



National Library
of Canada

Bibliothèque nationale
du Canada

Canadian Theses Service

Service des thèses canadiennes

Ottawa, Canada
K1A 0N4

NOTICE

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Previously copyrighted materials (journal articles, published tests, etc.) are not filmed.

Reproduction in full or in part of this microform is governed by the Canadian Copyright Act, R.S.C. 1970, c. C-30.

AVIS

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

Les documents qui font déjà l'objet d'un droit d'auteur (articles de revue, tests publiés, etc.) ne sont pas microfilmés.

La reproduction, même partielle, de cette microforme est soumise à la Loi canadienne sur le droit d'auteur, SRC 1970, c. C-30.

THE UNIVERSITY OF ALBERTA

A TECHNIQUE FOR AUTOMATED DETECTION OF BREAST TUMORS IN
MAMMOGRAMS

by

S.M. Lai

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES AND RESEARCH
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE
OF Master of Science

DEPARTMENT OF COMPUTING SCIENCE

EDMONTON, ALBERTA

FALL 1988

Permission has been granted to the National Library of Canada to microfilm this thesis and to lend or sell copies of the film.

The author (copyright owner) has reserved other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without his/her written permission.

L'autorisation a été accordée à la Bibliothèque nationale du Canada de microfilmer cette thèse et de prêter ou de vendre des exemplaires du film.

L'auteur (titulaire du droit d'auteur) se réserve les autres droits de publication; ni la thèse ni de longs extraits de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation écrite.

ISBN 0-315-45618-3

THE UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR : Shuk-Mei Lai

TITLE OF THESIS : A Technique for Automated Detection of Breast Tumors in
Mammograms

DEGREE : Master of Science

YEAR THIS DEGREE GRANTED : Fall 1988

Permission is hereby granted to THE UNIVERSITY OF ALBERTA LIBRARY to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purpose only.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

Shuk-Mei Lai

Permanent Address :

24B 3rd floor

Broadway

Mei Foo Sun Chuen

Kowloon, Hong Kong

Date :

Sept. 10, 1988

THE UNIVERSITY OF ALBERTA
FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled A Technique for Automated Detection of Breast Tumors in Mammograms submitted by Shuk-Mei Lai in partial fulfilment of the requirements for the degree of Master of Science.

Xiaobo Li
(Co-supervisor)

[Signature]
(Co-supervisor)

3 J Koles
[Signature]

W. A. Daves

Date : Sept. 10, 1998

ABSTRACT

This thesis attempts to solve the problem of breast tumor detection in mammograms. A technique for detecting one type of breast tumors, circumscribed masses, using a combination of detection criteria used by experts is presented. The criteria include the shape, brightness contrast and uniform density of tumor areas. This technique employs modified median filtering to enhance mammogram images and template matching to detect breast tumors. In the template matching step, suspicious areas are picked by thresholding the cross-correlation values and a percentile method is used to determine a threshold for each film. In addition, two tests are designed to remove false alarms from the resulting candidates. The results obtained by applying these techniques to some test images are described. The computational cost of the template matching procedure and its possible solutions are discussed. A coarse-fine approach is suggested for speeding up the procedure. In addition, a review on the general techniques found in the literature for noise cleaning and object detection is presented. Several noise filtering algorithms are implemented and their performance are compared.

Acknowledgements

I wish to express my profound gratitude to my supervisors, Dr. Xiaobo Li and Dr. B. Walter Bischof, for their painstaking efforts in guiding me through the writing and production process of this thesis. Thanks are also due to the thesis examining committee members, Dr. Wayne Davis, Dr. Zoley Koles and Dr. Mark Green for their time and suggestions; Dr. Bill Castor, director of the radiology department of the Cross Cancer Institute, for providing the mammograms; Mr. Jim Easton for helping in digitizing them. I am also indebted to Dr. Wlodek Dobosiewicz and Miss Meei Fen Teo for their proofreading.

Lastly, but not the least, I wish to express my appreciation to my family and Danny Huang for their patience, encouragement and support. Without their foremost patience and continuous encouragement, the completion of this thesis would not have been possible.

Financial support was given by the Department of Computing Science, University of Alberta, and the National Science and Engineering Research Council of Canada.

Table of Contents

Chapter	Page
1 Introduction	1
1.1 Background	2
1.2 Statement of the Problem	3
1.3 Conventions and Notations	5
1.4 Organization of Thesis	5
2 Basic Principles of Mammographic Diagnosis	7
2.1 Mammography and Tumors	7
2.2 Mammography Techniques	8
2.2.1 Mammography X-ray Unit	8
2.2.2 Mammography Recording Systems	9
2.3 Criterion in Detecting Circumscribed Tumors	10
2.4 Difficulties Using Computers to Detect Breast Tumors	11
2.5 Examples of Mammogram Images	11
3 Image Enhancement	15
3.1 Introduction	15
3.2 Survey of Existing Methods	15
3.2.1 Contrast Enhancement by Histogram Modification	17
3.2.2 Averaging Filters	20
3.2.2.1 Unweighted Averaging in the Spatial-Domain	20

3.2.2.2 Weighted Averaging	21
3.2.2.3 Averaging in the Frequency-Domain	22
3.2.3 Selective Averaging Filters	24
3.2.4 Median Filtering	32
3.3 Proposed Solution	35
3.3.1 Selective Median Filtering	36
3.3.2 Evaluations and Comparison	39
3.3.3 Performance of the Proposed Method	46
3.4 Summary	48
4 Object Detection	52
4.1 Introduction	52
4.2 Survey of Existing Methods	54
4.2.1 Template Matching	54
4.2.1.1 Similarity Measures	55
4.2.1.2 Criteria in Selecting Match Positions	57
4.2.1.3 Efficiency of the Matching Process	57
4.2.2 Image Segmentation	60
4.2.2.1 Pixel-Based Segmentation	61
4.2.2.2 Edge-Based Segmentation	63
4.2.2.3 Region-Based Segmentation	67
4.3 Proposed Solution	70
4.3.1 Template Matching Approach to Tumor Detection	70
4.3.2 Definition of the Tumor-like Template	71

4.3.3 Similarity Measure	73
4.3.4 Criteria in Selecting Suspicious Areas	73
4.3.5 False Alarm Tests	77
4.3.6 Experimental Results	82
4.3.7 Computation Consideration	88
4.3.7.1 Discussion	88
4.3.7.2 Experiments	88
4.4 Summary	90
5 Conclusion and Future Research	91
5.1 Conclusion	91
5.2 Directions for Further Research	92
Bibliography	94

List of Tables

Table	Page
3.1 Quantitative measurements of the enhanced images	40
4.1 Comparison of the computer's and the radiologist's interpretation on 7 mammograms	84
4.2 Comparison of the performance of the proposed method using different q values.....	83
4.3 Comparison of the computer's and the radiologist's interpretation on 17 mammograms	86
4.2 Summary of experimental results	87

List of Figures

Figure	Page
3.1 Histogram of the image in Plate 2.1	18
3.2 Histogram of the image in Plate 3.1	19
3.3 Examples of noise cleaning mask	23
3.4 Possible combinations of N_1 and N_2	27
3.5 Edge masks in four different orientations	30
3.6 Examples of straight edges	45
3.7 Examples of non-straight edges	45
4.1 A straight line in Cartesian coordinate space	65
4.2 A tumor-like template for matching with tumors of 5 pixels in diameter	72
4.3 A distribution curve of normalized cross-correlation values computed for the best case image (plate 3.24)	76
4.4 A distribution curve of normalized cross-correlation values computed for the worse case image (plate 3.24)	76
4.5 The normalized cross-correlation values computed around the center of a suspicious area	79
4.6 The normalized cross-correlation values computed around the center of a false alarm area	79
4.7 Grey-level histogram constructed in a suspicious area	81
4.8 Grey-level histogram constructed in a false alarm area	81
4.9 A distribution curve of normalized cross-correlation values computed for the unenhanced best case image (plate 2.1)	87

List of Plates

Plate		Page
2.1	Mammogram image 1.....	12
2.2	Mammogram image 2.....	13
2.3	Mammogram image 1 with suspicious areas marked.....	13
2.4	Mammogram image 2 with suspicious area marked.....	14
3.1	Result of applying histogram equalization to the test image.....	19
3.2	Result of applying unweighted averaging to the test image.....	21
3.3	Result of applying the edge preserving smoothing method to the test image.....	26
3.4	Result of applying the half neighborhood method to the test image.....	28
3.5	Result of applying the k-nearest neighbor method to the test image	29
3.6	Result of applying the directional smoothing method to the test image.....	31
3.7	Result of applying selective median filtering to the test image	38
3.8	Laplacian of the test image	42
3.9	Laplacian of the image in Plate 3.2a (3x3).....	42
3.10	Laplacian of the image in Plate 3.2b (5x5)	42
3.11	Laplacian of the image in Plate 3.3	42
3.12	Laplacian of the image in Plate 3.4a (3x3).....	43
3.13	Laplacian of the image in Plate 3.4b (3x5).....	43
3.14	Laplacian of the image in Plate 3.5a (3x3).....	43
3.15	Laplacian of the image in Plate 3.5b (5x5).....	43
3.16	Laplacian of the image in Plate 3.6a (3x3).....	44
3.17	Laplacian of the image in Plate 3.6b (5x5).....	44
3.18	Laplacian of the image in Plate 3.7a (3x3).....	44
3.19	Laplacian of the image in Plate 3.7b (5x5).....	44
3.20	Laplacian of the SMF-enhanced test image, T=3.....	50

Plate		Page
3.21	Laplacian of the SMF-enhanced test image, $T=5$	50
3.22	Laplacian of the SMF-enhanced test image, $T=7$	50
3.23	Laplacian of the SMF-enhanced test image, $T=10$	50
3.24	Result of applying selective median filtering to (a) mammogram image 1 and (b) mammogram image2, using a threshold value of 5, window size of 9 and five iterations.....	51
3.25	Result of applying selective median filtering to the original mammogram image 1, some false alarms are marked.....	51
4.1	Result of applying the proposed method to the enhanced mammogram images 1 and 2. All suspicious areas reported are marked by circles.....	85
4.2	Result of applying the proposed method to the original mammogram image 1. Some suspicious areas reported are marked by circles.....	85

Chapter 1

Introduction

This thesis attempts to search for techniques to solve a radiographic image analysis problem in a particular biomedical application – the detection of breast tumors in mammograms.

Breast cancer is not only a leading cause of death among all cancers for women of middle age and older [Tab85], but its incidence is rising. Primary prevention is not possible since the cause of this disease is still not understood. Once it has become disseminated, the chance for complete recovery is small, because the course of the disease at such a late stage has not been greatly affected by the tremendous efforts to improve treatment [Tab85].

– However, current methods of treatment are very effective against breast cancer in its early phase when the balance between the tumor and its host is more favourable [Str87]. Therefore, removal of the cancer while it is still in its early stages is the most promising way to achieve a significant change in the current breast cancer situation. Of all diagnostic methods currently available for this purpose, mammography is the most reliable method for the detection of early breast cancer [Zuc87]. Hence, a mass screening program utilizing mammography is obviously the best weapon against breast cancer.

A major problem expected with such a screening program would involve the interpretation of the large volume of images produced. In addition to a shortage of trained radiologists and the need to improve the cost benefit ratio of such a program, it is difficult for human radiologists to maintain interest in interpreting large numbers of images which show only a small number of abnormalities [Dha86]. Hence, the need to construct computer-aided systems to diagnose breast cancer in mammograms becomes

apparent.

1.1. Background

The processing of mammograms by computer can be roughly divided into three phases:

- (1) *Enhancement* : Enhancement of preselected features and removal of irrelevant details using application dependent techniques.
- (2) *Object Location* : Location of suspicious areas in mammograms.
- (3) *Classification* : Feature-based classification of these areas into non-tumors, benign or malignant tumor areas.

The most difficult of the three steps is object location. This is particularly relevant with respect to radiographic image analysis and the reasons are as follows [Hal 71]. First, radiographs, like mammograms, are characterized by their low resolution and low contrast ratios of small features superimposed onto non-uniform backgrounds. Second, the diagnostic information content in a radiograph is more semantic than statistical in nature. That is, the features essential for diagnosing a particular disease are determined both by the disease and the class of images. For example, heart size is relevant in the diagnosis of rheumatic heart disease, but is probably irrelevant in the diagnosis of pneumoniosis opacities. Third, the resolution required in a particular disease application may, for example, be an order of magnitude higher than that necessary for all general diseases using the same class of images.

A large repertoire of image processing techniques has been developed for image enhancement, object location and pattern classification. However, due to the problems of radiographic image analysis mentioned above, most image processing techniques which have been applied to biomedical situations have been found to be very application dependent. This phenomenon is analogous to the fact that radiologists adopt different strategies

to analyze different types of medical images. Thus, starting with a large repertoire of techniques, one may have to combine and modify some existing techniques to create the best technique for a particular application. Some examples are the histogram enhancement and zoom thresholding techniques developed by *Tou* [Tou79] for detecting lung tissue boundaries in lung tissue micrographs, and the region growing algorithm developed by *Li, Savol and Fong* [Li78] to detect pneumonconiosis opacities in chest X-rays.

In the past, several groups [Ack72, Win67], have demonstrated the potential use of computers in feature based classification of suspicious areas. Since human assistance is needed to locate these areas before the computer processes the images, these systems are not fully automated. *Hand et al.* [Han79] have developed a method for identifying and locating abnormal areas in xeromammograms. Their approach which utilizes fourteen texture parameters and two shape parameters together with a comparison of left and right breast images, shows a high false alarm rate of approximately fifty-three suspicious areas per xeromammogram. Since xeromammograms are produced using a different recording technique than that used for mammograms, the performance of *Hand et al.*'s method with mammograms cannot be directly determined.

In diagnosing breast cancer, radiologists use several indicators or "markers" in mammograms, all being defined by a set of criteria, such as area, brightness, contrast and shape. This thesis is the first stage of a long term research project which aims at implementing an expert system for breast cancer diagnosis based on the same set of markers and criteria the experts use.

1.2. Statement of the Problem

This thesis explores an approach to the detection of one category of breast tumor, *circumscribed masses*, using a combination of detection criteria considered relevant by experts. The criterion used for distinguishing "suspicious" from "clearly normal" regions

on a film mammogram is that a suspicious area is a bright (compared with surrounding tissue) and approximately circular area of uniform density, and of varying size [Tab85, Mar82].

The problem stated in this thesis can be summarized as follows : *Given an image of a film mammogram, the problem is to find a detection method which determines the location and size of all suspicious areas, if there are any.* It must be noted that a suspicious area in a mammogram is not necessarily a tumor area. It can be a benign tumor or an area which radiologists would choose to examine in greater detail. In the classification step, these suspicious areas produced by the method designed, are classified into non-tumor or tumor areas.

Locating suspicious areas in mammograms is difficult due to the poor image quality of mammograms. The approach described in this thesis takes this problem into account and existing techniques are modified to improve tumor detection. This approach employs *selective median filtering* to enhance the images and *template matching* to detect suspicious areas. Due to the fact that each mammogram may contain more than one suspicious area, our template matching approach uses a statistical decision criterion to select suspicious areas. In addition, since the candidate areas produced by this matching process will not be exclusively suspicious areas, two tests are designed to remove false alarms from the resulting candidates.

Some general approaches to image enhancement and object detection in radiographs are discussed. The results of applying some of these techniques on a mammogram image are discussed and compared with the techniques developed. To evaluate the effectiveness of the techniques described, no formal proof is given. Instead, the results obtained by applying these techniques to test images are reported.

