


# Feature Detection in Mammographic Image Analysis




Marius George Linguraru

*A thesis submitted to the Department of Engineering  
Science, the University of Oxford, in partial fulfilment of  
the requirements for the degree of Doctor of Philosophy.*

The University of Oxford  
Department of Engineering Science  
Medical Vision Laboratory

Michaelmas Term  
2002



*We know that one times one is one,  
but an unicorn times a pear have no idea what it is.*

Nichita Stănescu – “Another Kind of Mathematics”

# Feature Detection in Mammographic Image Analysis

Marius George Linguraru

A thesis submitted for the  
degree of Doctor of Philosophy  
at the University of Oxford

Medical Vision Laboratory  
Department of Engineering Science

Keble College  
Michaelmas Term, 2002

## Abstract

In modern society, cancer has become one of the most worrying afflictions due to the high and continually increasing death rate. The deep impact of the disease offers sufficient reasons for extensive research to be carried out in detecting and eradicating cancer of all forms. Breast cancer is one of the most common forms and approximately 1 in 9 women in the Western world will develop it over the course of their lives. Screening programmes have already proven that can reduce the mortality rate, but they introduce an enormous amount of information that must be processed by radiologists on a daily basis. Computer Aided Diagnosis (CAD) systems aim to assist clinicians in their decision-making process, for example by acting as a second opinion, with a view to improve the detection and classification ratios by spotting very difficult and subtle cases. This thesis presents results on detecting mammographic features in image analysis for improved effectiveness in cancer detection in screening programmes.

The detection of early signs of breast cancer is vital in exterminating such a fast developing disease with very poor survival rates. Some of the earliest signs of cancer in the breast are the clusters of microcalcifications. We propose a method based on image filtering comprising partial differential equations (PDE) for image enhancement. Subsequently, microcalcifications are segmented using characteristics of the human visual system, based on the superior qualities of the human eye in depicting localised changes of intensity in an image. We set the parameters according to the image attributes, which makes our method fully automated. Image normalisation is another key concept discussed in this thesis. As a step towards a more complete detection tool, we further investigate the detection of breast masses in temporal mammographic pairs. This latter algorithm is designed based on the detection sequence used by radiologists in clinical routine.

## Acknowledgments

Three years have past since my arrival in Oxford on the very same day of the year. Three years of great memories in the town of tradition and modernity, of history and contemporaneousness, of science and art. A place designed for studying was revealed to my eyes and senses, a place of intellectual challenges and stimulation, where inspiration can be found amidst great halls and weighing libraries, evergreen parks and narrow alleys, famous museum collections and historic houses and nevertheless amongst its people.

I like Oxford with its individual charm and I would need many more pages to eulogise it, but it is the people that I met here that I fell in love with. There is a long list of people that helped me during the years I spent in Oxford working on my thesis.

I start by thanking my supervisor, Mike Brady, a man of amazingly broad knowledge and fantastic support, the academic we would like to become one day. Thank you, Mike, for guiding my way from someone without background in medical vision to writing a thesis! You have a great lab!

The Medical Vision Laboratory is an excellent combination of scientific merit and friendly and helpful environment. I now have great friends and colleagues amongst its members. I am especially grateful to: Kostas, Maud, Rob, Margaret, Timor, Djamal, Alexis, Chris, Carolyn, Veit, Gerardo, Mike, Byong-Woo, Styliani, Vicky and Alison. The list goes on, thank you all! One further thankful word to the team from Mirada Solutions (cheers, Ralph) and the Churchill Hospital Breast Cancer Unit. I am also indebted to Professor Averil Cameron, the Warden of Keble College.

Many more friends from outside the lab offered me comfort and joy over the past few years. Your companion and care are inestimable and have engraved fantastic moments to my memory. Thank you, Manuela, Dan, Dana, John, Jen, Richard, Isabela, Camelia, Ruth, Marie-Josée, Steffen, Roland, Katja, Susana, Adam, Patrick, Izumi, Dana, Liana, Helen and Maria. Many other names add to them, especially the great party of people I know from back home.

My gratitude converges towards the ones so close to my heart: my mother Elena, my father Gheorghe and my sister Georgiana. The distance between us will never be large...

Marius George Linguraru  
Oxford, 28<sup>th</sup> October 2002



# Table of Contents

<b>1 Introduction .....</b>	<b>1</b>
1.1 The Prevalence of Breast Cancer .....	1
1.1.1 A Brief Anatomy of the Female Breast.....	3
1.1.3 The Pathology of Breast.....	5
1.2 Are Screening Programmes the Solution? .....	7
1.3 The Need for Image Segmentation.....	9
1.3.1 Image Quality .....	11
1.3.2 Future Trends.....	11
1.4 Hoping to “Make a Difference” .....	12
1.4.1 The Incentive of Work.....	12
1.4.2 Remarkable Achievements to Date.....	13
1.4.3 The Aim of the Thesis.....	15
1.5 Overview of the Thesis .....	18
<b>2 Diagnosing the Breast .....</b>	<b>21</b>
2.1 Imaging Modalities .....	22
2.1.1 X-ray Mammography .....	22
2.1.2 Digital Mammography.....	26
2.1.3 Magnetic Resonance Imaging .....	28
2.1.4 Ultrasound Imaging.....	32
2.1.5 Nuclear Imaging.....	34
2.1.6 Strengths and Weaknesses of the Breast Imaging Techniques .....	36
2.2 The Detection of Mammographic Anomalies .....	37
2.2.1 The Detection and Classification of Masses in X-ray Mammography .....	38

2.2.2 The Detection and Classification of Microcalcifications in X-ray Mammography.....	45
2.2.3 Yam <i>et al.</i> 's Physics Based Approach .....	50
2.3 Standard Mammogram Form.....	54
<b>3 Filtering <i>hint</i> Images .....</b>	<b>64</b>
3.1 Anisotropic Diffusion.....	65
3.1.1 The Diffusion Process .....	67
3.1.2 Nonlinear Anisotropic Diffusion.....	69
3.1.3 Discussion .....	71
3.2 Filter Model.....	74
3.2.1 Theory.....	75
3.2.2 Method.....	77
3.2.3 Results.....	83
3.2.3.1 Coarse Calcifications .....	86
3.2.3.2 Microcalcifications .....	87
3.2.4 Discussion .....	88
3.3 Conclusion.....	90
<b>4 Adapting Characteristics of the Human Visual System to Digital Mammography .....</b>	<b>92</b>
4.1 Pre-Processing.....	94
4.1.1 Shot Noise Removal.....	95
4.1.2 Curvilinear Structures Removal.....	96
4.2 Statistical Analysis and Image Enhancement.....	100
4.3 Foveal Segmentation .....	103
4.4 Discussion .....	115
4.4.1 Comparative ROC Analysis .....	116
4.4.1.1 CLS Removal .....	116
4.4.1.2 Image Diffusion .....	118
4.4.1.3 Perceptibility .....	119
4.4.1.4 Intensity Images .....	120
4.4.1.5 University of South Florida Database .....	121
4.4.1.6 Whole Mammograms.....	123
4.4.2 An Alternative CLS Removal .....	125
4.4.2.1 Theory.....	126
4.4.2.2 Initial Results .....	129
4.4.2.3 Discussion .....	131

4.5 The Detection of Microcalcifications using SMF .....	132
4.5.1 Comparative Results .....	132
4.5.2 Setting the Parameters .....	138
4.5.3 Discussion.....	143
4.6 Conclusion.....	147
<b>5 Temporal Comparison of Feature Enhanced Mammograms For Mass Detection.....</b>	<b>148</b>
5.1 Introduction.....	149
5.2 Diffusion and Masses.....	150
5.2.1 Defining Diffusion Parameters for Mass Detection .....	154
5.3 Method.....	158
5.3.1 Mammogram Registration.....	159
5.3.2 Anisotropic Diffusion of the Temporal Mammograms .....	160
5.3.3 Tissue Classification Based on Texture Analysis.....	161
5.4 Results.....	162
5.5 Conclusion.....	167
5.6 Future Work on Mass Characterisation Features.....	168
<b>6 Discussion And Future Work.....</b>	<b>172</b>
6.1 Summary and Discussion .....	172
6.1.1 Mammographic Image Filtering .....	172
6.1.2 Complex Pre-processing.....	173
6.1.3 Human Vision Based Segmentation.....	175
6.1.4 Temporal Mass Prompting .....	175
6.1.5 Conclusion.....	177
6.2 Future Work .....	179
6.2.1 Expectation Maximisation .....	179
6.2.1.1 Theory.....	180
6.2.1.2 Initial Results .....	182
6.2.2 Iso-Level Segmentation.....	184
6.2.2.1 Theory.....	184
6.2.2.2 Initial Results .....	186
6.3 Concluding Remark .....	188

<b>Phase Congruency and the Structure Multivector</b> .....	<b>189</b>
A.1 Local Energy and Phase Congruency (PC).....	190
A.2 Detecting Feature Type from Phase Angle.....	194
A.3 Extending PC to 2D.....	196
A.4 Multiscale Analysis.....	197
<b>Registration Framework for Mammography....</b>	<b>200</b>
B.1 Introduction.....	200
B.2 Partial Registration Using the Boundary.....	203
B.2.1 Breast Outline Detection.....	203
B.2.2 Curvature Analysis on the Breast Outline-consistent Landmarks.....	204
B.2.3 Anatomical Significance of Boundary Landmarks.....	205
B.2.4 Partial Registration from the Breast Boundary.....	205
B.3 Multi-scale Landmark Selection for Improved Registration.....	207
B.4 Landmark Matching and Registration Refinement.....	207
B.5 Results an Discussion.....	208
<b>Receiver Operating Characteristic (ROC) Curve Analysis .....</b>	<b>210</b>
<b>Bibliography .....</b>	<b>215</b>
<b>List of Figures .....</b>	<b>229</b>
<b>List of Tables.....</b>	<b>237</b>



## CHAPTER 1

### 1 Introduction

*It is the mark of an educated mind to be able to entertain a thought without accepting it.*

Aristotle

#### 1.1 The Prevalence of Breast Cancer

The incidence of cancer in the Western world is enormous and its threatening presence is an unfortunate reality in our living environment. The impact that this fierce affliction has for a large percentage of the population has become a cultural phenomenon. We grow and live with *the fear of cancer* invading our privacy and shadowing the existence of people dear to us. We read about the spreading of cancer, and not only in specialised literature, we see on TV how tumours are formed, we listen to stories about affected lives, we sense the pain of those suffering. *The hospital* lost its “copyright” for the term “cancer” and now shares it with general sciences and, through its massive connotation, the entire human society. While society reacts to the burden of cancer, science attempts to get a reaction from the disease and discover the long-searched for solution to improve worrying statistics.

When considering figures and statistics of breast cancer, we refer to studies conducted in developed countries of the Western world. This is mainly due to the lack of information on this subject from under-developed and developing countries. Although the number of studies

carried out in such regions of the world has increased, the lack of appropriate technical equipment to detect cancer, associated with the high costs of the procedure, makes the records of such studies quite imprecise. We should normally expect much higher rates of incidence than reported in the poorer areas of the world, due to the present lack of radiologists and detection equipment in those regions of the world. That will likely become more of an issue once people in undeveloped countries live long enough to die from cancer. The disease has a far higher incidence in Europe (especially Western Europe) and North America. In the Far East and parts of Africa, the mortality rate due to breast cancer is much lower, with an incidence about 5 times smaller than in the West, although there has been a substantial increase in the number of new cases. During the last few years, Japan has witnessed a growth of 10 times in the number of breast cancers. In the Western world, recent figures show that breast cancer accounts for a high percentage of the overall cancer incidence in women, approximately 24% of all cancer cases. Around the world, there are approximately 945,000 new cases of breast cancer every year, of which the Western world accounts for 437,000 and the European Community accounts for approximately 235,000 (figures published in 2001 [40]). Figures have changed rapidly, as in 1993 there were 570,000 cases world wide [34]. Amongst the developed countries, the UK is rated as one of the regions with the highest incidence in breast cancer, where 14,000 women died of the disease in 1995.

In a more general context, breast cancer is second after lung cancer (28% of cancer cases), the most feared form of cancerous death in women of all ages [186, 187]. Once cancer has been diagnosed, the chances of survival are reduced to just over 60% [34]. Two years ago, when I was writing my first year report, figures of the time (three year old) showed that an average of 1 in 12 women [65] in the western world develop breast cancer during the course of their lives. The astounding rise in the impact that cancerous diseases have especially over Western Europe and Americas modified these figures with incredible speed. Nowadays, it is estimated that approximately 1 in 10 women [15, 189] (1 in 9 in Britain and 1 in 8 in the USA [132]) will be victims of breast cancer (only 0.1% of the total incidence of breast cancer is attributed to men), showing an approximate increase of 0.55% per year in the number of women developing the


disease. If the rate of change remains constant, the number of cases of breast cancer will double by 2020, when almost a fifth of the female population will be exposed to the disease at some stage in their lives.

Why is cancer incidence growing rapidly in modern society? Could the causes be environmental? Researchers have tried to trace both genetic and environmental causes that lead to developing the disease; still, there is so far insufficient evidence to support theories that attribute unhealthy food, hormonal treatments and pollution as major factors in the expansion of the disease. Many attribute this to a change in the lifestyle of women, particularly those seeking a career other than homemaker, and in diet – more fatty foods. It is estimated that 70% of cancers have their origins in the foods we eat. New reports [68] emphasise the carcinogenic factors found in diet, smoking, alcohol consumption, sunbathing and sedentary lifestyle. However, our diet is full of surprises; while fruits, vegetables, fish and milk are some the most effective sources of protection against cancer, normal food is stripped down of its nutrients by modern farming and therefore becomes less efficient. Fruits and vegetables may contain a high level of pesticides and farmed fish is a source of toxins. Amongst the dangerous foods, red meat, salt, soy, baked, fried, grilled, barbequed food (containing acrylamide), sweeteners and thickeners found in processed food are the best-known triggers. Other well-established factors refer to family history, ethnic background, early menarche or late menopause, the absence of childbirth, obesity and there are increasingly worrying signs related to the use of hormone replacement therapy (HRT) [7, 110].

### **1.1.1 A Brief Anatomy of the Female Breast**

For a better understanding of the subject, an overview of the breast anatomy becomes necessary. Throughout the life of a woman, the breast goes through a set of continuous changes. The first major development occurs at teen-age, when the lactation system evolves. A second important stage is the menopause, when the milk-producing tissue changes into fat.


