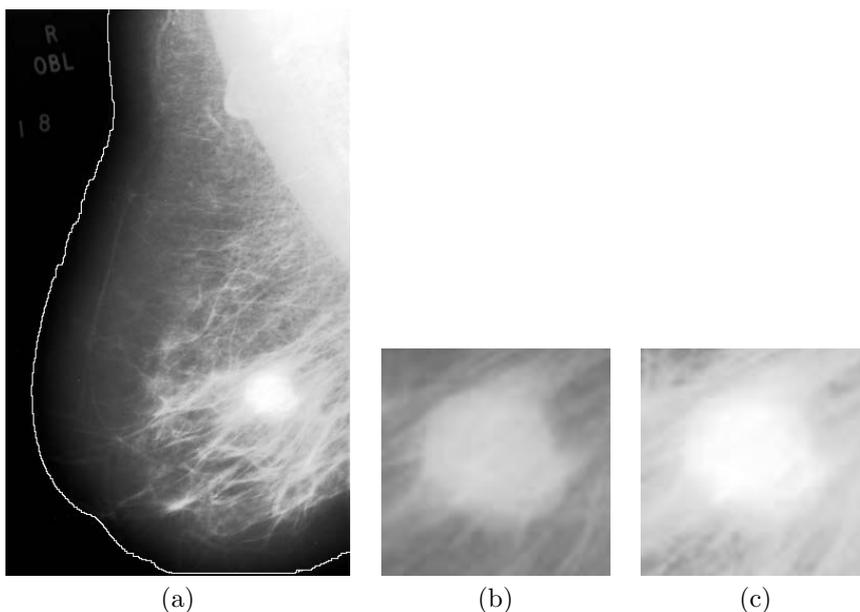


FIGURE 6.9: Contrast enhancement using histogram equalization. (a) Histogram equalization applied only to the breast region of MIAS image mdb028rl. (b) ROI of the original mass. (c) ROI of the enhanced mass.

6.6.1 Distribution Separation Measure (DSM)

The first measure proposed in [5] is *distribution separation measure* (DSM). To evaluate the performance of an enhancement technique, the separation of the grey-level distribution between the target and background regions before and after enhancement is measured. The original image is denoted as O . The target and background regions, T and B , are treated as two grey-level distributions in the histogram. The means of T and B are represented by μ_T^O and μ_B^O respectively; and the standard deviations of T and B are represented by σ_T^O and σ_B^O respectively. Likewise, for the enhanced image, E , the means and standard deviations of T and B are denoted as μ_T^E , μ_B^E , σ_T^E and σ_B^E . Then the best *decision boundary* for separating the two distributions in the original image is given by [5]

$$D_1 = \frac{\mu_B^O \sigma_T^O + \mu_T^O \sigma_B^O}{\sigma_B^O + \sigma_T^O} \quad ; \quad \sigma_B^O + \sigma_T^O \neq 0 \quad (6.45)$$

Similarly, the best decision boundary for separating the two distributions in the enhanced image is given by

$$D_2 = \frac{\mu_B^E \sigma_T^E + \mu_T^E \sigma_B^E}{\sigma_B^E + \sigma_T^E} \quad ; \quad \sigma_B^E + \sigma_T^E \neq 0 \quad (6.46)$$

Finally the distribution separation measure (DSM) is calculated as

$$\text{DSM} = (|D_2 - \mu_B^E| + |D_2 - \mu_T^E|) - (|D_1 - \mu_B^O| + |D_1 - \mu_T^O|) \quad (6.47)$$

If the image contrast has been enhanced, the DSM should be positive. The larger the value of DSM, the greater the separation between the distributions and therefore the better the image contrast enhancement. On the other hand, a negative DSM value indicates that the contrast of the enhanced image is worse than the original image.

6.6.2 Target to Background Contrast Ratio Using Standard Deviation (TBCs)

The second measure proposed in [5] is *target to background contrast ratio using standard deviation* (TBCs). (The term *variance* rather than *standard deviation* was actually used in [5], but the latter is more suitable to name this measure.) After contrast enhancement, the difference between intensity means of the target and background regions is expected to be greater, because μ_T^E is increased and μ_B^E is decreased. Therefore the ratio μ_T^E/μ_B^E should be greater than μ_T^O/μ_B^O . Furthermore, if the spread of the grey levels in the enhanced target region is reduced after enhancement, then the standard deviation σ_T^E should be greater than σ_T^O . Thus the target to background contrast ratio using standard deviation (TBCs) can be computed as

$$\text{TBCs} = \frac{(\mu_T^E/\mu_B^E) - (\mu_T^O/\mu_B^O)}{\sigma_T^E/\sigma_T^O} \quad (6.48)$$

where μ_B^E , μ_B^O , σ_T^E and σ_T^O are all non-zero.

Like the DSM, if the image contrast has been enhanced, the TBCs should be positive. The larger the value of the TBCs, the greater the difference between the means (or the standard deviations), and therefore the better the contrast enhancement. On the other hand, a negative TBCs value indicates that the contrast of the enhanced image is worse than the original image.

6.6.3 Target to Background Contrast Ratio Using Entropy (TBCe)

The third measure proposed in [5] is *target to background contrast ratio using entropy* (TBCe). The basis of this measure is similar to TBCs, except that image entropy is used here. The entropy of all the pixels in a region is computed as [54, p. 42]

$$\varepsilon = - \sum_{i=0}^{L-1} f(i) \log_2[f(i)] \quad (6.49)$$

where L is the number of grey levels in the histogram of that particular region, and $f \in [0, 1]$ is the normalized frequency or probability of the occurrence of each grey level, i.e.,

$$f(i) = \frac{h(i)}{\sum_{i=0}^{L-1} h(i)} \quad (6.50)$$

where h is the histogram of the region.

In this contrast enhancement measure, the entropy of the target region is expected to be reduced after image enhancement. Thus the target to background contrast ratio using entropy (TBCe) is computed as

$$\text{TBCe} = \frac{(\mu_T^E/\mu_B^E) - (\mu_T^O/\mu_B^O)}{\varepsilon_T^E/\varepsilon_T^O} \quad (6.51)$$

where ε_T^O and ε_T^E are the entropy of the target regions of the original and enhanced images respectively; and μ_B^E , μ_B^O , ε_T^E and ε_T^O are all non-zero.

Like the DSM and TBCs, if the image contrast has been enhanced, the TBCe should be positive. The larger the value of TBCe, the greater the difference between the means (or the entropies), and therefore the better the contrast enhancement. On the other hand, a negative TBCe value indicates that the contrast of the enhanced image is worse than the original image.

6.6.4 Combining the Quantitative Measures

After obtaining the three quantitative measures described above, they can be combined to produce the *final score* for contrast enhancement. First, each measure is normalized by linear scaling within the range $[0, 1]$, i.e., the minimum measured value is mapped to 0 and the maximum measured value is mapped to 1. Then the final score is computed by taking the *average* of all three normalized measures. Consequently, the final score that represents the overall performance of contrast enhancement is also in the range $[0, 1]$. This combining method was derived by the author, and it is *not* described in [5].

6.7 Experimental Setup

The modified texture demodulation algorithm that we have developed, and the two existing techniques, morphological enhancement and histogram equalization, were used to enhance all the masses from the MIAS database, on a total of 41 images with 22 circumscribed masses and 19 spiculated masses. The original $50 \mu\text{m}/\text{pixel}$ images were reduced in resolution to $400 \mu\text{m}/\text{pixel}$ by averaging 8×8 pixels to one. The original 8-bit grey levels $[0-255]$ was retained. The results of the three enhancement techniques were then evaluated by the three quantitative contrast measures.

For the quantitative contrast measures, target and background regions have to be defined on each enhanced image. However the exact boundaries of the masses are not provided with the MIAS database. The location and size of each mass is instead represented by a circular region which was drawn by radiologist. As a result, the target region T is approximated by such circular region, and its background B is defined as a larger concentric circle. Additionally, the areas of T and B should be the same. Thus, if the

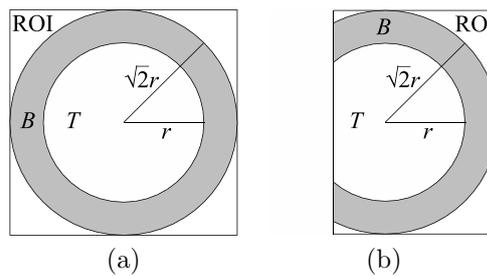


FIGURE 6.10: (a) Target and background regions, T and B respectively, defined in the ROI of a mass that is fully included on the mammogram. The areas of T and B are equal. (b) Target and background regions defined in the ROI of partially included mass.

radius of T is r and the radius of B is λr , then

$$\begin{aligned} \text{Area}(T) &= \text{Area}(B) \\ \pi r^2 &= \pi(\lambda r)^2 - \pi r^2 \\ 1 &= \lambda^2 - 1 \\ \lambda &= \sqrt{2} \end{aligned}$$

Therefore, the radius of B is $\sqrt{2}r$, as shown in Fig. 6.10.

The ROI of each mass was extracted from the whole image as a *square* of side length $2\sqrt{2}r$ pixels, centred on the target and background regions. Sometimes the mass was partially missing at the image edge and so was the pre-defined circular region. In this case the ROI was selected as a rectangular region adjacent to the image edge (see Fig. 6.10). The three quantitative contrast measures were then performed on the target and background regions inside the ROI.

6.8 Results

The results of the three quantitative contrast measures for all 41 images are shown in Fig. 6.11. The three measures for each image are graphically represented by three bars of different colours. For all three measures, the higher the value the better the performance. A negative value indicates that the contrast of the enhanced image is worse than that of the original image. In all three measures, the performance of texture demodulation was either comparable or better than that of morphological enhancement and histogram equalization. For DSM, texture demodulation and morphological operation were found to have around the same score for some images, and they both generally outperform histogram equalization. For TBCs and TBCe, the values for texture demodulation were predominantly greater than the other two, clearly demonstrating its excellent performance in contrast enhancement.

As mentioned before, the contrast measures were normalized from 0 to 1 and the aver-

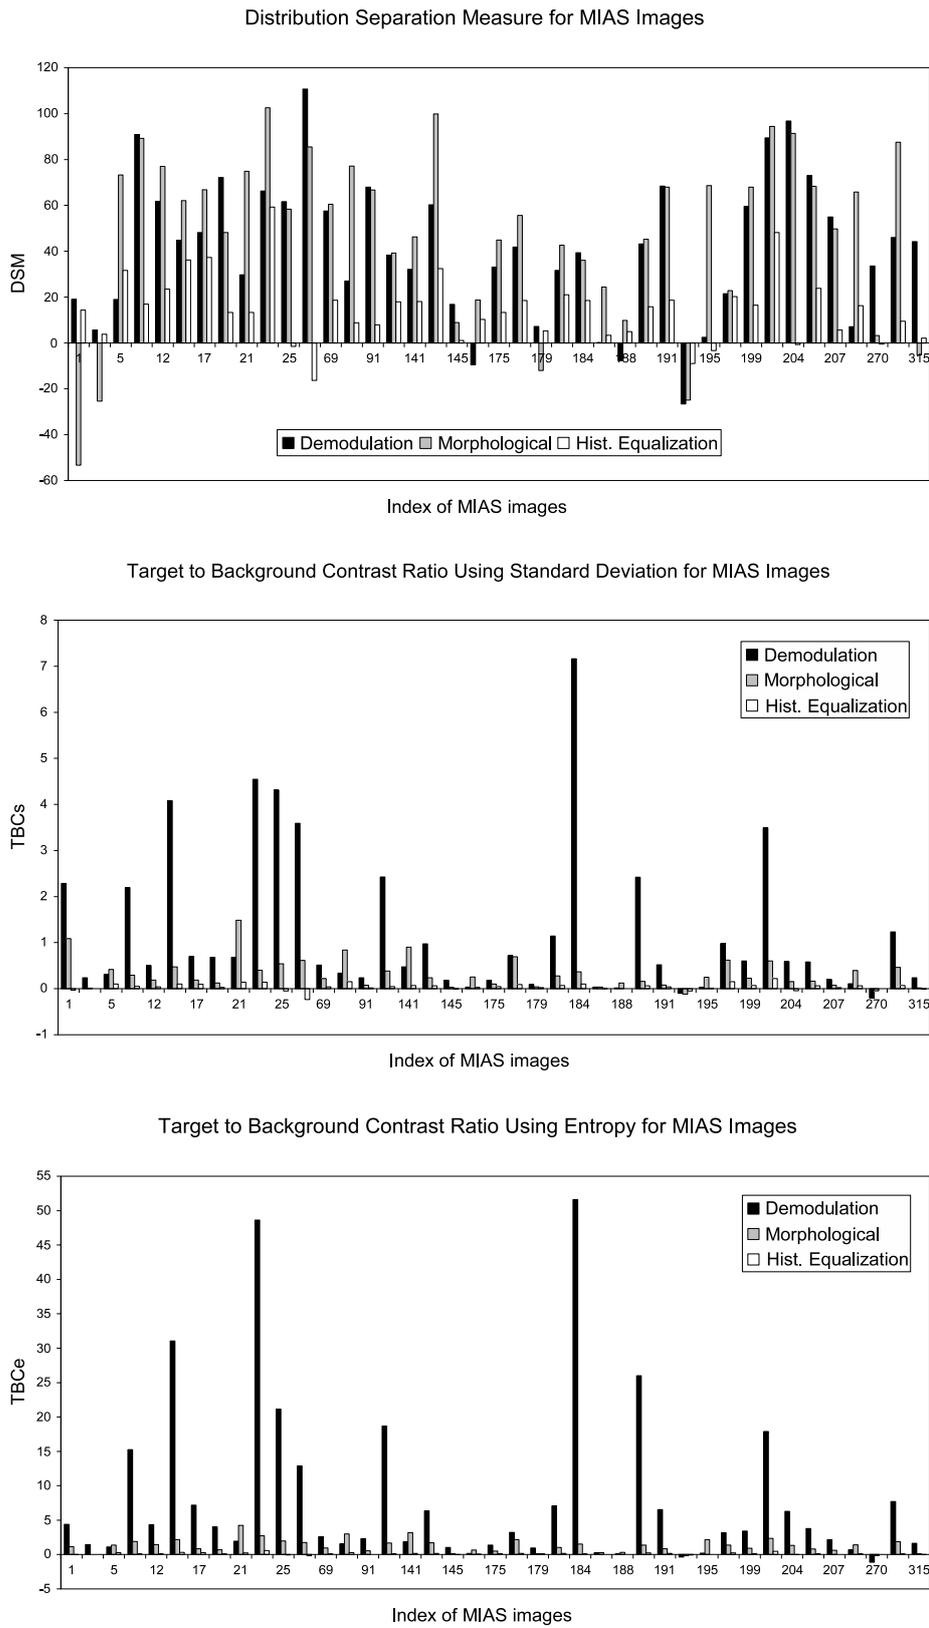


FIGURE 6.11: The three contrast enhancement techniques, texture demodulation, morphological operation and histogram equalization, are compared using three quantitative measures.

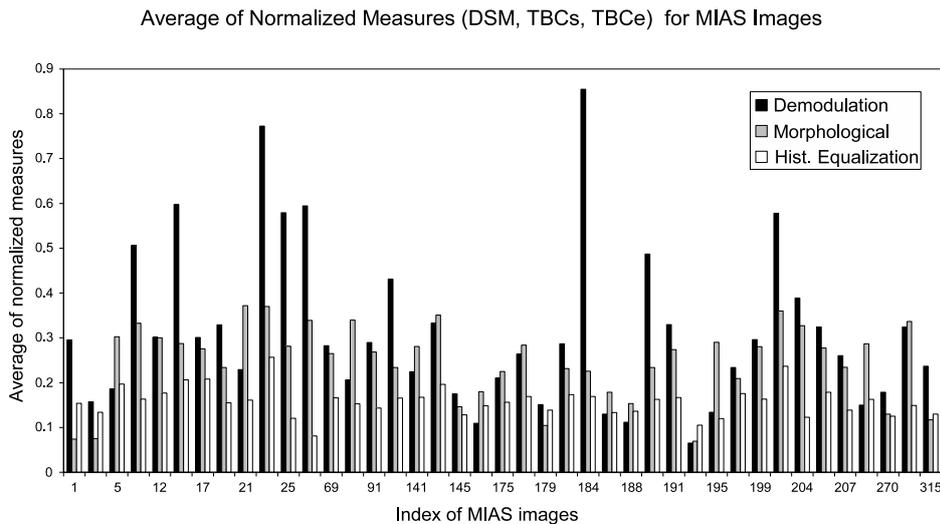


FIGURE 6.12: The average of the three normalized contrast measures (DSM, TBCs and TBCe), taken as the final score for the enhancement performance.

ages of the normalized measures were taken as the final scores for performance evaluation. The final scores of all the tested images are shown in Fig. 6.12. The higher the score the better the performance of the enhancement technique. The mean of the final scores for texture demodulation is 0.31, for morphological enhancement is 0.25, and for histogram equalization is 0.16.

One-tailed z -test was applied to the final scores to determine their differences statistically. The value of z was 2.20 and the P-value was 1.40×10^{-2} , when comparing texture demodulation and morphological enhancement. The same figures were 5.54 and 1.51×10^{-8} , respectively, when comparing texture demodulation and histogram equalization. We therefore conclude that the final scores of texture demodulation are statistically higher than those of the other two methods, to the significant level of 2%.

The number of highest scores achieved by each enhancement technique is listed in Table 6.1. A *highest score* is defined here as the largest measured value (in DSM, TBCs, TBCe or final score) among the three contrast enhancement techniques for a particular image. Texture demodulation achieved the highest scores for 16/41 images in DSM, 30 in TBCs and 31 in TBCe. Texture demodulation was also found to have the highest final scores on 27/41 images, while morphological enhancement achieved 13 and histogram equalization only achieved 1.

The Pearson correlation coefficients between the three contrast measures as well as the final scores are shown in Table 6.2. We found that DSM is not strongly correlated with TBCs and TBCe. This is expected since the formulation of DSM is so different from the other two—DSM is a distance measure but TBCs and TBCe are two ratios. Also they focus on different aspects of contrast enhancement. On the other hand, TBCs and TBCe were found highly correlated at 0.93, due to the fact that the only difference between these two measures is the use of standard deviation in one formula and entropy in the other.