1.3. Conventions and Notations

As used in this thesis, the term *mammogram image*, refers to a digital image of a mammogram film digitized by a TV camera.* It is considered as a two dimensional array of integers whose row and column indices identify a point in the image and the corresponding array element value identifies a gray level (brightness level) at that point. The elements of such a digital array are called picture elements or pixels.

Pixels are identified by their (i,j) positions with respect to the upper left-hand corner of the image, which has position $(0,0)$. The i -dimension increases to the right; the j -dimension increases downward. The intensities of each pixel are represented by numbers from 0 through 255, with 0 representing black, and 255 representing white.

1.4. Organization of Thesis

In Chapter 2, the basic principles of mammographic diagnosis are discussed. This includes a description of the mammographic techniques recently used and the criterion used by expert radiologists to detect circumscribed masses. Also, the difficulties in using a computer to detect this type of breast tumor are described.

Chapter 3 gives a survey of techniques for image enhancement. The two classes of techniques that will be discussed are histogram modification and image enhancement by smoothing. A modified median filter is proposed to enhance mammogram films. Results of applying these techniques to test images are given.

A review of several techniques developed for object detection is presented in Chapter 4. It includes techniques employing image segmentation and template matching. A method using template matching to detect suspicious areas in mammogram film is proposed. The design of templates for matching tumors, the similarity measure and the

If computer-aided diagnosis of breast cancer is realized, the mammogram images will be digitized directly from the machine that produces mammograms.

criteria used for interpreting this measure are discussed. Also, two tests designed to remove false alarms are presented. Results of applying these techniques on some test images are reported.

Chapter 5 presents the conclusion of this thesis and the direction for further research in the area of automated detection of breast tumors in mammograms.

Chapter 2

Basic Principles of Mammographic Diagnosis

2.1. Mammography and Tumors

At the present time the only reliable system for diagnosing early breast tumor is by means of routine X-ray mammography and the accurate interpretation of mammograms by expert radiologists [Her87]. Besides, to produce high quality mammograms with the least radiation dose, it is necessary to use machinery specifically designed for this purpose. Many mammography techniques have been developed to improve the performance of these machines. To design an automated system to diagnose breast cancer in mammograms, it is essential to gain a basic knowledge in mammography techniques that directly affect the quality of images produced. Also, the diagnostic procedures used by the expert radiologist in detecting breast tumors must be studied.

Mammography, the radiologic study of the breast, has changed dramatically during the past 20 years [Hau83]. The major hurdle in the development of mammography was the difficulty of obtaining sufficient contrast in the soft-tissues of the breast that would define the individual structures [Zuc87]. In addition, although mammography is presently accepted as a relatively low-risk method for detecting breast tumors [Str87], it was once viewed as a technique that increased the risk of breast cancer in women below the age of fifty. Controversy on the radiation risk of mammography gradually gave rise to further improvement of the technique during the 1970's and the emphasis of new techniques is placed on reducing the radiation dose while maintaining optimum diagnostic image quality [Zuc87].

The mammographic examination routinely includes two views of each breast in different view planes [D'Os83]. When an expert radiologist gets the four mammograms

of the same patient, the radiologist will guide the diagnostic workup which extends from physical examination through film reading to interventional procedures [Tab85]. The film reading procedure involves the interpretation of individual films and a comparison study of the four mammograms. Accurate diagnosis of breast tumors in individual films depends on the perception of the tumor site and the right classification of the site into a benign or malignant tumor. According to the mammographic appearance, the pathological lesions occurring in the breast can be grouped into four major categories: *circumscribed tumors*, *star-shaped lesions*, *calcifications* and pathological processes giving rise to the *thickened skin syndrome* [Mar82]. For each group there are a certain number of X-ray signs. This thesis is concerned with the detection of circumscribed tumors. The criterion in detecting this category of tumor will be discussed.

2.2. Mammography Techniques

Dedicated mammographic machines are built to provide special imaging requirements for mammography. The two main components of this unit are the X-ray unit and the recording system. The key factors that must be controlled in these two components in order to produce mammograms of good quality while minimizing radiation dose to the patient are discussed in this section.

2.2.1. Mammography X-ray Units

Subject contrast is defined as the ratio of the X-ray intensity transmitted through one part of the breast to that transmitted through a more absorbing adjacent area in the breast [Hau83]. Subject contrast is of vital importance in mammography, because it picks up the very small differences in the density between normal and pathologic soft tissues of the breast. It also helps to detect minute details such as tumors of extremely small size which are very often, the foci of attention in breast tumor diagnosis. Factors affecting subject contrast include absorption differences in the breast (thickness and density),

radiation quality and scattered radiation. The last two factors are the ones that can be controlled in the X-ray units [Hau83].

In addition, compression of the breast is crucial because it makes the breast thickness more uniform, thereby providing more even penetration by the X-ray and reduces the amount of scattered radiation [D'Or83]. Also, compression immobilizes the breast and reduces blurring caused by motion. However, the amount of compression is limited by the amount of pain a patient can endure. Various techniques for compressing breast tissue have been developed, ranging from a balloon fitted into the radiographic cone to a movable plastic shield. Very often, a poor quality image is caused by using improper compression equipment [D'Or83].

2.2.2. Mammography Recording Systems

The two recording systems in use today are film and xeroradiographic processing. In film processing, the use of direct exposure film with a fine grain, produces excellent radiographic details. However, the radiation dose to the breast is large, therefore, a better film-screen combination is used today [Hau83]. A film-screen combination makes use of a single high-definition screen in contact with a single emulsion film [Way79]. To ensure optimal film-screen contact, a vacuum exists inside the film cassette. This further reduces the radiation dose without sacrificing image quality [D'Or83].

Xeroradiography was first used as an alternate means of recording images in the 1970s. This method involves the production of an electrostatic image which is then coated with powder, placed on paper, and finally coated with plastic [D'Or83]. The advantages of xeroradiography include wide exposure latitude, increased small-area contrast and convenient interpretation (eg. no view box required) [Dod81].

Competition between film-screen mammography and xeromammograph to capture the market for mammographic equipment is intense [Gol83]. There are advocates for

each method who are adamant in their claims that one is superior to the other. In fact, while there are advantages to each technique, there does not appear to be significant difference in their ability to detect cancer [Pag80]. The success of either mammographic technique is, in fact, mostly dependent on the skill and experience of the radiologist [Gol83].

2.3. Criterion in Detecting Circumscribed Tumors

The process of mammographic interpretation begins with the search for lesions in the mammogram. Perception is influenced by the texture of normal breast tissue. The radiologist must ensure that all parts of each mammogram are carefully searched for the presence of lesions to help discover tumors which are partially obscured by overlying breast tissue [Tab85]. X-ray signs of circumscribed tumors described by radiologists are as follows [Tab85, Mar82]: Circumscribed masses have hazy and indistinct borders. Their shape is close to circular, but some masses have short spicules on their surface. Benign breast tumors are sharply outlined by their contour, and malignant tumors have an indistinct contour, sometimes with short spicules. Benign tumors usually have about the same attenuation as the surrounding breast tissue. Circumscribed tumors, which are partly or completely the same density as fat are benign. Circumscribed malignant tumors usually absorb more radiation than normal breast tissue. Because of their high-density radiopaque mammographic appearance, overlying structural elements such as trabeculae and blood vessels cannot be seen through them. Malignant circumscribed tumors containing mucin and relatively little fibroelastic components are of higher density on the mammogram than those which have the same density as fat. The size of a tumor and its orientation within the breast are of secondary importance, but are helpful in differential diagnosis.

2.4. Difficulties Using Computers to Detect Breast Tumors

In order to design an automated method to reliably detect likely tumor sites, the radiologists' descriptions of the X-ray signs of the tumor must be translated into more explicit terms. Based on the previous descriptions given by experts, a suspicious area is defined as an area which

- (1) has a higher density value (brighter) than its surrounding tissue,
- (2) has uniform density inside the area,
- (3) has an approximately circular shape and of varying size, and
- (4) has fuzzy edges.

Locating suspicious areas in mammograms is difficult for a number of reasons. The small differences in density between normal and tumorous tissues in human breasts create little contrast between a tumor area and its background in the image. This contrast is further reduced in the filming and digitization process of the mammogram images. In addition, the presence of noise and other anatomical structures, such as ducts and glands increases the background variations of tumor areas. The boundaries of tumor areas are fuzzy and in some instances, only partially visible. Together with the small size of early-stage tumors, makes any attempt to segment the image by global gray level thresholding technique very difficult. Also, some tumor areas lose some of the well-defined characteristics mentioned above. To tackle these difficulties, mammograms must be enhanced prior to applying any technique for tumor detection. A survey of existing techniques for image enhancement and a proposed method for enhancing mammogram images are presented in the next chapter.

2.5 Examples of Mammogram Images

Mammogram images 1 and 2 are shown in Plate 2.1 and 2.2 respectively. The suspicious areas in these mammograms are marked by circles in Plate 2.3 and Plate 2.4. The suspicious areas in image 1 is easier to detect than that in image 2. This is due to the fact that there are less background variations at the suspicious areas in image 1. In image 2, the background variations are due to the presence of gland and fatty tissue in the breast.



Plate 2.1 Mammogram image 1 with easily detectable suspicious areas. This image is the test image used for testing the image enhancement methods implemented in Chapter 3



Plate 2.2 Mammogram image 2 with a suspicious area which is difficult to detect



Plate 2.3 Mammogram image 1 with suspicious areas marked in circles




Plate 2.4 Mammogram image 2 with suspicious area marked in a circle

Chapter 3

Image Enhancement

3.1. Introduction

The main objective of enhancement techniques is to process a given image so that the result is more suitable than the original image for a specific application or for further processing. In general, this is achieved by either noise removal or by feature enhancement.

It is evident that much useful physical information is lost and noise is introduced during the imaging process. When a picture (or a film mammogram) is digitized, further information is lost due to quantization and sampling errors. In addition, the effect of equipment and other environmental factors also introduce noise in digitized images. In the case of mammogram images, the problem of information loss is worsened by the fact that during the X-ray imaging process, subject contrast is difficult to be retained due to scattered radiation, thickness of the breast, etc. (see Chapter 2).

There is a large repertoire of techniques available in the literature for image enhancement. However, no single technique is applicable to all classes of images and the choice of technique often depends on the specific application. Some simple, but general, techniques are studied in this chapter so that a suitable technique can be chosen for the purpose of enhancing mammogram images. Also, a method combining several techniques is proposed and its performance is evaluated.

3.2. Survey of Existing Methods

In general, existing image enhancement techniques can be divided into three categories :

- (1) contrast enhancement,
- (2) image smoothing or noise cleaning, and
- (3) image sharpening or deblurring.

Unlike image restoration, these techniques make little or no attempt to estimate the actual degradation process that has operated on the image. They do, however, take into account certain general properties of image degradation. Contrast enhancement tries to compensate for attenuation of the picture signal, image smoothing reduces noise from the image and image sharpening removes blurring which is introduced by the degradation process.

Contrast enhancement by histogram modification and various techniques for image smoothing are discussed in the following sections. Techniques for smoothing in both the spatial and frequency domain are considered and linear methods such as low-pass filtering and non-linear ones such as selective neighborhood averaging and median filtering are included. Image sharpening techniques will not be discussed here because these techniques are useful primarily as enhancement tools for highlighting edges in an image. However, pure edge information by itself is of little importance to our tumor detection method described in later sections.

Most of the techniques described in this section were implemented and were applied to mammogram image 1 (Plate 2.1). In the following discussion, this image is referred as the *test image*. The enhanced version of this test image is shown after the discussion of a specific method.

3.2.1 Contrast Enhancement by Histogram Modification

The grey level histogram (first-order probability density function) of a typical radiographic image is highly skewed towards the darker levels [Hal71]. Thus, details in the dark regions are often not perceptible and this can be observed in the test image. The histogram of this test image is shown in Figure 3.1. Histogram modification is a group of techniques that can be used to enhance this type of image by rescaling or translating the grey scale of the original image, so that the histogram of the enhanced image follows some desired form.

For many types of images, the "ideal" distribution of gray levels is a uniform distribution [Hal74]. This is due to the fact that a uniform distribution of grey levels makes equal use of each quantization level and tends to enhance low contrast information. Hence, histogram-enhanced images should possess more low contrast details than the original one. Histogram equalization [Hal74, Hum75] is a technique in which the histogram of the original image is mapped into another histogram which has a uniform distribution of grey levels. Algorithms for computing the mapping are given in [Hal71, Hum77, Ros82] and most of them are based on the selection of thresholds for the original histogram which map into equal areas ("vertical slices") of the new histogram.

The result of applying histogram equalization to the test image is shown in Plate 3.1 and its histogram is shown in Figure 3.2. The algorithm used is based on a single-valued transformation, therefore, gray-level bins of the original histogram were only merged together, but no bin was broken up. Thus, due to the large peaks in the original histogram, the transformed histogram (Figure 3.2) is only very approximately a flat histogram.

It should be noted that, although contrast enhancement by histogram modification is very effective for enhancing low-contrast detail, it does not discriminate between low-contrast information and noise. Also, in performing histogram equalization to an image,

we are assuming that the information in the sparsely populated grey level ranges is of no special significance, so that compressing them will not result in information loss [Ros82].

Another histogram modification technique which transforms the grey level distribution of an image into a hyperbolic distribution instead of a uniform distribution is suggested in [Fre77]. It is based upon the idea that the transformation which flattened the distribution of the outputs of the retinal receptors of the human visual system will provide the best visual enhancement. Then, by assuming a logarithm response of the human visual system, the resulting transformation transforms the histogram into a hyperbola. This method seemed to work well for satellite pictures with inherently poor contrast and it was claimed [Fre77] to be superior to a histogram equalization method.




Figure 3.1 Histogram of the test image shown in Plate 2.1




Figure 3.2 Histogram of the image in Plate 3.1




Plate 3.1 Result of applying histogram equalization to the test image

3.2.2. Averaging Filters

3.2.2.1. Unweighted Averaging in the Spatial-Domain

In the spatial-domain, unweighted averaging is a straight-forward technique for image smoothing. Given an image $f(x,y)$, every pixel (x,y) is replaced by the average of the grey level values of the pixels contained in a predefined neighborhood of (x,y) . In other words, the smoothed image $g(x,y)$ is obtained by calculating the average

$$g(x,y) = \frac{1}{N} \sum_{(i,j) \in W} f(x,y)$$

where W is the set of coordinates of pixels in the neighborhood of the pixel (x,y) and N is the total number of pixels defined in W .

The neighborhood of (x,y) can be defined by a window of various sizes and shapes. It is conventional practice to select a square array template of size 3×3 [Hal79]. However, intuitively, if the objects in the image are small relative to the smoothing neighborhood size, smoothing may blur the object. In the case of circumscribed masses, the smallest size they appear in mammogram images is around six pixels in diameter. Therefore, to preserve the edges of small masses, square windows of 3×3 and 5×5 pixels are used in implementing the method discussed in this section. With a window size of 5×5 pixels, the average is

$$g(x,y) = \frac{1}{25} \sum_{i=-2}^2 \sum_{j=-2}^2 f(x+i,y+j)$$

The effect of unweighted averaging is shown by applying the procedure to the test mammogram image. The smoothed image is shown in Plate 3.2a. The image is more diffused than the original and the edge of the suspicious area is blurred. This effect is increased with larger window size (see Plate 3.2b). Hence, when edge information must be preserved for future processing, unweighted averaging is not suitable for enhancing the

image.