Figure 1 shows a comparison between the appearance of a mammogram of a breast of a young woman and an image of a breast of a post-menopausal woman.



**Figure 1** The appearance of young versus menopausal breasts in mammography: (a) A pair of left and right cranio-caudal mammograms of a pre-menopausal woman with very dense appearance due to the presence of milk-producing tissue; (b) a pair of left and right cranio-caudal mammograms of a post-menopausal woman where there is a larger amount of fat-tissue, which makes the depiction of dense areas simpler.

The simplest portrayal of the mature female breast would have to enumerate the following types of incorporated tissues: epithelial or glandular tissue (the milk-producing tissue), which appears very dense in mammograms due to the high percent of calcium it contains; adipose tissue (fatty tissue, which is mainly transparent in X-ray), fascia (the connective tissue), muscles, ligaments and lymphatic and blood vessels [30, 53, 65]. Figure 2 presents the main anatomical features of the breast. The arborescent structure of the breast is nourished by acini (milk producing sacs inside a lobule) connected through terminal ducts. Lobules are fruits on the branches represented by subsegmental ducts that grow from the mammary ducts converging in the nipple.





**Figure 2** A brief anatomy of the breast showing the branching internal structure of ducts and lobules.

### 1.1.3 The Pathology of Breast

Breast cancer is an abnormally fast reproductive process of the epithelial cells in the lobular unit. It is also referred as *carcinoma* (the other type of cancer is *sarcoma*, which is much rarer and arises from a bone, muscle or other soft tissue [188]). The routes preferred by cancer in its spreading process are the blood and lymph vessels (which makes the axillary nodes an important feature in signalling metastatic diseases), but the direct invasion of the surrounding tissue may have the same effect. Therefore, one would naturally speak about spreading (*invasive* or *infiltrating*) and non-spreading (*in-situ* – which stay within the lobular or ductal

unit and has not gone through the basal membrane) forms of cancer. *In-situ* cancers are sometimes referred to as *pre-invasive* since further developments of the tumour may occur and invade the neighbouring tissue. Both *invasive* and *in-situ* breast carcinomas are commonly divided into lobular or ductal. Besides the ductal and lobular forms of cancer, other types of carcinoma may be medullar, tubular, papillary, cribriform and mucinous [53, 189].

Most ductal and lobular cancers lead to secretions that form calcifications. The smallest of these (under 1 mm in diameter) are called microcalcifications and represent some of the earliest signs of breast cancer. Microcalcification clusters may be the only indication of *in-situ* tumours. Approximately 80% of them are benign [65, 81, 166] their shape and topology discriminating them from the malignant type. The architectural distortions are another distinctive sign of breast cancer, when the tumour has no central mass [65].

Ductal Carcinoma in Situ (DCIS) is one of the most common types of *in-situ* cancers and can involve a large number of ductal structures. It is mainly associated with microcalcification clusters and corresponds to one of the earliest signs of malignancy. It appears in 40% of the screening detected cancers [95]. Its treatment may involve a partial mastectomy.

Infiltrating Ductal Carcinoma (DC) is the most usual invasive cancer (over 70% of tumours), a form of disease with very poor prognosis and which may require a total mastectomy, including the removal of the axillary lymph nodes and post-operative radiotherapy. It may be associated with microcalcifications.

Amongst the special types of cancer is the Phyllodes Tumour's [189] whose clinical behaviour is still not fully understood. Most of these tumours are benign, but there are also malignant forms. They can be cystic and sarcomatous and usually recur after initial excision. Their treatment is particularly difficult through the atypical behaviour of the tumour.

The treatment given to a woman diagnosed with breast cancer depends on the specific characteristics of that tumour. Moreover, we are dealing with different stages of the same type of disease and the treatment differs considerably from one stage to another. So does the prognosis, as a statistical measure of the chances of having a positive outcome of the treatment the patient is undergoing. The most relevant factors a doctor considers when deliberating on the

prognosis are the size and type of the tumour, the presence of metastasis, the stage of the disease, the status of the axillary lymph nodes and patient's age and medical condition. The prognosis is especially positive for tumours smaller than 2 cm without having lymph node involvement or other remote areas metastatic. A combination of surgery, radiotherapy and chemotherapy are typically necessary to eradicate advanced malignancy; still, the prognosis is rather poor.

Benign diseases may develop within or outside the ductal and lobular system and most are associated with microcalcification clusters. Adenosis, necrosis, hyperplasia, fibroadenomas and arterial calcifications are an important source of false positives (FP) in the classification of breast tumour.

Section 2.2 illustrates some typical cases of breast pathology in both the form of masses and microcalcifications.

## **1.2 Are Screening Programmes the Solution?**

The signs of breast cancer that appear in X-ray mammograms present a significant challenge to radiologists and they are generally difficult to distinguish in the highly textured breast anatomy. Breast screening programmes attempt to detect and eradicate cancer at the earliest possible stage to reduce the rate of mortality amongst women. From the first trials in USA and Canada in the sixties and its very first implementation in the seventies in Sweden, screening programmes were found to reduce mortality caused by breast cancer in women by nearly 30% [4, 29, 163].

Some results of the mass screening in UK show a very high rate of false negatives (FN), when *in situ* cancers are left to metastasise or invasive cancers are not detected, as well as a high rate of FP, where women are operated on without finding breast cancer. Hence, the early detection of breast carcinomas and subsequent classification as either benign or malignant is

subject necessitating further improvements with great impact on the fight against breast cancer and malignancy.

Breast cancer's incidence is low in women under 30 years old (although extremely aggressive when present) and thereafter increases with age. Between the ages of 40 and 50, women face a doubling of the rate of incidence which continues to increase over the age of 50, but more slowly. Younger women are encouraged to check the status of their breasts by simple palpation. Unfortunately, most women cannot reliably palpate a tumour smaller than 1 cm; therefore more thorough examinations are required. In the UK, the screening programme was arguably implemented for women between 50 and 64, since mammography has not to date been demonstrated to be clinically effective before menopause and this is the group of age mostly exposed to breast cancer. Other countries start the screening at ages below 50 and sometimes push the upper age limit to 75.

There are several criteria that need to be fulfilled before starting a screening programme. These include:

- the disease to be screened must be very common and a treatment must be available for it, since there are very high costs involved and there would be little point screening for a non-treatable affliction;
- the detection method must be robust and reliable and lead to good results for the overall screening process;
- it must have high specificity;
- patients must accept it, since the method would not be cost-effective without a large number of patients to be examined.

The key method used in screening programmes is X-ray mammography, as the most reliable process fulfilling the above criteria. If a mammogram presents any features that seem suspicious to the radiologist, the patient will be asked to attend an assessment clinic where more investigations will be performed by means of medical imaging and consulting. Magnetic Resonance Imaging (MRI) and ultrasound (US) are secondary imaging techniques used in the



triple assessment (palpation, imaging and core biopsy). Chapter 2 expands on the advantages and disadvantages of these imaging methods.

The screening programme should, according to optimistic statements, almost double the chances of survival in women that develop breast cancer. Studies have shown that 8% of women are recalled for further investigations [65], most of them not presenting any malignancy. There is intensive debate as to whether the breast screening clinical assessments should be performed every two years instead of the usual three-year period, as it has been noticed that the assessed women sometimes develop cancer over a period shorter than three years (“interval cancer”). While the UK has a screening interval of 3 years, other countries such as Sweden and Netherlands have already decreased the period to 2 years [29]. It is estimated that nearly 20% of cancers were visible in the mammograms previous to the current screening [150] and that interval cancers prefer to grow in the upper outer quadrant of the breast [17]. Segmentation methods may be more sensitive in detecting abnormalities at an earlier stage, but it is the radiologist whom will need to make the final decision. Even though screening programmes improve the results of detection, there is sufficient room for progress in the clinical performance to warrant further research.

### **1.3 The Need for Image Segmentation**

Image processing is a challenging but difficult task. Working with mammograms is especially complicated due to the complex appearance of the structures of interest in this particular type of image representation. Although a mammogram is a good “picture” of the breast, this is hardly sufficient when searching for small, subtle and complex anatomical parts, such as microcalcifications, masses or curvilinear structures (CLS) in the process of early detection of breast cancer. Statistics show that approximately 25% of cancers are missed and about 80% of biopsies are performed on benign cases [11, 169]. Such numbers draw our attention to the additional requirements of the diagnosis process. Besides saving lives, doctors are also

expected to find the least stressful and painful way to check the status of the disease, malignant versus benign. Regarding the unpleasantness of both mammogram and core biopsy acquisition, reducing the number of FP becomes as equally important as reducing the number of FN. Furthermore, it is estimated that about 22% of films are usually lost between visits and 5% of the mammograms have to be retaken.

The complex anatomy of the breast is the inevitable source of the highly textured structure of the mammograms. It provides a most difficult to analyse input for radiologists, who are expected to distinguish very subtle abnormalities out of this mass of structural ambiguity. The variability between any two cases adds to the difficult task that the human decision maker faces. The inter- and intra-radiologist variability of 30% emphasises the need for reliable image processing tools to assist the process of diagnosis. According to Krupinski [92], radiologists only investigate 87% of the mammogram area; in contrast, an automatic detection algorithm will not leave any area of the image unexamined.

With up to 3 million new mammograms to be analysed each year in the UK, the amount of information becomes overwhelming for radiologists, especially since a second opinion is requested before a decision is made in diagnosing a patient. Moreover, these figures are doubled as the previous mammograms of the patient are compared to the current ones at each screening session. Computer-aided diagnosis (CAD) can assist the radiologist at this point, by helping to balance the measure of confidence of the specialist and eliminate the second analysis. However, as Karssemeijer highlighted in [76], CAD systems can only be evaluated from a radiologist point of view:

- on annotated databases, in which case the effect of the CAD in practice cannot be predicted;
- on databases of known human reader, to compare their evolution with that of the clinician by receiver operating characteristic (ROC) curve analysis and predict the effect of the CAD in practice;
- by comparing the work of the radiologists with and without CAD;

- by doing prospective evaluations.

### **1.3.1 Image Quality**

The quality of the image depends on several physical factors. The most important are the time of exposure and the breast thickness along various other imaging factors. Since the X-ray dosage must be minimised for patient safety reasons, there is a compromise between dosage and the signal-to-noise ratio (SNR) of the mammogram. Reduction of X-ray exposure degrades the quality of the image, which already must reflect the superimposed structures of the breast – a three-dimensional (3D) structure - on a 2D projection. Noise further obscures these features. This presentation of the mammogram could be easily reduced to that of a large textured noisy image, which still represents the best tool for early cancer detection to be used around the world. The different breast deformation during the X-ray shot causes more difficulties in the detection process.

### **1.3.2 Future Trends**

In order to overcome the present limitations in medical image analysis, the need for computer aided image segmentation is most important. There are several implemented techniques that bring improvements in the field of medical imaging. Most of them have proved to be unsatisfactory for the purpose that they are used, such as automatic detection and managing mammograms. Moreover, although the resulting images may look quite impressive, it has not always been the case that radiologists worked better on the basis of such processed images. Therefore, there is sufficient room for improvements and further developments in image processing. Some present trends in the field include:

- the development of soft-copy reading workstations [12, 37, 142, 177], as the tool for the future use of digital mammograms;

- friendlier user-interfaces (touch-screen, automatic report generation, robust display) which would only require a minimal intervention from the human factor involved, the radiologist;
- the development of training-systems with immediate feedback and the use of larger databases;
- the development of real-time applications for making the best use of the image processing methods in clinical applications;
- finding more reliable, more robust and faster image processing algorithms;
- software integration of robust algorithms is an inherent condition in building strong performant medical systems.

## **1.4 Hoping to “Make a Difference”**

It has become clear that the earliest possible detection of signs of breast cancer is fundamental in making a difference in the lives of so many women having to face the disease. The challenges enumerated in the previous section combine to create a complex problem that has to be solved using basic science.

### **1.4.1 The Incentive of Work**

The research described in this thesis is motivated by the intrinsic facts of managing the detection of breast cancer. The first question arising is: why detect microcalcifications and not masses? The answer is quite straightforward; the presence of microcalcifications is one of the earliest indications of malignancy. Whether or not they appear in independent clusters or associated with masses, the existence of microcalcifications in a mammogram is often a clear warning of abnormality. They can be visible long before any palpable lesion has developed and



their early detection can indeed “make a difference” in the prognosis. The detection of masses is also investigated to a smaller extent in this thesis.

A second question arises at this stage: what type of medical imaging procedure should be investigated in the detection process? Although nuclear images of the breast are currently available, clinical practise reveals that only three techniques are of common practice in a breast assessment. These are: X-ray mammography, MRI and ultrasound (the main characteristics of each technique will be highlighted in Chapter 2). Unlike the first of these, the latter two (and most commonly only ultrasound) are only used to complete and reinforce the conclusion of triple assessment. Hence, the fundamental and primary procedure that “makes a difference” in the diagnosis remains X-ray mammography. As part of the screening programme, X-ray mammography provides the data that the radiologist will subsequently use to conclude: benign, malignant or the desired normal. From the point of view of detecting microcalcifications, only mammograms can provide the necessary spatial-resolution and SNR to distinguish microcalcifications from the background image. Both MRI and ultrasound are inadequate for this purpose.

#### **1.4.2 Remarkable Achievements to Date**

The task of CAD systems is to prompt suspicious regions in an image to the radiologist. The path that all methods must follow starts from image acquisition and visualisation and extends to quantitative and functional analysis. Directly digital mammography (cf. Chapter 2) will simplify the acquisition and will eliminate the digitisation stage. The results of analysis will tell the radiologist how big is the highlighted abnormality, what are its shape characteristics and how it behaves. Studies prove that the performance of radiologists is increased when a good segmentation or classification system is used [42, 48, 135, 152].

The mammogram is enhanced for better visualisation from improving the contrast and reducing the noise [65, 80] to compensating for breast margin and segmenting the pectoral

muscle [20, 79]. Two main branches in mammographic image segmentation are further developed: the detection and classification of microcalcification clusters and masses. Chapter 2 outlines a few notable methods.