TABLE 6.1: The Number of Highest Scores Achieved by the Enhancement Techniques

Measures	Demodulation	Morphological	Hist. Equalization
DSM	16	24	1
TBCs	30	9	2
TBCe	31	8	2
Final Score	27	13	1

TABLE 6.2: Pearson Correlation Coefficients between DSM, TBCs, TBCe and the Final Score

	DSM	TBCs	TBCe	Final Score
DSM	1.00	0.33	0.28	0.72
TBCs	0.33	1.00	0.93	0.88
TBCe	0.28	0.93	1.00	0.86
Final Score	0.72	0.88	0.86	1.00

Furthermore, because TBCs and TBCe represent $2/3$ of the final score, the correlation between each of these two measures and the final score is higher than that between DSM and the final score.

In addition, the processing time for the three enhancement techniques was collected. The average processing time for texture demodulation was 1.31 seconds per ROI. The actual processing time for each mammogram depends on the size of the ROI; the larger the mass the longer the time. The average processing time for morphological enhancement was much more at 475.40 seconds per image, because the top-hat operations were repeated on every pixel of the whole image. The actual processing time for each mammogram depends on the sizes of the structuring elements and the image dimensions. The average processing time for histogram equalization was the least at 0.36 second per image, since its computational complexity is far less than the other two methods. The actual processing time for each mammogram is quite consistent; it depends on the size of the breast region.

Finally, some enhanced images of the masses together with their final scores are shown in Fig. 6.13 for visual comparison.

6.9 Discussion

The modified demodulation algorithm was found to have very good sensitivity for highlighting circumscribed or spiculated masses on mammograms. With an accurate ROI provided, the contrast of the targeted mass and its background could be greatly enhanced, showing a clear outline of the mass.

One feature of our proposed method is that the background of the enhanced image is reduced to zero intensity, by image subtraction, with the surrounding glandular tissue suppressed. Its effect is even greater than the background correction used in morphological enhancement. The pure black background significantly improved the visualization of the

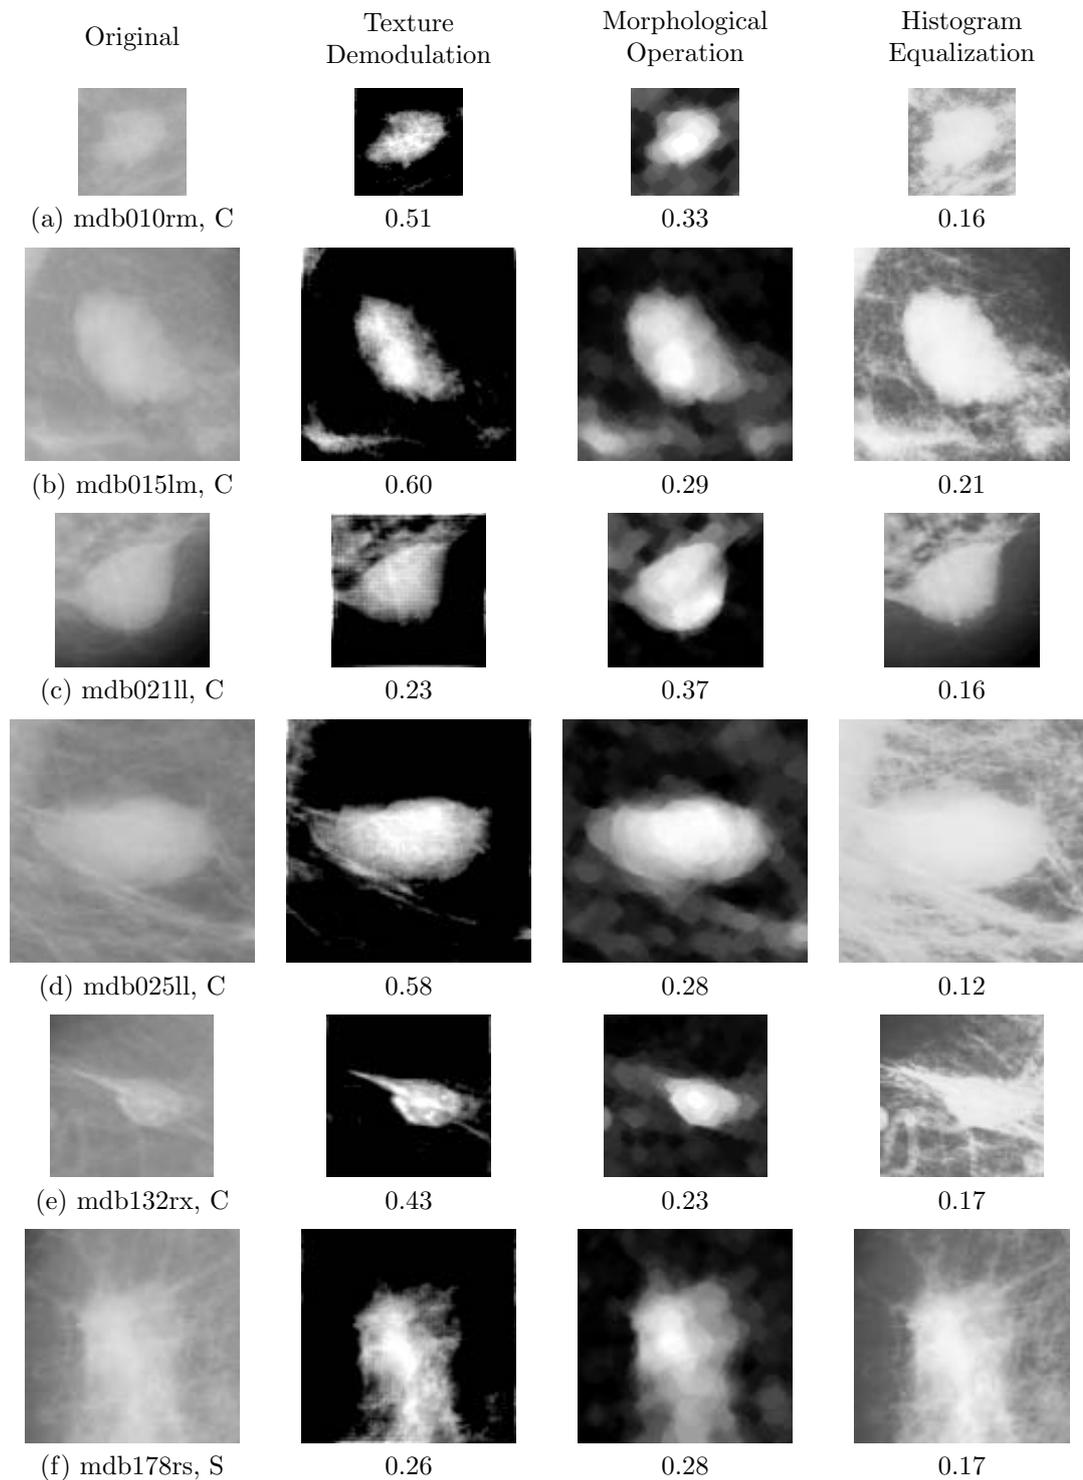


FIGURE 6.13: Some enhanced masses from the MIAS images. All images are on the same scale except for (h). The letter C stands for circumscribed, and S for spiculated. The number showing below each enhanced image is the final score, i.e., the average of the three normalized contrast measure.

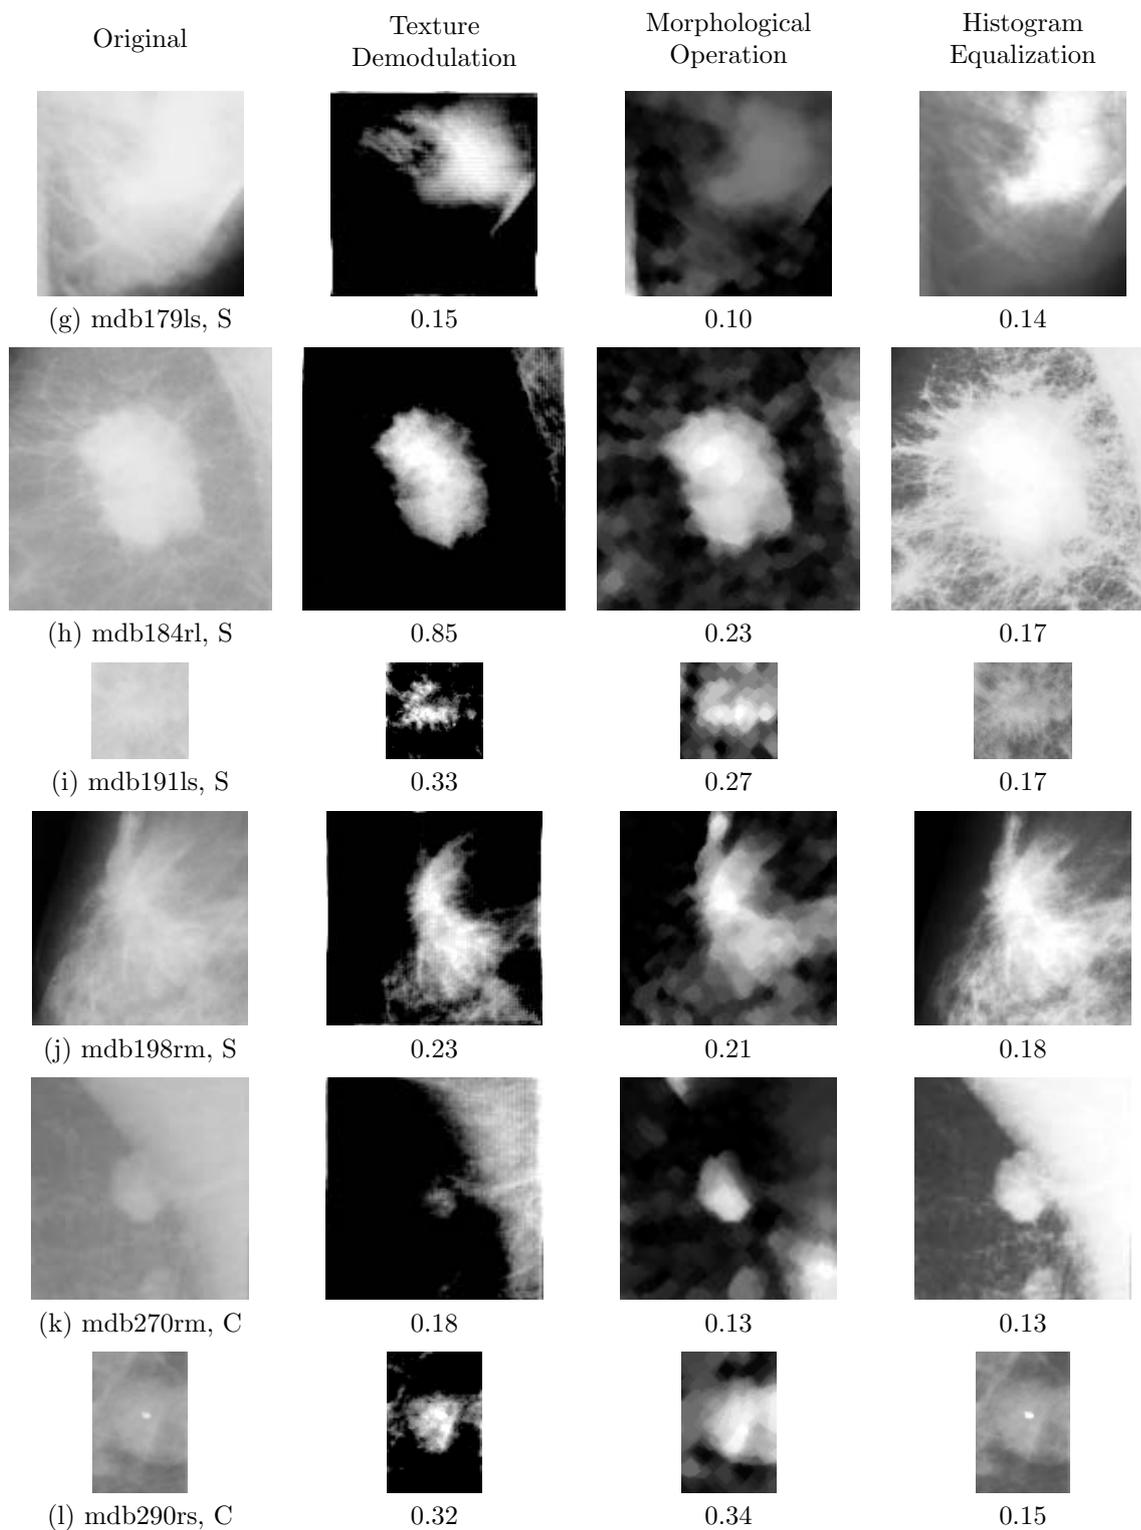


FIGURE 6.13: (Continued).

mass (see Fig. 6.13), and this is surely a very useful feature in mass segmentation. Masses superimposed by dense tissue are also highlighted as a solid pattern with most of the surrounding tissue removed, e.g. Fig. 6.13(g) and (j).

The enhancement effect of texture demodulation was maximized on images when the background intensity surrounding the mass was consistent. As shown in Fig. 6.13(c), only sub-optimal enhancement could be achieved on the mass that was near the breast edge, where the lower part of the background on the original image was darker. For the same reason, the demodulation method also failed to highlight masses that were superimposed by the pectoral muscle or that were partially missing at the image edge (see Fig. 6.13(k) and (l)). In these cases of uneven background intensity, morphological enhancement was found to give better results.

We have noticed that for a few images the quantitative contrast measures were negative, indicating a decrease in image contrast. However these negative measures were mostly caused by the inaccuracy of target and background regions defined in the ROI, and they did not necessarily reflect the change in actual image contrast. As explained in the previous section, the target region is only a circle enclosing the mass but not the exact boundary of the mass. Therefore some image background may have been included in the target region. Furthermore, some of the circles given in the MIAS database are clearly either not centred on the mass, or are larger than the visually apparent size of the mass, e.g. Fig. 6.13(g) and (k). These inaccurate target regions also affected the measurement of image contrast.

The only input parameters in the texture demodulation algorithm are the location and size of the mass. The ROI of the mass has to be extracted first, before the demodulation process is applied. This is very different from the morphological enhancement and histogram equalization which are applied to the whole image or breast region before extraction of the ROI. Unlike morphological enhancement, that was designed for lesion site selection [53], *texture demodulation is not suitable for mass detection where the mass location is unknown*. Instead, texture demodulation could be used to enhance suspicious masses that are already found by radiologist or computer, so that further processing such as mass segmentation and classification can be carried out more efficiently.

When comparing texture demodulation with morphological enhancement, we can see several advantages of our method. Despite the mathematical complexity of the demodulation algorithm, the processing time was reduced significantly when fast Fourier transform (FFT) was used to compute the frequency components. Also, the enhanced images produced by texture demodulation are not as blurred as those produced by morphological enhancement. Texture demodulation gives a clearer outline of the mass with some of the spicules around the mass retained, and may be useful in mass classification later on. Furthermore, the background texture surrounding the mass is more thoroughly removed using texture demodulation, so that standard deviation and entropy of the background is usually very small. This is reflected by the high contrast measures, TBCs and TBCe, of our method shown in Fig. 6.11.

6.10 Suggestions for Further Work

Although the experimental results have shown that the modified demodulation algorithm was successful in highlighting masses with high sensitivity, it did not appear to have good specificity. When it was applied to the ROIs of normal breast tissue where no mass lesions exist, the glandular tissue was itself highlighted rather than suppressed. As a result, some irregular and curly shapes were formed on the enhanced images of normal tissue. This effect was more significant for dense breast tissue (see Fig. 6.14).

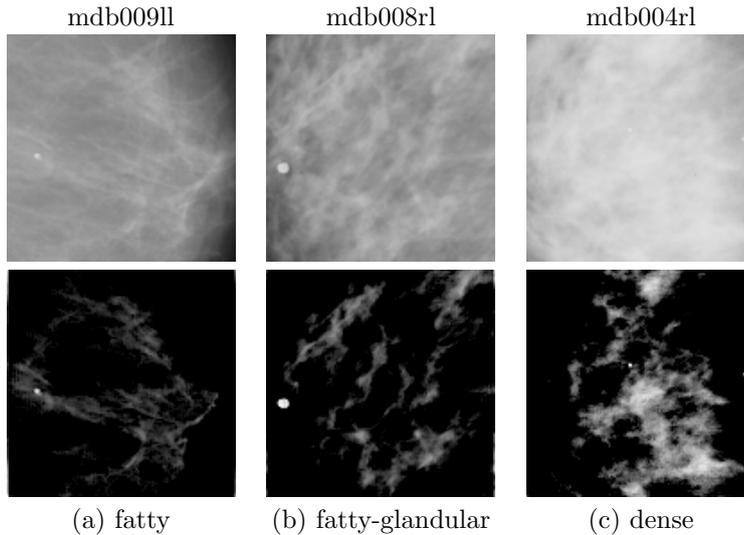


FIGURE 6.14: Original ROIs of normal tissue (above) and enhanced images by texture demodulation (below). The size of each ROI is 128×128 . The densities of breast tissue are (a) fatty, (b) fatty-glandular, and (c) dense.

However, the appearance of these highlighted normal tissue fractions are so different from the appearance of masses that it is not difficult to visually distinguish them from the enhanced images. The pattern of highlighted mass lesions is usually solid and circular in shape; whereas the pattern of enhanced normal tissue is likely to be incoherence and non-circular in shape.

In computer-aided diagnosis, texture demodulation may be used to highlight part of the breast tissue that contains a suspicious mass for the radiologist to make the final decision. We suggest that a shape analysis operator may be developed to distinguish circular and non-circular shapes and hence to improve the specificity of automatic mass detection. Furthermore, the highlighting of normal breast tissue may prove useful for modelling and subtracting out portions of the breast tissue for dense breasts, thus enhancing visualization.

The texture demodulation algorithm has been tested on both circumscribed and spiculated masses. We observed that some of the fine structures, like spicules, around the masses were removed on the contrast enhanced images. Due to the removal of these fine structures, the proposed method may not be suitable for mass classification. Further development is therefore required to preserve the spicules of masses on the enhanced images.