Plate 3.2 Result of applying unweighted averaging to the test image using window size of (a) 3×3 pixels and (b) 5×5 pixels

3.2.2.2. Weighted Averaging

The undesirable effect of an unweighted averaging operation can be prevented by using a weighted average, in such a way that it never averages across edges.

When a given pixel is on an edge, but not in the interior of a region, its neighbors which belong to the same region of the given pixel should take a heavier weight in the averaging operation, otherwise, the edge pixel will be "blurred". Based on this idea, a weighted averaging operation assigns different weight to the neighbors. The weight assigned to a neighbor represents the confidence that the neighbor and the given pixel come from the same region. A greater weight is given to neighbors that belong to the same region as the given pixel. Specifically, the weighted averaging operation [Lev77] can be defined as:

$$g(i,j) = \sum_{(i,j) \in A} w(i,j) f(i,j),$$

where $g(i,j)$ is the enhanced image and $f(i,j)$ is the original image, $w(i,j)$ is the weight given to neighbor (i,j) within the neighborhood defined by A .

Two methods of constructing the weights $w(i,j)$ are presented in [Lev77].

3.2.2.3. Averaging in the Frequency-Domain

The averaging operation in the spatial domain is equivalent to low-pass filtering in the frequency-domain. Before we define low-pass filtering, we first present the convolution theorem [Gon77] which is important for discussing frequency domain techniques. Let $g(x,y)$ be an image produced by the convolution of images $f(x,y)$ and $h(x,y)$, that is $g(x,y) = h(x,y) * f(x,y)$. Then, from the convolution theorem, we have $G(u,v) = H(u,v)F(u,v)$, where G , H and F are the Fourier transforms of g , h and f respectively.

Noise in an image generally has higher spatial frequency than the normal image

components because it lacks spatial correlation [Pra78]. Hence, this implies that noise attenuation can be achieved by attenuating a specific range of high-frequency components in the transform of a given image. We can treat $F(u,v)$ as the transform of the image we wish to smooth. Then, we have to select a function $H(u,v)$ which produces $G(u,v)$ by attenuating the high frequency components of $F(u,v)$. The inverse transform of $G(u,v)$ will then yield the desired smoothed image $g(x,y)$. Since high-frequency components are "filtered out", and information in the low-frequency range is "passed" without attenuation, this method is commonly referred to as "low-pass filtering" [Gon77]. The image $h(x,y)$ is often called a noise cleaning mask. Two examples [Pra78] of such masks are listed in Figure 3.3.

$$h_1(x,y) = \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} \quad h_2(x,y) = \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

Figure 3.3 Examples of noise cleaning mask

These two masks are normalized to unit weights so that the noise-cleaning process does not introduce a brightness bias in the smoothed image. Also, all the elements in the mask are positive numbers. These two facts hold true for all low-pass filter masks. It should be noted that a low-pass filter using $h_1(x,y)$ as mask is equivalent to the unweighted averaging process operating in the spatial domain, using a 3×3 square window. With large window sizes, it is computationally faster to do the operation in the frequency domain using Fast Fourier Transform algorithms.

3.2.3. Selective Averaging Filters

An alternative idea to using a weighted average is to use selected neighbors for averaging. This technique adapts to a given neighborhood by estimating the nature of the neighborhood and then selecting the best subset of neighbors to compute the average. As a result, averaging across edges can be avoided. These techniques can be viewed as a special class of weighted averaging techniques which assign a equal weight to all selected neighbors and zero to the rest, according to some selection scheme.

As discussed in [Lev77], better smoothing effects can be achieved by smoothing procedures if a larger neighborhood size is used. However, as mentioned earlier, due to the small size of circumscribed masses, the neighborhood size is limited to 5×5 pixels. One solution to this problem is to iterate the smoothing procedure over a small neighborhood in order to bring more pictorial context to bear on the smoothing of individual points [Dav78].

Four selective averaging algorithms are discussed in this section. All of them were implemented and were applied to the test image (see Plate 2.1). The enhanced version of the test image is shown after the discussion of each method. As constrained by the high computational cost in applying these techniques to the test image, each procedure was iterated five times maximum. If blurring occurred (detected by human visual judgement) before the fifth iteration, the procedure was terminated.

(1) Edge preserving smoothing

This scheme [Nag79] tries to search for a homogeneous neighborhood around each pixel in an image and averages on this neighborhood only. The search is done by rotating a *mask* or a window around a given pixel and the variance in these masks is used as a measure of homogeneity of an area. Nine masks are defined for each pixel. These masks are designed so as to smooth with a homogeneous neighborhood without blurring sharp

edges.

If the neighbors of pixel M are

A	B	C	D	E
F	G	H	I	J
K	L	M	N	O
P	Q	R	S	T
U	V	W	X	Y

the masks are defined as

G	H	I
L	M	N
Q	R	S

A	B
F	G
L	M

B	C	D
G	H	I
	M	

	D	E
H	I	J
M	N	

	I	J
M	N	O
	S	T

M	N
R	S
	T
X	Y

	M	
Q	R	S
V	W	X

	L	M
P	Q	R
U	V	

F	G
K	L
P	Q
	M

To apply this scheme, the variances of these nine masks are compared with each other at each image coordinate (x,y) and the average grey level of the smallest-variance mask is given to the pixel (x,y) .

The result of applying this technique to the test image is shown in Plate 3.3.



Plate 3.3 Result of applying the edge preserving smoothing method to the test image using a window size of 3×3 pixels.

(2) Half-neighborhood method

This method [Sch80] operates like unweighted averaging at pixels in the interior of a region. For pixels that lie on an edge, the five neighbors (in a 3×3 neighborhood) which lie in the same region as the given pixel are used to compute the average. This prevents averaging across edges.

This method assumes that an edge is straight and is represented by five consecutive pixels. The eight immediate neighbors of a given pixel are partitioned into two groups. Let N be the set of eight neighbors of pixel P , such that $N = N_1 \cup N_2$, where N_1 is a set of three consecutive neighbors and N_2 consists of the five remaining neighbors. The possible combinations of N_1 and N_2 are shown in Figure 3.4.




Figure 3.4 Possible combinations of N_1 and N_2 . Points in the shaded areas belong to the set N_1 and pixels in the blank areas belong to the set N_2 .

Let $\overline{N_1}$ and $\overline{N_2}$ denote the average grey levels of N_1 and N_2 respectively. Then the criterion used to determine if a pixel P sits on an edge is

- (i) To choose the strongest edge, we choose the N_2 for which $|\overline{N_1} - \overline{N_2}|$ is the greatest.
- (ii) If $|\overline{N_1} - \overline{N_2}| > T$ (where T is a predetermined threshold), then an edge is present.

This scheme was applied to the test image, the value used for T is 5 since the minimum contrast between the suspicious areas and their background was found to be around five grey levels. The result is shown in Plate 3.4.

(a)

(b)




Plate 3.4 Result of applying the half neighborhood method to the test image using window size of (a) 3×3 pixels and (b) 5×5 pixels

(3) The k-nearest neighbor method

This method [Dav78] is based on the idea that pixels in the same region should have similar grey values. Therefore, under this scheme, the average is taken using the k neighbors (e.g. $k=5$) whose grey levels are closest to that of the given pixel. This method is simple to implement and can be viewed as a generalization of the previous method. It makes no attempt to detect the presence of an edge and it does not require that the neighbors involved in the averaging are consecutive.

The result of applying this method to the test image is shown in Plate 3.5.

(a)

(b)




Plate 3.5 Result of applying the k-nearest neighbor method to the test image using a window size of (a) 3×3 pixels and (b) 5×5 pixels

(4) Directional smoothing

This method [Ros82] is similar to the half-neighborhood method in that it tries to detect if an edge is present. If this is the case, a directional average is taken on those neighbors that lie in the direction along the edge. If the given pixel sits in the interior of a region, averaging is done using all the neighbor pixels.

To decide if a pixel is on an edge, the range of the pixels values inside a neighborhood (e.g. in a 3×3 square) of a given pixel is checked, if it exceeds a given threshold, T , the pixel is treated as an edge pixel. The edge pixel is then replaced by the average of its two neighbors in the direction of minimum edge value. The edge values are obtained by convolving the image with a set of edge detection masks in four orientations. The masks are listed in Figure 3.5.




Figure 3.5 Edge masks in four different orientations

This method is easier to implement than the half-neighborhood method because the latter method needs a prespecified algorithm for partitioning the neighborhood which are obtained either by trial and error or by prior knowledge.

This method was applied to the test image using a threshold value of 5. The enhanced image is shown in Plate 3.6.

(a)

(b)




Plate 3.6 Result of applying the directional smoothing method to the test image using a window size of (a) 3×3 pixels and (b) 5×5 pixels

3.2.4. Median Filtering

Median Filtering is a non-linear filtering technique which has been successfully applied in many speech signal analysis (eg. [Rab75]) and image processing (eg. [Hua79]) tasks. Also, it has been found to be very powerful in removing noise from two-dimension signals without blurring edges [Bov87]. Due to this property, it is suggested in [Ioa84] that the median filter is particularly suitable for enhancing medical images. Since we use this filtering method for enhancing mammogram images, a detailed definition of median filters and a review of their properties and implementation aspects are presented in this section.

To apply median filtering to a digital picture, we replace the value at a pixel by the median of the values in a neighborhood of the pixel. Before giving a more precise definition of a median filter, we first give a basic definition on the order statistic, *median*, we are using.

Given a set of n numbers $\{x_1 \cdots x_n\}$ and we define the ordered set $\{x_1^* \cdots x_n^*\}$, then the median of the set is given by

$$\text{median} \{x_1 \cdots x_n\} = \begin{cases} x_{n/2}^* & n \text{ even} \\ x_{(n+1)/2}^* & n \text{ odd} \end{cases}$$

In one dimension, the median filter of size n on a sequence $\{x_i, i \in Z\}$ is defined as :

$$\hat{x}_i = \text{median} \{x_{i-v}, \dots, x_i, \dots, x_{i+v}\}, i \in Z,$$

where $v=(n-1)/2$ and v is odd.

Two-dimensional median filters can be defined in many different ways depending on both the size and shape of the filter window. Various shapes of filter windows can be used e.g.

line segments, squares, circles and crosses. The two-dimensional median filtering operation is defined as follows :

For a two-dimensional filter window $W(i,j)$ centered at image coordinates (i,j) of a picture $\{x_{ij} : (i,j) \in Z^2\}$, the median filtering output is

$$\hat{x}_{ij} = \text{median} \left\{ x_{r,s} : (r,s) \in N(i,j) \right\}, (i,j) \in Z^2$$

where $N(i,j)$ is the area in the image covered by window $W(i,j)$.

Similar to low-pass filters, the median filter removes noise by examining the trend of the data at different points over a small region, defined by a window, and uses a statistical estimator to replace the grey level value at the centre. While a low-pass filter usually uses the average (i.e. mean), a median filter uses the median as the estimator. It is noted that sometimes the median is a better estimator of the 'trend' than the average [Bed84], because the sample average tends to be pulled towards the direction of a single aberrant data value. For example, consider the average of ten numbers, nine of them are between zero and five, and the tenth is 1000. The average of these ten numbers surely does not reflect the 'trend' of the majorities, while the median does. As a result, low-pass filtering is not adequate for certain applications because it tends to blur the image and smear out the edges in the original image.

It is observed that median filters have several interesting properties that sometimes make them superior to low-pass filters[Kuh81]. If an image has impulse-like noise, median filtering can remove them without significantly modifying other components, and if an image contains edges, median filtering can preserve them. This can be explained by the fact that the median over a neighborhood does not "respond" to an edge when only a small fraction of the neighborhood overlaps the edge. In addition, Bovik *et al.* have successfully applied generalized filters based on linear combinations of order statistics to the image filtering problem, and demonstrated that among those the median filter is nearly

optimal for suppressing noise which is characterized by a large percentage of outliers, i.e. very heavy-tailed or impulsive noise [Bov83].

As mentioned earlier, the median filter is a non-linear filter. That is, in general,

$$\text{median} \left\{ f_{ij} \right\} + \text{median} \left\{ g_{ij} \right\} \neq \text{median} \left\{ f_{ij} + g_{ij} \right\}$$

where f_{ij} and g_{ij} are two different pictures.

Unfortunately, as is true for all non-linear operations, it is difficult to perform quantitative analysis on the behaviour of median filters. Nevertheless, for one-dimensional median filters, this has, to a large extent, been done by *Gallagher* and *Wise* [Gal81, Nod82]. They studied and defined various concepts associated with median filters. They formalized ideas such as "impulse-like", "edge" and "constant neighborhood". Also they proved that successive median filtering of a signal eventually reduces the original signal to an invariant signal called a *root* signal. This suggests the concept of a filter *passband* and *stopband*. Signals which do not reside entirely within the filter passband can be reduced to their passband component by repeated filtering. At its stopband, the filter is invariant to subsequent filtering. They showed that only signals composed of constant neighborhoods and edges are roots to the median filter. This implies that the smoothing power of median filters can be improved by iterating the filtering process and that edge information is not lost due to such iterations.

The extension of the above properties of one-dimensional median filters to two-dimensional filters is discussed in [Nod82]. It was noted that although preliminary studies showed that root (invariant) signals may also exist in two-dimensional median filtering, it is still not clear if or how most of the other properties can be extended to the 2D case. Part of the problem is that of defining terms, such as constant neighborhood in more than one dimension; also, two-dimensional windows can take on many different shapes, many of which have been found to be useful.

Furthermore, some statistical properties of one-dimensional median filters are analyzed in [Ata81]. *Kuhlmann* and *Wise* analyzed the output autocorrelation function and the output power spectrum of median filtered stationary sequences of independent data and they observed that one-dimensional median filters behave like a low pass filter when the filter size is small [Kuh81]. To improve the power of the median filter in removing non-impulsive noise components, a modification of the median filter has been developed in [Lee84] which combined desirable properties of both linear and non-linear filters.

For the implementation aspect of median filters, one criticism which is sometimes raised against median filters is that they are too slow for routine large data volume applications due to the large number of sorting involved. *Huang et al.* have introduced a fast algorithm [Hua79] for two-dimensional median filtering, the algorithm was reported as much faster than conventional sorting methods. For a window size of $n \times n$, the computing time required is $O(n)$. This algorithm is based on storing and updating the grey level histogram of the picture elements in the window and the amount of sorting required is minimized. For real time application, *Atama et al.* introduced a fast method [Aat80] for real time median filtering. A special hardware device for median filtering suitable for parallel processing is presented in [Dan81]. In addition, a separable median filter is introduced in [Nar81], although it is not a true two-dimensional filter, it was observed that its performance in image noise smoothing is close to that of the two-dimensional case. It was shown that the separable filter has a much simpler implementation in real-time hardware.

3.3. Proposed Solution

There are two possible approaches in enhancing mammographic features. One is to increase the contrast of suspicious areas and the other is to remove background noise.

Some techniques for contrast enhancement of film mammograms have been suggested earlier [Dha86, Gor84]. Their method is based on adaptive neighbourhood processing with a set of contrast enhancement functions or an optimal one to enhance the contrast of mammographic features. In our research, we take the other approach and enhance the images by removing background noise while preserving the edge information in suspicious areas in the images.