The ultimate goal of any method is to be robust enough for clinical application and provide reliable results when used in the hospital environment. A short review of the latest CAD systems is provided in [108]. A remarkable technique worth mentioning in this introduction was launched by R2 Technology with their *ImageChecker*<sup>®</sup>. This is a complex system that outputs both clusters of microcalcifications and dense regions and areas with radiating lines. The results are very promising, as it has been reported that using the *R2 ImageChecker*<sup>®</sup> microcalcifications are detected with 98.3% TP rate with 0.5 FP/image [141]. Other reports mention 100% TP rate on microcalcification detection with 2.2 FP/image, while 81.6% of masses are prompted [41]. However, although the diagnostic sensitivity of the clinicians rose when using the system, the positive predictive value of the clinician's interpretations worsened due to the high number of FP [42]. Furthermore, approximately half of the increase in the recall rate in screening programmes is due to the high number of FP in microcalcification detection [41].

The Standard Mammogram Form (SMF) is an image normalisation technique that eliminates the current limitations of the imaging process and relies only on anatomical breast structures. Chapter 2 will briefly overview the  $h_{int}$  model [65]. The SMF<sup>TM</sup> Workstation developed by Mirada Solutions embeds the quantification of the amount of non-fat and fat tissue for each pixel, temporal registration of the breast, reconstruction of the uncompressed breast and localising microcalcification clusters in 3D [180]. This system obtained a microcalcification cluster detection rate of 95% TP with 0.38 FP per image.

### 1.4.3 The Aim of the Thesis


The systems enumerated previously show impressive advances in detecting and diagnosing breast cancer, as there is considerable research ongoing in the field. Although figures almost reach 100% detection rate, the large number of women undergoing screening (about 3 million mammograms annually in UK only) mean that every percent that remains undetected represents a large number of women that will most probably have to confront cancer at an incurable stage. In Western Europe alone, 1% of the missed cancers would sentence approximately 2350 women each year to face long and painful treatments, both physically and psychologically, with extremely small chances of survival.

The motivation of this work is born from the impetuous necessity to develop detection techniques ready for clinical application. Such systems must prove sufficient robustness and trustworthiness to be used in hospitals in real-time. When dealing with particularly delicate problems, such as human lives, we must assure that results are optimal and accept full liability.


Many methods described in the literature attempt to improve the results of detection of breast abnormalities by tuning the variety of parameters used in the implementation of the algorithm to suit every single case studied. Although the outcome is impressive, the consistency and reproducibility of results is highly dependant on the operators and their capability to find the best parametrical configuration for the detection. The approach used in this thesis proposes a fully automated non-parametric method to detect microcalcifications using the SMF normalised representation of the breast. The aim is to overcome some of the current limitations in methods tackling this subject. The algorithm presented here can be similarly used for intensity images, since its implementation would follow the same logic on grey-level mammograms.

The first intrinsic objective is the removal of noise, as a major source of FP. The algorithm presented here considers several types of noise from quantum mottle to shot-noise and uses subsequent filters to eliminate it. Curvilinear structures (CLS) – ducts, blood vessels, ligaments


or tumour spiculations - proved to be equally important in the computation of the method specificity and a major challenge. A CLS removal step is embedded in the method, prior to the final segmentation. Built on a combination of partial differential equations (PDE), wavelets and statistics, the technique presents the user with a map of detected microcalcifications. The detection of these early signs of malignancy in the breast is meant to assist the radiologists in diagnosing breast cancer. The free-response receiver operating characteristic (FROC) curve of the microcalcification detection method is shown in Figure 3, along with some examples of detection that illustrate the sensitivity and specificity of the approach. The detection of masses in temporally registered enhanced mammograms is also investigated.




**Figure 3:** The FROC curve of the microcalcification-detection method.




**Figure 4:** Detection example 1; (a) the original contrast-enhanced SMF sample with a very large microcalcification cluster in a dense area of the breast; (b) the detection map of the detection method presented in this thesis depicting correctly the cluster.



**Figure 5:** Detection example 3; (a) the original contrast-enhanced SMF sample with a subtle microcalcification clusters in a dense breast area; (b) the detection map.



**Figure 6:** Detection example2; (a) The original contrast-enhanced SMF sample with a widespread microcalcification cluster; (b) the detection map.



**Figure 7** Fig Example4. (a) The original contrast-enhanced SMF sample with a cluster of very small microcalcifications in an area with several curvilinear structures; (b) the detection map.

## 1.5 Overview of the Thesis

The thesis aims to detect features signalling early development of breast cancer in mammographic images. We have reviewed some basic facts about breast cancer and its impact

in the Introduction. A brief anatomy of the breast is also presented here along with some of the most common mammographic anomalies. The incentive of the thesis is closely related to screening programmes; therefore, the screening programme is analysed along with future trends in mammography.

Chapter 2 begins with a comparative overview of current imaging modalities to detect breast cancer. Mammography (X-ray and digital), Magnetic Resonance Imaging (MRI), Ultrasound (US) and Nuclear Imaging (PET, SPECT) are described and analysed. With their strengths and weaknesses outlined, this is followed by a discussion of the detection of anomalies in mammograms. The detection and classification of both masses and microcalcifications using state-of-the-art algorithms are reviewed. The chapter ends with a description of the Standard Mammogram Form (SMF), the image normalisation technique at the heart of our algorithm evaluation.

With Chapter 3 we start presenting the results of our work. The beginning of the chapter introduces analytically the concept of diffusion and particularly anisotropic diffusion, a cornerstone idea in the development of the enhancing filter used throughout the thesis. Our work is the first attempt to use anisotropic diffusion to filter mammographic images and analyse them. This is a priori a reasonable thing to attempt as anisotropic diffusion smoothes the image, reduces noise (hence increases signal to noise, which is generally poor for mammographic images), and yet does a reasonable job at preserving image structure. The filter and its parameters are described and the chapter concludes with the presentation of the first results.

Chapter 4 presents the principal original contributions of the thesis. The removal of shot-noise and curvilinear structures (CLS) is introduced in the image pre-processing. We develop a statistical approach for deriving automatically the parameters of the enhancing anisotropic diffusing filter. The second original step is the development of a method for adaptively thresholding the filtered results in order to segment microcalcifications. This is based upon a model proposed originally to account for certain findings about the human visual system, though, of course, it can be judged on its own merits and by the results that using it leads to.

The alternation of filtering and segmentation steps employed contributes to the novelty of the method. We build several FROC curves for the validation of the detection algorithm. We compare detection methods on SMF images as well as the outcome of our algorithm on both SMF and intensity images.

Chapter 5 expands the use of the enhancing filter to the detection of masses. An original method to prompt mammographic masses is investigated; it uses image registration and enhancement as a pre-processing step, followed by texture analysis-based segmentation and visual comparison between temporal mammographic pairs. The chapter concludes with a discussion of features for mass classification.

The final chapter of the thesis, Chapter 6, lays down a general summary of the work presented in the thesis accompanied by discussion and conclusions. Future work ideas are also underlined and some initial examples are shown to the reader.



## CHAPTER 2

### 2 Diagnosing the Breast

*Real knowledge is to know the extent  
of one's ignorance.*

Confucius

The diagnosis of breast disease is based on a routinised process called triple assessment involving breast surgeons, histopathologists, radiologists, and oncologists. Medical imaging procedures form a key part of the evaluation. Generally, when speaking about medical vision, a variety of imaging techniques should be mentioned. Whether we refer to X-ray and computed tomography (CT), MRI, US, positron-emission tomography (PET) and scintimammography, tissue impedance imaging, infrared or optical imaging [170], there are three major problems to be dealt with:

- images tend to have poor signal-to-noise ratio (SNR);
- images are mostly highly textured and variable from one subject to another;
- clinically significant details are often subtle.

Each imaging technique has its advantages and disadvantages and the clinician would often need to use a combination of them for best results.

Breast imaging faces the same difficulties. While X-ray mammography is the primary imaging method used in screening programmes around the world, MRI and US have become auxiliary tools in the breast clinic triple assessment process. Nuclear imaging has also shown

rapid advances in breast cancer detection, but the cost of installing and maintaining PET systems has until recently slowed its adoption. This is likely to change rapidly over the next few years. The next few sections will expand on each of the techniques used in breast imaging and will be followed by a summary of their strengths and weaknesses.

## **2.1 Imaging Modalities**

### **2.1.1 X-ray Mammography**

X-ray mammography has been widely used to detect the earliest signs of breast cancer since the beginning of last century, due in part to the cost-effectiveness of the procedure relative to other imaging modalities [5]. Screening programmes throughout the world have proved its effectiveness to image non-palpable abnormalities in the breast with very high resolution (the equivalent of a 25 $\mu$ m pixel resolution [75]). The earlier the detection, the better the chance to cure the cancer and to date, mammography is considered to be the best modality to depict microcalcifications and even small tumours. Mammography has the best combination of sensitivity, specificity, low cost and short acquisition, as underlined in Section 2.1.6.


A mammogram is a two-dimensional X-ray image of the mammary gland produced by a radiation beam passing through the compressed breast. Photons are attenuated according to their initial intensity, the thickness of the tissue they pass through and the attenuation coefficient of the respective tissue. The breast must be initially compressed between two plates to even its thickness and spread out the breast tissue for the radiologist to detect density variations easier in a mammogram. This improves the appearance of the mammogram and the image quality is even across the breast, while exposing the patient to a lower radiation dose. Figure 8 shows the clinical imaging machine. From the X-ray source, the photons traverse the breast compressed between the superior and the inferior plates and their intensity will change. The beam eventually reaches the film cassette, which, in the case of a film-screen device,

contains the anti-scatter grid, the film and the intensifying screen. The photons are absorbed by the intensifying screen, which then emits light photons; these are recorded on the film, resulting in a mammogram.


The role of the anti-scatter grid is to absorb the majority of scattered radiation (the photons that have deviated from their initial path after attenuation and re-emission at a different angle). More specifically, the anti-scatter grid absorbs those photons arriving at a film location at a “low angle” determined by the dimensions of the grid. The advantage of an anti-scatter grid is that it removes much of the blur that can otherwise be seen in a mammogram. The disadvantage is that its use necessitates a significant increase in the X-ray dosage to the breast since scattered radiation can account for 40% of the total X-ray radiation exiting the breast [65]. Figure 9 shows the mammographic image formation process and the formation of scattered radiation. Mammographic applications must take into account:

- the relatively weak control of image formation;
- many non-linear effects, such as scattered radiation, time of exposure, breast compression;
- the variation in many image formation parameters between machines or within a single machine over time.

The Standard Mammogram Form (SMF) [65, 67] offers a model of estimating the scattered radiation that can eliminate the need for an anti-scatter grid.



**Figure 8:** The clinical mammographic film-screen machine. The X-rays pass through the compressed breast from the X-ray source towards the film cassette.



**Figure 9:** A representation of mammographic image formation and scattered radiation.


We noted that mammogram is a 2D image of the breast, but the anatomical information is 3D. Two views of each breast are taken: a medio-lateral-oblique (MLO) view (shoulder to opposite hip) and a cranio-caudal (CC) view (head to toe) [190]. In Figure 10 are shown all

four images from a single patient: MLO and CC views of the left and right breasts, as used by a radiologist in detecting abnormalities. Both views show the entire breast area from different angles. The MLO mammograms also image the axilla, a key in determining metastasis. Unfortunately, obtaining the four images can cause distress to the patient, not least because the procedure is rather painful due to the compression of the breast. (During the time that I have spent at the Breast Cancer Unit of the Churchill Hospital in Oxford I have repeatedly noticed patients complaining about pain.)

The 3D correlation between these views is still difficult due to the different compression factors and the considerable inter- and intra-doctor variability in finding the correspondences between regions in the MLO and CC views. Kita *et al.* introduced a correspondence between the MLO and CC in [86, 87] and computed a 3D model of the breast based on the acquired mammograms. This allows masses and microcalcifications to be viewed in the three dimensional framework of the breast, making it easier to diagnose and plan the eradication of tumours.

According to the radiation attenuation of breast tissue, there are three main anatomical categories in the breast, in the reverse order of their attenuation coefficient magnitude: calcifications, glandular tissue and fat tissue. In areas of lower exposure, the film will become brighter and the anomalies will be highlighted on a background of darker fat tissue. The drawback is the inefficiency of mammography in the detection of cancer for young women and women on hormone-replacement treatment (HRT), since tumours and glandular tissue have similar attenuation coefficients. Scars and angiogenesis (tumours show an increase in the local blood flow) are not always imaged in mammography.

Mammography is the only reliable method to visualise microcalcifications and tumour spiculations. For these reasons, mammography is currently the best imaging procedure to detect pre-invasive diseases and the most appropriate modality for screening.



**Figure 10:** The combination of four images used by the clinician in diagnosing the breast; (a) a pair of MLO images showing the breast, the pectoral muscle and axilla; (b) a pair of CC images from the same patient.

### 2.1.2 Digital Mammography

Currently, mammography is overwhelmingly film-based. Although many technical problems have retarded the uptake of mammographic image analysis of digitised films – in which X-ray photons are first converted to light, then expose the film, and then to electrons as the film is digitised –, directly digital mammography – in which X-ray photons are converted directly to electrons – is now available [12, 37, 142]. The digital signal is then stored on a computer in the

form of an image, which can be visualised on a workstation [106] or printed on a film. Full-field digital systems make a clear distinction between image acquisition, image storage and image visualisation. There are several clear advantages that digital imaging has to offer:

- *digital image acquisition* is expected to reduce the X-ray dose in the imaging process, (because of the higher sensitivity of the charged-couple device (CCD) versus film to radiation [158]) leading to less risk for the patient; it also improves the contrast resolution and the SNR will be higher as film granularity will no longer be a source of radiographic noise [134];
- *digital image storage* reduces the cost of the operation since it does not require film and chemicals, provides fast and reliable retrieval from the archive and allows the collection of large image databases. It also prevents films being lost between screening visits – currently 22% of films are lost;
- *digital image visualisation* enables the use of Computer Aided Diagnosis (CAD) for automatic detection, data documentation and image processing, unlike the rigid film. Given that the digitisation of A4 size X-ray films to 50 micron pixels is typically very slow (up to 10 minutes), directly digital mammography could well be the catalyst that leads to the widespread use of CAD.