Due to the small number of sample images, only quantitative measures [5] were used to evaluate the results in this chapter. For future work, the contrast enhancement methods described in this chapter should be further evaluated using receiver operating characteristic (ROC) analysis [50, 51]. By comparing the radiologists' performance in detecting masses on various contrast enhanced images, the usefulness of the texture demodulation algorithm could be determined.

6.11 Conclusions

In this chapter we have summarized the AMPM demodulation transform introduced by Daugman and Downing, and we have extended it to multi-carrier demodulation by partitioning the Fourier spectrum. Both the single- or multi-carrier demodulations were used to simplify the texture of some synthetic images, but they failed to simplify the texture of natural images, due to the fact that the Fourier spectrum of this kind of image is far from being well concentrated; the predictive power of the carrier is also accordingly low. AMPM demodulation with single or multiple carriers is therefore not very useful in the texture analysis of mammograms, which are also spectrally less concentrated.

However, we discovered that the AM component of such demodulation could be useful for image enhancement and hence we modified the existing demodulation transform. The modifications include: (1) zero-frequency carrier wave; (2) partial reconstruction; (3) image subtraction; and (4) contrast stretching. The modified algorithm was found to highlight mammographic masses with good contrast enhancement.

Two other contrast enhancement techniques, morphological operations [53] and histogram equalization, were implemented and compared with the new texture demodulation method. They were tested on 41 circumscribed and spiculated masses taken from the MIAS mammograms. The results were then evaluated using three quantitative measures suggested in [5], and these measures were combined to derive the final scores of contrast enhancement. The mean of the final scores for texture demodulation was 0.31, for morphological enhancement was 0.25, and for histogram equalization was 0.16. On average, the performance of texture demodulation was better than that of the other two techniques.

The texture demodulation method could be used to highlight masses on mammograms with good sensitivity. A significant feature is the pure black background produced on enhanced images. However, the method does not have good specificity as it also highlights some of the dense normal tissue. This shortcoming could be improved by the use of shape classification. Nevertheless, it may prove useful in the analysis of dense mammographic parenchyma.

This concludes all the algorithms developed for the applications of mammographic segmentation and analysis. The next chapter is the conclusion of this thesis.

Chapter 7

Conclusion

In this thesis, a mammogram analysis system has been presented for the attribute-driven segmentation and analysis of mammograms. In this system, mammogram attributes are derived progressively from low-level to high-level in a cyclic procedure. An attribute dependence graph has been drawn for the acquisition of mammogram attributes; a standard format has been specified for the normalization of mammogram attributes; and an XML database with the support of XML Schema has been built for the storage of mammogram attributes. These three major components of the system are discussed in Chapter 3, and they constitute the first part of the original work of this thesis.

New algorithms have been proposed for mammogram segmentation and analysis, following the framework described in Chapter 3. These algorithms were implemented as computer applications and integrated into the mammogram analysis system. They include:

1. automatic pectoral muscle segmentation (Chapter 4);
2. adequacy assessment of positioning and exposure (Chapter 5); and
3. contrast enhancement of mass lesions (Chapter 6).

The detail of each of these three applications constitutes the second part of the original work of this thesis.

7.1 Summary of Research Outcomes

The research outcomes of this thesis are summarized in the following sections:

7.1.1 Mammogram Analysis System

In Chapter 3, a solid framework is provided for the attribute-driven segmentation and analysis of mammograms. A set of mammogram attributes are defined in three categories: image data, visual features and analytical information. All the knowledge about the image that we have at each stage of image processing is represented by the whole set of

mammogram attributes accumulated in the system. As the set of attributes expands, the knowledge about the image also increases. Hence more complex algorithms can be developed using the existing mammogram attributes to generate even more accurate results. This is the fundamental concept of the entire framework.

The mammogram attributes were organized in an attribute dependence graph with priorities assigned to them, the sequential acquisition order of mammogram attributes was then determined. The dependence graph also showed that all the adequacy measures are independent of each other, so they can be acquired in any order. By changing the attribute priorities in the graph, we have shown that adequacy assessment and mass lesion detection can be performed independently of each other, but they both depend on the anatomical features given by mammogram segmentation.

A standard format has been specified for the normalization of mammogram attributes. This format was required so that information exchange between applications could be carried out easily. All the common aspects of attribute normalization have been considered during the process of standardization for mammogram attributes.

A mammogram attribute database has been built using XML for data storage, and XML Schema for data definition. An open-ended and application-independent element structure of the mammogram attributes was specified in the database schema. Mammogram attributes were divided into four categories: breast anatomy, adequacy, abnormality and texture. Each of these categories can be extended easily by attaching new elements at the end of the element tree. By comparing the dependence graph and the element structure of mammogram attributes, we have shown that the attributes in the former graph are tightly coupled, whereas the attributes in the latter one are loosely coupled. So the attribute acquisition order is much more restrictive than the attribute insertion order to the database. This is a characteristic of the attribute-driven system.

7.1.2 Automatic Pectoral Muscle Segmentation

In Chapter 4, a new method has been developed for automatic segmentation of the pectoral muscle on mammograms of the MLO view. In this method, the pectoral muscle is first approximated by a straight line using iterative threshold selection and gradient test. The position of the straight line is validated according to the end-points of the breast border. This straight line is then refined into a curve using iterative cliff detection, in which the pectoral edge is located by fitting intensity profiles to a sigmoid model. Finally, an enclosed region is extracted to represent the pectoral muscle. This method has been tested on 322 mammograms from the MIAS database, and on average 83.9% of the curve segmentations were rated adequate or better by two radiologists. The results have also shown that the segmentation accuracy was significantly improved after refining the straight line into a curve. The straight line approximation and the curve delineation of the pectoral muscle were both used in the automatic assessment of breast positioning.

7.1.3 Adequacy Assessment of Positioning and Exposure

In Chapter 5, it has been demonstrated that the adequacy of breast positioning and exposure of mammograms may be automatically assessed by computer. The purpose of breast positioning assessment is to ensure that the optimal amount of breast tissue is included, while the purpose of exposure assessment is to guarantee that most of the breast tissue is visualized in the useful optical density range. Algorithms have been derived for these two assessments, based on several quality criteria that are stated in the radiological literature. In these algorithms, some anatomical features extracted from mammogram segmentation are used, including the breast border, the nipple and the pectoral muscle. Thus the accuracy of adequacy assessment is affected by the accuracy of the preceding segmentation steps. This is reflected in the dependence graph of mammogram attributes that adequacy measures are dependent on anatomical features. The algorithms have also been tested on 322 mammograms from the MIAS database. Only 4.7% of these images were found adequate in positioning and 53.4% exhibited proper exposure. We therefore conclude that very few mammograms from the MIAS database are diagnostically adequate. The results also showed that breast tissue exclusion, overexposure and poor positioning could be a source of errors for mammogram segmentation algorithms.

7.1.4 Contrast Enhancement of Mass Lesions

In Chapter 6, a novel technique has been developed to highlight mass lesions on mammograms. It is a modification of the AMPM demodulation transform (developed by Daugman and Downing), based on texture demodulation, partial reconstruction and image subtraction. This contrast enhancement technique was found to have good sensitivity for highlighting masses on mammograms, but insufficient specificity, as it also enhances some of the dense normal tissue. One significant feature of this technique is that it increases contrast by producing a pure black background in the enhanced image. The new contrast enhancement technique was compared with two existing methods: morphological enhancement and histogram equalization. All three techniques were applied to 41 mammograms with circumscribed and spiculated masses from the MIAS database. The enhancement results were then evaluated using three quantitative measures, and these measures were combined into one final score of contrast enhancement for each image. The mean of the final scores was 0.31 for the new technique, 0.25 for morphological enhancement, and 0.16 for histogram equalization, indicating that the new technique had the best performance on contrast enhancement.

7.2 Suggestions for Further Work

A hierarchical framework has been proposed in this thesis for the development of a mammogram analysis system. The research described in this thesis forms part of the proposed system. While the segmentation algorithms may be further refined, the work on the ad-

equacy assessment and lesion enhancement hold great potential for future development and applications.

Other applications of mammographic image analysis, such as density classification, bilateral comparison, and other abnormality detection, may also be included in the mammogram analysis system. The mammogram attributes described in this thesis by no means cover all the possible attributes that can be derived from a mammogram. In fact, they form only an initial subset of useful mammogram attributes. More attributes should be included in the future as new applications are developed and integrated into the system.

An XML database was constructed for the storage of mammogram attributes. The syntax of XML is straight forward and simple, and its hierarchical structure is particularly useful in describing medical data. Currently there are efforts directed at the development of mammographic ontology, which is similar to our attribute database. Several modelling languages have been suggested in [84] for defining the Breast Cancer Imaging Ontology (BCIO), e.g., Resource Description Framework (RDF), RDF Schema, DARPA Agent Markup Language and Ontology Inference Layer (DAML+OIL), and Ontology Web Language (OWL). These languages provide a means to formally define the semantics of a specific domain. All of them conform to the syntactic constraints of XML. They should be investigated in the future work and may be used to replace the current XML schema for mammogram attribute specification. In addition, BCIO was created using the terminology of the Breast Imaging Reporting and Data System (BI-RADS), which should be included into our set of mammogram attributes.

At the current stage, most of the anatomical features on a mammogram can be identified and segmented out automatically. The algorithms used in the work of this thesis, for segmenting the breast border, nipple and pectoral muscle, are accurate enough for most of the subsequent image analysis applications. Further work has to be carried out on segmenting the fibroglandular tissue on mammograms. This anatomical feature is particularly important in breast positioning assessment and breast tissue density classification. One way to extract this feature is to use Gaussian-mixture models and the Expectation-Maximization (EM) algorithm as proposed in [17]. Another possibility would be the use of edge detector and texture segmentation as suggested in [19].

Because the method proposed in this thesis for pectoral muscle segmentation refines a straight line into a curve, the curve segmentation could be improved by giving it a more accurate straight line approximation. Expanding window thresholding proposed in [32] may be used to replace the iterative threshold selection method described in this work, since the former one was evaluated to be more accurate [32]. Furthermore, a texture edge detector may be combined with iterative cliff detection (which mainly detects intensity edges) so that the texture edge at the lower portion of the pectoral muscle can be more accurately delineated.

For the work on adequacy assessment, two quality criteria of positioning related to the posterior glandular tissue were not evaluated, because the boundary of the fibroglandular

region had not been outlined, as accurate automatic segmentation of fibroglandular tissue was not available. Besides the adequacy of positioning and exposure, other quality attributes should be included in the assessment. Some quantitative attributes, such as contrast, compression, sharpness and artefacts, are likely to be measured on individual mammograms using the computer. Automatic adequacy assessment should also be extended to evaluate mammograms of the CC view, since it is a standard mammographic view as well. Once this is achieved, the posterior nipple lines on the MLO and CC views should be compared for breast positioning assessment.

Contrast enhancement of mass lesions is only part of the whole process of lesion detection on mammograms. Unlike other enhancement techniques, the texture demodulation method proposed in this thesis requires the ROI of the mass to be extracted first. Thus methods for suspicious lesion site selection, such as [53], should be investigated in the further work. The two subsequent steps of mass enhancement are lesion segmentation and classification; these should be further investigated as well. In addition, to improve the specificity of the proposed texture demodulation method, a shape analysis operator may be used to classify circular and non-circular patterns highlighted on the enhanced image, so that mass lesions can be distinguished from normal breast tissue.

Finally, it is hoped that the generic abstract system model presented in Appendix A of this thesis contributes not only to the field of mammographic image analysis, but also find application in other domains as well.

Appendix A

Abstract System Model

A.1 Introduction

The problems of image processing in the real world are usually very complex. They are unlikely to be solved by a single, simple, generic technique. Therefore, in many areas of image processing, most of the problems have been successfully solved by a combination of different techniques, that were specifically designed and developed for a particular class of images. Although these techniques are not generic, we believe that from the abstract point of view, there exists a generic framework that may be applicable to the segmentation and analysis of any class of images.

In this appendix, an *abstract system model* is introduced for automatic, attribute-driven segmentation and analysis of images. This generic model has been designed to integrate an arbitrary number of special-purpose image processing techniques into one system, in which the hierarchical segmentation process is driven by the attributes of the image. Different computer applications communicate with each other through the exchange of attributes, and the information about the image is centralized in an attribute database.

Note that the proposed abstract system model has not been fully evaluated but only tested on the segmentation and analysis of mammograms. Further development of this system model is outside the scope of this thesis, so it is left to future work.

A.2 Image Attributes

In an attribute-driven system for image segmentation and analysis, all the information that we have gathered from an image is represented by a set of *image attributes*. This set of image attributes can be considered, in the broad sense, as a *knowledge base* [55, p. 25–28], which interacts with image processing applications in the system. Image attributes are not only simple data items, but they are also often meaningful and useful information belonging to a specific domain. There is no generic set of attributes that is suitable for any class of images. Instead, a specific set of attributes has to be chosen carefully,

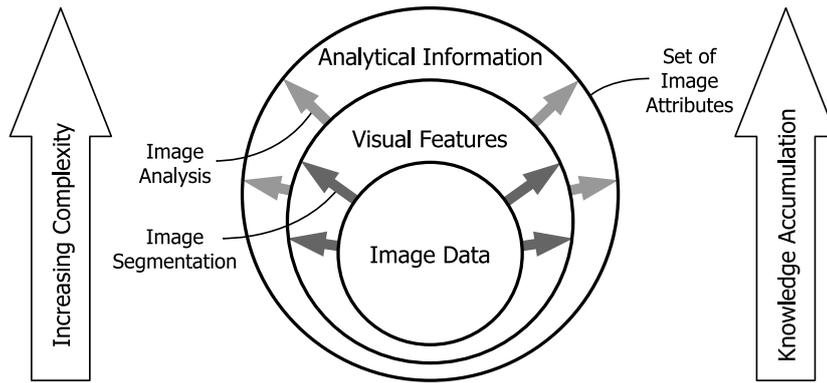


FIGURE A.1: The set of image attributes expands from image data to visual features when image segmentation is applied; and it expands from visual features to analytical information when image analysis is carried out. As the knowledge about the image increases, more complex algorithms may be performed on the image.

according to the purpose of analysis, for a particular class of images. Image attributes can be categorized into three types: *image data*, *visual features* and *analytical information*.

Image data includes the pixel values of the digital image that is being analyzed, and all the image properties, such as dimensions, resolution and bit-depth. Image data can usually be located in the header of an image file. When the image is first read by the system, the image data constitutes the *first* set of attributes, before any segmentation is performed on the image. This first set of attributes is most frequently accessed by subsequent applications in the system.

Visual features include the description of an object on the image, such as colour, size and shape. They are also used to describe the structural relationship among objects or within parts of an object, as well as the texture of different regions on the image. Visual features are usually extracted *directly* from the image by image segmentation and texture analysis. They are *low-level image attributes* that are usually obtained in the earlier stages of image processing. Based on these attributes, more refined segmentation and sophisticated image analysis can be performed in the later stages.

Analytical information consists of all the *high-level image attributes* that are derived from the low-level ones. It is the result of image analysis, which is normally carried out in the later stages of image processing. Analytical information may include, for example, the classification of an object, the clinical findings on a medical image, or the identification of a person on a photograph. External data that is not derived from the image, such as the age of a patient, can also be considered as part of analytical information.

During the process of image segmentation and analysis, image attributes are acquired progressively in a step-by-step fashion. From the attributes of image data, visual features are extracted by image segmentation. Analytical information is then derived from image analysis. With each successive stage of attribute acquisition, the knowledge that we have about the image is assumed to increase monotonically as more attributes are derived, leading to refinement of accuracy. We assume that there is no reduction in information,

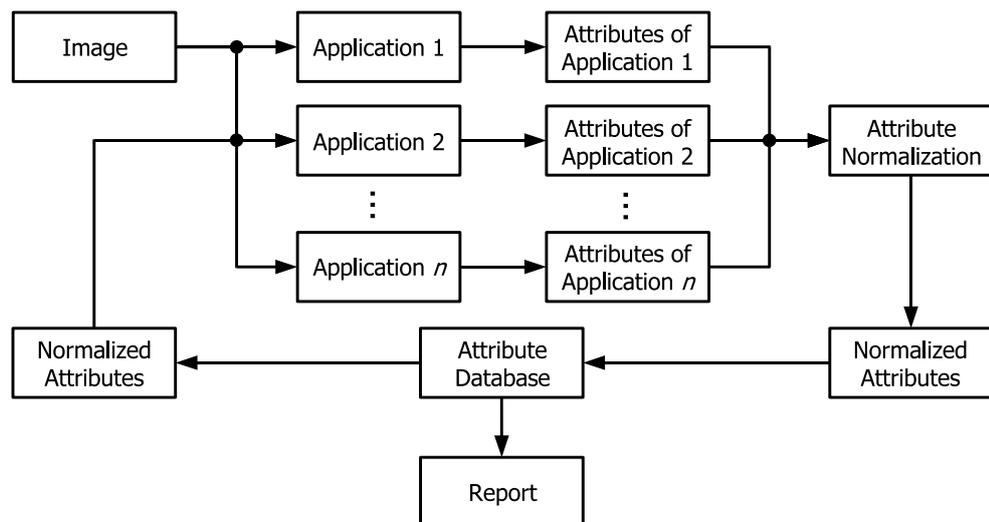


FIGURE A.2: A system model of attribute-driven image segmentation and analysis. The input to the system is an image and the output is a report. The n number of applications are executed in a pre-defined sequential order, from 1 to n , to extract attributes from the image. The image attributes are normalized and stored in an attribute database. They can be retrieved and used by any subsequent application.

and the correlations between the set of image attributes, the knowledge about the image and the complexity of algorithms are all positive. This fundamental concept of *knowledge accumulation* can be visualized as a growing set of attributes, illustrated in Fig. A.1.

Attributes are usually extracted automatically from the image using computer applications. Although these applications may operate with data formats that are incompatible with each other, attributes can still be exchanged between applications if they are *normalized* to a standard format. Normalized attributes are stored persistently in a database, and they can be retrieved from the database by any application in the system. The workflow of the system is described in Section A.3. The relationship between *attribute acquisition*, *normalization* and *storage* are discussed in Section A.4.