As in most medical images, the structure of the signal in mammogram images is not well defined, so that the design of a more specific filter to remove noise in mammogram images is not feasible. Noise reduction and edge preservation are the two most important requirements in choosing a filtering method. The ability of the median filter to provide both noise reduction and edge preservation makes it promising for this application.

3.3.1. Selective Median Filtering

Investigation shows that the edge preservation power of a pure median filter is not sufficient for enhancing mammogram images due to the fuzziness of the boundaries of suspicious areas. In order to further preserve the boundaries of suspicious areas, a modification of the median filter is introduced and is termed Selective Median Filter (SMF). The two-dimensional Selective Median Filter is defined as follows :

For a window $W(i,j)$ centered at image coordinates (i,j) , the selective median filtering output is

$$\hat{x}_{ij} = \text{median} \left\{ x_{r,s} : (r,s) \in N(i,j) \text{ and } |x_{r,s} - x_{ij}| < T \right\}, (i,j) \in Z^2$$

where $N(i,j)$ is the area in the image covered by window $W(i,j)$ and T is a threshold.

Thus, in computing the median, the set of pixels is restricted to those with a difference in gray level no greater than T . By adjusting T , the amount of edge smearing

can be controlled. This modification of the median filter is related to selective averaging schemes developed for linear filters [Ros82] that show good results in improving the edge preserving power of linear low-pass filters.

In general, to achieve sufficient noise suppression, one needs either a filtering technique allowing a large window size, or the filter has to be applied repeatedly to the data [Ioa84]. Median filters act as low-pass filters in homogeneous areas and as their window size increases, they respond with increasingly narrow pass-bands [Gal81]. Also, it was shown in [Gal81] that as the window size increases, noise is reduced but distortion is introduced into the actual signal. As discussed earlier, iteration of median filters will not remove edge information from the image. Therefore in using the SMF, we choose to iterate the filtering operation to achieve sufficient noise reduction.

The pure median filter has no design parameters other than the window size and cannot be adjusted to the given signal and noise characteristics. Such filtering may degrade important information in some type of images or possibly introduce artifacts [Ioa84]. The selective median filter, SMF, has three design parameters that can be varied to adjust the filter to the noise characteristics of all mammogram images. The parameters are the window size W , the gray level difference threshold T and the number of iterations. The effect of the SMF is controlled mainly by the threshold T . If T is small, the edge preserving power of SMF is strong, but its smoothing effect will be small. If T is large, the SMF behaves the other way round.

SMF using a T value of 5 was applied to the test image, the result is shown in Plate 3.7.

(a)

(b)




Plate 3.7 Result of applying selective median filtering to the test image using a window size of (a) 3×3 pixels and (b) 5×5 pixels

3.3.2. Evaluations and Comparison

In this section, the performance of the neighborhood averaging methods discussed in Section 3.3 and the proposed method are compared and evaluated. In enhancing mammogram images, we require an enhancement technique which clears noise while preserving edges of suspicious areas.

To evaluate the performance of the various techniques, by following the literature [Hal71], some quantitative measurements were computed for the enhanced images. They are

- (1) the zero-shift correlation,
- (2) the mean and variance,
- (3) the mean error and mean-square error.

In computing the zero-shift correlation, mean error and mean square error, the reference image is the original image. The correlation value measures the amount of change in the picture structure of the enhanced image, an enhanced image which has little resemblance to the original image will produce a small correlation value. The other measurements measure the amount of change of the overall grey level values in the enhanced images. For instance, a large mean-square error implies a great change. The results are tabulated in Table 3.1.

From Table 3.1, we can observe that the SMF-enhanced images are the only ones that have larger mean, variance and mean error values than those of the original image. This implies that the SMF tends to pull up the intensity of an image slightly in its smoothing process. In addition, the correlation measures of all the enhanced image are slightly less than one. This shows that the original picture structure is not altered severely by the smoothing process.

	Mean	Variance	Correlation	Mean Error	MSE*	Time (sec)
Original Image	34.82	834	1.000	0.000	00	-
SMF (3)	35.60	850	0.997	0.782	187	6
SMF (5)	34.94	838	0.996	0.128	333	24
Edge Preserving (5)	33.76	820	0.995	-1.053	467	24
Half-neighborhood (3)	34.09	819	0.997	-0.727	222	7
Half-neighborhood (5)	33.99	812	0.990	-0.821	1724	19
K-nearest neighbor (3)	33.39	804	0.997	-1.426	171	7
K-nearest neighbor (5)	34.32	826	0.999	-0.454	37	35
Directional (3)	34.14	827	0.999	-0.316	09	4
Directional (5)	34.16	827	0.999	-0.363	19	7
Unweighted (3)	34.45	820	0.998	-0.364	68	2
Unweighted (5)	34.39	815	0.996	-0.429	336	4

Table 3.1. Quantitative measurements of the enhanced images
*mean square error reduced by 4 orders of magnitude

These quantitative measurements cannot measure the amount of semantic structure in an image. For example, they cannot tell if important information, such as edges of suspicious areas, are enhanced or removed in an image. In fact, at present, no objective method is available to judge which image is best enhanced. In my own view, the most effective method in comparing images is still subjective human evaluation. To facilitate the evaluation of the edge preserving power of the various enhancement techniques, the boundary information of the enhanced images was extracted for human evaluation. A high-pass (Laplacian) filter was used for this purpose, the high-pass filter output is defined as

$$g(x,y) = \left| \sum_{i=-1}^1 \sum_{j=-1}^1 f(x+i,y+j) - 9f(x,y) \right|$$

where $f(x,y)$ is an enhanced image.

The high-pass filtered image was then converted to a binary image using a fixed threshold. The binary high-pass filtered images of the original test image and of the enhanced images produced by the various techniques are shown in Plate 3.8 through Plate 3.19. By examining these images, it is observed that among the six techniques implemented, the unweighted averaging method has the best noise-removal power, but it blurs edges. Blurring increases as the window gets larger. Therefore, this technique is not suitable for our application.




Plate 3.8 Laplacian of the test image in Plate 2.1



Plate 3.9 Laplacian of the image in Plate 3.2a (smoothed by unweighted averaging using a window size of 3x3)




Plate 3.10 Laplacian of the image in Plate 3.2b (smoothed by unweighted averaging using a window size of 5x5)




Plate 3.11 Laplacian of the image in Plate 3.3 (smoothed by the edge-preserving smoothing method)



Plate 3.12 Laplacian of the image in Plate 3.4a (smoothed by the half-neighborhood method using a window size of 3×3)




Plate 3.13 Laplacian of the image in Plate 3.4b (smoothed by the half-neighborhood method using a window size of 5×5)




Plate 3.14 Laplacian of the image in Plate 3.5a (smoothed by the k-nearest neighbor method using a window size of 3×3)




Plate 3.15 Laplacian of the image in Plate 3.5b (smoothed by the k-nearest neighbor method using a window size of 5×5)




Plate 3.16 Laplacian of the image in Plate 3.6a (smoothed by direction smoothing using a window size of 3x3)




Plate 3.17 Laplacian of the image in Plate 3.6b (smoothed by direction smoothing using a window size of 5x5)




Plate 3.18 Laplacian of the image in Plate 3.7a (smoothed by selective median filtering using a window size of 3x3)




Plate 3.19 Laplacian of the image in Plate 3.7b (smoothed by selective median filtering using a window size of 5x5)

The directional smoothing method cannot preserve edges well when compared with the other methods and it tends to preserve straight edges better. Intuitively, this can be explained by the fact that this method assumes that edges are quite straight, such as those in Figure 3.6. However, in mammogram images, the shape of suspicious areas is approximately circular, therefore, edges do not always lie in a straight line (see Figure 3.7). In this case, no matter which direction the average is computed, the edge pixel is always being averaged with a pixel which does not belong to the same region as the edge pixel. As a result, the edges of suspicious areas are blurred. As the window size increases, this effect is more prominent.

1	1	0
1	1	0
1	1	0

1	1	1
1	1	1
0	0	0

Figure 3.6 Examples of straight edges

1	1	0
1	1	0
0	0	0

1	1	1
1	1	0
0	0	0

Figure 3.7 Examples of non-straight edges

Although the k-nearest neighbor method is easy to implement, it can not preserve edges as well as the half-neighbor or SMF method. The k-nearest neighbor method assumes that the five nearest neighbors (in grey level) must belong to the same region of the centre pixel in a window. This is not always true and the five nearest neighbors may include pixels that lie outside the region. The half-neighborhood method preserves edges better because it treats edge pixels differently and restricts that the five neighbors selected for averaging must be consecutive pixels. This latter restriction reduces the chance of averaging an edge pixel with pixels belonging to another region. By using a threshold to select pixels from which the median is chosen, the SMF achieves the same effect.

In comparing the high-pass results of the various enhanced images, it was noted that only the edge preserving smoothing method (Plate 3.11) and the SMF method (Plate 3.19) preserve the edges of the suspicious areas well. In both plates, the edges of the suspicious areas appear as approximately closed rings. However, the noise cleaning power of the SMF method is better than that of the edge preserving smoothing method. This can be illustrated by the fact that there are less noise edges in Plate 3.19 than that in Plate 3.11. Based on these observations, it was concluded that the SMF method with a 5×5 window size enhances the mammogram images better than the other methods. In addition, although Table 1 shows that the computation time of these two techniques is close, the fast algorithm designed for median filtering (see section 3.2.4) makes SMF more attractive (in terms of computation time) than the other methods.

3.3.3 Performance of the Proposed Method

To further evaluate the performance of the SMF, a set of 24 mammogram films were used. They were randomly selected from mammogram files maintained by the Radiology Department of the Cross Cancer Institute. They were diagnosed by an expert radiologist who circled the suspicious areas with wax pencil in each test case. His mark-

ing were then translated into x, y coordinates representing the approximate center and radius measures.

The SMF has three parameters that can be adjusted to adapt the filter to the noise characteristics of the mammograms. The three parameters are the threshold T , the number of iterations and the window size W . To set the values of these three parameters, a training set of seven mammograms are used. The other seventeen mammograms are grouped into a test set for testing the performance of the proposed method after the values of the parameters are determined.

To estimate the maximum value for the parameter T , the contrast between pixels in the suspicious areas (of the training set images) were checked. It was noted that the contrast ranges from 5 to 15 grey levels. Using different threshold values within this range, the SMF was applied to the training set. The high-pass result of the enhanced versions of the test image (Plate 2.1) are shown in Plates 3.20 through 3.23. By observing the amount of noise edges and "relevant" edges in these plates, it was concluded that although the use of a large threshold value, such as 10, can remove noise more effectively (as shown in Plate 3.23), a smaller threshold value, such as 5, can preserve edges of all suspicious areas better (Plate 3.22). Therefore, to preserve the edges of all suspicious areas and to remove a moderate amount of noise, a threshold value of 5 was selected for the SMF.

In addition, it was noted that the performance of SMF does not improve much after the 5th iteration if the other two parameters (W and T) are kept constant. This is in agreement with Gallager and Wise's [Gal81] observation that every signal can only be reduced down to a certain point no matter how many times we apply the median filter. At its stopband the filter is invariant to subsequent filtering.

Best performance of the SMF was found for window size $W = 9 \times 9$, number of itera-

tions = 5 and edge threshold $T = 5$. Using these parameter values, the SMF was applied to the mammogram image 1 (Plate 2.1) and 2 (Plate 2.2) which are members of the training set. The filtered images are shown in Plate 3.24. In comparison to the original images, the background variation is reduced in the filtered images while the boundaries of all suspicious areas are preserved. However, it was noted that the filtered images have a mottled appearance and this is caused by the iteration of the filter. This effect is not desirable and it may produce false alarms in the tumor detection process. This problem can be easily overcome by applying a false alarm test using the original unfiltered image. This test is discussed in detail in Chapter 4. In Plate 3.25, areas that may produce this kind of false alarms in the images shown in Plate 3.24 are marked by circles.

In applying this SMF on the test set of seventeen mammogram images, the performance of the SMF was found to be satisfactory. The background variations in the filtered images is reduced and the boundaries of all suspicious areas are preserved.

3.4. Summary

In this chapter, several general image enhancement techniques are discussed and a proposed method, Selective Median Filtering, is presented. Five existing image smoothing techniques were implemented and their performance is compared with the proposed method. These five techniques are unweighted averaging, edge preserving smoothing, half-neighborhood smoothing, k-nearest neighbor smoothing and directional smoothing.

In enhancing mammogram images, an effective enhancement technique is required to remove noise and preserve edges of suspicious areas. But these two requirements are, very often, conflicting. Hence, in choosing an optimum method, we prefer a method which works moderately well in both noise clearing and edge preserving, rather than one which removes noise effectively but does not preserve edges or one that does the opposite. Based on this evaluation criterion, we found that the Selective Median Filtering

method is more effective than the other five techniques in enhancing mammogram images.

A training set of 7 mammogram images were used to set the parameters of this Selective Median Filtering method. It was then applied to 17 test images and its performance is satisfactory.



Plate 3.20 Laplacian of the SMF-enhanced mammogram image 1, threshold $T = 3$




Plate 3.21 Laplacian of the SMF-enhanced mammogram image 1, threshold $T = 5$




Plate 3.22 Laplacian of the SMF-enhanced mammogram image 1, threshold $T = 7$





Plate 3.23 Laplacian of the SMF-enhanced mammogram image 1, threshold $T = 10$

(a)



(b)

— Plate 3.24 Results of applying selective median filtering to (a) mammogram image 1 and (b) mammogram image 2, using a threshold value of 5, window size of 9×9 pixels and five iterations




Plate 3.25 False alarms caused by iterating the SMF are marked in the enhanced images

Chapter 4

Object Detection

4.1. Introduction

Object detection is a typical image analysis problem which involves the extraction of objects from their background. In three-dimensional scenes, object detection is complicated by the requirement of an understanding of the imaging geometry, that is, how three-dimensional scenes are mapped into two-dimensional images. In addition, besides the problem of translation, rotation and scale distortion, the image obtained from a three-dimensional scene may suffer from perspective distortion as well. This thesis is concerned with detection of objects in mammogram films which are practically in a two-dimensional medium. Hence, the discussion following is limited to object detection techniques applied to two-dimensional scenes only.

Ideally, an object detection technique should possess two characteristics. First, the technique should be able to detect an object in an image independent of the transformation (position, orientation and size) of the object relative to those specified. Second, it should give positive response only to truly matched objects. In other words, the false alarm rate of the technique must be kept minimal. However, in reality, none of the existing object detection techniques possesses all these characteristics. In addition, to solve an object detection problem in a particular domain, one would have to combine and tailor some existing general techniques according to the image characteristics and nature of the object. To gain some knowledge on the techniques available for tumor detection, a brief survey of existing object detection techniques is presented in this chapter.

Many techniques have been developed for object detection, the two most general approaches are *template matching* and *image segmentation*. Template matching is the

simplest approach to the problem of locating an object in an image. In this approach, a *template* which is a standard representative pattern of the object is compared to all locations in a data image. The degree of match at each image location is represented by some similarity measure. Then by interpreting this metric, the location of the match position, if any, is known. The advantages of this approach are that detection is not affected by translation of the object and the shape of the object can be integrated into the criterion for matching very easily, though the match may be disturbed by shape distortion. Also, the location of the detected object is easily obtained from the output of the process. On the other hand, the main disadvantage of the template matching approach is that the detection is sensitive to rotation and size of the object. Modifications of the method are required to detect objects that will appear in different orientations and sizes in the data image. A detailed discussion of this approach is given in Section 4.2.1.