Early detection can be accomplished only with high quality detection and processing tools in the different levels of the screening programmes. Directly digital mammography is the new trend in X-ray mammography, which should bring notable improvements in the development of screening programmes. The *soft-copy reading environment*, a computer workstation that displays digital mammograms for the radiologist to read, is a tool that will make digital mammography available in hospitals for better results in detecting the structures of interest within the breast. Figure 11 shows the digital mammography machine developed by Siemens. Recent evaluations [60, 61, 136] proved that radiologists have a similar performance to using film-screen mammography when working on soft-copy machines.



**Figure 11:** A soft copy environment produces digital mammograms, which are stored directly into a computer. This figure shows the Opdima System from Siemens, which was designed for near real-time computer guided biopsy.

### 2.1.3 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is a complementary imaging technique to X-ray in the diagnosis of breast cancer. X-ray mammography has evolved significantly over the past few decades and is an extremely useful imaging modality for screening programmes, but needs to be seconded by MRI and US for best clinical results.


As its name implies, MRI results from magnetising the body tissue of the patient. The patient is introduced into the scanner (Figure 12), where she must lie with as little movement as possible for 20-30 minutes. An external magnetic field is applied to the organ to be imaged. The water protons (which are abundant because of the preponderance of water in human physiology) enter a higher energy state when a radio-frequency pulse is applied and this energy is re-emitted when the pulse stops. A coil is used to measure this energy, which is proportional to the quantity of hydrogen and therefore specific to the tissue type. The fat and dense tissues of the breast thus give different intensities in the final MR image [27]. However, to date, not pulse sequence has been found which demonstrates a clear differentiation between carcinomas and benign tissue changes.






**Figure 12:** The MRI scanner with a patient in the right position for an MRI of the breast.

MR images offer a good 3D representation of the organ; they are a collection of successive slices of the region of interest (ROI), a profitable way to visualise 3D data. Images can be axial, as in Figure 13 (slices perpendicular transversal to the body, similar to the CC view in mammography), sagittal (slices perpendicular along the body, closer to the MLO view in mammography) and coronal (slices parallel to the body through the breast). The information gathered from any of the 3 views can be used to build a 3D representation of the breast [148]. Also, fully 3D MR sequences are available, in which data is captured in a 3 dimensional 3D Fourier space, rather than the data being captured separately in a set of individual slices in a 2D Fourier space and then stacked to make a 3D volume.



**Figure 13:** An axial T1-weighted MRI slice of the breast (using a gradient echo sequence), the closest view to the X-ray CC image. In this image, fatty tissue appears brighter, as its T1 value (around 200 ms) is considerably less either that that of normal healthy tissue (typically 700-1000 ms) and of cancerous tissue (typically 1500 ms).

Knowing that tumours have an increased localised blood flow, called neoangiogenesis, researchers embedded some dynamic information of the breast tissue in MRI. Using contrast agents, the examination of the angiogenesis of the breast can discriminate between malignant tumours and the surrounding tissue (Figure 14). The contrast agent used in clinical practice is Gd-DTPA (gadolinium diethylene triamene pentaacetic acid) [65], which is paramagnetic due to Gd's unpaired electrons. Its intravenous injection will enhance the blood vessels in the MR image and highlight the regions with high concentration of blood supply and vessel leakage, such as carcinomas. Malignant tumours have elevated vascularity both at the edges and within the tumour [157], unlike benign anomalies and can be identified in a sequence of pre- and post-enhancement MR images [59, 160].



**Figure 14:** A sequence of contrast-enhanced MR images of the breast; (a) is the MR slice before the injection of contrast agent; (b) is the MR slice after contrast enhancement, where the tumour is highlighted due to its higher vascularity.


In addition to X-ray mammography, MRI has shown great results in detecting tumours in pre-menopausal women, where mammograms cannot distinguish between parenchymal tissue, scars and carcinoma. Women with implants or undergoing HRT treatment also benefit from MRI. Although there is no harmful radiation involved in the process, Gd-DTPA is actually quite toxic and some patients have adverse reactions. There is no breast compression in breast MRI, but the time of acquisition is much longer. Furthermore, very small pre-invasive lesions as well as microcalcifications are not visible in MRI [126, 168] and the spatial resolution is much lower than in a mammogram ( $1 \text{ mm}^3$  voxels, while the most common resolution in mammography is 50 micron pixels). The currently high capital and recurrent costs of MRI, the cost of the contrast agent, and the time taken for the procedure make MRI unsuitable for screening programmes. However, the UK Government has started a pilot study aimed at MRI screening of young women who are epidemiologically at highest risk of contracting breast cancer.

The primary drawback of MRI is the inability to image microcalcifications and there is extensive debate ongoing on this subject [63]. Until researchers can overcome this inadequacy, there is substantial work done in data fusion. A combination of X-ray and MRI can combine the early visualisation of anomalies with the availability of 3D data and angiogenesis.

### 2.1.4 Ultrasound Imaging

Ultrasound (US) is the third imaging modality currently used frequently in clinical assessments. When something unusual is noticed in the mammogram of a patient, it is commonly the case that an US image is subsequently taken and analysed. To date, there is no cheaper and simpler method to visually investigate the breast than US. For this reason, and its efficiency to differentiate between soft and hard tissue, US is used as a complementary detection and evaluation method next to X-ray mammography.


US images of the breast are formed from the reflections produced by high frequency acoustic waves that reach the breast tissue [170]. The response of the different tissue depths and densities will return a map of densities. This is the main utility of US, that of distinguishing between soft and hard tissue and therefore, to some extent, between benign and malignant masses. Figure 15 shows the HP Sonos 5500 echographic machine.



**Figure 15:** An ultrasound machine; as the probe is swept over the patient's body, the clinician can visualise in real-time the US images on the machine monitor.

US images show strong echoes for calcium and skin, but because of the small size of microcalcifications and the poor SNR, only macrocalcifications may be visible. Then there are

weaker echoes from the parenchyma and finally, almost no reflection from fat and tumours [105]. Figure 16 shows the representation of a cyst and a cancer in US. The cyst appears as a dark consistent region surrounded by glandular and fat tissue, slightly squeezed by the pressure of the probe. The cancerous tumour is more vertically elongated with fuzzier edges and there is a shadow following it on the lower part of the image (opposite the US probe). The reduced SNR and the low lateral resolution mean that their margins are poorly defined.



**Figure 16:** Two US images of the breast; (a) the image of a cyst, a compact dark area squeezed by the probe; (b) the image of a tumour, an elongated dense area with less well-defined margins.


Ultrasound images can be obtained using a free-hand probe that is swept over the ROI. Using 3D reconstruction, the area of the breast can be visualised three-dimensionally [176]. Due to the limited dynamic resolution, US does not provide a full 3D reconstruction of the breast, but the information is essential in determining the extent of the tumour. Contrast agents are also used in US as well to enhance the acoustic impedance of tissue [82, 151]. Furthermore, an elastographic approach is used by clinicians when palpating the breast with a US probe [18, 31, 46]. The displacement of the ROI is estimated and the larger that it is the higher the chances to have detected a benign tumour. Other features used by radiologists to classify lesions include the margin definition, echogenic texture, shadowing and lesion shape [71].

Besides being the best detector of cysts, the real-time response of the procedure makes it valuable in performing image-guided minimal-invasive biopsy. It is also applicable to patients of all ages and, from the patient's perspective, it is the most comfortable breast imaging procedure. The changes caused by HRT do not represent a problem when searching for fibrocystic diseases. US has its disadvantages though: it has the poorest resolution and SNR of the modalities that we are considering and does not image microcalcifications and very small lesions. Although US remains unsuitable for breast cancer screening, it is a field that attracts a great deal of interest from researchers and many people regard it as the next major step in screening programmes.

One notable achievement related to US is the vibro-acoustic tissue mammography proposed by Fatemi *et al.* [38] in which we can visualise microcalcification in an US image for the first time. They use low frequency US to make the tissue vibrate and utilise the response to image the tissue hardness. The method is at its very early stages and is not as reliable in imaging microcalcifications as the X-ray mammography, but opens a new route for breast imaging.


### **2.1.5 Nuclear Imaging**

Nuclear medicine is a relatively new branch of breast imaging. Positron emission tomography (PET) and single-photon emission computed tomography (SPECT) have recently begun to be more widely used in breast diagnosis. The PET scanner appears similar to an open MRI scanner, as in Figure 17. Scintimammography (SPECT of the breast) has proved to have good sensitivity for detecting metastasis and palpable cancers [19, 128]. Images are acquired in multiple cross-sectional slices of the breast by moving a single planar detector around the organ or by using multiple detectors around the body of the patient [170], offering a good 3D localisation of the tumour.



**Figure 17:** A PET scanner.

Nuclear medicine uses radiopharmaceuticals to differentiate between tumours and normal tissue. Such agents are injected into the body, and since tumours tend to have a much greater uptake of the agent, they are identified by a gamma camera [19]. Scintimammography uses a very small dosage of radioisotope with low residual concentration. As Webb observes [170], there is no measurable toxic effect with the use of radiopharmaceuticals since the emitted radiation is not very strong from a small mass of isotope. Figure 18 shows some examples of scintimammograms and the sensitivity of the technique to multi-focal cancers.



**Figure 18:** An example of scintimammograms showing the sensitivity of SPECT images to multi-focal tumours.

Nuclear medicine is expensive, although scintimammography is less expensive than MRI. The cost of a PET scan is currently about £700, compared to about £350 for MRI and less than £50 for mammography and ultrasound. Its sensitivity is high, but it fails to detect small non-palpable tumours. PET actually has the best sensitivity and specificity because of the cell uptake of sugar and the subsequent positron emission from the radio tag. Scintimammograms are feasible in detecting cancer on younger women and dense breasts, but have low resolution and poor SNR. Its spatial resolution will soon rival that of MRI, since it has improved substantially over the past few years. The utility of nuclear medicine is also evident in tracing the effects of chemotherapy faster than MRI, by imaging changes in the local angiogenesis. The axilla are also well imaged, but the acquisition time of PET is about 60 minutes.

### **2.1.6 Strengths and Weaknesses of the Breast Imaging Techniques**

We have briefly reviewed the imaging modalities used today in the diagnosis of breast cancer. Clinical practice has shown that none of the techniques suffices for every patient, so in general a combination is required. The triple assessment associated with screening programmes uses mammography as the main and US (and sometimes MRI) as the complimentary imaging procedures. MRI and nuclear imaging are used to a more limited extent in the detection of breast cancer. Table 1 compares the strengths and weaknesses of each assessed method.

Reviewing the advantages and disadvantages of mammography, MRI, US and nuclear medicine, it becomes evident that at present mammography offers the best compromise between costs and performance. For its good specificity, sensitivity and low cost, mammography was chosen as the main screening modality. Its widespread use and obvious utility are evident reasons for extensive research to be carried out in improving the impact mammography has on the detection of breast anomalies.



**Table 1:** The strengths and weaknesses of breast imaging procedures

	<b>Mammography</b>	<b>MRI</b>	<b>US</b>	<b>Nuclear Imaging</b>
<b>Tumour sensitivity</b>	High	Poor to pre-invasive tumours	Low	Poor to pre-invasive tumours
<b>Microcalcifications/ lymph nodes/ spiculations</b>	Yes/Yes/Yes	No/Yes/No	No/Yes/No	No/Yes/Yes
<b>Multi-focal diseases</b>	Poor	Yes	Poor	Yes
<b>Young + HRT women</b>	No	Yes	Yes	Yes
<b>Resolution/ SNR</b>	High/High	Medium/High	Poor/Poor	Poor/High
<b>Costs</b>	Cheap	Expensive	Cheap	Expensive
<b>Acquisition time</b>	Short, but unpleasant	20-30 minutes	Real-time	60 minutes
<b>3D information</b>	No depth info, 3D applications	Yes	Yes	Yes
<b>Angiogenesis</b>	No	Yes	No	Yes
<b>Toxicity/ compression</b>	Yes/Yes	Yes/ natural deformation	No/ deformation	Yes/ natural deformation

## 2.2 The Detection of Mammographic Anomalies

The development of Computer Aided Diagnosis (CAD) systems has reached the point where they offer extremely valuable information to the clinician in the detection and classification of abnormalities. So far, they can only assist the medical staff in making a decision. To date, a CAD system performs about as well as a radiologist, but the combination can perform better than either alone [48, 76, 135, 152]. In X-ray mammography, CAD systems hope to assist the clinician in diagnosing breast cancer at the earliest possible stage. Their main use to date is in screening programmes, where the large number of mammograms to be processed requires a large number of radiologists and the difficulty of their interpretation necessitates robust and reliable assistance. Furthermore, the rapid development of digital mammography increases the

utility of CAD in everyday image processing (e.g. zooming, contrast enhancement, edge detection, image registration and subtraction) and fully automated detection methods.


There are two groups of major anomalies in the breast: masses and microcalcifications; detection and classification methods usually tackle one of the two. The following two sections review the state-of-the-art in detecting and classifying masses and microcalcifications in X-ray mammography. The first complex CAD system detecting both masses and microcalcifications approved by the Food and Drug Administration (FDA) in the USA was the R2 Technology *ImageChecker*<sup>®</sup> [141, 152].

### **2.2.1 The Detection and Classification of Masses in X-ray Mammography**

Masses in the breast take their name from the characteristic well-defined mammographic appearance. They tend to be brighter than their surrounds due to the high density within their boundaries [95]. In fact, they do not always have clearly defined edges, this definition being typically applicable to benign masses. Most malignant lesions have rather ill-defined forms, since they infiltrate into the surrounding tissue and may present radiating spiculations (stellate lesions) [127, 159]. Malignant microcalcification clusters might also develop in association with cancers. Figure 19 illustrates the variety in the appearance of breast tumours. Some masses appear to create a ‘halo of security’, which contradicts the ideal mass appearance and which may correspond to the angiogenesis region surrounding the dense necrotic tissue [65].

Mass detection algorithms aim to depict tumourous regions in mammograms by differentiating abnormal tissue from the fatty background and parenchyma. The task is extremely challenging due to the very broad variety of masses and the subtle appearance of some of them. Furthermore, parenchyma has similar density to that of tumours at the passage of X-ray through the breast and represents the main source of false positives (FP) in mass detection algorithms. The 2D highly textured appearance of mammograms with overlapped tissues contributes to the significant percentage of errors encountered in tumour detection.

Using information from both MLO and CC views can clarify some of the ambiguities. Unfortunately, although recent methods tend to equal the radiologist's performance, no single algorithm has yet been developed to function well on all mammographic cases.