A.3 System Model

The abstract system model of attribute-driven image segmentation and analysis is illustrated in Fig. A.2. This system model is generalized so that it may be applied to any class of images. The input to the system is a digital image. The attributes of the image are extracted by a number of *applications*, which are the computer programs of specific segmentation and analysis techniques. These applications may be written in different programming languages and be run on different platforms. The attributes generated by each application could possibly be in a data format that is not compatible with other applications. Therefore it is necessary to *normalize* the attributes to a standard format. The normalized attributes are stored in an *attribute database* that is designed to be application- and platform-independent. An open-ended data structure is used in the database so that

it can be extended easily when more applications and attributes are added into the system.

Figure A.2 also illustrates the *cyclic* nature of attribute-driven segmentation and analysis. All the applications in the system are executed in a *pre-defined sequential order*. First, Application 1 reads the image and generates some attributes. These attributes are normalized and stored into the database. Then, Application 2 requests and retrieves the normalized attributes produced by Application 1 from the database. In Application 2, these normalized attributes are converted to the data format that is compatible with its implementation. Application 2 then performs further segmentation or analysis and generates additional attributes. The same procedure is repeated until all the subsequent applications in the system have been executed. Ultimately, an analytical report with annotated images can be produced using the attributes in the database, and presented to the user at a workstation.

A.4 Three Major Components

The system model for attribute-driven segmentation and analysis may be decomposed into three major components, which are:

1. attribute acquisition,
2. attribute normalization, and
3. attribute storage.

These three components are linked to each other to form a feedback loop in the system, as shown in Fig. A.3. The detailed description of these three components is given in Sections A.5, A.6 and A.7 respectively.

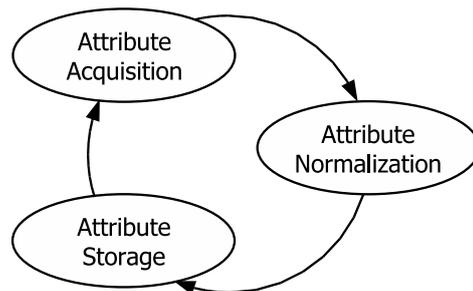


FIGURE A.3: The three major components comprising the system model of attribute-driven image segmentation and analysis.

In the system model, attribute acquisition is handled by a series of applications; attribute normalization is specified by the standard format of attributes; and attribute storage is achieved by the attribute database. Among these three components, attribute acquisition is the most difficult to implement, since the domain-specific applications usually involve complex, non-trivial image processing techniques. In contrast, attribute normalization and attribute storage are relatively easy to accomplish by specifying a standard

format for the attributes and by designing and building a database. Consequently, when developing the actual system for the analysis of a particular class of images, more time and effort are expected to be spent on the attribute acquisition component, rather than on the other two.

A.5 Acquisition of Attributes

As previously described in Section A.2, the acquisition of some attributes requires the input of other attributes from earlier stages. The high-level attributes from image analysis rely on the low-level attributes from image segmentation. The attributes themselves are sometimes closely coupled with each other, especially in hierarchical image segmentation. Since all the attributes required by an application have to be available when the application is executed, it is particularly important to work out the *data dependence* of all the possible attributes that would be stored in the database. To illustrate the dependence relationship between the attributes, an *attribute dependence graph* is therefore created.

A.5.1 Attribute Dependence Graph

In the dependence graph, image attributes are shown as *nodes* and the dependence relations are represented by *directed edges*. If attribute B is dependent on attribute A , then an edge is directed from node A to node B (see Fig. A.4(a)). An attribute may be dependent on more than one attribute; and a particular attribute may be required by more than one other attribute (see Fig. A.4(b)).

However, the dependence graph must be *acyclic*. It is not allowed to form a path in the dependence graph that starts and ends on the same node. For example, attribute B in Fig. A.4(c) can never be generated because it virtually depends on itself through the *cycle* $\{B \rightarrow C \rightarrow D \rightarrow B\}$. With the use of a dependence graph, we can keep track of the attributes available in the current stage and find out which attributes may be derived in the next stage of segmentation. Furthermore, most importantly, we can use the dependence graph to determine the *order of attribute acquisition*.

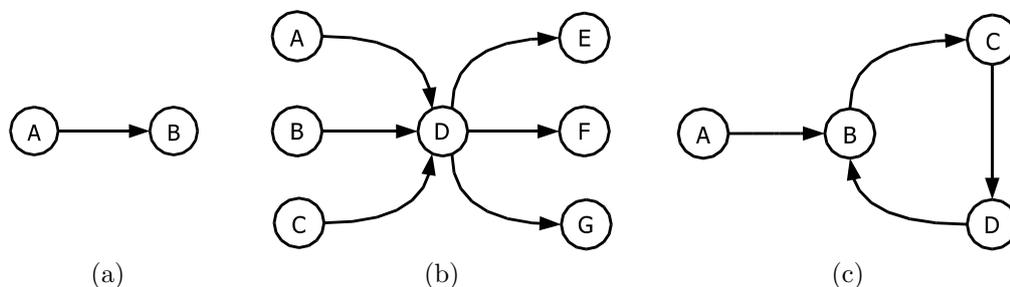


FIGURE A.4: Dependence graphs. (a) Attribute B depends on attribute A . (b) Attribute D depends on attributes A , B and C , and it is required by attributes E , F and G . (a) and (b) are both acyclic directed graph. (c) Cyclic dependence graph where attribute B virtually depends on itself through a cycle; attributes B , C and D can never be acquired by the proposed scheme.

A.5.2 Sequential Attribute Acquisition

In the attribute-driven image analysis system, the applications are executed sequentially. This means only one of the n applications is executed at each cycle and only the output attributes of that particular application are normalized and stored into the database. By examining the dependence graph of all the attributes and assigning a *priority* to each attribute, the *sequential order* of attribute acquisition can be determined. The algorithm that we use to assign priorities to attributes is called *attribute priority assignment*.

Attribute Priority Assignment

This algorithm starts from the attributes that we want to acquire, and iteratively searches for other attributes that are required in the attribute acquisition. In each iteration of this algorithm, attribute priorities are assigned to the nodes in the dependence graph, from the lowest priority to the highest one.

1. Let $G = (V, E)$ be the dependence graph with the node set V and the edge set E . G is an acyclic directed graph with two maps, *init*: $E \rightarrow V$ and *ter*: $E \rightarrow V$, assigning to every edge $e \in E$ an *initial node*, $\text{init}(e)$, and a *terminal node*, $\text{ter}(e)$ [60, p. 25]. So the edge e is directed from $\text{init}(e)$ to $\text{ter}(e)$. Each node $v \in V$ can be *labelled* with a *priority number*, which is a non-negative integer. The priority of a node v is given by $\text{pri}(v)$. Initially, all nodes in G are *unlabelled*, i.e., no priorities assigned. Also, the attributes that we want to acquire are called *desired nodes*, denoted by $W \subseteq V$. If all the attributes in the dependence graph are to be acquired, then they are all desired nodes, i.e., $W = V$. If we only want to obtain a single attribute (which may depend on other attributes), then only this particular attribute is selected as the desired node. In addition, let $U \subseteq V$ be a set of *flagged nodes*, $u \in U$, and $R \subseteq V$ be a set of *required nodes*, $r \in R$.
2. Label every desired node $w \in W$ with the lowest priority zero, i.e., $\text{pri}(w) := 0$. Flag all the desired nodes, so we set $U := D$. We also empty the set of required nodes $R := \emptyset$.
3. For each flagged node $u \in U$, search for all the nodes that u depends on, i.e., the required nodes, and put them in R . Also, assign a higher priority to the required nodes. Thus,


```

for all  $u \in U$  {
  for all  $e \in E$  {
    if (  $\text{ter}(e) \in U$  ) {
      add  $\text{init}(e)$  into  $R$ 
      if ( (  $\text{pri}(\text{ter}(e)) + 1 > \text{pri}(\text{init}(e))$  ) or (  $\text{init}(e)$  is unlabelled ) ) {
        set  $\text{pri}(\text{init}(e)) := \text{pri}(\text{ter}(e)) + 1$  } } }
      
```

4. If no required nodes are found for every flagged node, i.e., $R = \emptyset$ at the end of step 3, we have reached the nodes with the highest priority, so go to step 5 to finish the assignment. Otherwise set $U := R$, then set $R := \emptyset$, and repeat step 3.
5. The priority assignment is completed. The order of attribute acquisition for all the desired nodes is determined by the priorities of all the labelled nodes in the graph. The attribute with the highest priority must be acquired first. The desired attributes have priority 0 so they are acquired last. If there are two or more attributes with the same priority, they are not dependent on each other, therefore they can be acquired in any order. Any node that is not labelled with any priority is not required by the desired nodes; hence it is not necessary to include unlabelled nodes in the acquisition order.

Example 1: Single Desired Node

Suppose we have six attributes, A to F , organized in a dependence graph as shown in Fig. A.5. We want to acquire a single attribute F , so node F is the only desired node. Using attribute priority assignment, the attribute acquisition order is derived as $\{A \rightarrow B \rightarrow D \rightarrow E \rightarrow F\}$, and attribute C is not required by F .

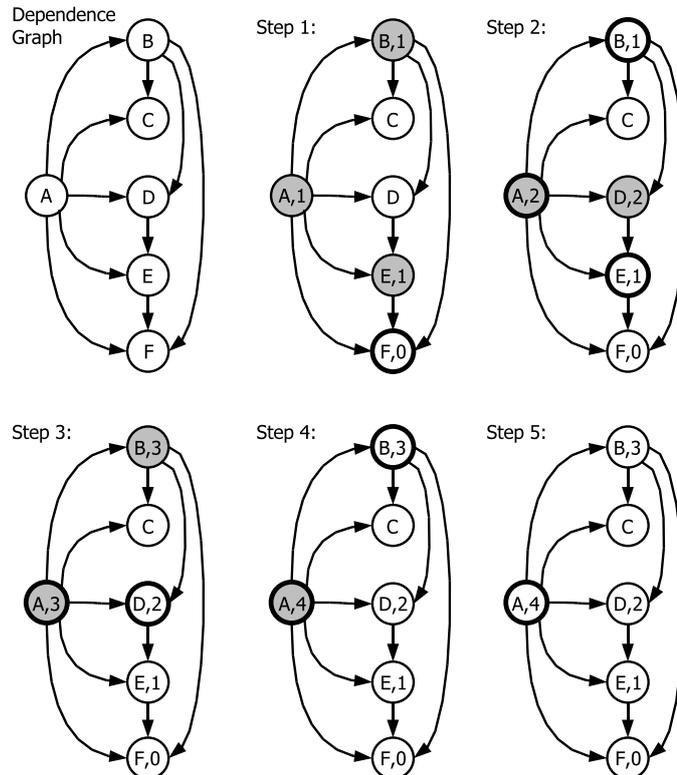


FIGURE A.5: A step-by-step illustration of attribute priority assignment in a dependence graph. The priority of the node is given as a number inside the circle. The single *desired* node, F , initially has the priority 0. In every step, the *flagged* nodes are represented by thickened circles; and the *required* nodes are represented by shaded ones.

Example 2: Multiple Desired Nodes

Another dependence graph is shown in Fig. A.6. All the attributes in the dependence graph are considered the desired nodes, hence they are all labelled with priority 0 and flagged in the first step. The priorities of all attributes may therefore be determined using attribute priority assignment. The attribute acquisition order is $\{A \rightarrow B \rightarrow D \rightarrow E \rightarrow C \rightarrow F\}$ or $\{A \rightarrow B \rightarrow D \rightarrow E \rightarrow F \rightarrow C\}$. In either case all the attributes will be generated correctly.

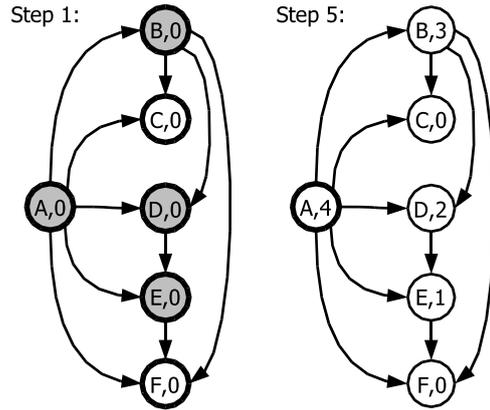


FIGURE A.6: The first and last steps of the derivation of attribute priorities in a dependence graph. All the attributes are desired nodes.

A.5.3 Parallel Attribute Acquisition

From the previous example, there are two ways to set the attribute acquisition order and both of them are valid. The two attributes with priority 0, C and F , are not dependent on each other. So in *sequential* attribute acquisition, C and F are generated one before the other, but the order does not matter. However, in *parallel* attribute acquisition, it is possible to compute C and F *at the same time*, as soon as the attribute with priority 1 is available.

In fact, it is not unusual to have more than two attributes with the same priority in the dependence graph. Since *all the attributes with the same priority can be acquired in parallel*, the corresponding applications can be executed in parallel as well. Simultaneous application processing can then be performed using multiple computers or CPU's. For example, in Fig. A.7, the priorities of all the attributes in the dependence graph are determined using the same method, attribute priority assignment. Among them, three attributes have priority 2 and three attributes have priority 1. Thus, B , C and D can be acquired in parallel; whereas E , F and G can also be generated simultaneously. If we use $//$ to denote the attributes that can be acquired in parallel, then the attribute acquisition order in this case is $\{A \rightarrow B//C//D \rightarrow E//F//G \rightarrow H\}$.

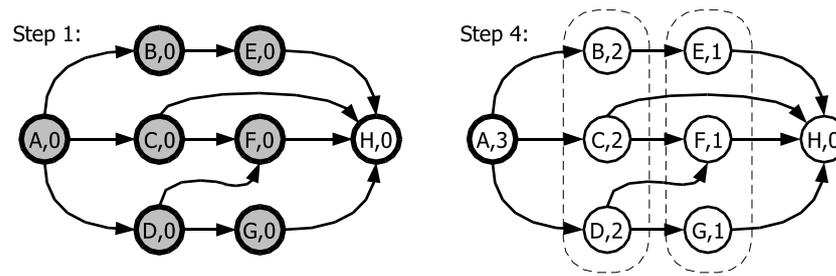


FIGURE A.7: The first and last steps of the derivation of attribute priorities in a dependence graph where parallel attribute acquisition is possible. All the attributes with the same priority can be acquired in parallel.

A.6 Normalization of Attributes

The derived attributes may vary in scale and format due to different segmentation and analysis approaches, and different application implementations. Attribute normalization is required to convert the attributes to a standard format, before they are inserted into the database. Normalized attributes ensure that data exchange between applications is carried out in a standardized manner. In this way, the application developer only needs to deal with one standard data format and not worry about the data format used by other applications. The application may convert or scale the normalized input attributes to its own data format during its execution, and then normalize its output attribute to the standard format.

When the format of image attributes is being standardized, the following aspects are usually considered:

- image orientation
- co-ordinate system
- Spatial resolution
- measurement units
- representations of object boundary

A.6.1 Image Orientation

The orientation of an image is very important in image segmentation because most of the algorithms require the image to be placed in a certain orientation, in order to identify and segment the targeted objects. Different segmentation algorithms may use various image orientations, and the attributes produced by these applications may refer to their own image orientations rather than the original one. However, the normalized attributes must refer to the *same* image orientation defined in the standard format. Without the knowledge of image orientation, the segmented boundaries cannot be reproduced correctly on the image and they become useless to other applications in the system.

An image may be oriented in eight different ways as shown in Fig. A.8. In the top row of the figure, the image orientations are (a) correct, (b) rotated 90° , (c) rotated 180° ,

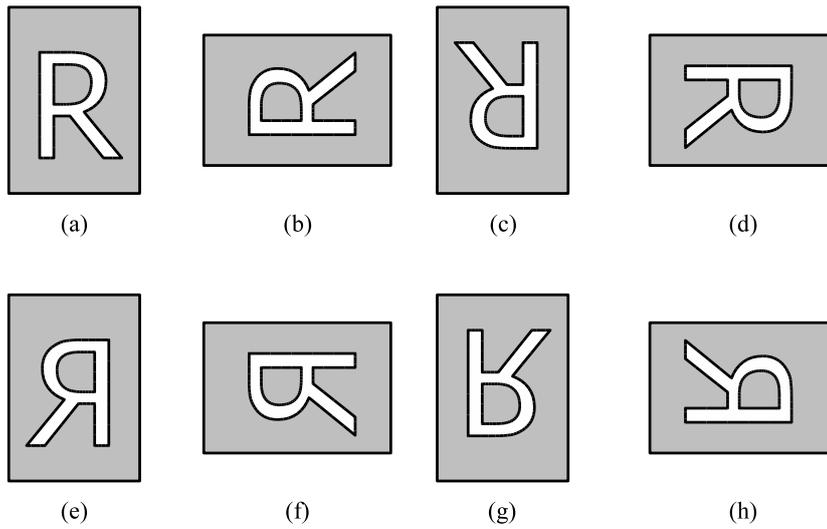


FIGURE A.8: Eight possible orientations of an image.

and (d) rotated 270° . In the bottom row, the images are flipped around the vertical axis, hence the image orientations are (e) flipped, (f) rotated 90° and flipped, (g) rotated 180° and flipped, and (h) rotated 270° and flipped. One of these orientations should be selected as the standard image orientation in attribute normalization. Only quadrantal changes of image orientation are considered here, and other image orientations are not recommended to be the standard image orientation.

A.6.2 Co-ordinate System

The x - y co-ordinate system used in the image domain should also be standardized for attribute normalization. Some possible co-ordinate systems used in image processing are shown in Fig. A.9. In the figure, (a) The origin is at the image top-left corner; x -axis is horizontal and y -axis is vertical. This co-ordinate system is commonly used to specify the pixels on a digital image, but it does not comply with the right-hand rule of the Cartesian co-ordinate system. (b) The origin is also at the image top-left corner, but x -axis is vertical and y -axis is horizontal. This co-ordinate system complies with the right-hand rule. (c) The origin is at the image bottom-left corner; x -axis is horizontal and y -axis is vertical. This co-ordinate system also complies with the right-hand rule. (d) The origin is at the centre of the image; x -axis is horizontal and y -axis is vertical. This is another co-ordinate system that complies with the right-hand rule.