The image segmentation process partitions an image into disjoint regions which usually correspond to meaningful units. In solving the problem of object detection, ideally, the segmentation process partitions the image into object(s) and background segments. Since the orientation, size and shape of the object seldom affects the partition decision, this approach to object detection has the advantage of being insensitive to rotation, size and shape distortion of the object. However, the disadvantage of this approach is that a segmented image provides no direct information on the size, shape and location of the object segment. Very often, further processing is required to extract them. Image segmentation involves a large range of techniques that are developed for image analysis. But not all of them are applicable to the problem of object detection. No attempt is made in this chapter to review all segmentation techniques but only those which are generally applicable are included in the survey presented in Section 4.2.2.

In the case of mammograms, tumor detection is relatively difficult because the low quality mammogram images are noisy and the tumors are not well defined. A method for

breast tumor detection is proposed in Section 4.3. This method uses a template matching technique to detect suspicious tumor areas in mammogram images. Despite the disadvantages of the template matching method mentioned above, the proposed method is insensitive to size and orientation of the tumor. Also, to reduce the false alarm rate, two false alarm tests are designed to discriminate noise from true suspicious areas. The performance of this method is evaluated by applying it to 24 test images. The result is presented in Section 4.3.6.

4.2. Survey of Existing Methods

4.2.1. Template Matching

In this approach, object detection is achieved by a window search process in which a template is matched with all locations in a data image to find a match location, if any. The object to be detected is usually known *a priori* so that the template can be constructed based on representative features of the object. This method produces another image of the same dimension as the data image in which each pixel value gives an indication of the degree of match of the template in its neighborhood. In interpreting these match values or similarity measures, one can pick the location with the peak value as the match position, or some thresholding method can be used to select a set of candidate points which are possible locations of the object.

The disadvantage of this approach is that the match is sensitive to the size and orientation of the object. In general, if a geometrically distorted copy of the template is present in the data image, it will be left undetected due to the low similarity measure produced in its immediate vicinity. To circumvent this problem, multiple templates are used.

Other important issues of this technique are:

- (1) the similarity measure between image and template,
- (2) the criteria in selecting match positions,
- (3) the efficiency of the matching process.

These issues are discussed in the following sections.

4.2.1.1. Similarity Measures

Major similarity measures used in template matching are the absolute differences and the cross-correlation coefficient. Other similarity measures have also been used in the window search process such as invariant moment [Won78a] and Walsh-Hadamard transform coefficients [Sch80b]. In this section, the most widely used similarity measure, cross-correlation, is discussed.

- Let S be the image, an $L \times L$ array of pixels each taking one of K gray levels; W be the template, which is $M \times M$ with $M \ll L$. Each $M \times M$ subimage of S can be uniquely referenced by its upper left corner coordinates (i, j) . There are $(L-M+1)(L-M+1)$ such (i, j) 's.

The unnormalized cross-correlation between image and template is defined as

$$R(i, j) = \sum_{k=1}^M \sum_{m=1}^M W(k, m) S(i+k-1, j+m-1)$$

and locations can be defined as peaks of $R(i, j)$ over the whole image area.

One disadvantage of using unnormalized cross-correlation is that it is sensitive to local properties of the image, such as its average brightness. As a result, the maximum of the unnormalized correlation does not always yield the real location with the exact match [Bal82]. Improved detection performance is achieved by using normalized cross-correlation, defined as [Ros82]:

$$R(i,j) = \frac{\sum_{k=1}^M \sum_{m=1}^M \left\{ \left[W(k,m) - \mu_w \right] \left[S(i+k-1, j+m-1) - \mu_s(i,j) \right] \right\}}{\sqrt{\sum_{k=1}^M \sum_{m=1}^M \left[W(k,m) - \mu_w \right]^2 \sum_{k=1}^M \sum_{m=1}^M \left[S(i+k-1, j+m-1) - \mu_s(i,j) \right]^2}}$$

where μ_w is the mean of the template and μ_s is the mean of the subimage centered at image point (i,j) .

This normalized cross-correlation has an absolute value less than or equal to one. It takes on a maximum value of one where the template and the subimage are identical. In addition, it is less dependent on the local properties of the subimage than is the unnormalized correlation, but it is more sensitive to noise that may mask the peak correlation [Li86]. This may lead to multiple match candidates. However, this merely introduces false alarms, but does not discard true target areas.

Note that for a given window W , the term $\sum_{k=1}^M \sum_{m=1}^M \left[W(k,m) - \mu_w \right]^2$ is a constant.

Therefore, if only one template is used, the computation can be reduced to

$$R(i,j) = \frac{\sum_{k=1}^M \sum_{m=1}^M \left\{ \left[W(k,m) - \mu_w \right] \left[S(i+k-1, j+m-1) - \mu_s(i,j) \right] \right\}}{\sqrt{\sum_{k=1}^M \sum_{m=1}^M \left[S(i+k-1, j+m-1) - \mu_s(i,j) \right]^2}}$$

The process of cross-correlating a template with an image can be justified by the *matched filter theorem* [Ros82]. This theorem states that under certain assumptions, the best filter to use for finding matches between a template f and an image g is f itself. Furthermore, the cross-correlation operation is equivalent to convolving g with f rotated by 180 degrees.

4.2.1.2. Criteria in Selecting Match Positions

A template match is rarely very exact due to image noise and *a priori* uncertainty as to the exact shape and structure of an object to be detected. Consequently, to choose the best matched subimage, one of the following can be used :

- (1) pick the location with the maximum cross-correlation value,
- (2) pick the locations with the maximum cross-correlation value for each of a number of templates used, if more than one template is applied,
- (3) pick all the locations with cross-correlation value exceeding a threshold.

The choice of a selection criterion depends on the nature of the object detection problem. In a situation where a unique match is required, (1) can be used. But in other situations, such as tumor detection in mammogram images, where the object may be found in more than one location in the an image, (2) or (3) should be considered.

4.2.1.3. Efficiency of the Matching Process

One disadvantage of template matching using cross correlation measure is its high computational cost. For an image of size $L \times L$ and a template of size $M \times M$, the number of operations (addition, multiplication, division and comparison) required is proportional to $M^2(L-M+1)^2$.

To speed up the matching process, the two general approaches are:

- (1) speed up the computation at each test location, and
- (2) reduce the number of test locations to be fully checked.

For the first approach, a typical solution is to compute the cross-correlations in the frequency domain via a Fast Fourier Transform algorithm. By the convolution theorem, cross-correlating a template with an image in the spatial domain is equivalent to multiplication in the frequency domain. However, the time saving is significant only if the

template size is large. The computation requirement is the same for normalization. The trade off of this method is the large amount of memory space required.

Barnea and Silverman proposed a *sequential decision method* [Bar72] which is based on the idea that since the matching process spends most of its time in checking locations that do not match the object, much time can be saved if we can discover a mismatch earlier (before all locations in the template are tested). Their method starts by calculating the measure of match (sum of absolute differences) using corresponding pairs of sample elements out of the subimage and the template in pseudo-random order. If at any time, the measure becomes non-optimal according to their decision criteria, the calculation is discontinued and the area is discarded as non-matching. Otherwise, the calculation is continued until the whole area has been included. They claimed that this method was a factor of 50 faster than the FFT correlation method.

In the above methods, although less time is needed for the computation at each test location, all locations in the image are involved in the matching process. In fact, a lot of effort is wasted due to the large amount of local coherence existing in most images. That is, in most images, an area centered at one pixel does not usually differ greatly from an area centered at a neighboring pixel. Due to this reason, the second approach to speed up the template matching process leads to very efficient algorithms.

A *two-stage template matching* approach [Van77] proposed by *Vanderburg and Rosenfeld* is a typical example of methods using the second approach. Their method divides the process into two stages. In the first stage, a subtemplate which is a subset of template pixels, is used for matching. In the second stage, the entire template is matched only at those positions where the subtemplate has matched well. In [Ros77], the same group proposed a similar approach termed *Coarse-Fine template matching*. In the first stage of this method, a coarse template (ie., reduced-resolution template) is matched with a coarse image (ie., reduced-resolution image). Only those positions where the coarse

template has matched well are used to match with the full resolution template in the second stage.

The computational savings of these two methods are due to the reduction in the number of test locations needed to be checked thoroughly. It should be noted that although they both divide the matching process into two stages, they construct the subtemplate in different ways. A reduced-resolution subtemplate does include the information of every pixel in the original template, but it is not true in the case of the *two-stage template matching* approach because only a subset of template pixels is used. This difference may have impact on the performance of these methods.

Wong and Hall proposed a even more efficient *hierarchical scene matching* algorithm [Won78b]. This algorithm first creates two "pyramids" which are a set of search area images and a set of templates in decreasing resolution. The search for a matching position starts at the lowest resolution level, that is, the top of the pyramid and then it proceeds to the next higher resolution level. At each search level k , a search is made only at test locations where a sufficient match has been recorded in the $(k+1)$ th level. As k decreases, fewer and fewer test locations are checked. The most promising locations are examined at level 0. At the lowest resolution level, the number of possible test location is reduced to $\left[\frac{L}{2^k} - \frac{M}{2^k} + 1 \right]^2$ where k is the total number of search level, L and M are the dimension of the search area and the template respectively. Comparing this to the possible test locations of $(L-M+1)^2$ at the highest resolution when $k=0$, there is a reduction of nearly 2^{2k} in possible location of computation. To further minimize the computation, the search at each level of resolution can be guided by the sequential testing and detection techniques used in [Bar72].

4.2.2. Image Segmentation

The image segmentation process partitions an image into disjoint regions such that picture elements with similar properties are grouped into the same region. Each region usually corresponds to a meaningful unit and the definition of its meaning is domain dependent. In the case of object detection, the resulting segments should correspond to the object(s) and the background. One drawback of this approach to object detection is that it is difficult to extract geometrical information, such as location, shape and size, of the detected object from a segmented image. Further processing is often required to complete the job. For most segmentation techniques, it is difficult to use the size, orientation or shape of an object as a criterion for segmenting an image. As a result, the segmentation process can operate successfully without regard to the size and orientation of the object in an image. Also, it can stand large shape distortion of the object.

The main issue for this technique is the property used for partitioning an image. Based on this issue, segmentation techniques can be categorized into three types :

- (1) *Pixel-based* – global techniques which segment an image according to the property of picture elements without regard to their positions or any local conditions in the image.
- (2) *Edge-based* – techniques which detect an object by outlining the object boundary within an image.
- (3) *Region-based* – techniques which segment an image according to local regional properties.

These three types of techniques are discussed in the following sections.

4.2.2.1. Pixel-based Segmentation

Pixel-based segmentation techniques attempt to group picture elements having similar properties, such as brightness or color, into regions without regard to their position in the image. The simplest method of this kind, *grey-level thresholding*, partitions an image based on the grey level value of each image point. Assuming that the object is a bright area in a dark background, this approach segments the object from its background by selecting a threshold grey level value, T , such that pixels with grey level larger than T belong to the object region and the rest are in the background region. Then the position of the regions can be shown by forming a new image $g(x,y)$, whose grey level corresponds to the region number. The decision rule [Hal79]

$$g(x,y) = \begin{cases} 1 & \text{if } f(x,y) \geq T \\ 0 & \text{if } f(x,y) < T \end{cases}$$

applied to each point of the original image, $f(x,y)$, would produce the new image.

The effectiveness of this approach depends on the selection of an appropriate threshold which correctly separates the object grey levels from the background. In a simple situation, such as an image with homogeneous objects in a homogeneous background of sharp contrast, the threshold can be determined by choosing the point that separates the two peaks in the bimodal intensity histogram of the image [Ros82, pp.62]. Usually the histogram is first smoothed and then an exhaustive search is made for the peaks and valleys. This approach assumes that the object grey levels are similar and the size of the object is big enough to create a significant peak in the histogram. As a result, the histogram contains peaks corresponding to grey levels of objects, while the valleys are assumed to result from edge pixels that are much less numerous.

If there is more than one object in the image and their brightness is not the same, a single threshold may not be enough to segment the image and indeed, the intensity histogram may contain more than two peaks. In such cases, several thresholds may have to be

used [Ros82, pp. 66].

In the case of radiographs, threshold selection is complicated by the fuzziness of the object boundary, the small size of the object and the varying grey levels of the background. In most cases, the intensity histogram is not bimodal, and exhibits only one distinct peak [Tou79]. Hence, the simple thresholding technique mentioned above fails to generate a solution. Some modifications of the grey level thresholding method have been developed to circumvent these problems.

For example, Tou and Liu introduced the concept of *intensity histogram enhancement* to produce a bimodal histogram and a *zoom thresholding technique* to select an optimum threshold [Tou79]. Instead of using only global information, their method uses edge points, i.e. local information, in an image. To generate a bimodal histogram, only a fixed percentage of pixels with high Laplacian values are selected to generate the enhanced histogram. Hence, with a suitable percentile value, only points that lie in the immediate vicinity of the object boundary are used to plot the histogram. Among these points, the probability that a point lies in the background is about equal to the probability that it lies in the object. Thus, the enhanced histogram should be bimodal.

The *zoom-thresholding technique* determines an optimal threshold by minimizing the criterion

$$\Delta A(t) = \left| \sum n(t) - \sum n(t+1) \right|$$

where $n(t)$ is the number of pixels with intensity equal to or greater than the threshold t and $\Delta A(t)$ denotes the rate of change of the object area using the threshold t .

This technique is based upon the idea that in the neighborhood of the true boundary of the object, the rate of change of the object area, A , is minimal. Therefore, the threshold, t , "zooms" to the true boundary threshold when $\Delta A(t)$ reaches a minimum. These techniques have been applied to the problem of determining lung tissue boundaries in

lung tissue micrographs. The result was found to be satisfactory based on two test cases [Tou79].

Another optimum threshold selection technique is developed by Chow and Kaneko for outlining boundaries of the left ventricle in cardioangiograms.¹ In their application, a single global threshold is inadequate due to contrast differences throughout the image, therefore, a *dynamic thresholding technique* is used [Cho72]. Their technique divides the image up into partially overlapping rectangular subimages and computes a threshold for each subimage. If a subimage fails to have a threshold (ie. the histogram is not bimodal), it receives interpolated thresholds from neighboring subimages that are bimodal. Finally, the entire picture is thresholded using separate thresholds for each subimage. In the last step of the process, the thresholds are selected dynamically; the value of the threshold at a point depends on its proximity to boundary points which have bimodal histograms.

From these examples, we can see that the power of pixel-based segmentation technique can be improved by considering global information as well as local information. A detailed discussion of the various extensions of this segmentation technique is presented in [Ros82, pp.61-83].

4.2.2.2. Edge-Based Segmentation

Boundaries are regions between objects and the background, therefore, the object detection problem can be considered as a problem of finding boundaries between these two regions. Edge detection or edge enhancement is a widely studied topic, involves techniques which aim at extracting primitive edge elements from an image. But these edge elements are often disconnected and do not yield enough information for object detection without further processing. The goal of edge-based segmentation is to connect individual local edge elements to produce meaningful boundaries. In object detection,

¹ Cardioangiograms are x-ray pictures of a heart which have been injected with a dye.

the segmentation is considered correct, if the resulted boundaries correspond to object boundaries.

In the following discussion, the edge detection step is treated as a preprocessing step necessary for most edge-based segmentation techniques. Since, it is not the focus of this survey, the various techniques developed for edge detection are not discussed. A detailed survey of these techniques can be found in [Bal82, Ros82].