**Figure 19:** Some examples of masses in X-ray mammography; (a) a benign cyst will well-defined boundaries and an ellipsoidal self-contained shape; (b) a stellate mass with spiculations radiating from the central mass into the surrounding tissue; (c) an ill-defined mass with low contrast at the boundary, which can be easily confused with the neighbouring tissue.

The classification of tumours as either benign or malignant is performed by taking into account several mass features. The main features used are the shape of the lesions (spiculated or circumscribed), the texture of the mass and the contrast around the edges.

When related to masses, automated image analysis relies on three main applications:

- Locating abnormal regions in a single mammogram, when, as a result of several features calculated (i.e. enhancement, likelihood measure), possibly pathological regions of the breast are extracted and pointed out to the radiologist [14, 16, 24, 49, 52, 79, 88, 92, 94, 115, 145, 146, 184, 185]. Such methods are usually designed on different types of mammographic lesions.
- Matching bilateral pairs, making use of the same-view left and right breast mammograms of the same woman at the same time, acknowledging the approximate symmetry of the two breasts [33, 78, 182]. Areas of high asymmetry are labelled as

suspicious in both CC and MLO views and referred to further examination. Previous image registration is required.

- Matching temporal pairs, using the same-view of the same breast mammograms of the same woman at two different times, searching the major changes that appear between the two mammograms [89, 97, 104, 109, 110, 111]. All regions showing changes over time must be thoroughly evaluated, excepting the involution of dense tissue into fat. The breasts of women on HRT have denser appearance in mammograms over time, an exception to the general rule. The registration of the temporal pairs is also necessary before analysing the images.

The first approach to detecting masses is probably the most elaborate, since it does not look for bright and unmatched regions, but it searches for sets of features that describe the appearance and properties of a breast mass. We shall further expand on the most popular approaches leading to the most recent and relevant achievements in detecting breast lesions in single mammograms.

Most algorithms for detection and characterisation of masses follow a three-step process. They start with a pre-processing step, which aims to filter the image and remove some possible FP (i.e. noise or background removal). The second step prompts suspicious areas, which form a set of mass candidates. The selection is done after computing statistical features of the mammogram at pixel level, sometimes aided by learning, for example by neural networks [15, 49, 79, 147, 175]. One of the most basic methods to segment lesions is based on region growing, labelling connected similar pixels from a seed [91, 98]. The final step is the characterisation of the depicted masses as either malignant or benign, based on the analysis of shape, sharpness, size and texture (Chan *et al.* [24] extract a total of 320 spatial grey level dependence texture features).

The detection step mostly follows along conventional lines. One of the common approaches preferred in identifying masses is the convolution of the breast image with a second order derivative (Laplacian of the Gaussian), a zero mean filter with a positive centre surrounded by negative boundaries [16, 147]. For such algorithms, the most relevant feature of

a mass is its local brightness. The method will prompt lesions with a bright central mass, but may fail in detecting more subtle masses. Another approach is template matching using a model of a mass to find resemblances between the model and areas of the breast [94, 123]. This approach performs better on faint masses, since the correlation is a normalised intensity measure, but the wide variety of masses will still confine it to the detection of central masses. Furthermore, it will depend on the shape of the chosen template.

The evolution of mass detection methods lead to the development of more sophisticated algorithms, since the above mentioned approaches do not have the necessary flexibility to adapt to the variety of mass types in mammograms. Statistical analysis of the mass area can bring valuable information, since the orientation of the gradients will be radial and pointing towards the centre of the mass [16]. Groshong *et al.* [52] adapted the concept to a Hough transform. A signature based on a recursive median filter at various orientations on a pixel was used to determine the shape of a structure as either blob or line by Zwiggelaar *et al.* [184, 185]. Miller and Ramsay [115] perform a multiscale non-linear analysis of maximum entropy, while Li *et al.* [97] use morphological filtering. In another approach, the distance from the nipple to the mass in both MLO and CC images is evaluated to discard FP [33].

Chan, Sahiner *et al.*, developed an original method in their ROC study to detect and characterise masses [24, 134, 145, 146, 147]. The performance of the classifier was evaluated using ROC analysis, after a previous training step performed on 238 mammograms. A background correction is applied to the original images and three images are subsequently generated for each one of them: a median-filter smoothed image and two high-frequency enhanced images which are used in a clustering algorithm to classify the pixels as either mass or background. The outcome of the segmentation of the mass provides a mass region smaller than the real mass and then using a rubber-band-straightening transform (RBST), a 40-pixel wide region around the tumour, the mass is transformed into a rectangular shape. The RBST transforms the band of pixels surrounding the mass into a Cartesian plane (the RBST image) and has the advantage of making the mass margins almost parallel, with perpendicular spicules to the length of the rectangle.

Chan, Sahiner *et al.*'s classification algorithm relies on two types of texture features: measures from the spatial grey-level dependence matrices of the RBST image containing information about image characteristics (i.e. homogeneity, contrast, etc. resulting in 320 features); and features computed from the run-length statistics matrices of the horizontal and vertical gradient images of the RBST image (20 features). Using statistical analysis, the most effective features are retained, eventually reducing to 41, and these are used as input for Fisher's linear discriminant classifier [10] to compute the relative malignancy rating of masses.

The likelihood of a mass to be malignant is very high when spiculations are present. Besides detecting the central mass, there has been considerable interest shown in the detection of spicules. They can be associated with a central mass, although there are architectural distortions that can be prompted only by detecting the radiating spicules. Spicules typically radiate from the centre of the mass (their pixels are directed to the centre of the mass [14,16]) and have the histogram of orientations flatter than normal tissue [83]). While a normal mammogram has a radial pattern of ducts converging towards the nipple (especially at lower resolution), a spiculated area introduces another radiating centre. Kegelmeyer [83] first noticed the difference between the largely homogeneous orientations of normal edges (in small windows) and the variety in orientations of suspicious edges. He proposed to use this histogram of their orientations to differentiate between them. A model for spicules was built in [130]. Curvilinear structures (CLS) may sometimes be confused with mass spiculations. We shall describe the detection and removal of CLS in Chapter 4.

Karssemeijer and te Brake [14, 15, 16, 78, 79] developed an elaborate algorithm to detect stellate lesions without relying on the presence of a central mass. The method is laboriously tested on different sets of images with a multiscale use of parameters. It makes use of familiar pre-processing steps, including noise equalisation [80], pectoral muscle removal from MLO images and breast edge correction. The results converge to an approximate fraction of 90% TP at 1 FP per image.

Spiculated lesions are detected more easily than architectural distortions. If an increase in pixels oriented towards a certain central region is found, then that area may be abnormal. The

best results are obtained by combining the line-based pixel orientation map with features signalling the presence of a central mass. Also, the multi-scale approach improves the detection figures, but is very dependent on the choice of the optimal scale. A comparison is made between three methods for mass detection: Laplacian filtering, template matching and gradient orientation analysis with clear general best results provided by the latter method.

The automated classification of masses in digital mammograms to a level that is clinically acceptable remains an open subject. There are no clinically robust methods to differentiate between benign and malignant lesions. There are several features that most papers take into account: the size of the lesion, its shape (round for benign and jagged for malignant), the sharpness of the intensity transformation at the edges (well-defined for benign, vague and stellate for malignant), texture and contrast [24, 55, 120, 139]. Some original work can be found in [73] on the use of a radial edge gradient and [137] on looking at the lesions external roughness.

Kita *et al.* [86, 87] found a spatial correspondence between the MLO and CC views of the same breast. The result of this was a 3D reconstruction of the breast with a model of the detected mass within it. The authors computed a model of the decompressed breast from both views, after having the mass prompted, and add the recomputed mass coordinates to it. There are some obvious difficulties related to the MLO-CC correspondence of the deformed breast. A considerable number of approximations are employed in this method and they result in obtaining a model of an “idealised” breast, hardly matching the large variety of shapes encountered in screening programmes. Some of the factors contributing to this are: the estimation of the rotation angle of the X-ray source to acquire the MLO view; the modelling of the expanded breast area (when squeezed between the compression plates) by dilation proportional to the breast thickness (later improvement introduced by Yam in [180]); the correct detection of the nipple position; the ellipsoidal approximation of the cross-section of the breast. The authors acknowledge the insufficient deformation data that can be extracted from just two views and introduce some approximations regarding tissue movements. Although the 3D reconstruction of the breast from its two breast views is rather non-deterministic, the

method computes an estimation of the tumour or cluster of calcification location and situates the abnormality in an actual 3D context.

Highnam and Brady [65] comment on the advantages of using the  $h_{int}$  representations, noting that a mass would correspond to a hill-like structure surrounded by a smoother region corresponding to fibroglandular tissue. The ultimate aim of the detection of masses is their interpretation; therefore, the shape and arrangement of the salient region are relevant features in this process. Section 2.3 of this Chapter will introduce the Standard Mammogram Form, as image representation of  $h_{int}$ .

A potential problem that arises in evaluating detection algorithms is the use of the same database to both train and evaluate the algorithm. Since such detection algorithms need to adjust their particular parameters in the training stage, there is a clear risk that they might end up being tuned to a specific image database and their performance on a different set of images remains uncertain.

Increasingly, detection algorithms are evaluated on a set of publicly available databases including that provided by the University of Nijmegen (50 $\mu$ m, but no longer available), MIAS (50 $\mu$ m, <http://www.wiau.man.ac.uk/services/MIAS/MIASweb.html>, available free of charge at small resolution) or University of South Florida (42-50 $\mu$ m, <http://marathon.csee.usf.edu/Mammography/Database.html>) databases. Although this provides a good basis for algorithm performance comparison (when several methods are compared on the same collection of images), there also involves the danger of developing worldwide detection algorithms that may only work for certain image acquisition characteristics. For example, there is a huge difference between the same mammogram (film) digitised on a laser scanner (e.g. Lumisys) or CCD (e.g. Canon). Rarely, images from more than one database are used in the evaluation of methods (i.e. the Nijmegen group uses images collected from two databases), but no explicit comparison is made between results on the different sets of images. Image normalisation is the solution to these problems.



## 2.2.2 The Detection and Classification of Microcalcifications in X-ray


### Mammography

Microcalcifications represent one of the earliest signs of breast cancer and they account for half of the non-palpable lesions that appear in mammograms. Calcifications are small deposits of calcium (and related) salts representing either warnings of malignancy or just benign formations. They are encountered in approximately 25% of mammograms and appear as bright spots or clusters of such spots, due to the high X-ray attenuation factor of calcium [65]. Figure 20 shows some examples of samples of mammograms containing calcifications. To differentiate malignant from benign microcalcifications, radiologists believe that they take into consideration several criteria, such as their shape, size and arrangement of calcifications as clusters.

According to their size, calcifications can be classified into macrocalcifications or coarse calcifications (usually interpreted as meaning that their size exceeds 1 mm) and microcalcifications. While it is believed that large, single, regular-shaped calcifications are benign, small clustered whorled calcifications are more likely to signify malignancy, as noticed by Caseldine *et al.* [22] and Le Gal *et al.* [44]. Although the detection and classification of calcifications are two fields that have improved significantly in recent years, there is still no robust differentiation between benign and malignant calcifications. However, each incremental improvement in the detection rate has a potentially significant impact on breast cancer screening.

The aim of automatic detection is to find a sufficiently reliable algorithm to be used in clinical practice. Such algorithms are meant to assist the radiologist in making decisions and to improve the overall sensitivity (the ratio of TP) and specificity (1 minus the proportion of FP, c.f. Appendix C) of the detection process. Some of the factors that drastically influence the TP and FP figures are:

- the variability of the anatomy of the breast; every mammogram has different features related to different tissue types and correspondingly variable brightness in the mammographic appearance;
- the imaging conditions, such as shot noise, quantum mottle, patient movement, low contrast in mammograms due to low X-ray dosage and glare;
- faint microcalcifications lost in a dense background, the superposition of certain breast structures (such as CLS).



**Figure 20:** Some examples of microcalcifications in X-ray mammography; (a) an isolated large calcification; (b) a group of two isolated macrocalcifications; (c) a subtle cluster of microcalcifications following the shape of a duct (ductal carcinoma *in-situ*); (d) a compact malignant microcalcification cluster.

There are many aspects of microcalcifications that need to be understood before attempting to detect them. Lefebvre discusses them carefully in building a simulation model of calcification clusters [96]. In order to overcome the above-mentioned compromising factors,

most conventional detection algorithms consist of three main stages, similar to the approach taken in detecting masses:

- a preprocessing step based on filtering the image; the filter is meant to detect and remove noise and enhance the structures of interest, the microcalcifications; Chan *et al.* [25] and Nishikawa *et al.* [124] use a difference image, the result of subtracting a signal-enhanced image from a signal-suppressed image;
- a detection and segmentation step based on adaptive thresholding or local contrast, due to the bright spot appearance with very high contrast of the microcalcifications; in [124] both global thresholding and morphological erosion are used, in [153] the segmentation is done by a region growing algorithm, in [28] an adaptive thresholding is combined with feature analysis, while in [8] morphological erosion and dilation are used to detect microcalcifications;
- a clustering step using a fixed size kernel to eliminate noise points and isolated calcifications and identify clusters (formed of more than 3 calcifications).

Although most conventional detection methods are based on the above three-step algorithm, the literature includes some novel detection methods, which are further discussed. These new approaches tend to overcome some of the problems that conventional algorithms have with faint microcalcifications and image noise.

One approach uses automatic neural network classifiers in subtracting calcification candidates from mammograms and then building clusters with them [85, 143]. As expected, the classifiers need to be trained and may become adapted to certain databases. Neural networks are further used to discard FP [32] and classify [56]. Aghdasi [1] uses a neural network classifier after using an adaptive wavelet transform to enhance the signal. Wavelet filters are also used in [93] and [113] to highlight the high frequency signal from the background. Noise is also high frequency and therefore remains the main problem to be solved.