Generally, any setting of the co-ordinate axes may be chosen as the standard co-ordinate system for attribute normalization. Different co-ordinate systems may be used in the algorithms of attribute extraction, but all output attributes must be normalized to the standard co-ordinate system. The conversion of the data values from one 2D co-ordinate system to the standard one normally just involves swapping the x and y values, and shifting the origin. If it is necessary, a 3D co-ordinate system or a polar co-ordinate

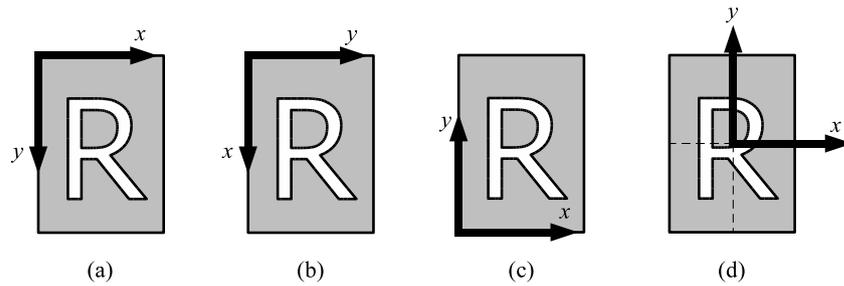


FIGURE A.9: Some possible co-ordinate systems used in image processing.

system may also be defined in the standard attribute format.

A.6.3 Spatial Resolution

The original image is often subsampled and low-pass filtered to a lower resolution when extracting large features on the image. The subsampling operation filters out noise and unwanted image texture. It improves the accuracy of feature extraction and significantly reduces the processing time. Consequently, the attributes extracted from an image at lower resolution are also reduced in scale, since the subsampled image is smaller in size than the original one.

In attribute normalization, all attributes are scaled back to the highest spatial resolution, which corresponds to the original image. For example, if the original image resolution is $50 \mu\text{m}/\text{pixel}$ and the subsampled image resolution is $400 \mu\text{m}/\text{pixel}$, then the scaling ratio is $400/50 = 8$. If the attribute to be normalized is a straight line described by two points, then the x and y values of each point are multiplied by the scaling ratio. The normalized straight line will be drawn correctly on the original image if it is reused. On the other hand, if the attribute to be normalized is a continuous boundary composed of an array of points, then after scaling, the resulting boundary is no longer continuous and the points are evenly spread. Some kind of interpolation must then be used to link up the points when the boundary is reproduced on the original image.

A.6.4 Measurement Units

When the Euclidean distance between two points on the image is calculated using the x - y co-ordinate pairs, the result is usually expressed in pixels. However, pixels are not the preferred measurement unit for the normalized attributes, since they depend on the image resolution. A measurement given in pixels can be converted to micrometres (μm) by multiplying it with the image resolution ($\mu\text{m}/\text{pixel}$).

Therefore, all attributes that involve measurement should be described by standard measurement units, which should be meaningful and suitable to the object that is measured. For example, the diameter of a lesion is normally measured in millimetres rather than metres. The units and prefixes defined in the *International System of Units* (SI) are

preferable.

A.6.5 Representations of Object Boundary

A segmented object or region is normally represented by a set of pixels on the image. The boundary of the object is formed by all the pixels between the object and the background. When the object boundary is described as an attribute in the database, the boundary pixels are represented by a sequence of points. The sequence of points are called *boundary points*, which should be described by the standard co-ordinate system on standard image orientation and resolution. The normalization methods defined in previous sections should therefore be used to transform the boundary pixels to boundary points.

Sometimes it is not necessary to include every segmented boundary pixel as a boundary point. When the image resolution is high and the object is large, using each boundary pixel may produce a sequence of points that is too long for storage. To save storage space, the segmented boundary may be reduced in spatial resolution so that the number of boundary points are also reduced. However, after the reduction, the original boundary can no longer be reproduced exactly on the image. One type of interpolations (possibly linear) has to be used to connect up the boundary points.

In another case, if the image resolution is reduced during the segmentation of the object, then the segmented boundary is also smaller in scale. After scaling the boundary points back to the original resolution, they would be disconnected and evenly spread out on the image. Again, interpolation may be used to link them up.

In order to perform interpolation on the boundary points, they must be arranged in a predefined sequential order. If the boundary is closed (a loop), then the boundary points should be described in clockwise or anti-clockwise direction. If the boundary is open (an edge), then the boundary points should be described in a single direction, e.g., from top to bottom, or from left to right. With this kind of consistency, the subsequent applications can reproduce and reuse the boundary more easily.

Furthermore, the approximate location and size of an object can be represented by some simple shapes, e.g., a circle, a rectangle or a polygon. The advantage of using simple shapes is that they can be described by a few parameters and hence with less data. As with the boundary points, these parameters must be normalized to a standard format too. In addition, chain codes are also widely used in different areas of image processing [56, pp. 219–220], and they may be considered for efficient boundary representation as well.

A.7 Storage of Attributes

The normalized attributes are permanently stored in an attribute database for any subsequent application in the system to use them again. The attributes can be stored in any data format, as long as the value of each attribute remains unchanged when it is retrieved from the database. All the approximation, reduction and conversion of the at-

tribute values should not be carried out within the database; they are instead part of the normalization process. Only the integrity of data items and the structure of data entries are verified by the database. The correctness and accuracy of the attributes are not part of its verification.

The database should be independent of the applications that are used for attribute acquisition, and should be portable to a wide range of computer platforms. It should also support some degree of scalability and extensibility so that further development of the system can be implemented without significant changes to the existing database design. If it is necessary, online access of the database should be enabled for remote applications.

Generally, any type of database can be used to store the normalized attributes. However, depending on the hierarchical structure of the attributes, one type of database may be more suitable to capture the data than other databases. Typically, Extensible Markup Language (XML) database is preferable for storing the attributes of any class of images, due to a number of reasons [85, p. 85], [59, p. 10]:

- XML has the expressive power to describe a complex data model;
- XML is useful for data exchange between applications; and
- XML can be viewed and manipulated with a wide range of tools.

The pros and cons of XML database are discussed in the following section. It is also compared with relational databases and object-oriented programming languages.

A.7.1 XML Database

XML is a meta-markup language that defines a syntax in which field-specific markups can be used to represent data as a string of text, or as a document [85, p. 3]. An XML database stores data in a collection of XML documents. Each XML document is simply a text file, containing XML elements and XML attributes¹, that can be easily accessed and manipulated with an editor or programming interface.

XML elements are the basic building blocks of XML. An element may contain an arbitrary number of subelements, and each subelement may contain other sub-subelements, and so on. An XML document must have exactly one root element, that completely contains all other elements, at the very top level. Thus the nested elements actually form a tree hierarchy in the document and data information is represented in a hierarchical style. The nested element structure can be extended by attaching new elements to the existing hierarchy. An element may also contain an arbitrary number of XML attributes, which are name-value pairs in string format. Unlike the XML elements, XML attributes cannot be nested and therefore they are not extensible.

XML is a self-describing language with great expressive power and flexibility. One of its advantages is the ability to markup structured and unstructured data, both in the

¹The definition of XML attributes should not be confused with that of image attributes.

same document. It can be used to capture a complex data model with a mixture of data and text. Most languages are designed to describe either well-structured data or free-form text, but XML can integrate both in the same document [59, p. 9].

The use of XML should be seriously considered when data exchange is required between multiple applications, particularly when the structure of the data changes [59, p. 10]. Due to the flexible and extensible data representation of XML, *the element structure of an XML document can be iteratively refined as the data model changes* [59, p. 9]. This makes XML particularly useful for areas that consist of complex relationships and are likely to undergo frequent revision [59, p. 9].

The syntax of XML is very well defined in the specification [86] from World Wide Web Consortium (W3C). The standard is supported by a range of generic tools and XML parsers [85]. As a plain text file, an XML document can be viewed and edited using a simple text editor. It can also be manipulated using a programming interface that parses the document and converts it into a tree of objects. Hence XML has been designed to be a powerful language that is easy to read and write by human beings and by computer applications.

Comparison with Relational Databases

Compared with the relations used in relational databases, the structure of XML elements is certainly more flexible and expressive. A *relation* is often visualized as a table with rows and columns. Each relation has an unordered collection of tuples (or rows), where each tuple has a fixed number of characteristics (or columns) [59, p. 22]. A relation cannot contain other relations and the number of characteristics has to be the same for all tuples in the relation. Because a relation is not very flexible, it is only suitable to capture data that is rigid and regularly structured; whereas XML has the expressive power to describe both structured and unstructured data.

One advantage of using a relational database is that complex queries can be expressed in Structured Query Language (SQL). SQL is a well developed and widely used query language that is supported by many database products. It provides a means of creating a relational database, as well as searching, inserting, updating and deleting data in the database. Currently, the query languages of XML, such as XPath and XQuery (both are W3C working drafts), are still under development. Most of them only support data searching in XML databases. Other data operations, like insert, update and delete, have to be done by other means, but they will be supported in the future.

As mentioned previously, the database of image attributes has to support some degree of scalability. The size of the database is likely to scale up in two situations: (i) when new image attributes are added to the existing set; and (ii) when more images are analyzed and their attributes are inserted into the database. As the size of the database increases, the search time of data items in the database scales up also. In terms of data storage, the scalability of XML database is as good as that of relational database. More XML elements

and attributes can be added to the XML document for new image attributes; and more XML documents can be created for new images.

However, in terms of data searching, the scalability of XML database is relatively poor when compared with relational database. An XML database stores data in XML documents. If the search is carried out within a single document, the processing time is still reasonable, since the data item is only searched for in a single hierarchy of elements. On the other hand, if the search is performed across multiple documents, perhaps across all the documents in the database, then it is a time-consuming process. The same search operation is faster in a relational database because a particular characteristic can be centralized in a single column of a relation.

In summary, different implementation approaches have to be taken for different system requirements. If the attribute-driven system is built solely for the purposes of image segmentation and analysis, then an XML database is sufficient since image attributes are usually retrieved from a single XML document that contains all the attributes of a particular image. On the other hand, if the system is designed for the purpose of image similarity matching or image retrieval based on image attributes, then a relational database could be used for data storage and searching, together with an XML interface for data exchange between applications.

The main topic of this thesis, segmentation and analysis of digitized mammograms, corresponds to the former case described above.

Comparison with Object-Oriented Programming Languages

Some features of XML are similar to those of object-oriented programming languages, such as C++ and Java. The features provided by XML, such as element types, named attributes and hierarchical structure, are also supported by the *objects* used in object-oriented programming languages [59, p. 16]. Due to these similarities, an XML document can be directly transformed into a tree of objects using an application programming interface (API), such as Document Object Model (DOM, a W3C recommendation). Computer programs can therefore be used to manipulate the data stored in XML documents.

Nevertheless, XML elements are not totally the same as objects. XML is not related to the concept of encapsulation, that is fundamentally defined in object-oriented programming languages. For objects, all internal structure is hidden from other objects; whereas for an XML element, all internal structure is transparently exposed [59, p. 16]. Objects also have the programming features, like function and attribute inheritance, that are not supported by XML. Therefore, while XML is excellent for data representation and data exchange between applications, a programming language is still required for data manipulation in applications, i.e., for editing the content of XML documents [59, p. 16].

A.7.2 XML Schema

XML Schema is a schema definition language (W3C recommendation). The element structure and data values that may occur in the XML documents are defined and constrained by an XML schema. XML Schema also provides a mechanism to enable inheritance for element, attribute, and datatype definitions [87, p. 11]. By defining tighter and looser constraints to formalize the document structure and datatypes, only a restricted class of XML documents are allowed to be created. This formalization is a necessary step for software design [88, p. 1]. Applications can be developed to deal with all the possible structures of an XML document. In many ways, schemas serve as design tools, establishing a framework on which implementations can be built [88, p. 1].

XML Schema can be interpreted by computer because its definition language is written in XML. An XML schema is primarily used to validate an XML document, i.e., to ensure that the document contents conform to the data definition described in the schema. The schema serves as a firewall, between the applications and the database, against the diversity and extensibility of XML. Due to the powerful validation features provided by XML Schemas, more validation work that has traditionally been performed by application code can be taken care of by defining a schema [87, p. 228]. The more constraints we put in the schema, the tighter its restrictions on the class of documents, and the less validation code we have to write in applications [87, p. 228].

A.8 Design Strategy for Actual Systems

The abstract system model presented in this appendix is a generic scheme that may be applied to the attribute-driven segmentation and analysis of a wide range of image classes. When designing the actual system for a particular class of images, the following steps are usually involved:

1. Identify the three types of image attributes, i.e., image data, visual features and analytical information. These attributes may be determined by an exhaustive search, experimental methods, past experience, or other manual processes.
2. Draw the dependence graph of the image attributes and show the hierarchical structure of low-level and high-level attributes.
3. Determine the attribute priorities, and hence the attribute acquisition order, either sequential or parallel.
4. Develop the computer application for acquiring the attribute with the highest priority.
5. Specify a standard format for the attribute if it has to be normalized.
6. Define the XML representation of the attribute in the database schema.

7. Repeat steps 4 to 6 for all other attributes from highest to lowest priorities.

In the beginning of a research project, it may not be possible to identify all the useful image attributes for a particular class of images, as we only have little knowledge of how to segment and analyze the images automatically using computer. Therefore the first draft of the dependence graph may not include all the image attributes; it may only contain the image data and some low-level attributes of image segmentation. Low-level attributes are likely to have high priorities, so they can be acquired in the early stages. Computer applications can be developed for the low-level attributes without considering any of the high-level attributes. After the low-level attributes are successfully generated using computer, and they are normalized and stored in the database in the standard format, high-level attributes of analytical information may then be identified and included in the dependence graph. Hence the applications of image analysis can be developed for the high-level attributes, which may depend on some of the low-level attributes in the database.

From the observation above, we notice that many components of the system are continuously evolving during the system development. The set of image attributes and the dependence graph are expanding as more attributes are included. New applications are implemented for the new attributes. The standard attribute format and the database schema are also modified to include the new attributes. One of the advantages of our attribute-driven system model is that it can be extended easily. The entire system may be built up incrementally, application by application, until we have reached our goals of image analysis. New applications can be developed separately without knowing the internal implementation of other applications; and new applications can be integrated into the system without influencing the existing applications.

Three examples are outlined below to demonstrate the possible ways of designing the actual systems of different classes of images.

A.8.1 Example: Bank Cheques

The information on a bank cheques may be located, extracted and processed automatically using computer. The first set of image attributes acquired from the cheque are likely to be the locations of several hand-written items, including the legal amount in words, the numeric amount, the date and the signature [89]. These hand-written items may be written in words or numbers with various styles. They can be located on the cheque by detecting the horizontal baselines under the items [89]. The regions of interest around the baselines are then extracted, and the hand-written data can be segmented out. Other low-level image attributes of the cheque could be the visual features of the segmented hand-written data. To extract the visual features, the region of each hand-written item may be divided into words, and also subdivided each word horizontally into 5 smaller regions, corresponding to the main body and the baseline of the word [90]. Within each of these five regions, finer features of the word may be generated, such as ascender, descender, loops,

word length, as well as horizontal, vertical and diagonal strokes [91]. They are all image attributes of the cheque. After all the low-level attributes are obtained and stored into the database, applications of information processing, such as handwriting recognition and signature matching, may be carried out automatically. The results of these applications are the high-level image attributes of the cheque.

A.8.2 Example: Human Faces

The visual features of the face provide very important information in a face recognition system. These low-level image attributes may be obtained progressively using hierarchical face segmentation. Locating the face region is arguably the first step of any system performing face recognition. The head of the targeted person is first segmented from the image background to locate the frontal view of the face. From the face region, several features can be extracted, such as skin colour, face shape, and the ratio of face width and height [92]. Then the face may be horizontally partitioned into five distinctive regions, from top to bottom: forehead, eyes, nose, mouth and chin [93]. Within each of these five regions, smaller facial features can be detected, such as hair colour, hair style, lengths of the eyebrows, eye colour, shape of the eyes, size of the nostrils, and width of the mouth [92]. After collecting all the visual features in the database, high-level analysis, primarily face recognition and identification, may be performed.

A.8.3 Example: Mammograms

An automatic image analysis system may be developed for breast cancer detection on mammograms. The mammogram can be hierarchically segmented into several distinctive regions according to the anatomical structure of the breast [61]. The low-level attributes of a mammogram are therefore its anatomical features. The breast region is first segmented from the image background, then the nipple is located on the breast border, and the pectoral muscle is delineated inside the breast region. The fibro-glandular region may be segmented from the breast region as well. Using the anatomical features, several types of image analysis can be performed, such as image adequacy and quality assessment, breast density classification, and abnormality detection. This analytical information constitutes the high-level attributes of mammograms. More detail of the mammogram analysis system is described in the Chapter 3.

A.9 Conclusions

In this appendix, the abstract view of an original, attribute-driven system is presented for automatic image segmentation and analysis. The system model has been generalized for most classes of images, but it can be customized for a particular class of images through the design and development of domain-specific image processing techniques.

There are generally three types of image attributes: image data, visual features and analytical information. Based on image data, visual features can be derived from image segmentation, and hence analytical information obtained through image analysis. The knowledge about the image is captured by the whole set of attributes in the database. The set of attributes expands as we gather more information about the image, leading to more accurate segmentation and analysis results. This iterative refinement is supported by the cyclic process of attribute-driven segmentation and analysis that comprises three major components: attribute acquisition, normalization and storage.

Attribute acquisition is achieved by developing a series of domain-specific applications for a particular class of images. The priorities of attributes may be determined from the attribute dependence graph introduced in this appendix. The applications may then be executed in a pre-defined order so that the attributes with higher priorities are generated first.

In order to formalize the data exchange between applications, the image attributes produced by different applications need to be normalized to a standard format, which is defined by the system developer. Aspects associated with attribute normalization, such as image orientation, co-ordinate system, image resolution, measurement units and representation of object boundary, are discussed in this appendix.