Existing techniques for edge-based segmentation vary in the amount of prior knowledge incorporated into the segmentation process. Some techniques require prior knowledge such as an estimates of object location, while others make very little assumption on the image content and are generally applicable. Three generally applicable approaches are discussed in the following :

(1) Hough Transform

The classical Hough technique for curve detection is applicable if the location of the object boundary is not known, but the shape is known and can be described as a parametric curve, such as a straight line or conic. Its main advantage is that it is relatively unaffected by gaps in curves and by noise [Bal82].

This method [Dud72] involves the transformation of a curve in Cartesian coordinate space to a point in polar coordination space. Consider the detection of straight lines in an image. A straight line shown in Figure 4.1 may be parametrically described by the angle θ of its normal and its algebraic distance ρ from the origin. After the Hough transform, the equation of this straight line becomes

$$\rho = x \cos \theta + y \sin \theta$$

The important property of the Hough transform is that this line can be defined by a single point (ρ, θ) in the θ - ρ plane. Also, a family of lines that pass through a common point in the x - y plane is transformed to a sinusoidal curve in the θ - ρ

plane. Hence, it is easy to show that the curves in the θ - ρ plane which correspond to a set of colinear points in the x - y plane will have a common point of intersection.




Figure 4.1 A straight line in Cartesian coordinate space

Suppose we have a set of edge points and we want to find a straight line that fits them. Based on the above properties of Hough transform, all points are first transformed into sinusoidal curves in the θ - ρ plane. Then the problem can be solved by finding concurrent curves in the θ - ρ plane. The implementation of this method is discussed in [Dud72].

The generalization of this technique to other curves is simple. For example, if the object shape is circular, then we will like to transform edge points into the (a, b) space where in the x - y plane, the equation of the circle is

$$(x-a)^2 + (y-b)^2 = a^2 + b^2.$$

However, for other conics which have many parameters, the Hough technique becomes impractical due to the large amount of memory required and the increased computation time.

This technique has been used successfully in several medical applications which involve radiographs, such as in the detection of tumors [Sk179] and ribs [Wec77] in chest radiographs.

(2) Boundary Following

Boundary following is also known as boundary tracking. When the shape of the object is not known *a priori*, this method is commonly used to obtain global line or curve-information from local edge elements. It operates by beginning at a point where a local edge occurs and searches for a succession of each points that lie along a curve. This approach usually finds the curves in the image one at a time. Many tracking algorithms have been developed, a detailed survey can be found in [Ros82, pp.133-137].

The main issues of this method are the search mechanisms used to search for successive points that lie along a curve and the tracking criterion for acceptance of the next point of interest. Classical tracking techniques use a search mechanism which proceed along the best path (curve) at each step and never reverse a decision once made. An alternative to this is to allow backup, ie, if an acceptance decision seems to be leading to a series of poor subsequent acceptance, one can go back and alter the decision. In fact, the tracking problem can be viewed as one of searching for optimum curves that stretches across an image [Ros82]. Optimality might be defined in terms of maximizing an evaluation function which is an estimate of the value of the best curve in the picture that passes through a point. Good evaluation functions are discussed in [Bal82, pp.133].

To find optimum curves, heuristic search techniques can be used. A heuristic technique would not only evaluate a given point on the basis of how it can be reached from the starting point, but would also attempt to look ahead and estimate

how the goal (a curve that crosses the image) can be reached from the given point. Based on this technique, *Ramer* suggested a system that used heuristic search to organize bottom-up edge following in an image [Ram75]. In medical applications, *Greer* used a limited look-ahead heuristic search to successfully locate lung boundaries in tomographs [Gre78].

The problem of finding optimum curves can also be solved using a *dynamic programming technique* [Mon71]. In this approach, the criterion for selecting the best curve is embedded in a merit function which is to be maximized in the optimization process. This approach is computationally costly, but it is found to be effective in some difficult situations. *Ballard and Sklansky* successfully applied this technique for tumor detection in chest radiographs [Bal73].

4.2.2.3. Region-Based Segmentation

Region growing is a typical approach to region-based segmentation. It involves the process of joining neighboring points into larger regions and the goal is to partition the entire image into disjoint regions such that they have no two dimensional overlaps and no pixel belongs to the interior of more than one region. Early efforts in developing region growing algorithms involved the merging of regions with similar properties. Later, refinements of these methods were obtained by combining splitting and merging mechanisms to produce more reliable segmentations. Both approaches are discussed in the following sections.

(1) Merging techniques

The most primitive region growing algorithms use only aggregates of properties of local groups of pixels to determine regions. It starts with a point that meets a detection criterion and then all its neighboring points which meet a growth acceptance criterion are merged into the same region. The acceptance criterion may

depend on some statistical measures of a region. For example, it can be based on the similarity in average grey level of the region and the grey level of the candidate point. When a new point is accepted, it becomes part of the region and the process is repeated with the resulting region. For this algorithm, the results of segmentation always depend on the choice of starting points and the order in which points are added to the region. Optimization or heuristic search techniques can also be used to control the process [Ros82]. In simple situations, if the starting points and the number of regions are known *a priori*, this method can produce good segmentation results.

To improve the performance of this algorithm, *Brice and Fennema* use a more global approach by developing heuristics which evaluate parameters depending on more than one region. In this method [Bri70], the picture is initially partitioned into "atomic" regions of constant grey level by inserting an elementary boundary segment between pixels of different intensity. These large number of atomic regions are then combined by the successive application of two heuristics which are based on the properties of the edge boundaries between regions. The first heuristic ("phagocyte") merges two adjacent regions if the boundary between them is weak and the resulting new region has a shorter boundary than the previous two. That is, two adjacent regions R_i and R_j have a weak boundary and are merged if

$$W/\min(P_i, P_j) > \theta$$

where θ is a suitably chosen threshold, W denotes the length of the common boundary and P_i denotes the perimeter of region R_i .

In this heuristic the threshold θ controls the merging. If it is small, regions are merged frequently, otherwise, few regions are merged and in fact one region would practically have to surround another for merging to happen.

The second heuristic ("weakness") merges two regions if the weak portion W of their common boundary is above some predetermined percentage of their total shared boundary. This heuristic is applied to refine the results of the phagocyte heuristic.

This method of region growing provides reasonably accurate segmentation of simple scenes with few objects and little texture [Bri70], but does not perform well on complex scenes. *Benbaun* and *Weyl* [Ten75] utilized semantic knowledge for merging guidelines. Their extension provide improved results, but it is quite complex.

(2) Merging and Splitting technique

In this approach to region growing, splitting of non-homogeneous regions is incorporated into the merging algorithm. The general idea of this approach is to start with a given initial partition, then the process proceeds by merging adjacent regions if the resulting new region is sufficiently homogeneous and splitting a region if it is not homogeneous enough. After a region is split, its parts become candidates for merging with adjacent regions. The homogeneity is measured by the closeness to which a region can be approximated, usually by a function of a given type.

An algorithm based on this approach was suggested by *Horowitz* and *Pavlidis*. This algorithm use a pyramidal data structure which is a stack of pictures decreasing in resolution, with the smallest being a single point. Beginning with the original image in the bottom, the next picture higher in the stack is obtained from the previous one by averaging every four adjacent points in the lower picture to produce one point in the next higher picture. The split and merge operations are then defined in terms of these ascending collections of square regions. Based on this structure, the merge and split region-growing algorithm operates as follows: Given a homogeneity property H , for any region such that H is false, it is split into four subre-

gions. For any *four* appropriate regions such that H is true, they are merged into a single region. This process continues until no region can be further split or merged. Then for any *two* neighboring regions such that H is true, they are merged. Details of this algorithm are presented in [Hor74]. As for the performance of this method, *Rosenfeld* comments that in the absence of special knowledge about the desired segmentation, this method appears to be the preferred approach to image segmentation,

4.3. Proposed Method

4.3.1. Template Matching Approach to Tumor Detection

In examining criteria used by radiologist to extract suspicious areas from film mammograms, we notice that the approximately circular shape and the brightness homogeneity of tumor areas are important characteristics used in detecting suspicious areas. Among the general techniques developed for object detection, pixel-based and region-based segmentation techniques can make use of the brightness homogeneity property of regions to segment the image, but such detection processes will be blind to the shape of tumors. Since brightness homogeneity is not a unique characteristic of tumor areas, it is very difficult to use this criterion alone to segment the image successfully into tumor and background segments.

Among other techniques, two approaches which make use of the shape characteristic of the target object are possible solutions to the problem of tumor detection. One is to use the region information by applying a template matching technique to search for a homogeneously-dense region with a circular shape. The other approach is to use the boundary information by applying an edge-based segmentation technique on a high-pass filtered image and to detect ring like structures. The edge-based segmentation approach requires a preprocessing step to extract edge elements from the mammogram images. Since edge detection techniques make use of local information to detect the presence of

edges, a tremendous amount of noise edges will be extracted from those mammogram images which are rich in texture. These noise edges may make successful boundary tracking impossible. In addition, if segmentation techniques are used to detect breast tumors, further processing is necessary to extract the locations and sizes of the suspicious areas from the segmented images.

A method based on the template matching approach is proposed in this section. With this approach, the shape and homogeneity characteristics of tumor areas can be used as the match criterion by defining them in the templates. The location and size of the detected suspicious area(s) can be easily obtained from the output of the matching operation. The problem of object orientation variance is solved by the fact that the shape of breast tumors is approximately circular. To cope with the object size variance, multiple templates can be used.

A general discussion of the template matching approach is presented in Section 4.2.1. In the case of tumor detection, the search area of the matching operation is an SM-filtered mammogram image and the template is an integer array which represents an 'ideal' tumor. The important issues of the proposed method are discussed in the following sections.

4.3.2. Definition of the Tumor-like Template

A tumor like template is defined based on three characteristics of tumor areas and they are :

- (1) brightness contrast (ie. a bright object in a dark background),
- (2) uniform density,
- (3) approximately circular shape.

The template used to match tumors with a diameter of 5 pixels is shown in Figure

4.2. The templates are designed so as to ascribe varying weights to points within the templates. The circular patch of 1's in the center of the template represents a tumor area having uniform density. To allow tumor shape to deviate slightly from a perfect circle, the patch of 1's is bounded by a ring of 0's. This is a "don't care" area in the match. The background of the patch is filled with -1's instead of 0's, because we are looking for a light object on a dark background. The size of the ring of 0's and the background in each template increase in proportion to the size of tumor the template represents. In addition, the shape of the template is circular instead of square, so as to increase the sensitivity of the match. By using a circular template, all the neighboring pixels which locate evenly around the tumor are checked in the matching process.

```

      -1 -1 -1
    -1 -1 0 -1 -1
  -1 -1 0 1 0 -1 -1
-1 -1 0 1 1 1 0 -1 -1
-1 0 1 1 1 1 1 0 -1
-1 -1 0 1 1 1 0 -1 -1
  -1 -1 0 1 0 -1 -1
    -1 -1 0 -1 -1
      -1 -1 -1

```

Figure 4.2 A tumor-like template for matching with tumors of 5 pixels in diameter

4.3.3. Similarity Measure

To measure the similarity between a true suspicious area and the template, we need a similarity measure that is not sensitive to properties of the mammogram that vary with the offset, such as its average brightness. In addition, the measure produced by templates of different sizes must be comparable. Hence, the following normalized cross-correlation measure is used :

$$R(i, j) = \frac{\sum_{k=1}^M \sum_{m=1}^M \left\{ \left[W(k, m) - \mu_w \right] \left[S(i+k-1, j+m-1) - \mu_s(i, j) \right] \right\}}{\sqrt{\sum_{k=1}^M \sum_{m=1}^M \left[W(k, m) - \mu_w \right]^2 \sum_{k=1}^M \sum_{m=1}^M \left[S(i+k-1, j+m-1) - \mu_s(i, j) \right]^2}} \quad (1)$$

notation and indices of this equation are the same as those used in section 2.4.1.1.

As mentioned earlier, one disadvantage of using template matching is that the orientation and size of the target object may disturb the match. However, the shape of the tumor area is circular, therefore, the match is orientation invariant. To cope with size variations of the object, multiple templates are used. Investigations showed that the size of breast tumors ranges from three to fourteen pixels in radius. Therefore, twelve templates are used to match against each image. The circle in the smallest template has a radius of three pixels, while the largest template has a radius of fourteen pixels.

4.3.4. Criteria in Selecting Suspicious Areas

The template matching operation produces twelve output images in which each pixel value is the result of cross-correlating the template and the subimage centered at the point. It is important to devise an appropriate method to interpret these values so that all suspicious areas are detected and non-suspicious areas are excluded.

Due to the design of the templates, a circular dark object on a bright background

will produce a large negative cross-correlation value. Hence, locations in the images which have negative cross-correlation values are not considered as suspicious areas.

An effective criterion for selecting suspicious areas must be able to solve two problems. First, the most suspicious areas have the maximum cross-correlation value when being matched with one of the twelve templates, but we have no guarantee that it is always true. Second, we do not have prior knowledge of the size and the number of tumors in a mammogram film. In fact, a mammogram film may contain more than one suspicious area, and several suspicious areas in the same mammogram film may have the same size. It is obvious that criteria which look for locations with peak cross-correlation value are not suitable for this application, because they either assume that there can be at most one tumor in each mammogram image or that no two tumors of the same size will appear in a mammogram image. A more relaxed criterion must be used. The thresholding criterion seems to be a possible candidate, but some modifications must be made.

This criterion requires a threshold to be found for each image. However, a single threshold for all images is not sufficient, because of the following reason. It is observed that the breast image in some mammograms, such as the one shown in Plate 2.2, contain a lot of glands and fatty tissue which produce a very rich texture in the images. As a result, these images create many false suspicious areas that have large cross-correlation values. Some of these values are even greater than those of the suspicious areas in other images. Therefore, a single threshold will either produce too many false alarms or miss some true suspicious areas.

A *percentile* method [Doy62] is used to determine a threshold for each image. We assume that suspicious areas only occupy a fixed percentage of the breast area. Hence, by analyzing the normalized cross-correlation distribution of an image, we choose a threshold value, R , such that q percent of the locations in the image having a correlation

larger than R are considered as suspicious areas. This corresponds to mapping q percent of the locations into suspicious areas and the value of q is constant for all images.

The template matching process only considers local information, therefore it can not adjust to the global texture of each image. This percentile method improves the template matching step by taking into account the global information in an image. In general, if many locations in an image (with rough texture) produce large cross-correlation values, a large threshold will be selected to minimize the number of non-suspicious areas being considered as suspicious areas. In the case of an image with smooth texture, a smaller threshold is used to ensure the detection of suspicious areas. The cross-correlation distribution curves built for the images shown in Plate 3.24 are displayed in Figures 4.3 and 4.4. Note that in using the same q value, a large threshold (0.67) was selected for the image with rough texture while a smaller threshold (0.58) was selected for the smoother image.

To implement this percentile method, the largest among the twelve cross-correlation values (corresponding to the twelve templates) computed for each image coordinate was stored in a two-dimensional array and the corresponding radius was stored as well. To save processing time, locations with a cross-correlation value less than 0.2 were discarded as insignificant locations. Then the center of each "significant" location in the array was determined by a clustering operation based on the given radius and cross-correlation value. Among these locations, q percent, having the largest cross-correlation value were picked as suspicious areas in each image. The x, y -coordinates and radius of each suspicious area were stored in a list for further examination.