Yam *et al.* present a physics-based approach to detecting microcalcifications in [178, 179, 180]. The algorithm uses both grey level and SMF images and is based on two of the major characteristics of calcifications in each of the image types:

- on a grey level image, calcifications have an X-ray attenuation that is about 26 times higher than that of fat or dense tissue; therefore, on a 3D plot of a mammogram they “grow” higher and more abruptly than the surrounding tissue;
- on an  $h_{mt}$  image, the estimated volume of the interesting tissue corresponding to the region where a calcification is detected must exceed the estimated volume of the real 3D model of that region;

The detection method is described in more detail in Section 2.2.3

Most existing classifiers for microcalcification clusters consider 2D features from the mammogram. A sense of 3D information is implied in the algorithms that utilise both CC and MLO views. Yam *et al.*'s 3D reconstruction of microcalcifications builds a fully three-dimensional model of the cluster and allows a thorough visual and statistical analysis of the microcalcifications. The improvements brought by using information from CC and MLO views are also underlined in [144]. Taylor *et al.* originally investigate the use of computerised decision-making support to classify microcalcifications [2, 161, 162].

Another innovative detection method is introduced by Karssemeijer in [80, 81] and further developed in [165, 166, 167] and [113] to reduce false positives. An adaptive noise equalisation algorithm was developed to deal with the variation of noise characteristics in an image to make the detection algorithm less dependant on image acquisition. Karssemeijer builds his noise equalisation model by considering the strong dependency of image noise on signal intensity. He relates the term “noise” to the standard error of feature values. The rescaling of data is then based on a high-pass filtered image representing local contrast, since local features will only depend on high frequency noise components.

The segmentation uses Bayes' rule for labelling and a Markov Random Field (MRF), a combination widely applied in the enhancement of noisy images and subsequent classification. These are the first results obtained from applying simulated annealing in mammography. The

method uses an extended number of parameters making its performance extremely dependent on them. The parameters being very specific, the algorithm is oriented towards detecting only certain types of calcifications, as emphasised in [72]. Poissonnier and Brady also comment on the ineffectiveness of the noise equalisation approach in images that have the relation between intensity and density different than on the Nijmegen databases [138]. Nevertheless, it should be acknowledged that it is an elegantly complex innovative approach with very promising results in detecting microcalcifications and reveals a different perspective in approaching mammography.

Reports [42, 141] show figures of 98% TP with the R2 Technology *ImageChecker*<sup>®</sup>, a great improvement in the sensitivity of breast cancer diagnosis with a number of 0.5 FP/image as the state-of-the-art in the detection of microcalcifications. All the clusters are reported detected at 2.2 FP/image [41]. The results are excellent, but there is still room for improvement. Since perfect detection figures could not be achieved, the ideal numbers should tend towards 99% TP with 0.1 FP/image. The work of this thesis aims to take a next step towards achieving such results.

The reduction of the number of FP remains the major problem that researchers have to solve in the detection of microcalcifications. In [164], a Hough transform is used to detect calcified vessels and discard them, while Edwards [32] uses a Bayesian Neural Network to eliminate FP. An important source of FP is the presence of film-screen artefacts. Highnam and Brady [3, 66] use a model-based approach of the blurring functions in the X-ray imaging process to detect dust and dirt on the film. The  $h_{int}$  normalised representation of the breast imbeds the film screen artefact removal and offers a ‘cleaner’ image for subsequent detection algorithms.

The classification of microcalcification clusters as either benign or malignant is based on a number of cluster features describing the shape of the cluster, the shape of the individual microcalcifications and their distribution. The resulting set of computed characteristics is input to a neural network or pattern recognition system. In [153], a back-propagation neural network on three layers is reported, while in [129] the classification is performed by a k-nearest

neighbour classifier. Artificial neural networks are also used in [56] and [74]. Patrick [131] describes an expert system subdivided in five sub-systems using clinical data and the CC and MLO views. Both views are also used in [165] in a k-nearest-neighbourhood classification.

### 2.2.3 Yam *et al.*'s Physics Based Approach


Yam *et al.* [179] proposed a method to detect microcalcifications using both grey-level and Standard Mammogram Form (SMF) images [65] (see Section 2.3). The values corresponding to the  $h_{int}$  of pixels corresponding to calcifications do not represent the real thickness of the interesting tissue of the breast in that specific area. Therefore, the physics-based approach makes two major assumptions:

- microcalcifications have an X-ray attenuation that is 26 times higher than normal tissue; in SMF microcalcifications appear as towers – this is approximately true for a range of X-ray photon energies as is shown on page 40 of Highnam and Brady's book [65] and reproduced here in Table 2;
- the volume of the detected microcalcifications must exceed the estimated volume of the 3D model of that region;

The first step of the method described by Yam *et al.* is the region extraction. A 12-bit image with grey levels between 0 and 4095 is thresholded at every 32 grey level between levels 2000 and 4000, resulting in layers of surface regions surrounded by contours of the same intensity. This is essentially a watershed algorithm. Only those regions with an area corresponding to the size of calcifications are selected. An area change constraint is then imposed to the candidate regions. All candidates are discarded unless for every  $r_G$ :

$$\Delta A(r_G, r_{G+32}) \leq \Delta A_{thresh} \quad , \quad (1)$$

where  $r_G$  is the candidate region,  $G$  represents the grey level,  $\Delta A_{thresh}$  is a preset threshold and  $\Delta A(r_1, r_2) = (A_1 - A_2) / (A_1 + A_2)$ , as seen in Figure 21.



**Figure 21:** The extraction of candidate microcalcification regions in Yam's algorithm.


The second stage of the detection algorithm computes the volume of the interesting tissue of the extracted regions. The corresponding blob's volume  $v_{int}^{blob}$  is calculated according to equation (3), where  $v_{int}^{blob+surr}$  is the sum of the  $h_{int}$  values of all pixels within the extracted region and  $v_{int}^{surr}$  is the volume of the interesting surrounding tissue. Figure 22 shows the interesting volume computation.  $N$  is the number of pixels within the blob region,  $p$  is the pixel size,  $r_d$  and  $r_e$  represent the extracted region and the region after dilation respectively and  $n$  is the number of pixels in  $r_d \setminus r_e$ .

$$h_{int}^{surr} = \sum_{i \in r_d \setminus r_e} h_{int}(i) / n \quad (2)$$

$$v_{int}^{blob} = v_{int}^{blob+surr} - v_{int}^{surr} \quad (3)$$

$$v_{int}^{blob+surr} = \sum_{i \in r_e} h_{int} \times p^2 \quad (4)$$

$$v_{int}^{surr} = h_{int}^{surr} \times N \times p^2 \quad (5)$$



**Figure 22:** Computing the interesting tissue volume (removing the background) in Yam’s algorithm.

The other relevant volume for the method is the volume of the extracted blob estimated from the image. The blobs are assumed to have ellipsoidal shape and therefore the corresponding 3D volumes are estimated by ( 6), where  $a$  and  $b$  are half of the length of the sides of a rectangular box that encloses the region.

$$v_{3D}^{blob} = \frac{4}{3} \pi a^2 b \quad (6)$$

The interesting tissue composition is then defined as  $v_{int}^{ratio}$ , as in equation ( 7). Finally, the ratio is thresholded at some preset value and the regions with  $v_{int}^{ratio}$  above the chosen value are marked as calcifications.

$$v_{int}^{ratio} = \frac{v_{int}^{blob}}{v_{3D}^{blob}} \quad (7)$$

The described method stands for the segmentation step in the detection algorithm for microcalcifications. Yam *et al.* introduced a de-noising approach for  $h_{int}$  images built upon a Wiener-type filter in [178].

Yam uses a simplified Wiener filter, which has the following transfer function in the frequency domain:

$$G(u, v) = \frac{P_s(u, v)}{P_s(u, v) + P_n(u, v)} \quad (8)$$



$P_s(u,v)$  is the power spectrum of the signal, while  $P_n(u,v)$  is the power spectrum of the noise. The two power spectra are approximated, within a small local neighbourhood, by their local variances  $\sigma_s^2$  and  $\sigma_n^2$ . The filtered signal through an adaptive Wiener becomes:

$$f(x, y) = m_s(x, y) + \frac{\sigma_s^2(x, y)}{\sigma_s^2(x, y) + \sigma_n^2(x, y)} (h(x, y) - m_s(x, y)) \quad (9)$$

$h(x,y)$  is the original degraded signal and  $m_s(x,y)$  is the signal mean within a local neighbourhood of  $(x,y)$ . The mean of the signal is approximated by that of the original signal  $m_h(x,y)$ . Radiographic noise is additive, therefore

$$\sigma_s^2(x, y) = \begin{cases} \sigma_h^2(x, y) - \sigma_n^2(x, y) & \text{if } \sigma_h^2(x, y) > \sigma_n^2(x, y) \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

$$\sigma_h^2(x, y) = \frac{1}{M} \sum_{(x,y) \in \delta(x,y)} (h(x, y) - m_h(x, y))^2 \quad (11)$$

where  $M$  is the number of pixels in the neighbourhood  $\delta$  of  $(x,y)$ .  $\sigma_n^2$  is estimated from the computation of quantum mottle (caused by the spatial fluctuations in the number of X-ray photons absorbed per unit area of the intensifying screen and the variation in light photon emission per absorption event) and film granularity (fluctuations in the number of silver halide grains per unit area of the film emulsion). The computation of both these variances has been proposed in literature using characteristics of the film, intensifying screen and light photon energy.


The robustness of the physical-based approach is improved and the filtering step achieved. It should be noted that this method uses both grey-level images and  $h_{int}$  representations. However, the basic Wiener filter may alter the results in a way in which both small spots of noise and very fine microcalcifications may be overlooked during the filtering process. The method has been tested on both isolated calcifications and clusters with impressive results (see Figure 76). Yam *et al.* algorithm offers state-of-the-art results (93% TP fraction at 0.16

FP/image) with a novel physical interpretation of mammograms. The results of their algorithm are in the range of those obtained by commercial detection toolboxes. An implementation of the algorithm on SMF images only is tested and the results are presented in Section 3.2.3.

## 2.3 Standard Mammogram Form

The concept of a Standard Mammogram Form (SMF) has emerged from the core problems encountered in mammography. A normalisation of mammograms is necessary since the image brightness combines image-specific and anatomical information, making their appearance dependant on the imaging process. Furthermore, the breast is compressed during the X-ray shot and even if the density is constant, the deformation induces changes in area measures.

The  $h_{int}$  representation is a physics-based approach to mammographic analysis, an image normalisation method based on a complete understanding of the imaging process. Since the quality of mammograms is so highly dependent on the imaging conditions, the  $h_{int}$  model is an alternative *quantitative* representation of the breast tissue. Figure 23 shows a depiction of the  $h_{int}$  surface of a breast.



**Figure 23:** The  $h_{int}$  surface; (a) a mammogram presenting a lump; (b) the SMF that is generated from the mammogram where the ducts become ridges, and the mass a mountainous area.

The intensity or attenuation value of a pixel in a mammogram is determined by the amount of the X-ray photons absorbed by the tissue present between the X-ray source and the respective pixel. Table 2 shows the linear coefficients for various tissue types. The  $h_{int}$  value of a pixel represents the thickness of the breast tissue of interest (in mm) underlying between the X-ray source and the actual pixel. By interesting tissue one must consider the non-fatty tissue present in the breast, such as glandular, cancerous and fibrous tissue, which have high attenuation. Hence, the  $h_{int}$  representation (or Standard Mammogram Form when visualised as an image) is not dependent on the imaging procedure the same way the intensity value is. Other types of tissue present within the breast are the fatty tissue with low attenuation and the calcifications with very high attenuation, since they contain a concentrated level of calcium. They lead to the definition of  $h_{fat}$  and  $h_{calc}$ , the thickness of the fatty tissue and calcification. Calcifications are very small anatomical features of the breast and are considered an exception in the generation of the  $h_{int}$  representation; therefore their height will be considerably larger. The total thickness of the compressed breast can be computed with:

$$H = h_{int} + h_{fat} \quad (12)$$

**Table 2:** The linear coefficients for various tissue types reported by Highnam and Brady [65] after Johns and Yaffe. The coefficients of fibrous tissue and tumour overlap, while that of fat is clearly distinctive. Microcalcifications also have different attenuation coefficients, much higher than that of fibrous tissue.

Tissue type	No. patients		$\mu$ ( $\text{cm}^{-1}$ ) at energy (keV)		
			18	20	25
Fat	7	Minimum	0.538	0.441	0.317
		Mean	0.558	0.456	0.322
		Maximum	0.585	0.476	0.333
Fibrous (Glandular) (Parenchymal)	8	Minimum	1.014	0.791	0.499
		Mean	1.028	0.802	0.506
		Maximum	1.045	0.816	0.516
Infiltrating duct carcinoma	6	Minimum	1.061	0.826	0.519
		Mean	1.085	0.844	0.529
		Maximum	1.137	0.884	0.552

Although formula ( 12) has a very simple appearance, the computation of  $h_{int}$  is a rather complex process. It is based on an analysis of the mammographic imaging process from X-ray generation to film exposure.

It starts in the X-ray tube, more precisely in the anode, where accelerated electrons are converted into X-ray photons (with an energy spectrum that is quite complex, but which is then filtered to cut out photons of high energy that could be harmful to tissue. The result is a spectrum of photons that have energies between about 17 and 32 keV). The filtered X-ray photons then form the beam that passes through the breast. The variation of the tube voltage in time has a peak called kVp. Depending on the size and density of the breast, different times of exposure ( $t$ ) are chosen to assure a good exposure of the whole area of the breast. From the tube current ( $I$ ) and the photon output ( $f$ ), we can compute the photon flux at position  $(x,y)$ , as:

$$\phi(V_t, x, y) = f(V_t) \times I_t(V_t), \text{ where } f(V_t) \text{ is the photon output when kVp} = V_t \quad (13)$$

We can further compute the incident radiation at a certain kVp over a small area  $A$ , where  $N_{rel}$  is the relative number of photons of energy  $\varepsilon$  and  $\varepsilon_{max} = V_t$ .

$$E(x, y) = \phi(V_t, x, y) A t \int_0^{\varepsilon_{max}} N_{rel}(V_t, \varepsilon) \varepsilon d\varepsilon \quad (14)$$

The higher the kVp, the higher the average energy of the emitted photon, the lower the patient dose, but also the lower the image contrast. To reduce the amount of extra-focal radiation (photons that reach the breast from other directions than the focal-spot) a collimator is used (see Figure 9).