The normalized attributes are stored persistently in a database. Any type of database can be used in the implementation, but it has to be application- and platform-independent. The database should support some degree of scalability so that the attributes of a large number of images can be stored in the database without a significantly increase in retrieval time. The database should also be extensible so that new image attributes can be added to the system without significant changes to the existing database design. XML database and XML schema have been suggested for the storage of image attributes, and their properties have been discussed.

Finally, a design strategy has been given to layout the steps to be taken for customizing the abstract system model into an actual system. This design strategy provided a picture of the project life cycle that may be undertaken in the development of the attribute-driven system. Most importantly, it also revealed that the abstract system model is not only a generic scheme for automatic segmentation and analysis, but it also serves as a development platform for us to implement the entire system incrementally, application by application, until our end goals of image analysis have been achieved. Three examples have also been given to illustrate the possible applications of the abstract system model to some domain-specific problems.

Appendix B

XML Schema of Mammogram Attributes

The complete XML schema of the mammogram attribute database described in Chapter 3 is given in this appendix. The basic elements used in the schema are first described. The element structure of the mammogram attributes is then presented, followed by a complete code listing of the XML Schema file.

All diagrams in this appendix were generated by the XMLSpy Schema Editor¹, and hence its graphical representation is adopted.

B.1 Basic Elements

Basic elements can be included and re-used in the definitions of other elements in the same schema.

B.1.1 Point

A particular point on the image is described by a `<Point>` element with four XML attributes: `x`, `y`, `z`, and `i`, where `x` and `y` define the co-ordinate pair; `z` can be used to represent any quantity, such as pixel value; and `i` is the index number.



FIGURE B.1: Graphical representation of the `<Point>` element.

¹XMLSpy is available at <http://www.xmlspy.com>.

B.1.2 Line

An element of `lineType` describes a straight line on the image with two end-points, which are defined by `<Point>`.

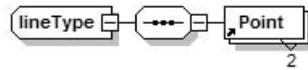


FIGURE B.2: Graphical representation of the element of `lineType`.

B.1.3 Array

A boundary or curve on the image is described by an element of `arrayType`, which consists of an array of `<Point>` elements in sequential order. The number of points in an array is specified by the XML attribute `array_length`.

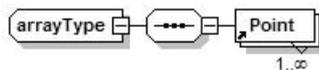


FIGURE B.3: Graphical representation of the element of `arrayType`.

B.1.4 Image

The file location of an image is described by an `<Image>` element with three XML attributes: `id`, `location` and `resolution`, where `id` is the identification number of the image; `location` contains the URL of the image file; and `resolution` is the image resolution.



FIGURE B.4: Graphical representation of the `<Image>` element.

B.2 XML Elements of Mammogram Attributes

B.2.1 Mammogram

All mammogram attributes are described under a single element `<Mammogram>`, which is the root of the element tree. `<Mammogram>` has three subelements: `<Patient_Info>`, `<Digital_Image>` and `<Comment>`. The elements of Mammogram attributes are divided into four categories under `<Digital_Image>`. The four categories are `<Breast_Anatomy>`, `<Abnormality>`, `<Adequacy>` and `<Texture>`. The first three are explained in the following sections.

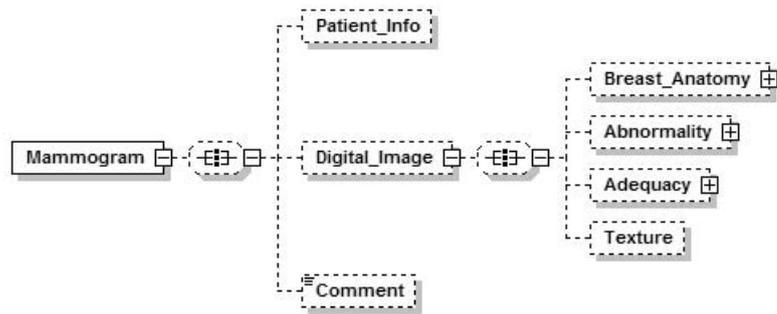


FIGURE B.5: Graphical representation of the <Mammogram> element and its subelements.

B.2.2 Breast Anatomy

The anatomical features of the mammogram are described in the element <Breast_Anatomy>. The <Breast_border> element is of arrayType. The <Nipple> element is a <Point> and one of its XML attributes indicates whether the nipple is in profile. The <Pectoral_muscle> is described by a <Straight_line> and a <Curve>. The <Boundary> of the <Parenchyma> is of arrayType, and its <Density> is represented by a string datatype.

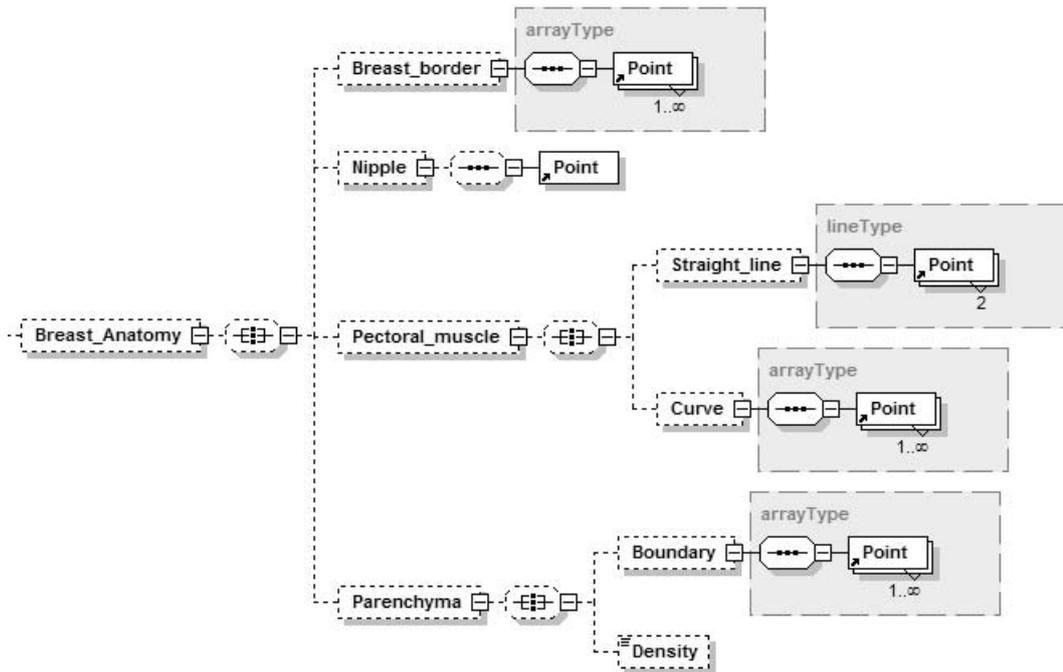


FIGURE B.6: Graphical representation of the <Breast_Anatomy> element and its subelements.

B.2.3 Adequacy

The adequacy measures of the mammogram are described in the element `<Adequacy>`. The exclusions of breast tissue at the image edges are described as elements of `lineType` in the `<Tissue_exclusion_group>`. The pectoral muscle extension to the level of the nipple, the angle and the percentage of convexity of pectoral muscle are described as XML attributes in `<Pectoral_muscle>`. The posterior `<Nipple_line>` is defined by the nipple location and the other `<End_point>`. The location of `<Inframammary_fold>` is a `<Point>`. The mean optical density and the percentage of breast tissue in the useful density range are described by the XML attributes of `<Exposure>`. Other adequacy measures are not within the scope of this thesis and they are not defined here.

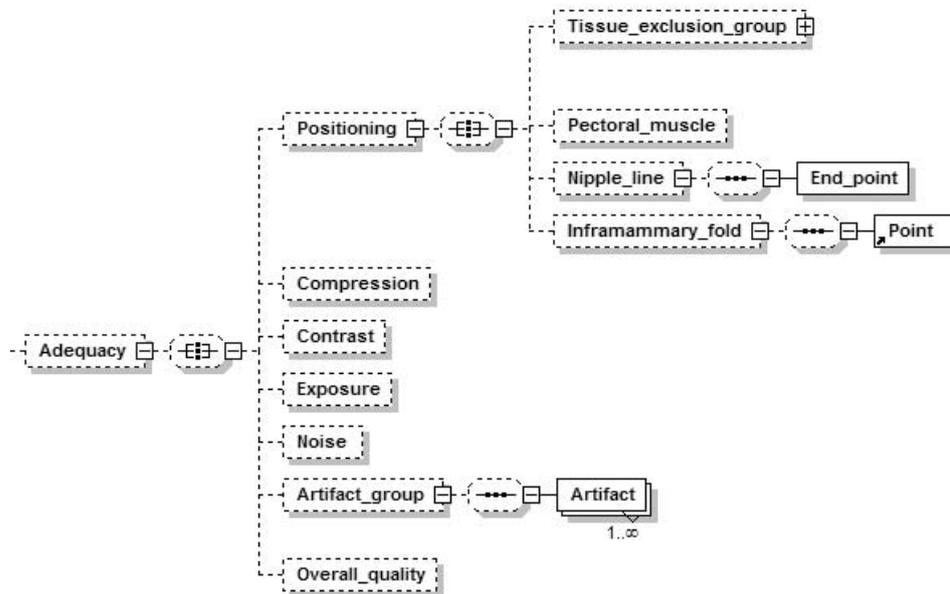
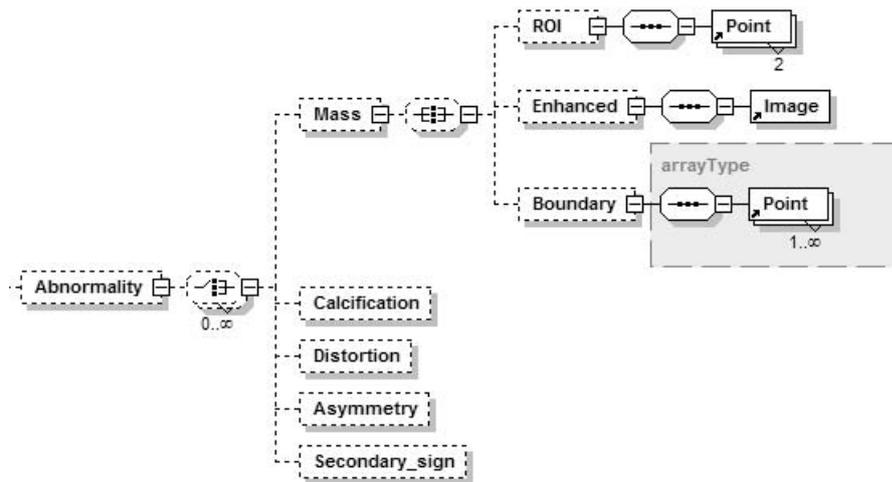


FIGURE B.7: Graphical representation of the `<Adequacy>` element and its subelements.

B.2.4 Abnormality

The descriptions of abnormalities found on the mammogram are included in the element `Abnormality`. The ROI of a `<Mass>` lesion is defined by two `<Point>`s; one for the top left corner of the window and the other for the bottom right corner. The `<Enhanced>` image of the mass is stored as a separate image file, referenced by `<Image>`. The `<Boundary>` of the mass is of `arrayType`. Other abnormalities are not within the scope of this thesis and they are not defined here.

FIGURE B.8: Graphical representation of the `<Abnormality>` element and its subelements.

B.3 XML Schema File

The text version of the XML Schema file, corresponding to the above graphical representation, is shown below.

```

<?xml version="1.0" encoding="UTF-8"?>
<!-- edited with XMLSPY -->
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema" elementFormDefault="qualified"
  attributeFormDefault="qualified">
  <xs:element name="Mammogram">
    <xs:complexType>
      <xs:all minOccurs="0">
        <xs:element name="Patient_Info" minOccurs="0">
          <xs:complexType>
            <xs:attribute name="id" type="xs:ID" use="required"/>
            <xs:attribute name="age" type="xs:nonNegativeInteger" use="optional"/>
            <xs:attribute name="gender" use="required">
              <xs:simpleType>
                <xs:restriction base="xs:string">
                  <xs:enumeration value="M"/>
                  <xs:enumeration value="F"/>
                </xs:restriction>
              </xs:simpleType>
            </xs:attribute>
            <xs:attribute name="weight" use="optional">
              <xs:simpleType>
                <xs:restriction base="xs:decimal">
                  <xs:minInclusive value="0"/>
                </xs:restriction>
              </xs:simpleType>
            </xs:attribute>
          </xs:complexType>
        </xs:element>
        <xs:element name="Digital_Image" minOccurs="0">
          <xs:complexType>

```

```

<xs:all minOccurs="0">
  <xs:element name="Breast_Anatomy" minOccurs="0">
    <xs:complexType>
      <xs:all minOccurs="0">
        <xs:element name="Breast_border" type="arrayType" minOccurs="0"/>
        <xs:element name="Nipple" minOccurs="0">
          <xs:complexType>
            <xs:sequence minOccurs="0">
              <xs:element ref="Point"/>
            </xs:sequence>
            <xs:attribute name="profile" type="xs:boolean" use="optional"/>
            <xs:attribute name="radius" use="optional">
              <xs:simpleType>
                <xs:restriction base="xs:decimal">
                  <xs:minInclusive value="0"/>
                </xs:restriction>
              </xs:simpleType>
            </xs:attribute>
          </xs:complexType>
        </xs:element>
        <xs:element name="Pectoral_muscle" minOccurs="0">
          <xs:complexType>
            <xs:all minOccurs="0">
              <xs:element name="Straight_line" type="lineType" minOccurs="0"/>
              <xs:element name="Curve" type="arrayType" minOccurs="0"/>
            </xs:all>
          </xs:complexType>
        </xs:element>
        <xs:element name="Parenchyma" minOccurs="0">
          <xs:complexType>
            <xs:all minOccurs="0">
              <xs:element name="Boundary" type="arrayType" minOccurs="0"/>
              <xs:element name="Density" minOccurs="0">
                <xs:simpleType>
                  <xs:restriction base="xs:string">
                    <xs:enumeration value="F"/>
                    <xs:enumeration value="G"/>
                    <xs:enumeration value="D"/>
                  </xs:restriction>
                </xs:simpleType>
              </xs:element>
            </xs:all>
          </xs:complexType>
        </xs:element>
      </xs:all>
    </xs:complexType>
  </xs:element>
  <xs:element name="Abnormality" minOccurs="0">
    <xs:complexType>
      <xs:choice minOccurs="0" maxOccurs="unbounded">
        <xs:element name="Mass" minOccurs="0">
          <xs:complexType>
            <xs:all minOccurs="0">
              <xs:element name="ROI" minOccurs="0">
                <xs:complexType>
                  <xs:sequence>
                    <xs:element ref="Point" minOccurs="2" maxOccurs="2"/>
                  </xs:sequence>
                </xs:complexType>
              </xs:element>
            </xs:all>
          </xs:complexType>
        </xs:element>
      </xs:choice>
    </xs:complexType>
  </xs:element>

```

```

        </xs:complexType>
    </xs:element>
    <xs:element name="Enhanced" minOccurs="0">
        <xs:complexType>
            <xs:sequence>
                <xs:element ref="Image"/>
            </xs:sequence>
        </xs:complexType>
    </xs:element>
    <xs:element name="Boundary" type="arrayType" minOccurs="0"/>
</xs:all>
<xs:attribute name="radius" use="optional">
    <xs:simpleType>
        <xs:restriction base="xs:decimal">
            <xs:minInclusive value="0"/>
        </xs:restriction>
    </xs:simpleType>
</xs:attribute>
<xs:attribute name="type" use="optional">
    <xs:simpleType>
        <xs:restriction base="xs:string">
            <xs:enumeration value="CIRC"/>
            <xs:enumeration value="SPIC"/>
            <xs:enumeration value="MISC"/>
        </xs:restriction>
    </xs:simpleType>
</xs:attribute>
<xs:attribute name="severity" use="optional">
    <xs:simpleType>
        <xs:restriction base="xs:string">
            <xs:enumeration value="BEN"/>
            <xs:enumeration value="MAL"/>
        </xs:restriction>
    </xs:simpleType>
</xs:attribute>
    <xs:attribute name="highlight_image" type="xs:IDREF" use="optional"/>
</xs:complexType>
</xs:element>
<xs:element name="Calcification" minOccurs="0">
    <xs:complexType>
        <xs:attribute name="radius" use="optional">
            <xs:simpleType>
                <xs:restriction base="xs:decimal">
                    <xs:minInclusive value="0"/>
                </xs:restriction>
            </xs:simpleType>
        </xs:attribute>
        <xs:attribute name="severity" use="optional">
            <xs:simpleType>
                <xs:restriction base="xs:string">
                    <xs:enumeration value="BEN"/>
                    <xs:enumeration value="MAL"/>
                </xs:restriction>
            </xs:simpleType>
        </xs:attribute>
        <xs:attribute name="count" type="xs:positiveInteger" use="optional"/>
    </xs:complexType>
</xs:element>

```

```

    <xs:element name="Distortion" minOccurs="0"/>
    <xs:element name="Asymmetry" minOccurs="0"/>
    <xs:element name="Secondary_sign" minOccurs="0"/>
  </xs:choice>
</xs:complexType>
</xs:element>
<xs:element name="Adequacy" minOccurs="0">
  <xs:complexType>
    <xs:all minOccurs="0">
      <xs:element name="Positioning" minOccurs="0">
        <xs:complexType>
          <xs:all minOccurs="0">
            <xs:element name="Tissue_exclusion_group" minOccurs="0">
              <xs:complexType>
                <xs:sequence minOccurs="0">
                  <xs:element name="Tissue_exclusion" type="lineType"
                    maxOccurs="unbounded"/>
                </xs:sequence>
              </xs:complexType>
            </xs:element>
            <xs:element name="Pectoral_muscle" minOccurs="0">
              <xs:complexType>
                <xs:attribute name="convex" type="xs:boolean" use="optional"/>
                <xs:attribute name="angle" use="optional">
                  <xs:simpleType>
                    <xs:restriction base="xs:decimal">
                      <xs:minInclusive value="0"/>
                      <xs:maxInclusive value="90"/>
                    </xs:restriction>
                  </xs:simpleType>
                </xs:attribute>
                <xs:attribute name="to_level_of_nipple" type="xs:decimal"
                  use="optional"/>
                <xs:attribute name="percentage" type="xs:decimal" use="optional"/>
              </xs:complexType>
            </xs:element>
            <xs:element name="Nipple_line" minOccurs="0">
              <xs:complexType>
                <xs:sequence minOccurs="0">
                  <xs:element name="End_point">
                    <xs:complexType>
                      <xs:attribute name="x" type="xs:decimal" use="required"/>
                      <xs:attribute name="y" type="xs:decimal" use="required"/>
                    </xs:complexType>
                  </xs:element>
                </xs:sequence>
                <xs:attribute name="length" use="required">
                  <xs:simpleType>
                    <xs:restriction base="xs:decimal">
                      <xs:minInclusive value="0"/>
                    </xs:restriction>
                  </xs:simpleType>
                </xs:attribute>
              </xs:complexType>
            </xs:element>
            <xs:element name="Inframammary_fold" minOccurs="0">
              <xs:complexType>
                <xs:sequence minOccurs="0">