Figure 4.3 A distribution curve of normalized cross-correlation values computed for the best case image (Plate 3.24a). Using a q value of 2.5, the threshold R of this image is 0.58




Figure 4.4 A distribution curve of normalized cross-correlation values computed for the worst case image (Plate 3.24b). Using a q value of 2.5, the threshold R of this image is 0.67

4.3.5 False Alarm Tests

The percentile method produces a set of candidate points which will not be exclusively suspicious areas. To improve the performance of this tumor detection technique, we need some selection methods to remove as many false alarms as possible. One important requirement of a false alarm test is that they must not respond to true suspicious areas because if it does, it will degrade the overall performance of the proposed method. After analyzing the characteristics of the many false alarm areas, two tests are designed. They examine a mammogram image at the locations reported as suspicious by the template matching procedures and try to discriminate false alarms from the suspicious areas.

(1) Neighborhood Test

This test is constructed based on the fact that at the center of a suspicious area, the cross-correlation value will not fall off significantly in its immediate neighborhood. The normalized cross-correlation values computed around the center of a suspicious area and a false-alarm area are shown in Figures 4.5 and 4.6 respectively. It can be observed that there is a cluster of large correlation values at the center of the suspicious area, but such a cluster does not exist at the false alarm area. Recall that the boundary of a suspicious area is fuzzy and the shape of a suspicious area is circular. Therefore, a perfect match between a template and a suspicious area will not be confined to a single pixel sitting at the center of the suspicious area. Instead, the pixels in the immediate neighborhood of the center pixel will also produce high cross-correlation value when being matched with a template having the same size as the suspicious area. On the other hand, in false alarm areas where a bright, homogeneous and circular area does not exist, the cross-correlation value may fall off quickly away from the center.

Moreover, as mentioned earlier, some false alarms are due to an artificial phenomenon created by the filtering process. The neighborhood test also aims at removing this type of false alarm by applying the test to the original image instead of the filtered image. This is based on the belief that when the matching is applied to the original image, a cluster of high cross-correlation values should not exist in these false alarm areas because they should appear as normal tissue in the unfiltered image.

To implement this test, the immediate neighborhood of the center of a suspicious area is defined by a 3×3 cross-shaped window. To detect the cluster at the center of each suspicious area, the window is centered at the area in the original image and the cross-correlation value of the pixels inside the window are calculated using the known tumor radius. Then the average of these cross-correlation values are taken. This average value reflects the "strength" of the cluster. If the cross-correlation value fall off significantly in the immediate neighborhood of the center of the suspicious area, the average value will be small. Otherwise, the average value will be large and in fact, should be larger than the threshold, R , selected for the particular image by the percentile method. Therefore, to determine if a cluster exist, this average value is compared with the threshold R . If it is less than the threshold, R , of the particular image, the suspicious area is discarded as a false alarm.

16	26	35	38	36	31	24	17	11	08	03
12	24	34	41	42	39	33	26	20	14	08
07	21	34	42	47	47	41	34	27	19	12
05	15	32	45	53	55	48	42	32	23	16
01	13	30	46	54	60	54	47	38	26	20
01	13	29	43	55	62	58	50	43	32	27
02	11	24	38	49	54	52	47	44	38	34
04	14	22	30	39	44	45	43	43	42	40
04	14	22	28	29	33	32	33	39	40	40
02	10	17	22	23	24	21	24	31	35	38
00	06	13	17	19	18	17	20	25	28	32

Figure 4.5 The normalized cross-correlation values computed around the center of a suspicious area located in the image shown in Plate 3.24a

23	26	28	37	39	40	39	36	32	28	24
29	32	37	40	42	43	40	36	30	26	23
29	33	37	40	42	42	40	35	29	24	20
27	31	34	38	40	40	38	34	28	22	18
25	28	30	34	36	38	36	32	43	34	27
24	26	29	30	33	61	59	53	44	38	32
25	27	28	30	30	56	56	49	43	39	36
26	27	29	29	29	51	50	45	39	39	38
27	28	28	28	27	46	44	40	37	36	37
25	28	28	27	25	23	38	38	36	35	35
22	25	27	27	25	22	34	34	33	34	34

Figure 4.6 The normalized cross-correlation values computed around the center of a false alarm area located in the image shown in Plate 3.24a

(2) Histogram Test

It is observed that many of the false alarms are caused by the presence of noise and anatomical structures in the image. The contrast of these areas is very small when compared with that of the true suspicious areas. A second test is designed to discriminate these false alarms. This test examines each suspicious areas and checks if there are two regions corresponding to the tumor area and its surrounding region. To define a region, the homogeneity property used is its intensity. Those areas which have only one region are non-suspicious areas.

To implement this test, an approach based on the *grey-level thresholding* method (refer to section 4.2.2.1) is used. A gray level histogram is constructed in each suspicious site using the known tumor radius reported by the template matching process. After subsequent smoothing of the histogram, the number of peaks in the histogram is checked. Areas which have a single-peaked histogram are rejected as possible suspicious areas. This test relies heavily on the structure of the gray level histogram, which contains peaks and valleys corresponding to gray level sub-populations of the image. A tumor and its surrounding area (represented by histogram peaks) are assumed to differ significantly in average gray level. Hence, a true suspicious area must have at least two peaks in its gray level histogram and the peaks should remain after subsequent smoothing of the histogram. The smoothed grey-level histograms constructed in a true suspicious area and in a false alarm area are shown in Figures 4.7 and 4.8 respectively. Note the number of peak in each histogram.




Figure 4.7 Grey-level histogram constructed in a suspicious area located in the image shown in Plate 3.24a




Figure 4.8 Grey-level histogram constructed in a false alarm area located in the image shown in Plate 3.24a

4.3.6. Experimental Results

To determine the value of the parameter, q , used in the percentile method, the procedures described were applied to the training set of seven mammogram images (please refer to Section 3.3.3) using different q values. The three different q values used are 0.5, 2.5 and 5. These values are chosen because they are commonly used as percentile values in solving statistical problems.

In comparing the diagnostic results produced by the computer and the radiologist, one could not expect that the location and size of a suspicious area determined by the radiologist and the computer were exactly the same, some criterion for determining a positive detection was established. A suspicious area was considered as detected, if the circle placed around a suspicious area by the radiologist and the nearest area mapped by the computer overlap by at least 50%. If the overlap was less than 50%, it was counted as a miss. The result of this comparison based on the training set is shown in Table 4.1, the q value used is 2.5. The table lists the size and location of the suspicious area(s) diagnosed by the radiologist along with the areas located by computer. Also included are the total number of suspicious areas reported in each test case.

The result of comparing the performance of the proposed method using different q values are shown in Table 4.2. This table shows that as the q value increases, both the total number of suspicious areas detected and the false alarm rate increase. A high false alarm rate will affect the overall accuracy of the proposed method and it will result in decreased computational efficiency. Hence, a q value smaller than 5 should be used. On the other hand, although the false alarm rate is zero when a q value of 0.5 is used, the hit rate produced falls to 54%. As a result, a q value of 2.5 is chosen for the percentile method because, as shown in Table 4.2, it produces a hit rate of 100% and a moderate false alarm rate of 1.1 per mammogram. In Plate 4.1, the suspicious areas reported by the proposed method are circled in the mammogram images 1 and 2.

The result of applying the proposed method to the test set of seventeen images (SMF-enhanced images) is shown in Table 4.3. The data shown in this table is summarized in Table 4.4. Although more test images were used, the hit rate of the proposed method in detecting suspicious areas is retained at 100%, that is, all suspicious areas detected by the radiologist were detected by the proposed method. The false alarm rate is 1.7 per mammogram.

In addition, to investigate the importance of the image enhancement step, we applied template matching to the original images. We found that many false alarms of a small radius size were produced and some suspicious areas were not detected by the template matching technique. The false alarms detected in the original mammogram image 1 are marked by circles in Plate 4.2. The cross-correlation distribution curve built for this image is shown in Figure 4.9. Note that there are more locations in this image that produce large cross-correlation values when compared with that in the enhanced image (see Figure 4.3). This result reflects that the SMF filtering process effectively removes noise from the image and this process is significant to the proposed tumor detection technique.

q	Hit-rate*	Total number of false alarm detected	Average number of false alarm per film	Average number of false alarm per pixel
0.5	54%	0	0	0.00%
2.5	100%	8	1.1	0.01%
5.0	100%	47	6.7	0.04%

Table 4.2 Comparison of the performance of the proposed method using different q values

* In the training set, 11 suspicious areas are reported by the radiologist

case number	Coordinates of suspicious areas						Total number of areas detected
	Radiologist			Computer*			
	X	Y	radius (mm)	X	Y	radius (mm)	
136801	170	116	6	166	116	5	2
151138	180	176	11	181	178	10	4
	100	186	8	102	185	8	
	150	93	10	150	96	9	
	186	208	8	184	206	6	
006356	200	149	4	199	149	3	2
63387a	34	165	4	35	165	4	1
119343	160	118	12	159	119	11	2
142322	188	192	9	187	194	8	5
89993a	134	128	8	135	128	8	3
	192	101	5	194	103	4	1

Table 4.1. Comparison of the computer's and the radiologist's interpretation on 7 mammograms in the training set ($q=2.5$).
 * Coordinates of suspicious area most closely correlating with radiologist.




Plate 4.1 Result of applying the proposed method to the enhanced mammogram images 1 and 2. All suspicious areas reported are marked by circles




Plate 4.2 Result of applying the proposed method to the original mammogram images 1. Some suspicious areas reported are marked by circles

case number	Coordinates of suspicious areas						Total number of areas detected
	Radiologist			Computer*			
	X	Y	radius (mm)	X	Y	radius (mm)	
89993b	171	140	6	174	140	6	5
124133a	101	126	4	104	126	4	4
	122	134	6	122	136	5	
	175	159	8	172	158	7	
124133b	130	42	10	131	44	10	3
129194a	88	112	10	88	110	9	3
	90	150	6	92	149	7	
	198	70	4	200	71	5	
129194b	198	208	4	196	208	5	4
145079	200	64	3	199	64	3	3
40274a	125	194	4	124	194	4	5
63387b	38	162	3	38	164	3	1
40274b	118	68	7	118	66	9	1
125758	56	22	5	55	22	5	1
069591	192	164	10	192	164	8	5
152506	-	-	None	-	-	-	1
004948	103	71	4	104	70	5	0
004703	100	44	4	100	46	5	7
129194	138	76	13	140	78	12	2
012914	125	73	7	124	71	6	2
933302							2

Table 4.3. Comparison of the computer's and the radiologist's interpretation on 17 mammograms in the test set ($q=2.5$).

* Coordinates of suspicious area most closely correlating with radiologist.

Number of films tested	= 17
Total number of true suspicious areas	= 19
Average number of true suspicious areas per film	= 1.1
Hit-rate of the technique	= 100%
Average number of false alarms found per film	= 1.7
Average number of false alarms found per pixel	= 0.026%

Table 4.4 Summary of experimental results




Figure 4.9 A distribution curve of normalized cross-correlation values computed for the original mammogram image 1. Using a q value of 2.5, the threshold R of this image is 0.66

4.3.7. Computation Consideration

4.3.7.1. Discussion

The computational cost of this template matching method is high. In order to match tumors of varying sizes, twelve templates have to be used. The sizes of the templates range from 10×10 to 45×45 pixels. Using a mammogram image size of 256×256 , the number of operations required to match the smallest template with the image is proportional to $10^2(256-10+1)^2 = 6.0 \times 10^6$. Experimental results show that approximately 15 minutes are required to complete the matching operation on one image using all twelve templates on a VAX 11/780 processor.

To improve the speed of the template matching process, one solution is to compute the cross-correlation in the frequency domain via a Fast Fourier Transform algorithm (FFT). The numerator of equation (1) in Section 4.3.3 can be treated by FFT, with a certain amount of time saving. The computational complexity of computing the numerator is proportional to $256^2 (\log_2 256) = 5.2 \times 10^5$.

Another way to speed up the computation is to reduce the number of locations to be fully checked. Methods similar to the two-stage template matching approaches mentioned in Section 4.5.1 can be used.

4.3.7.2. Experiments

A coarse-fine template matching method was designed and implemented. In the first stage of this method, a reduced-resolution template is matched with a reduced-resolution version of the image. Only those locations which produces a match value larger than a predetermined threshold are used in the second stage where a full resolution template and image are matched. Normalized cross-correlation is again used as the similarity measure. The coarse templates and images are obtained by averaging four pixels in

a higher resolution level to produce one pixel in a low resolution level.

In this coarse-fine template matching process, the expected computational cost [Ros77] at each image point is of the form $c + pd$, where

c = cost of applying the coarse template in the first stage,

p = probability of an above-threshold match to the coarse template,

d = cost of applying the fine (ie., full-resolution) template in the second stage.

In general, costs c and d are proportional to the number of pixels in the coarse template. If the coarse template has very few points, c will be small; but this will lead to a higher probability of above-threshold match, so that pd may be large.

In this implementation, the cost of the first stage, c , is about one sixteenth of the original "one-stage" template matching, since the image and template are reduced to a quarter of their original sizes. The cost of the second stage, pd , depends on the threshold. If the threshold is high, the cost of the second matching stage is low, but some suspicious areas may have a correlation value smaller than the threshold and therefore will not be used in the second stage matching. This is particularly true for small size tumors. They are small and have low contrast with their surroundings; therefore, they are often smeared in low resolution images. As a result, they produce correlation values smaller than the threshold in the first stage. This leads to a higher miss rate. On the other hand, if the threshold is small, the saving in computational time will become insignificant.

Investigations show that to produce approximately a 30% computational time saving, this method has to have a miss rate of 9% in comparison to 0% in the original method, ie. the suspicious areas were not detected in two out of the twenty-four test images.

4.4. Summary

In this chapter, some general techniques for object detection in two-dimensional scenes are discussed. They include methods based on the image segmentation approach and the template matching approach. In addition, a tumor detection method based on the template matching approach is proposed. The value of the parameter used is determined by applying the technique to a training set of seven mammogram images and the performance of this method is evaluated by applying it to 17 test images.

The circular shape, the brightness homogeneity and the contrast of tumor areas are important characteristics for detecting suspicious areas in mammogram images. The proposed method uses these three characteristics as the criterion for matching and it successfully detected all suspicious areas in the 17 test images. The false alarm rate reported is 1.7 per mammogram image.

The computation consideration of the proposed method is discussed and it is shown that the computational cost could be reduced by using either a Fast Fourier Transform algorithm in the frequency domain or a coarse-fine template matching approach.

Chapter 5

Conclusion and Future Research

5.1. Conclusion

The purpose of this thesis is to design a method to solve the problem of breast tumor detection in mammograms. We first started with a study of the basic principles of mammographic diagnosis. The criteria used by expert radiologists to detect one type of breast tumor, the circumscribed masses, are that the suspicious area is a bright area with uniform density and an approximately circular shape of varying sizes. The poor image quality of mammograms makes it difficult to use a computer to detect this type of breast tumor.

The first step towards tumor detection is image enhancement by noise removal. Several algorithms which attempt noise cleaning in digital images were studied. Some techniques based on selective averaging were implemented due to their simplicity and their promising properties for edge preserving and noise removal. The proposed method for mammogram enhancement is termed as selective median filtering and is a modification of the non-linear median filtering method. A training set of seven mammogram images were used to adjust the parameter values used in this method. This technique was then applied to a test set of 17 mammograms in which suspicious areas were identified by a radiologist. The Selective Median Filtering method was found to be more effective than the other techniques implemented in enhancing mammogram images.