The X-ray photons that pass through the compressed breast and the compression plates and are attenuated. Those coming from the focal-spot and passing undeflected form the primary radiation, while those re-emitted in different directions are part of the scattered radiation. The majority of scattered radiation can be removed using an anti-scatter grid, but the pay-off is the increase of radiation dose to compensate for the loss of primary radiation. According to Beer's law, the number of exiting photons ( $p_{out}$ ) after travelling through a material of thickness  $h$  and attenuation coefficient  $\mu(\varepsilon)$  is proportional to the number of incident photons ( $p_{in}$ ) as below:

$$p_{out} = p_{in} \times e^{-h\mu(\varepsilon)} \quad (15)$$

From Equation ( 14) and knowing the thickness and attenuation coefficients of the lucite plates and breast, we can compute the primary energy exiting the breast:

$$E_p^{exit}(x, y) = \phi(V_t, x, y) At \int_0^{\epsilon_{max}} N_{rel}(V_t, \epsilon) \epsilon e^{-\mu_{luc}(\epsilon)h_{plate}} e^{-\mu(\epsilon)h} d\epsilon \quad (16)$$

Once the radiation has passed through the breast, it reaches the film-screen cassette, where the intensifying screen and the X-ray film are placed. Curiously (for that is the way current scintillators work), the photons pass through the film, then the screen absorbs the X-ray photons and emits light photons back in the direction of the incoming X-ray photon beam that expose the film. The intensifying screen reduces the radiation dose due to its amplifying property, but also induces a blur on the film, which is called glare. The primary energy imparted to the intensifier screen becomes a function of the screen absorption value  $S(\epsilon)$  and the anti-scatter grid transmission ratio  $G(\epsilon)$ :

$$E_p^{exit}(x, y) = \phi(V_t, x, y) At \int_0^{\epsilon_{max}} N_{rel}(V_t, \epsilon) \epsilon S(\epsilon) G(\epsilon) e^{-\mu_{luc}(\epsilon)h_{plate}} e^{-\mu(\epsilon)h} d\epsilon \quad (17)$$

After reviewing briefly the mammographic image formation, we will focus on the different stages that lead to the generation of SMF images.


Here are the steps to be followed to generate SMF, based on [65]:

- Convert the pixel value  $P(x,y)$  into film density  $D(x,y)$ ; this is achieved by considering the linear relationship between  $P(x,y)$  and  $D(x,y)$  as in ( 18), where  $m$  and  $c$  are constants obtained from the digitiser calibration data.

$$P(x, y) = mD(x, y) + c \quad (18)$$

- Convert the film density  $D(x,y)$  into energy imparted to the intensifying screen  $E_{pse}^{imp}(x,y)$ ; the film-screen characteristic curve is relevant for this purpose, which is found by exposing a lucite step wedge phantom. The energy imparted image appears as an inverted version of the density image, where dark parts correspond to high energies;

- Compensate  $E_{pse}^{imp}(x,y)$  for the intensifying screen glare by using the point-spread function of the intensifying screen. The energy imparted is deconvolved using a weighting mask  $w(x,y)$  constituting the number of glare photons. They are emitted from various positions and depths of the intensifying screen and impress the film at position  $(x,y)$ . The algorithm estimates the thickness of one of the  $n$  layers of the  $t_p$  thick screen as being  $dt_p = t_p/n$ . The glare reaching the film at  $(x,y)$  from  $(x_c, y_c)$  on the layer at depth  $z$  on the screen is computed according to Figure 24, where  $\mu_{light}$  is the linear attenuation value of light photons and  $z/\cos\theta$  is the path they have to travel.



**Figure 24:** The glare process.

The energy imparted at  $(x_c, y_c)$  is the difference between the attenuation at depth  $z$  and the layer below, where  $E^{in}$  is the energy reaching the intensifying screen:

$$E_z^{imp}(x_c, y_c) = E^{in}(x_c, y_c) e^{-\mu_{phosphor}^{xray} z} - E^{in}(x_c, y_c) e^{-\mu_{phosphor}^{xray} (z+dt_p)} \quad (19)$$

A symmetry assumption is made in order to compute the weighting masks at each depth  $z$  and they are scaled to sum to 1, while the glare is computed as:

$$glare_{(x,y)}^{(x_c, y_c)z} = \theta e^{-\mu_{light} \frac{z}{\cos\theta}} E_z^{imp}(x_c, y_c) \quad (20)$$

- Compensate  $E_{pse}^{imp}(x,y)$  for the anode-heel effect and diverging X-ray beam, by taking into account the variation between the incident photon flux between different spatial locations on the film. A blank film is first exposed without any breast present so that no scatter is present and the primary energy imparted is simplified:

$$E_{blank}^{imp}(x, y) = \phi(V_t, x, y)At \quad (21)$$

The pixel on the film with greatest exposure is  $(x_a, y_a)$  and is positioned under the anode:

$$E_{blank}^{imp}(x_a, y_a) = \phi(V_t, x_a, y_a)At \quad (22)$$


Therefore:

$$E_{corrected}^{imp}(x, y) = \frac{\phi(V_t, x_a, y_a)}{\phi(V_t, x, y)} E^{imp}(x, y), \text{ where:} \quad (23)$$

$$\frac{\phi(V_t, x_a, y_a)}{\phi(V_t, x, y)} = \frac{E_{blank}^{imp}(x_a, y_a)}{E_{blank}^{imp}(x, y)} \quad (24)$$

A limitation of the correction for the anode-heel effect is the lack of notion about the variation of the X-ray energy spectrum across the beam.

- Estimate the scattered radiation  $E_s^{imp}(x, y)$ , since this component of the imparted energy contains no information about the breast tissue, but influences the neighbourhood of the pixel. Once more, a weighting mask  $w(x, y)$ , which estimates the relative scatter reaching the central pixel  $(x_c, y_c)$  from pixel  $(x, y)$ , is computed. The initial estimate is made without the presence of an antiscatter grid. A lucite cylinder that approximates 50% fat and 50% dense tissue (of thickness  $H$ , inner radius  $r$  and outer radius  $r+dr$ ) is used for the estimation, as in Figure 25.



**Figure 25:** The cylinder used in scatter estimation.

To compute the scatter coming from the cylinder  $dE_s^{imp}(r)$ , we use the scatter-to-primary ratio for the cylinder ( $s/p(r)$ ), which becomes approximately constant after a certain value of the cylinder radius:

$$dE_s^{imp} = E_s^{imp}(r+dr) - E_s^{imp}(r) = E_p^{imp}(r) \left( \frac{s}{p}(r+dr) - \frac{s}{p}(r) \right) \quad (25)$$

If  $R$  is the stopping radius, the total scatter becomes:

$$E_s^{imp}(R) = E_p^{imp}(R) \frac{s}{p}(R), \text{ and} \quad (26)$$

$$\frac{dE_s^{imp}}{E_s^{imp}(R)} = \frac{\frac{s}{p}(r+dr) - \frac{s}{p}(r)}{\frac{s}{p}(R)} \quad (27)$$

is the proportion of total scatter for each cylinder. As in Figure 25, the cylinder is divided into  $n$  horizontal slices of thickness  $dh=H/n$  to estimate the scatter from different positions along the height of the cylinder,  $p(r,h)$ :

$$p(r,h) = \frac{dE_s^{imp}(r)}{E_s^{imp}(R)n} \quad (28)$$

In order to estimate the scatter when an anti-scatter grid is used, we must know the grid relative transmission ratio for s photon with incident angle  $\theta$ ,  $t(\theta)$ :

$$p_g(r,h) = p(r,h)t(\theta), \text{ where} \quad (29)$$

$$\theta = \tan^{-1} \left( \frac{r}{h} \right) \quad (30)$$

As in the estimation of glare, assuming azimuthal symmetry, a weighting mask  $w(x,y)$  can be computed to represent the relative scatter. This mask is convolved with the total energy imparted  $E_{pse}^{imp}$  for three example cases (just dense tissue, just fat tissue, half fat and half dense) to approximate the scatter function  $s$ . Due to the linearity of  $s$ , it can be combined with  $w(x,y)$  to obtain the scatter component  $E_s^{imp}$  of  $E_{pse}^{imp}$  by direct convolution.



The deconvolution of scatter radiation is one of the major sources of error in noise estimation in the SMF generation, as deconvolution is an intrinsically ill-conditioned problem. The original technique by Highnam and Brady uses the standard method that operates in Fourier domain. Recent unpublished work by Ancelin (by personal communication) shows that alternative deconvolution techniques can significantly improve results of estimating scatter.


- Estimate the extra-focal radiation  $E_e^{imp}(x,y)$  component, which is relevant at the curved breast edge where photons arriving with low angles can reach the intensifying screen. It is assumed that the extra-focal radiation is constant ( $E_c$ ) over the image if there is no breast present. The breast edge is represented semi-circularly, as in Figure 26. The extra-focal energy is computed along the inner projected breast edge  $\Omega$ , where there is only fat tissue. Hence, along  $\Omega$ :

$$E_e^{imp}(\Omega) = E_{pse}^{imp}(\Omega) - E_s^{imp}(\Omega) - E_p^{imp}(\Omega) \quad (31)$$

$E_c$  is approximated from the average  $E_e^{imp}$  along  $\Omega$  over the total attenuation of the extra-focal photons under the assumption of symmetry:

$$E_c = \frac{\overline{E_e^{imp}(\Omega)}}{2\pi \int_0^{\pi/2} e^{-\mu_{fat} h_{fat}(x,y,\theta)} t(\theta) d\theta}, \text{ where } 0 < \theta < \pi/2 \quad (32)$$

From the value of  $E_c$  we can now compute  $E_e^{imp}$  by multiplying it by the attenuation of the extra-focal photons along the fatty tissue inside the breast edge.



**Figure 26:** Modelling the breast for the estimation of extra-focal radiation component.

- Compute the primary radiation  $E_p^{imp}(x,y)$ ;

$$E_p^{imp}(x,y) = E_{pse}^{imp}(x,y) - E_e^{imp}(x,y) - E_s^{imp}(x,y) \quad (33)$$

- Convert  $E_p^{imp}(x,y)$  into  $h_{int}$  using the conversion equations developed by Highnam and Brady. From equation (16) and estimating the breast thickness  $H$ , we can write:


$$h\mu = h_{int}\mu_{int} + h_{fat}\mu_{fat} = h_{int}(\mu_{int} - \mu_{fat}) + H\mu_{fat} \quad (34)$$

The resulting  $h_{int}(x,y)$  is a float value representing the thickness of dense breast tissue at the pixel location.

Examining the  $h_{int}$  surface, we observe significant amounts of high-frequency noise, a clear impediment in image analysis and detection of abnormalities, notably microcalcifications. The SMF image is noisy, since the removal of intensifying screen glare [63] amplifies the image noise. In [174] an estimation of the radiographic noise is computed. A Wiener filter (c.f. Section 2.2.3) is applied to the original images before the  $h_{int}$  generation, which improves the signal-to-noise-ratio of the de-noised SMF image.

The  $h_{int}$  representation is a robust and reliable method resulting in a floating-point form that corresponds to the thickness of interesting tissue in the breast. By removing most of the unwanted effects of the imaging process, such as glare, scatter radiation, anode-heel effect and extra-focal radiation, the output of the method presents a much more adequate representation of the real anatomical structure of the breast. Hence, the  $h_{int}$  representation of the breast is mainly a 3D surface built from the  $h_{int}$  values of the image pixels. By removing the image parameters, the  $h_{int}$  images stand as normalised images of the breast.

A  $h_{int}$  representation can be easily visualised as an image in SMF, since the  $h_{int}$  values are in floating-point format, where brighter parts correspond to thicker parts of the breast or calcifications as in Figure 27. Importantly, the depicted surface of a  $h_{int}$  representation of the breast can show important anatomical features, such as masses as hills in a less dense background, while the background is mainly flat.



**Figure 27:** The Standard Mammogram Form image of a breast; (a) an MLO digitised intensity mammogram; (b) the correspondent SMF image.

## CHAPTER 3

### 3 Filtering *hint* Images

*'Contrariwise', ..., 'if it was so, it might be, and if it were so, it would be; but as it isn't, it ain't. That's logic!'*

Lewis Carroll - "Alice in Wonderland"

The field of Medical Vision is constantly concerned with developing novel approaches of turning a representation of parts of the human body into valuable information for clinicians in order to get a better view and understanding of the structural anatomy of the region of interest. Unfortunately, this process is far from simple as medical images, in general, have quite poor signal-to-noise ratios, and so they need to be enhanced in order to become useable. Of particular importance in what follows is the fact that the SNR of medical images tends to be an order of magnitude poorer than for regular visible images, for which most image enhancement schemes have (naturally) been developed. Few image enhancement schemes have been developed specifically for medical images, nor have they been adequately tested on such images. It has repeatedly been demonstrated that filtering methods can substantially improve the quality of the image of interest by means of eliminating artefacts and reducing the weight of unwanted information in the original image, but also simplifying the appearance of otherwise complicated anatomical structures.

The notion of image filtering has already been discussed in the previous Chapter as an image pre-processing step used by several calcifications and mass detection algorithms. Conventional filtering methods include a background smoothing stage (e.g. convolution with a low pass filter) followed by enhancement of the structures of interest (e.g. high pass filtering) and the subtractions of the two newly obtained images. A single such method is not capable of dealing with the large variability of the anatomical features that must be considered in practice. There are numerous algorithms based on evolving partial differential equations (PDE) for noise removal and image enhancement, but, as noted above, few of them have been tested thoroughly on medical imaging. Ultrasound images are largely acknowledged for their very noisy appearance and there have been many attempts to develop noise removal filters in medical ultrasound imaging [117]. Several other algorithms deal with the application of diffusion tensors in MRI [140, 149, 173]. Since the detection of microcalcifications is influenced significantly by the presence of noise of similar shape and magnitude, a PDE filter for noise removal in X-ray mammography is developed in this thesis. The next section introduces the basic theory behind anisotropic diffusion, a particularly important example of such a PDE filter.

### 3.1 Anisotropic Diffusion

Anisotropic diffusion has its origins in the classical nonlinear diffusion filter developed by Perona and Malik in 1987 [133], which is based on a PDE in divergence form. It is the cornerstone for new developments in multi-scale image analysis aiming to simplify the image appearance while enhancing structures of interest, such as edges or coherent structures. Its name is derived from the classical diffusion (or heat) equation.