```

```

        <xs:element ref="Point"/>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:all>
<xs:attribute name="adequate" type="xs:boolean" use="optional"/>
<xs:attribute name="score" type="xs:decimal" use="optional"/>
</xs:complexType>
</xs:element>
<xs:element name="Compression" minOccurs="0">
  <xs:complexType>
    <xs:attribute name="adequate" type="xs:boolean" use="optional"/>
    <xs:attribute name="score" type="xs:decimal" use="optional"/>
  </xs:complexType>
</xs:element>
<xs:element name="Contrast" minOccurs="0">
  <xs:complexType>
    <xs:attribute name="adequate" type="xs:boolean" use="optional"/>
    <xs:attribute name="score" type="xs:decimal" use="optional"/>
    <xs:attribute name="average" use="optional">
      <xs:simpleType>
        <xs:restriction base="xs:decimal">
          <xs:minInclusive value="0"/>
          <xs:maxInclusive value="100"/>
        </xs:restriction>
      </xs:simpleType>
    </xs:attribute>
  </xs:complexType>
</xs:element>
<xs:element name="Exposure" minOccurs="0">
  <xs:complexType>
    <xs:attribute name="adequate" type="xs:boolean" use="optional"/>
    <xs:attribute name="score" type="xs:decimal" use="optional"/>
    <xs:attribute name="avg_optical_density" use="optional">
      <xs:simpleType>
        <xs:restriction base="xs:decimal">
          <xs:minInclusive value="0"/>
        </xs:restriction>
      </xs:simpleType>
    </xs:attribute>
    <xs:attribute name="percentage" type="xs:decimal" use="optional"/>
  </xs:complexType>
</xs:element>
<xs:element name="Noise" minOccurs="0">
  <xs:complexType>
    <xs:attribute name="adequate" type="xs:boolean" use="optional"/>
    <xs:attribute name="score" type="xs:decimal" use="optional"/>
  </xs:complexType>
</xs:element>
<xs:element name="Artifact_group" minOccurs="0">
  <xs:complexType>
    <xs:sequence minOccurs="0">
      <xs:element name="Artifact" maxOccurs="unbounded"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
<xs:element name="Overall_quality" minOccurs="0">
  <xs:complexType>

```

```

        <xs:attribute name="adequate" type="xs:boolean" use="optional"/>
        <xs:attribute name="score" type="xs:decimal" use="optional"/>
    </xs:complexType>
</xs:element>
</xs:all>
</xs:complexType>
</xs:element>
<xs:element name="Texture" minOccurs="0">
    <xs:complexType/>
</xs:element>
</xs:all>
<xs:attribute name="location" type="xs:anyURI" use="optional"/>
<xs:attribute name="width" type="xs:positiveInteger" use="required"/>
<xs:attribute name="height" type="xs:positiveInteger" use="required"/>
<xs:attribute name="resolution" use="required">
    <xs:simpleType>
        <xs:restriction base="xs:decimal">
            <xs:minInclusive value="0"/>
        </xs:restriction>
    </xs:simpleType>
</xs:attribute>
<xs:attribute name="bit-depth" type="xs:positiveInteger" use="required"/>
<xs:attribute name="min_optical_density" use="optional">
    <xs:simpleType>
        <xs:restriction base="xs:decimal">
            <xs:minInclusive value="0"/>
        </xs:restriction>
    </xs:simpleType>
</xs:attribute>
<xs:attribute name="max_optical_density" use="optional">
    <xs:simpleType>
        <xs:restriction base="xs:decimal">
            <xs:minInclusive value="0"/>
        </xs:restriction>
    </xs:simpleType>
</xs:attribute>
</xs:complexType>
</xs:element>
<xs:element name="Comment" type="xs:string" minOccurs="0"/>
</xs:all>
<xs:attribute name="id" type="xs:ID" use="required"/>
<xs:attribute name="view" use="required">
    <xs:simpleType>
        <xs:restriction base="xs:string"/>
    </xs:simpleType>
</xs:attribute>
<xs:attribute name="side" use="required">
    <xs:simpleType>
        <xs:restriction base="xs:string">
            <xs:enumeration value="L"/>
            <xs:enumeration value="R"/>
        </xs:restriction>
    </xs:simpleType>
</xs:attribute>
<xs:attribute name="pair_id" type="xs:ID" use="optional"/>
<xs:attribute name="film_density" use="optional">
    <xs:simpleType>
        <xs:restriction base="xs:decimal">

```

```
        <xs:minInclusive value="0"/>
      </xs:restriction>
    </xs:simpleType>
  </xs:attribute>
  <xs:attribute name="exposure_time" use="optional">
    <xs:simpleType>
      <xs:restriction base="xs:decimal">
        <xs:minExclusive value="0"/>
      </xs:restriction>
    </xs:simpleType>
  </xs:attribute>
</xs:complexType>
</xs:element>
<xs:element name="Point">
  <xs:complexType>
    <xs:attribute name="x" type="xs:nonNegativeInteger" use="required"/>
    <xs:attribute name="y" type="xs:nonNegativeInteger" use="required"/>
    <xs:attribute name="z" type="xs:double" use="optional"/>
    <xs:attribute name="i" type="xs:nonNegativeInteger" use="optional"/>
  </xs:complexType>
</xs:element>
<xs:complexType name="lineType">
  <xs:sequence>
    <xs:element ref="Point" minOccurs="2" maxOccurs="2"/>
  </xs:sequence>
</xs:complexType>
<xs:complexType name="arrayType">
  <xs:sequence>
    <xs:element ref="Point" maxOccurs="unbounded"/>
  </xs:sequence>
  <xs:attribute name="array_length" type="xs:positiveInteger" use="required"/>
</xs:complexType>
<xs:element name="Image">
  <xs:complexType>
    <xs:attribute name="id" type="xs:ID" use="required"/>
    <xs:attribute name="location" type="xs:anyURI" use="required"/>
    <xs:attribute name="resolution" type="xs:decimal" use="required"/>
  </xs:complexType>
</xs:element>
</xs:schema>
```

Appendix C

Iterative Threshold Selection

The iterative threshold selection method described in Chapter 4 was claimed to be “the best possible solution for dividing the histogram while preserving the image average luminance” by Magid *et al.* in [68]. Their mathematical explanation is given below.

Let T be the threshold. A binary image can be created by replacing all pixels having a darker grey-level than T by the grey-level A , and by replacing those brighter than T by the grey-level B . Then an error function e^2 can be defined to evaluate the sum of the difference squared between the pixel values of the binary image and those of the original image (using integral expressions for convenience):

$$e^2 = \int_0^T (i - A)^2 h(i) di + \int_T^N (B - i)^2 h(i) di \quad (\text{C.1})$$

where $h(i)$ is the grey-level histogram of the original image.

We have to prove that the solution given by iterative threshold selection minimizes the error function e^2 by showing that A and B are the means of the pixel values below and above the threshold T , respectively, in the histogram.

The error function is minimized with respect to T by differentiating (C.1) with respect to T and equating to 0, i.e.,

$$\begin{aligned} \frac{d(e^2)}{dT} &= 0 \\ \frac{d}{dT} \int_0^T (i - A)^2 h(i) di + \frac{d}{dT} \int_T^N (B - i)^2 h(i) di &= 0 \\ \frac{d}{dT} \int_0^T (i - A)^2 h(i) di - \frac{d}{dT} \int_T^N (B - i)^2 h(i) di &= 0 \\ (T - A)^2 h(T) - (B - T)^2 h(T) &= 0 \\ (T - A)^2 &= (B - T)^2 \\ T - A &= B - T \\ T &= \frac{A + B}{2}. \end{aligned} \quad (\text{C.2})$$

So T is the average of A and B .

By replacing the value T in the integral limits in (C.1) with (C.2), and minimizing the error function with respect to A , we obtain

$$\begin{aligned} \frac{d(e^2)}{dA} &= 0.5[((A+B)/2) - A]^2 h((A+B)/2) \\ &\quad + \int_0^{(A+B)/2} (-2)(i-A)h(i)di \\ &\quad - 0.5[((A+B)/2) - B]^2 h((A+B)/2) = 0. \end{aligned} \quad (C.3)$$

The first and third parts cancel each other, so

$$\begin{aligned} \int_0^T (i-A)h(i)di &= 0 \\ \int_0^T i \cdot h(i)di - \int_0^T A \cdot h(i)di &= 0 \\ \int_0^T i \cdot h(i)di &= A \int_0^T h(i)di \\ A &= \frac{\int_0^T i \cdot h(i)di}{\int_0^T h(i)di}. \end{aligned} \quad (C.4)$$

which shows that A is the mean of all the pixel values below the threshold T . In the same way, we can show that when the error function is minimized, B is the mean of all the pixel values above the threshold T .

Thus, we have proven that by selecting the threshold as the average of the means of pixel values above and below the threshold in the histogram, the error function given in (C.1) is minimized, hence image average luminance is preserved.

Bibliography

- [1] Australian Institute of Health and Welfare, “BreastScreen Australia monitoring report 2000–2001,” Australian Institute of Health and Welfare, Canberra, Australia, Report AIHW cat. no. CAN 20, December 2003, Cancer Series no. 25.
- [2] S. K. Moore, “Better breast cancer detection,” *IEEE Spectrum*, vol. 38, no. 5, pp. 50–54, May 2001.
- [3] R. Highnam and M. Brady, *Mammographic Image Analysis*. Dordrecht: Kluwer Academic Publishers, 1999.
- [4] J. G. Daugman and C. J. Downing, “Demodulation, predictive coding, and spatial vision,” *Journal of the Optical Society of America A*, vol. 12, no. 4, pp. 641–660, Apr. 1995.
- [5] K. J. Bovis and S. Singh, “Enhancement technique evaluation using quantitative measures on digital mammograms,” in *IWDM 2000: 5th International Workshop on Digital Mammography*, M. J. Yaffe, Ed. Madison, WI, USA: Medical Physics Publishing, 2001, pp. 547–553.
- [6] P. Miller and S. Astley, “Automated Detection of Breast Asymmetry Using Anatomical Features,” in *State of the Art in Digital Mammographic Image Analysis*, ser. Machine Perception and Artificial Intelligence, K. W. Bowyer and S. Astley, Eds. Singapore: World Scientific, 1994, vol. 9, pp. 247–261.
- [7] U. Bick, M. L. Giger, R. A. Schmidt, R. M. Nishikawa, D. E. Wolverton, and K. Doi, “Automated Segmentation of Digitized Mammograms,” *Academic Radiology*, vol. 2, no. 1, pp. 1–9, Jan. 1995.
- [8] J. Suckling, D. R. Dance, E. Moskovic, D. J. Lewis, and S. G. Blacker, “Segmentation of mammograms using multiple linked self-organizing neural networks,” *Medical Physics*, vol. 22, no. 2, pp. 145–152, Feb. 1995.
- [9] A. J. Méndez, P. G. Tahoces, M. J. Lado, M. Souto, J. L. Correa, and J. J. Vidal, “Automatic detection of breast border and nipple in digital mammograms,” *Computer Methods and Programs in Biomedicine*, vol. 49, no. 3, pp. 253–262, May 1996.

- [10] M. Abdel-Mottaleb, C. S. Carman, C. R. Hill, and S. Vafai, "Locating the boundary between the breast skin edge and the background in digitized mammograms," in *Digital Mammography '96: Proceedings of the Third International Workshop on Digital Mammography, Chicago, USA, 9–12 June 1996*, ser. Excerpta Medica International Congress Series, K. Doi, M. L. Giger, R. M. Nishikawa, and R. A. Schmidt, Eds., vol. 1119. Amsterdam, The Netherlands: Elsevier Science, 1996, pp. 467–470.
- [11] R. Chandrasekhar, "Systematic Segmentation of Mammograms," Ph.D. Thesis, Centre for Intelligent Information Processing Systems, Department of Electrical and Electronic Engineering, The University of Western Australia, Nedlands, WA 6907, Australia, Oct. 1996.
- [12] R. Chandrasekhar and Y. Attikiouzel, "Gross Segmentation of Mammograms using a Polynomial Model," in *Bridging Disciplines for Biomedicine: Proceedings of the 18th Annual International Conference, IEEE Engineering in Medicine and Biology Society, October 31–November 3, 1996, Amsterdam, The Netherlands*. Amsterdam, The Netherlands: IEEE, Oct. 1997, pp. 1056–1058.
- [13] R. Chandrasekhar and Y. Attikiouzel, "A Simple Method for Automatically Locating the Nipple on Mammograms," *IEEE Transactions on Medical Imaging*, vol. 16, no. 5, pp. 483–494, Oct. 1997.
- [14] N. Karssemeijer, "Automated classification of parenchymal patterns in mammograms," *Physics in Medicine and Biology*, vol. 43, no. 2, pp. 365–378, Feb. 1998.
- [15] S. R. Aylward, B. M. Hemminger, and E. D. Pisano, "Mixture modeling for digital mammogram display and analysis," in *Digital Mammography: Nijmegen, 1998*, ser. Computational Imaging and Vision, N. Karssemeijer, M. Thijssen, J. Hendriks, and L. van Erning, Eds. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1998, vol. 13, pp. 305–312.
- [16] R. Chandrasekhar and Y. Attikiouzel, "Segmenting the Breast Border and Nipple on Mammograms," *Australian Journal of Intelligent Information Processing Systems*, vol. 6, no. 1, pp. 24–29, 2000.
- [17] R. J. Ferrari, R. M. Rangayyan, J. E. L. Desautels, and A. F. Frère, "Segmentation of mammograms: identification of the skin boundary, pectoral muscle, and fibroglandular disc," in *IWDM 2000: 5th International Workshop on Digital Mammography*, M. J. Yaffe, Ed. Madison, WI, USA: Medical Physics Publishing, 2001, pp. 573–579.
- [18] R. Chandrasekhar and Y. Attikiouzel, "Automatic breast border segmentation by background modeling and subtraction," in *IWDM 2000: 5th International Workshop on Digital Mammography*, M. J. Yaffe, Ed. Madison, WI, USA: Medical Physics Publishing, 2001, pp. 560–565, Proceedings of the Workshop, June 11–14, 2000, Toronto, Canada.

-
- [19] R. Chandrasekhar and Y. Attikiouzel, “New range-based neighbourhood operator for extracting edge and texture information from mammograms for subsequent image segmentation and analysis,” *IEE Proceedings—Science, Measurement and Technology*, vol. 147, no. 6, pp. 408–413, Nov. 2000.
- [20] F. Georgsson, “Algorithms and Techniques for Computer Aided Mammographic Screening,” Ph.D. dissertation, Department of Computer Science, Umeå University, Sweden, 2001.
- [21] M. Yam, M. Brady, R. Highnam, C. Behrenbruch, R. English, and Y. Kita, “Three-dimensional reconstruction of microcalcification clusters from two mammographic views,” *IEEE Transactions on Medical Imaging*, vol. 20, no. 6, pp. 479–489, June 2001.
- [22] R. Chandrasekhar and Y. Attikiouzel, “Segmentation of the pectoral muscle edge on mammograms by tunable parametric edge detection,” in *Advances in Signal Processing and Computer Technologies*, G. Antoniou, N. Mastorakis, and O. Panfilov, Eds. Athens, Greece: World Scientific and Engineering Society WSES Press, 2001, pp. 55–60.
- [23] M. Adel, M. Rasigni, G. Rasigni, B. Seradour, and P. Heid, “An automatic scheme to breast region segmentation,” in *Digital Mammography: IWDM 2002 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Berlin: Springer, 2003, pp. 161–163.
- [24] B. H. Allen, M. E. Oxley, and M. J. Collins, “A universal segmentation platform for computer-aided detection,” in *Digital Mammography: IWDM 2002 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Berlin: Springer, 2003, pp. 164–168.
- [25] R. Zwiggelaar, P. Planiol, J. Martí, R. Martí, L. Blot, E. R. E. Denton, and C. M. E. Rubin, “EM texture segmentation of mammographic images,” in *Digital Mammography: IWDM 2002 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Berlin: Springer, 2003, pp. 223–227.
- [26] M. A. Wirth and A. Stapinski, “Segmentation of the breast region in mammograms using active contours,” in *Visual Communications and Image Processing*, T. Ebrahimi and T. Sikora, Eds., vol. 5150, Lugano, Switzerland, 2003, pp. 1995–2006.
- [27] R. J. Ferrari, R. M. Rangayyan, J. E. L. Desautels, R. A. Borges, and A. F. Frère, “Automatic identification of the pectoral muscle in mammograms,” *IEEE Transactions on Medical Imaging*, vol. 23, no. 2, pp. 232–245, February 2004.
- [28] M. Masek, Y. Attikiouzel, and C. J. S. deSilva, “Skin-air interface extraction from mammograms using an automatic local thresholding algorithm,” in *Proceedings of the*

- 15th Biennial International Conference Biosignal 2000*, Brno, Czech Republic, June 2000, pp. 204–206.
- [29] M. Masek, Y. Attikiouzel, and C. J. S. deSilva, “Combining data from different algorithms to segment the skin-air interface in mammograms,” in *Proceedings of the World Congress on Medical Physics and Biomedical Engineering*, Chicago, USA, July 2000, CD-ROM, 4 pages.
- [30] M. Masek, R. Chandrasekhar, C. J. S. deSilva, and Y. Attikiouzel, “Spatially based application of the minimum cross-entropy thresholding algorithm to segment the pectoral muscle in mammograms,” in *ANZIIS 2001: Proceedings of the Seventh Australian and New Zealand Intelligent Information Systems Conference*. Perth, Australia: ARCME, The University of Western Australia, Nov. 2001, pp. 101–106.
- [31] M. Masek, C. J. S. deSilva, and Y. Attikiouzel, “Comparison of local median with modified cross-entropy for pectoral muscle segmentation in mammograms,” in *Proceedings of the 16th Biennial International Conference Biosignal 2002*, Brno, Czech Republic, June 2002, pp. 320–322.
- [32] M. Masek, “Hierarchical segmentation of mammograms based on pixel intensity,” Ph.D. Thesis, Centre for Intelligent Information Processing Systems, School of Electrical, Electronic and Computer Engineering, The University of Western Australia, Crawley, WA 6009, Australia, Feb. 2004.
- [33] S. H. Heywang-Köbrunner, D. D. Dershaw, and I. Schreer, *Diagnostic Breast Imaging: Mammography, Sonography, Magnetic Resonance Imaging, and Interventional Procedures*, 2nd ed. Stuttgart, Germany: Georg Thieme Verlag, 2001.
- [34] G. W. Eklund, G. Cardenosa, and W. Parsons, “Assessing Adequacy of Mammographic Image Quality,” *Radiology*, vol. 190, no. 2, pp. 297–307, Feb. 1994.
- [35] G. W. Eklund and G. Cardenosa, “The art of mammographic positioning,” *The Radiologic Clinics of North America*, vol. 30, no. 1, pp. 21–53, Jan. 1992.
- [36] V. F. Andolina, S. Lillé, and K. M. Willison, *Mammographic Imaging: A Practical Guide*. Philadelphia: J B Lippincott Company, 1992.
- [37] L. W. Bassett, I. A. Hirbawi, N. DeBruhl, and M. K. Hayes, “Mammographic positioning: Evaluation from the view box,” *Radiology*, vol. 188, pp. 803–806, 1993.
- [38] S. H. Taplin, C. M. Rutter, C. Finder, M. T. Mandelson, F. Houn, and E. White, “Screening mammography: Clinical image quality and the risk of interval breast cancer,” *American Journal of Roentgenology*, vol. 178, pp. 797–803, Apr. 2002.
- [39] C. Olsén, “Automatic determination of mammogram adequacy,” Master’s thesis, Department of Computing Science, Umeå University, Sweden, September 2002.