After the mammograms are enhanced, the next step is concerned with object detection. Techniques based on image segmentation and template matching were reviewed. The technique chosen for tumor detection is based on template matching. This proposed method is capable of detecting suspicious areas in mammograms independent of their sizes, orientations and positions. However, many false alarms would be produced by this

matching process. To reduce the false alarm rate, two false alarm tests were designed to examine the output of the template matching process. The tests can discriminate most noise areas from true suspicious areas. These routines were applied to 17 test mammogram images and the results indicate that the routines correctly identified all suspicious areas presented although the false alarm rate reported is 1.7 per mammogram image.

One disadvantage of the proposed template matching method is the high computational cost incurred. Several methods for reducing this cost were discussed. A coarse-fine template matching method was implemented and the results show that the computational cost of the proposed method can be reduced by such an approach.

5.2. Directions for Further Research

The results obtained with the approach presented are quite encouraging. By combining three criteria, namely the contrast, the uniform density and the circular shape of tumor areas, the detection algorithm is capable of locating all tumor areas in the test set of 24 images. Nevertheless, to test the stability of this algorithm, future work would involve testing on a larger number of samples. In addition, testing of the method using a large sample size can provide more complete knowledge for setting the two parameters: the threshold T of the selective median filter and the percentile q of the percentile method.

In diagnosing breast tumors, radiologists first interpret a single mammogram film based on the criterion mentioned and then a comparison of left and right breast areas is often conducted to confirm the decision. Therefore, it is anticipated that the false alarm rate of this algorithm can be further reduced by comparing the left and right breast areas. If a suspicious area is found in the same location in both breasts, there is a high chance that it is an anatomical structure instead of a breast

Recall that the expert radiologists use four criteria to detect suspicious areas, but the proposed method only uses three of them. The one which is not used is that a suspicious area has fuzzy boundary. This fuzziness can be measured by the directional gradient at the boundary of a suspicious area. Hence, this suggests another metric that can be used for detecting false alarms which do not have fuzzy boundaries.

Bibliography

- [Aat80] Ataman, E., Aatre, U.K. and Wong, K.M., "A fast method for real time median filtering", *IEEE Trans. Acoust., Speech, Signal Processing*, ASSP-28, August, 1980.
- [Ack72] Ackerman, L.V. and Gose, E., "Breast lesion classification by computer and xeroradiograph," *Cancer*, pp. 1025-1035, 1972.
- [Bal82] Ballard, D.H. and Brown, C.M., *Computer Vision*, Prentice-hall, Inc., New Jersey, 1982.
- [Bal73] Ballard, D. and Sklansky, J., "Tumor detection in radiographs", *Computers and Biomedical Research*, vol. 6, pp. 299-321, 1973.
- [Bar72] Barnea, D.I. and Silverman, H.E., "A class of algorithm for fast digital image Registration," *IEEE Trans. Comput.*, C-21, pp. 179-186, February, 1972.
- [Bed84] Bednar, J.B. and Watt, T.L., "Alpha-trimmed means and their relationship to median filters", *IEEE Trans. on Acoust., Speech, Signal Processing*, ASSP-32, No. 1, 1984.
- [Bov87] Bovik, A.C., Huang, T.S. and Munson, D.C., "The effect of median filtering on edge estimation and detection," *IEEE Trans. Pattern Anal. and Machine Intell.*, PAMI-9, No. 2, pp. 181-194, March 1987.
- [Bov83] Bovik, A.C., Huang, T.S. and Munson, D.C., "A generalization of median filtering using linear combinations of order statistics", *IEEE Trans. Acoust., Speech, Signal Processing*, ASSP-31, pp. 1342-1350, 1983.
- [Bri70] Brice, C. and Fennema C., "Scene analysis using regions", *Artificial Intelligence*, Vol. 1, no. 3, pp. 205-226, 1970.
- [Cho72] Chow, C.K. and Kanedo, T., "Automatic Boundary Detection of the Left Ventricle from Cineangiograms", *Computer and Biomedical Research*, vol. 5, pp. 388-410, 1972.
- [Dan81] Danielsson, P.E., "Getting the median faster", *Computer Graphics and Image Processing*, vol. 17, pp. 71-78, 1981.
- [Dav78] Davis, L.S. and Rosenfeld, A., "Noise cleaning by iterated local averaging", *IEEE Trans. System, Man and Cyber.*, SMC-8, No. 9, pp. 705-710, 1978.
- [Dha86] Dhawan, A.P., Buelloni, G. and Gordon, R., "Enhancement of mammo-graphic features by optimal adaptive neighborhood image processing," *IEEE Trans. Medical Imaging*, MI-5, No. 1, pp. 108-120, March, 1986.

- [Dod81] Dodd, G.D., "Radiation detection and diagnosis of breast cancer", *Cancer*, vol. 37, pp. 1766-1769, 1981.
- [DOr83] D'Orsi, C.J. and Smith, E.H., "Breast imaging", in *Carcinoma of the Breast, Diagnosis and Treatment*, D'Orsi, C.J. and Wilson, R.E. (eds), Little, Brown and Company, Boston/Toronto, pp. 95-162, 1983.
- [Doy62] Doyle, W., "Operation useful for similarity-Invariant pattern recognition," *J. Assoc. Comput. Mach.*, vol. 9, pp. 259-267, 1962.
- [Dud72] Duda, R.O. and Hart, P.E., "Use of the Hough Transform to detect lines and curves in pictures", *Commun. ACM*, vol. 15, No. 1, pp. 11-15, 1972.
- [Fre77] Frei, W., "Image enhancement by histogram hyperbolization", *Computer Graphics and Image Processing*, Vol. 6, pp. 286-294, 1977.
- [Gal81] Gallaager, N.C. and Wise, G.L., "Theoretical analysis of the properties of median filters," *IEEE Trans. Acoust., Speech Signal Processing*, ASSP-29, No. 6, pp. 1136-1141, December 1981.
- [Gol83] Gold, R.H. and Bassett, L.W., "Mammography : History and state of the art", in *Breast Carcinoma, Current Diagnosis and Treatment*, Feig, S.A. and McLelland, R. (eds), Masson Publishing USA, Inc., pp. 95-98, 1983.
- [Gon77] Gonzalez, R.C. and Wintz, P., *Digital Image Processing*, Addison-Wesley Publishing Company, Inc., 1977.
- [.Gor84] Gordon, R. and Ranagayyan, R.M., "Feature enhancement of film mammograms using fixed and adaptive neighborhoods," *Appl. Opt.*, 23, No. 4, pp. 560-564, 1984.
- [Gos84] Goshtasby, A., Gage, S.H., and Bartholic, J.F., "A Two-stage cross correlation approach to template matching", *IEEE Trans on Patt. Anal. and Machine Intell.*, PAMI-6, No. 3, pp. 374-378, 1984.
- [Gra62] Graham, R.E., "Snow removal - a noise-stripping process for picture signals", *IRE Trans. Infor. Theory*, IT-8, pp. 129-144, 1962.
- [Gre78] Greer, D.S., "Use of heuristic search to find lung boundaries in conventional tomographs", *Proceedings of Pattern Reg. Image Proc.*, pp. 62-65, 1978.
- [Hal71] Hall, E.L., Kruger, R.P., Dwyer, S.J., *et al.*, "A survey of preprocessing and feature extraction techniques for radiographic images", *IEEE Trans. Comp.*, C-20, No. 9, pp. 1032-1044.
- [Hal74] Hall, E.L., "Almost uniform distributions for computer image enhancement", *IEEE Trans. Comput.*, C-23, No. 2, pp. 207-208, 1974.

- [Hal79] Hall, E.L., *Computer Image Processing and Recognition*, Academic Press, Inc., (London) Ltd., 1979.
- [Han79] Hand, W., Semmlow, J.L., Ackerman L.V., *et al.*, "Computer screening of xeromammograms : A technique for defining suspicious areas of the breast," *Computers & Biomed. Research*, vol. 12, pp. 445-460, 1979.
- [Hau83] Haus, A.G., "Physical principles and radiation dose in mammography", in *Breast Carcinoma, Current Diagnosis and Treatment*, Feig, S.A. and Mclelland, R. (eds), Masson Publishing USA, Inc., pp. 99-114, 1983.
- [Hor74] Horowitz, S.L. and Pavlidis, T., "Picture segmentation by a directed split-and-merge procedure", *Proc. 2nd Int. Joint Conf. Patt. Reg.*, pp. 424-433, August 1974.
- [Hua79] Huang, T.S., Yang, G.J. and Tang, G.Y., "A Fast Two-Dimensional Median Filtering Algorithm," *IEEE Transaction On Acoustic, Speech Signal Processing*, ASSP-27, No. 1, February 1979.
- [Hum75] Hummel, R., "Histogram modification techniques", *Computer Graphics and Image Processing*, Vol. 4, pp. 209-224, 1975.
- [Hum77] Hummel, R., "Image enhancement by histogram transformation", *Computer Graphics and Image Processing*, Vol. 6, pp. 184-195, 1977.
- [Ioa84] Ioannidis, A., Kazakos, D. and Watson, D.D., "Application of median filtering on nuclear medicine scintigram images," *Proceedings of Pattern Reg. and Image Proc.*, No. 1, pp. 33-36, 1984.
- [Kuh81] Kuhlmann, F. and Wise, G.L., "On second moment properties of median filtered sequences of independent data", *IEEE Trans. Commu.*, COM-29, pp.1374-1379, 1981.
- [Lai88] Lai, S.M., Li, X., Bischof, W.F., "Automated detection of breast tumors", *Proceedings Vision Interface, Edmonton, Alberta*, pp. 35-40, June 1988. Revised version in *Computer Vision and Shape Recognition*, Krzyzak, A., Kasvand. T., and Suen. C. (eds), (to appear), 1988.
- [Lee84] Lee, Y.H. and Kassam, S.A., "Application of non-linear filter for image enhancement," *Proceedings of Pattern Reg. and Image Proc.*, pp. 930-931, 1984.
- [Lev77] Lev, A., Zucker, S.W. and Rosenfeld, A., "Iterative enhancement of noisy images", *IEEE Trans. on System, Man and Cyber.*, Vol. SMC-7, No. 6, pp. 435-444, 1977.
- [Li78] Li, C.C., Savol, A.M. and Fong, C.P., "An improved algorithm for the detection of small, rounded pneumonconiosis opacities in chest X-Rays,"

- IEEE Trans. Medical Imaging*, pp. 370-395, 1978.
- [Li86] Li, X., Shanmugamani, C., Wu, T., *et al.*, "Correlation measures for corner detection", *Proceedings of Comput. Vision Pattern Reg.*, pp. 643-647, 1986.
 - [Mar82] Martin, J.E., *Atlas of Mammography, Histologic and Mammographic Correlations*, Williams and Wilkins, Baltimore/London, pp. 65-101, 1982.
 - [Mon71] Montanari, U., "On the optimal detection of curves in noisy pictures", *Commun. ACM*, vol. 14, No. 5, pp. 335-345, 1971.
 - [Nag79] Nagao, M. and Matsuyama, T., "Edge preserving smoothing", *Computer Graphics and Image Processing*, Vol. 9, pp. 394-407, 1979.
 - [Nar81] Narendra, P.M., "A separable median filter for image noise smoothing", *IEEE Trans. Pattern Analy. Machine Intell.*, PAMI-3, No. 1, 1981.
 - [Nod82] Nodes, T.A. and Gallagher, N.C., "Median filters: some modifications and their properties", *IEEE Trans. Acoust., Speech, Signal Processing*, ASSP-30, No. 5, 1982.
 - [Pra78] Pratt, W.K., *Digital Image Processing*, New York: Wiley-Interscience, 1978.
 - [Pag80] Pagani, J.J., Bassett, L.W., Gold, R.H., *et al.*, "Efficacy of combined Film Screen/Xeromammography", *Am. J. Radiol.*, vol. 135, pp. 141-146, 1980.
 - [Rab75] Rabiner, C.R., Sambur, M.R., and Schmidt, C.E., "Applications of a non-linear smoothing algorithm to speech processing", *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-23, pp. 552-557, 1975.
 - [Ram75] Ramer, U., "Extraction of line structures from photographs of curved objects", *Computer Graphics and Image Processing*, vol. 4, pp. 81-103, 1975.
 - [Ros82] Rosenfeld, A. and Kak, A.C., *Digital Picture Processing*, Vol. 1 and 2, 2nd edition, New York: Academic, 1982.
 - [Ros77] Rosenfeld, A. and Vanderfrug, G.J., "Coarse-fine template matching", *IEEE Trans. Syst., Man, Cybern.*, SMC-7, pp. 104-107, Feb. 1977.
 - [Sch80a] Scher, A., Velasco, F.R.D., and Rosenfeld, A., "Some new image smoothing techniques", *IEEE Trans. System, Man and Cyber.*, Vol. SMC-10, No. 3, pp. 153-158, 1980.
 - [Sch80b] Schutte, H., Frydrychowicz, S. and Schroder, J., "Scene matching with translation invariant transforms", *Proc. 5th Int. Joint Conf. Pattern Recognition*,

pp. 195-198, 1980.

- [Skl79] Sklansky, J. Sankar, P., Katz, M. *et al.*, "Toward computed detection of nodules in chest radiographs", *Proceedings of Pattern Reg. Image Proc.*, pp. 484-487, 1979.
- [Str87] Strax, P., "Mass screening of asymptomatic women", in *Breast Cancer, Diagnosis and Treatment*, in Ariel, I.M. and Cleary, J.B. (eds), McGraw-Hill Book Company, pp.145-151, 1987.
- [Tab85] Tabar, L., Dean, P.B., "Basic principles of mammographic diagnosis", *Diagn. Imag. clin. Med.*, vol 54, pp. 146-157, 1985.
- [Ten75] Tenenbaum, J.M. and Weyl, S., "A region analysis subsystem for interactive scene analysis", *Proceedings 4th Inter. Joint Conf. on Artificial Intelligent*, September, 1975.
- [Tou79] Tou, L.T., "Zoom-thresholding technique for boundary determination", *Int. J. Comp. and Infor. Sci.*, Vol. 8, No. 1, pp. 3-8, 1979.
- [Van77] Vanderbrug, G.J. and Rosenfeld, A., "Two-stage template matching", *IEEE Trans. Comput.*, C-26, No. 4, pp. 384-393, 1977.
- [Way79] Wayrynen, R.E., "Fundamental aspects of mammography", in *Reduced Dose Mammography*, Logan, W.W. and Muntz, E.P. (eds), New York, Masson Publishing USA, pp.512-528, 1979.
- [Wec77] Wechsler, H. and Sklansky, J., "Finding the rib cage in chest radiographs", *Pattern Recognition*, vol. 9, pp. 21-30, 1977.
- [Won78a] Wong, R.Y. and Hall, E.L., "Scene matching with invariant moments", *Comput. Graphics Image Processing*, vol. 8, pp. 16-24, 1978.
- [Won78b] Wong, R.Y. and Hall, E.L., "Sequential hierarchical scene matching", *IEEE Trans. Comput.*, C-27, No. 4, pp. 359-366, 1978.
- [Win77] Winsberg, F., "Detection of radiographic abnormalities in mammograms by means of optical scanning and computer analysis," *Radiology*, 89, pp. 211-215, 1967.
- [Zuc87] Zuckerman, H.C., "The role of mammograph in the diagnosis of breast cancer", in *Breast Cancer, Diagnosis and Treatment*, in Ariel, I.M. and Cleary, J.B. (eds), McGraw-Hill Book Company, pp. 152-172, 1987.