- [40] J. Law, "Consistency of film optical density in mammographic screening programmes," *The British Journal of Radiology*, vol. 69, no. 820, pp. 306–310, April 1996.
- [41] J. Law, "Improved image quality for dense breasts in mammography," *The British Journal of Radiology*, vol. 65, pp. 50–55, 1992.
- [42] K. C. Young, M. G. W. Wallis, and M. L. Ramsdale, "Mammographic film density and detection of small breast cancer," *Clinical Radiology*, vol. 49, pp. 461–465, 1994.
- [43] A. P. Dhawan, G. Buelloni, and R. Gordon, "Enhancement of Mammographic Features by Optimal Adaptive Neighbourhood Image Processing," *IEEE Transactions on Medical Imaging*, vol. MI-5, no. 1, pp. 8–15, Mar. 1986.
- [44] A. P. Dhawan and E. Le Royer, "Mammographic feature enhancement by computerized image processing," *Computer Methods and Programs in Biomedicine*, vol. 27, pp. 23–35, 1988.
- [45] R. H. Behrman, R. G. Zamenhof, and K. M. Blazo, "Evaluation of a Commercial Mammography Image-Enhancement System," *Journal of Digital Imaging*, vol. 2, no. 3, pp. 163–169, Aug. 1989.
- [46] R. P. Highnam, "Model-Based Enhancement of Mammographic Images," Ph.D. dissertation, Oxford University Computing Laboratory, Programming Research Group, Oxford, UK, 1992.
- [47] A. F. Laine, S. Schuler, J. Fan, and W. Huda, "Mammographic Feature Enhancement by Multiscale Analysis," *IEEE Transactions on Medical Imaging*, vol. 13, no. 4, pp. 725–740, Dec. 1994.
- [48] R. Gupta and P. E. Undrill, "The use of texture analysis to delineate suspicious masses in mammography," *Physics in Medicine and Biology*, vol. 40, pp. 835–855, 1995.
- [49] W. M. Morrow, R. B. Paranjape, R. M. Rangayyan, and J. E. L. Desautels, "Region-based contrast enhancement of mammograms," *IEEE Transactions on Medical Imaging*, vol. 11, no. 3, pp. 392–406, Sept. 1992.
- [50] R. M. Rangayyan, L. Shen, R. B. Paranjape, J. E. L. Desautels, J. H. MacGregor, H. F. Morrish, P. Burrowes, S. Share, and F. R. MacDonald, "An ROC Evaluation of Adaptive Neighbourhood Contrast Enhancement for Digitized Mammography," in *Digital Mammography: Proceedings of the Second International Workshop on Digital Mammography, York, England, 10–12 July 1994*, ser. Excerpta Medica International Congress Series, A. G. Gale, S. M. Astley, D. R. Dance, and A. Y. Cairns, Eds., vol. 1069. Amsterdam, The Netherlands: Elsevier Science, 1994, pp. 307–313.

- [51] R. M. Rangayyan, L. Shen, Y. Shen, J. E. L. Desautels, H. Bryant, T. J. Terry, N. Horeczko, and M. S. Rose, "Improvement of sensitivity of breast cancer diagnosis with adaptive neighborhood contrast enhancement of mammograms," *IEEE Transactions on Information Technology in Biomedicine*, vol. 1, no. 3, pp. 161–170, September 1997.
- [52] N. Petrick, H.-P. Chan, B. Sahiner, and D. Wei, "An adaptive density-weighted contrast enhancement filter for mammographic breast mass detection," *IEEE Transactions on Medical Imaging*, vol. 15, no. 1, pp. 59–67, February 1996.
- [53] H. Li, Y. Wang, K. J. R. Liu, S.-C. B. Lo, and M. T. Freedman, "Computerized radiographic mass detection—Part I: Lesion site selection by morphological enhancement and contextual segmentation," *IEEE Transactions on Medical Imaging*, vol. 20, no. 4, pp. 289–301, Apr. 2001.
- [54] A. K. Jain, *Fundamentals of Digital Image Processing*, ser. Prentice Hall Information and System Sciences Series. Englewood Cliffs, NJ, USA: Prentice-Hall, 1989.
- [55] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall/Pearson Education International, 2002.
- [56] I. N. Bankman, Ed., *Handbook of Medical Imaging: Processing and Analysis*. San Diego: Academic Press, 2000.
- [57] M. Masek, C. deSilva, and Y. Attikiouzel, "Automatic breast orientation in mediolateral oblique view mammograms," in *Digital Mammography: IWDM 2002 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Berlin: Springer, 2002, pp. 207–209.
- [58] M. J. Homer, *Mammographic Interpretation: A Practical Approach*. New York: McGraw-Hill, 1991.
- [59] M. Graves, *Designing XML Databases*. Upper Saddle River: Prentice Hall PTR, 2002.
- [60] R. Diestel, *Graph Theory*, 2nd ed. New York: Springer, 2000.
- [61] R. Chandrasekhar and Y. Attikiouzel, "Mammogram-attribute database: A tool for mammogram segmentation and analysis," in *Proceedings of the IASTED International Conference: Signal Processing, Pattern Recognition, and Applications (SP-PRA), July 3–6, 2001, Rhodes, Greece*, M. H. Hamza, Ed. Calgary, Canada: ACTA Press, 2001, pp. 143–148.
- [62] S. M. Kwok, R. Chandrasekhar, Y. Attikiouzel, and M. T. Rickard, "Automatic pectoral muscle segmentation on mediolateral oblique view mammograms," *IEEE Transactions on Medical Imaging*, vol. 23, no. 9, pp. 1129–1140, September 2004.

- [63] Y. Hatanaka, T. Hara, H. Fujita, S. Kasai, T. Endo, and T. Iwase, "Development of an automated method for detecting mammographic masses with a partial loss of region," *IEEE Transactions on Medical Imaging*, vol. 25, no. 12, pp. 5209–1212, Dec. 2001.
- [64] P. K. Saha, J. K. Udupa, E. F. Conant, D. P. Chakraborty, and D. Sullivan, "Breast tissue density quantification via digitized mammograms," *IEEE Transactions on Medical Imaging*, vol. 20, no. 8, pp. 792–803, Aug. 2001.
- [65] R. Chandrasekhar, S. M. Kwok, and Y. Attikiouzel, "Automatic evaluation of mammographic adequacy and quality on the mediolateral oblique view," in *Digital Mammography: IWDM 2002: 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Heidelberg, Germany: Springer-Verlag, 2003, pp. 182–186, Proceedings of the Workshop, June 22–25, 2002, Bremen, Germany.
- [66] P. K. Saha, J. K. Udupa, E. F. Conant, and D. P. Chakraborty, "Near automatic quantification of breast tissue glandularity via digitized mammograms," in *Medical Imaging 1993: Xmage Processing*, ser. SPIE Proceedings, K. M. Hanson, Ed., vol. 3661. San Diego, CA, USA: SPIE, 1999, pp. 266–276.
- [67] J. Suckling, J. Parker, D. R. Dance, S. Astley, I. Hutt, C. R. M. Boggis, I. Ricketts, E. Stamatakis, N. Cerneaz, S.-L. Kok, P. Taylor, D. Betal, and J. Savage, "The Mammographic Image Analysis Society Digital Mammogram Database," in *Digital Mammography*, ser. Excerpta Medica International Congress Series, A. G. Gale, S. M. Astley, D. R. Dance, and A. Y. Cairns, Eds., vol. 1069. Amsterdam, The Netherlands: Elsevier Science, 1994, pp. 375–378.
- [68] A. Magid, S. R. Rotman, and A. M. Weiss, "Comment on 'Picture thresholding using an iterative selection method'," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 20, no. 5, pp. 1238–1239, September/October 1990.
- [69] M. Sonka, V. Hlavac, and R. Boyle, *Image Processing, Analysis, and Machine Vision*, 2nd ed. Pacific Grove, CA, USA: PWS Publishing, 1999.
- [70] S. Siegel and N. J. Castellan, Jr., *Nonparametric Statistics for the Behavioral Sciences*, 2nd ed. New York: McGraw-Hill, 1988.
- [71] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. Belmont, Calif: Wadsworth International Group, 1984.
- [72] S. M. Kwok, R. Chandrasekhar, and Y. Attikiouzel, "Automatic pectoral muscle segmentation on mammograms by straight line estimation and cliff detection," in *Proceedings of the Seventh Australian and New Zealand Intelligent Information Systems Conference*. Perth, Western Australia: ARCME, The University of Western Australia, Nov. 2001, pp. 67–72.

- [73] L. Blot, A. Davis, M. Holubinka, R. Marti, and R. Zwiggelaar, “Automated quality assurance applied to mammographic imaging,” *EURASIP Journal of Applied Signal Processing*, vol. 7, pp. 736–745, 2002.
- [74] J. A. Thomas, K. Chakrabarti, R. V. Kaczmarek, and A. Romanyukna, “Automated method for measuring contrast-to-noise-ratio (CNR) and signal-to-noise-ratio (SNR) on a full field digital mammography system (FFDMS),” in *Digital Mammography: IWDM 2002 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Berlin: Springer, 2002, pp. 128–132.
- [75] K. R. Bijkerk, L. J. Oostveen, R. E. van Engen, M. M. J. Swinkels, and W. J. H. Veldkamp, “Interpretation of contrast detail images in digital mammography: Human observer versus computer detection,” in *Digital Mammography: IWDM 2002 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Berlin: Springer, 2002, pp. 133–135.
- [76] S. Muller, C. Ciofalo, G. Peter, A. Noël, and J. Stinès, “Automated scoring for CD-MAM phantom on digital mammography,” in *Digital Mammography: IWDM 2002 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Berlin: Springer, 2002, pp. 136–138.
- [77] B. Ancelin, M. Yam, R. P. Highnam, and J. M. Brady, “Film-screen artefact detection in mammograms,” in *Digital Mammography: IWDM 2002 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed. Berlin: Springer, 2003, pp. 174–178.
- [78] M. Masek, Y. Attikiouzel, and C. J. S. deSilva, “Automatic removal of high-intensity labels and noise from mammograms,” in *IWDM 2000: 5th International Workshop on Digital Mammography*, M. J. Yaffe, Ed. Madison, WI, USA: Medical Physics Publishing, 2001, pp. 587–592, Proceedings of the Workshop, June 11–14, 2000, Toronto, Canada.
- [79] J. Daugman and C. Downing, “Demodulation: A new approach to texture vision,” in *Spatial Vision in Humans and Robots: The Proceedings of the 1991 York Conference on Spatial Vision in Humans and Robots*, L. Harris and M. Jenkin, Eds. New York: Cambridge University Press, 1993, ch. 5, pp. 63–87.
- [80] P. Brodatz, *Textures : A Photographic Album for Artists and Designers*. New York: Dover Publications, 1966.
- [81] S. M. Kwok, R. Chandrasekhar, and Y. Attikiouzel, “Adaptation of the Daugman-Downing texture demodulation to highlight circumscribed mass lesions on mammograms,” in *DSP 2002: Proceedings of the 14th International Conference on Digital Signal Processing, July 1–3, 2002, Santorini, Greece*, A. N. Skodras and A. G. Constantinides, Eds., vol. 1. Santorini, Greece: IEEE, July 2002, pp. 449–452.

-
- [82] J. S. Lim, *Two-dimensional signal and image processing*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1990.
- [83] P. Soille, *Morphological Image Analysis: Principles and Applications*, 2nd ed. Berlin: Springer, 2003.
- [84] S. Dasmahapatra, D. Dupplaw, B. Hu, H. Lewis, P. Lewis, and N. Shadbolt, "Facilitating multi-disciplinary knowledge-based support for breast cancer screening," *International Journal of Healthcare Technology and Management*, 2004, Available from <http://www.ecs.soton.ac.uk/people/sd>.
- [85] E. R. Harold, *XML Bible*, 2nd ed. New York: Hungry Minds, 2001.
- [86] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau, "Extensible Markup Language (XML) 1.0 (Third Edition)," Published electronically, February 2004, W3C Recommendation. Available from <http://www.w3.org/TR/2004/REC-xml-20040204>.
- [87] K. Cagle, J. Duckett, O. Griffin, S. Mohr, F. Norton, N. Ozu, I. Stokes-Rees, J. Tension, and K. Williams, *Professional XML Schemas*. Birmingham: Wrox Press, 2001.
- [88] E. van der Vlist, *XML Schema*. Sebastopol: O'Reilly, 2002.
- [89] L. Lam, C. Y. Suen, D. Guillevic, N. W. Strathy, M. Cheriet, K. Liu, and J. N. Said, "Automatic processing of information on cheques," in *International Conference on Systems, Man and Cybernetics*, Vancouver, October 1995, pp. 2353–2358.
- [90] D. Guillevic and C. Y. Suen, "HMM word recognition engine," in *Proceedings of the International Conference on Document Analysis and Recognition*, vol. 2, Ulm, Germany, 1997, pp. 544–547.
- [91] D. Guillevic and C. Y. Suen, "Cursive script recognition applied to the processing of bank cheques," in *Proceedings of the International Conference on Document Analysis and Recognition*, Montreal, Canada, August 1995, pp. 11–14.
- [92] H.-S. Oh, J.-S. Park, and D.-H. Chang, "Development of the content based image retrieval system using object information," in *SPIE Conference on Multimedia Storage and Archiving Systems IV*, Boston, USA, September 1999, pp. 172–180.
- [93] F. Samaria, "Face segmentation for identification using hidden markov models," in *Proceedings of the 4th British Machine Vision Conference*. Guilford: Springer-Verlag, 1993, pp. 399–408.

List of Publications

The work presented in this thesis has resulted in the following publications:

1. S. M. Kwok, R. Chandrasekhar, Y. Attikiouzel, and M. T. Rickard, “Automatic pectoral muscle segmentation on mediolateral oblique view mammograms,” *IEEE Transactions on Medical Imaging*, vol. 23, no. 9, pp. 1129–1140, September 2004.
2. S. M. Kwok, R. Chandrasekhar, and Y. Attikiouzel, “Automatic pectoral muscle segmentation on mammograms by straight line estimation and cliff detection,” in *Proceedings of the Seventh Australian and New Zealand Intelligent Information Systems Conference*. Perth, Western Australia: ARCME, The University of Western Australia, Nov. 2001, pp. 67–72.
3. R. Chandrasekhar, S. M. Kwok, and Y. Attikiouzel, “Automatic evaluation of mammographic adequacy and quality on the mediolateral oblique view,” in *Digital Mammography: IWDM 2002: 6th International Workshop on Digital Mammography*, H.-O. Peitgen, Ed., Heidelberg, Germany, 2003, pp. 182–186, Proceedings of the Workshop, June 22–25, 2002, Bremen, Germany.
4. S. M. Kwok, R. Chandrasekhar, and Y. Attikiouzel, “Adaptation of the Daugman-Downing texture demodulation to highlight circumscribed mass lesions on mammograms,” in *DSP 2002: Proceedings of the 14th International Conference on Digital Signal Processing, July 1–3, 2002, Santorini, Greece*, A. N. Skodras and A. G. Constantinides, Eds., vol. 1. Santorini, Greece: IEEE, July 2002, pp. 449–452.
5. S. M. Kwok, R. Chandrasekhar, and Y. Attikiouzel, “A mammogram-attribute database in XML format for data-driven segmentation and image analysis,” in *IFMBE Proceedings: World Congress on Medical Physics and Biomedical Engineering (WC2003)*, vol. 4. Sydney, Australia: IFMBE and IOMP, Aug. 2003, 4 pages, CD-ROM, ISSN: 1727-1983.
6. S. M. Kwok, R. Chandrasekhar, and Y. Attikiouzel, “Automatic assessment of mammographic positioning on the mediolateral oblique view,” in *Proceedings of the IEEE International Conference on Image Processing*, Singapore, October 24–27, 2004, pp. 151–154.