Anisotropic diffusion is the solution we adopt for its inherent properties of smoothing and edge enhancement, but other methods are reported in literature with good results in image regularisation. Amongst them, we mention morphological methods [125], which have a

geometrical interpretation of images and are based on finding specific spatial characteristics in an image. Grey-scale morphology decomposes the image into its level sets [172]. The SUSAN noise removal and feature detector [155] is a nonlinear approach to feature detection that uses a circular mask for local measurements (excluding the central pixel). Those parts of the local image that are similar to each central pixel are used to compute the pixel's value. Wavelet packages [107] are also used to filter out high frequency noise through subspace decomposition of a characteristic function. Wavelets enable filters to be constructed with well defined spatial and frequency attributes obtaining a set of characteristics that can be used to detect specific image features.

A rather different approach to image smoothing is based on Markov Random Fields (MRF) [47], which allow stochastic modelling of images ending with a maximum a posteriori solution. The value of each pixel is probabilistically conditioned by the pixel values in its neighbourhood. Particular models inevitably impose specific constraints on the neighbourhood suited to the application. Any model requires that its parameters be estimated, therefore a parameter evaluation stage is also necessary for MRF, as for anisotropic diffusion. MRF usually having a large number of parameters that must be estimated from a training set, the computational time increases accordingly. Crucially, for example, Karssemeijer's [80, 81] original image filtering technique was based on Markov Random Fields; several papers have drawn attention to the many parameters it necessitated, its lack of robustness, and its poor convergence properties.


When the diffusivity function is a constant, the approaches and goals of both anisotropic diffusion and MRF are very similar. The energy is shown differently. The local energy function in anisotropic diffusion is infinitely differentiable and very smooth. Hence images diffuse slowly and anisotropic diffusion works well for removal of high frequency noise. MRF have the local energy expressed as a step function that is not continuous. That makes this technique appropriate for image segmentation. In general, the basic difference between anisotropic diffusion and the MRF is that in the first the local linearity is explicit in the structure of the filter, but in MRF the local linearity is determined very locally through propagation of

probabilities across cliques. Depending on the definition of its cliques, the MRF can be isotropic or anisotropic.

### 3.1.1 The Diffusion Process

One of the most commonly used methods for smoothing an image  $f: R^2 \rightarrow R$  is by convolving it with a Gaussian with standard deviation  $\sigma$  (35). The effect of the Gaussian kernel will be to blur the central point (considered to be the origin) into the neighbourhood, as shown in Figure 28.

$$K_{\sigma}(x) = \frac{1}{2\pi\sigma^2} \cdot \exp\left(-\frac{|x|^2}{2\sigma^2}\right) \quad (35)$$



**Figure 28:** The plot of the 2D Gaussian, where the central point (the top of the hill-like shape) will be gradually smoothed into the background.

The image  $f$  is transformed into a family of gradually smoother versions over an often large repeated convolution with a Gaussian (later referred as number of iterations  $t > 0$ ). An increasing scale will simplify the appearance of the original image. There are several limitations to this method, as observed in [171, 172]:

- although convolution with a Gaussian reduces noise, it also blurs important anatomical structures in the image, such as edges;
- linear smoothing dislocates edges when changing from finer to coarser scales.

The nonlinear diffusion process was proposed as an alternative to smoothing images by a Gaussian kernel, which does not preserve edges. Since it derives from a process of equilibrating concentration differences, it can be expressed through a continuity equation of Fick's law [171]:

$$J = -D \cdot \nabla u \quad (36)$$

$$\delta_t u = \text{div}(D \cdot \nabla u) \quad (37)$$

$D$  is called the *diffusion tensor*, a positive definite symmetric matrix that represents the relation between the concentration gradient ( $\nabla u$ ) and the flux ( $J$ ) that aims to compensate for this gradient (36). In image processing, the concept of concentration is replaced by that of grey level. The diffusion tensor may be replaced by a positive scalar-valued *diffusivity*  $g$ . If  $J$  and  $\nabla u$  are parallel, the diffusion is called *isotropic*. In the *anisotropic* case,  $J$  and  $\nabla u$  are not parallel. Equation (37) is called the *diffusion equation*. If the diffusion tensor is space-dependent, then the diffusion is called *inhomogeneous*, while a constant diffusion tensor is related to a *homogeneous* diffusion.

In order to overcome the scale correspondence problem (the coarse-to-fine tracking difficulties), the inhomogeneous linear diffusion filtering introduces  $|\nabla f|$  (the gradient of the original image) as edge detector to preserve different entities in the image. High values of the detector indicate the presence of edges in the image. The diffusivity function  $g$  was set to:

$$g(|\nabla f|^2) = \frac{1}{\sqrt{1 + |\nabla f|^2 / k^2}} \quad (k > 0) \quad (38)$$

and the diffusion equation reduces to:

$$\delta_t u = \text{div}\left(g(|\nabla f|^2) \cdot \nabla u\right) \quad (39)$$

Introducing feedback into the diffusion process, by adapting the diffusivity to the gradient of  $u(x,t)$  - the actual image - rather than the original  $f(x)$ , the diffusion equation becomes nonlinear and therefore the diffusion filtering becomes nonlinear and isotropic:



$$\delta_t u = \text{div}\left(g\left(|\nabla u|^2\right) \cdot \nabla u\right) \quad (40)$$

The Perona-Malik model [84, 132] represents the first nonlinear diffusion filter. It provides stable edges over a large number of iterations based on a rapidly decreasing diffusivity, but will only enhance those edges for which the gradient is larger than the contrast parameter  $k$ .

$$g\left(|\nabla u|^2\right) = \frac{1}{\sqrt{1 + |\nabla u|^2 / k^2}} \quad (k > 0) \quad (41)$$

Catté *et al.* [23] introduced the Gaussian convolution of  $u$ :  $u_\sigma = K_\sigma * u$  and the result of it was:

$$\delta_t u = \text{div}\left(g\left(|\nabla u_\sigma|^2\right) \cdot \nabla u\right) \quad (42)$$

This new form of the diffusion equation solved the spatial regularisation problem of the inhomogeneous filtering, meaning that the solution of the nonlinear filtering method of images aims to achieve a steady state. Moreover, a new parameter is introduced in the process, *the scale parameter*  $\sigma$ . As a result, the process is now controlled by three parameters,  $t$  (time),  $k$  (contrast) and  $\sigma$  (scale), which substantially reduce the impact of the choice of diffusivity over the whole process and make the use of it more flexible and robust. Although the contrast parameter works similarly to the Perona-Malik model, the scale parameter makes the filter less sensitive to small-size structures, such as noise, by increasing  $\sigma$ , the kernel of the Gaussian.

### 3.1.2 Nonlinear Anisotropic Diffusion

The main improvement introduced by nonlinear anisotropic filters is smoothing along the isophote and, when the value of the gradient is large, not across it [54]. While for low gradients smoothing is performed in the usual way, diffusion is inhibited at edges. Weickert [171, 172] introduces a system of eigenvectors  $v_1, v_2$  of the diffusion tensor  $D$ .  $v_1$  and  $v_2$  are orthonormal (see Figure 29) and

$$v1 \parallel \nabla u_\sigma \quad (43)$$

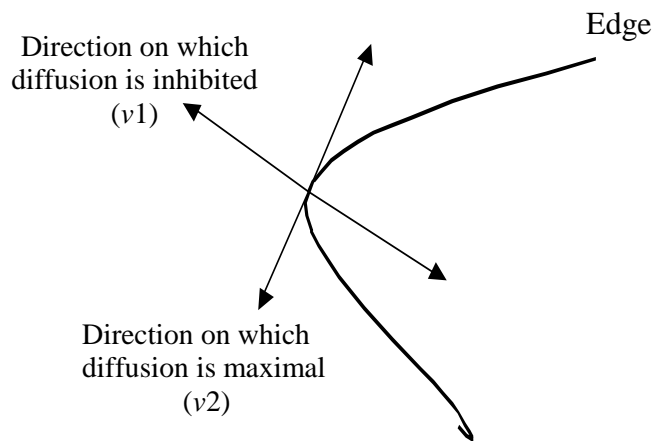
$$v2 \perp \nabla u_\sigma \quad (44)$$

The corresponding eigenvalues are:

$$\lambda_1 = g(|\nabla u_\sigma|^2) \quad (45)$$

$$\lambda_2 = 1 \quad (46)$$

In general  $\nabla u_\sigma$  is not parallel to one of the eigenvectors of  $D$  for  $\sigma > 0$  and Weickert's model behaves highly anisotropically. As  $\sigma$  tends to 0, the process tends to behave like the original Perona-Malik model.



**Figure 29:** The diffusion tensor eigenvectors;  $v1$  is parallel with the edge gradient and the smoothing is inhibited across the edge;  $v2$  is orthonormal to  $v1$  and the diffusion is permitted along the edge.

The edge-enhancing diffusion model proposed by Weickert and described above is the one that is used in our initial experiments; it gave the best results, although we found in practice that the value of the constant  $-3.31488$  (47) is not crucial. The diffusion across edges is performed according to the following eigenvalue:

$$\lambda_1 = \begin{cases} 1 & |\nabla u_\sigma| = 0 \\ 1 - \exp\left(\frac{-3.31488}{(|\nabla u_\sigma|/k)^8}\right) & |\nabla u_\sigma| > 0 \end{cases} \quad (47)$$


### 3.1.3 Discussion

Anisotropic diffusion overcomes some major limitations of linear and nonlinear isotropic filters (see Figure 30 for some comparative results):

- enhances noisy edges and flow-like structures (this might, for example, be useful in the detection of curvilinear structures, as proposed for future work in Chapter 6);
- inhibits diffusion at edges;
- is more flexible due to the larger number of parameters, but not so much so as to alter the robustness and accuracy of the method;

The use of anisotropic diffusion, as observed in [172], ranges from computer-aided quality control to post-processing fluctuating data, target tracking in infrared images and blind image restoration, to enumerate just a few of the applications. However, most applications [117, 140, 149, 172, 173] concerned with filtering medical images were mainly developed for ultrasound and magnetic resonance images.

However, it is not sufficient to use “blind” filtering methods when the features we need to preserve in an image are so precise and specific. The use of *a priori* or even *a posteriori* knowledge in the diffusion process (i.e. image characteristics) must be embedded in the use of PDE-based filters, a field that is evolving rapidly and which has to bring many advantages in the overall development of medical vision.



**Figure 30:** Some comparative diffusion results. (a) the original image; (b), (c), (d) the smoothed image with nonlinear isotropic diffusion (Perona-Malik) after 20, 40 and 100 iterations; (e), (f), (g) the smoothed image with nonlinear anisotropic diffusion (Weickert) after 20, 40 and 100 iterations; (h), (i), (j) the smoothed image with linear diffusion after 20, 40 and 100 iterations.

Anisotropic diffusion is one particular technique of nonlinear diffusion, from the larger family of methods referred to collectively as scale-space theory. It can be thought of in terms of the application of partial differential equations (PDE) to image analysis, which adds a-priori knowledge to the classical scale-space evolution. The regularisation by convolution with a Gaussian builds an edge detector that does not depend on noise smaller than the order of the Gaussian kernel and ensures the uniqueness of results. The filter class with diffusion tensor has a unique solution, dependant on the original image, which is infinitely differentiable for  $t > 0$ . [172]. This guarantees that the diffused image has similar properties to the input image. For  $t \rightarrow \infty$  the result converges to the most simplified version of the input, namely a constant image with the same average grey level as the original. This is an important issue in medical imaging, where grey level has physical meaning. But for finite, although large, values of  $t$ , the result shows enhanced contrast at edges. Weickert also notes that such filters respect an extremum principle and non-enhancement of local extrema.

If we choose a fast decreasing diffusivity, as in our application, the reasoning of employing anisotropic diffusion becomes obvious. As Weickert points out, since diffusion is much stronger at both sides of an edge than at the edge itself, the contrast of the edge becomes enhanced. The filter acts like a backward diffusion at edges, while smoothing between them. This is precisely why we chose anisotropic diffusion for our application. After a few iterations, the less important features will be filtered out of an image, while the salient ones will persist over time. Therefore, it is important to include the appropriate set of saliency descriptors into the diffusion model for good filtering results.

A practical problem inevitably encountered using anisotropic diffusion is its parametric nature. Although it has far fewer parameters than probabilistic methods, typified by Markov Random Fields, it is difficult to find a general solution for a wide range of images. The parameters must be tuned for robust results and this will limit their relevance in the broad-spectrum situation. Note however that this is true of every technique that has been developed for image analysis. Nonlinear anisotropic diffusion adds two further parameters to the number

of iterations ( $t$ ) found in standard linear diffusion, namely the contrast  $k$  and the scale  $\sigma$ , but as argued in this Chapter 4, there are rational reasons to inform the choice of parameter values for this particular application. It is sometimes argued that another major disadvantage of anisotropic diffusion is its computational complexity, as it is an iterative solution to a PDE. However, this is no longer a valid objection as research over the past ten years have shown how semi-implicit and parallel implementations can reduce the computational complexity to the point where anisotropic diffusion systems nowadays run in real time on moderately powerful workstations.

Some specific concern in using anisotropic diffusion in mammography is the spatial scale of diffusion compared with thin structures in the breast. When  $\sigma \rightarrow 0$ , the filter becomes isotropic ( $\nabla u$  becomes eigenvector of the diffusion tensor  $D$ ), hence extremely small structures cannot be preserved. The resolution of mammograms becomes relevant to the type of filter we need to build. While a  $50\mu\text{m}$  resolution seems sufficient for structures under a quarter of a millimetre in width,  $100\mu\text{m}$  in resolution would be insufficient. Thus, in the case of mass detection, where a resolution of  $100\mu\text{m}$  is common, some fine spicules may be overlooked. Usually, extremely thin lines are corrupted because of the high gradients on both sides of the line. Very small microcalcifications face the same problem, as well as corners, which become rounded in time.


### 3.2 Filter Model

A novel approach to filtering mammographic images for detecting microcalcifications will be introduced in this section. Consistent with the overall aims of the thesis, the method has been implemented and tested primarily on images in Standard Mammogram Form; but it is not confined to this image format. The theoretical foundation of the algorithm uses anisotropic diffusion. The aim of our approach is to “clean” the noisy  $h_{int}$  images while preserving structures of interest, specifically calcifications and the portions of noise of the signal that

could help in the detection of false positives. The diffusion process becomes a method of detecting both microcalcifications and noise in X-ray mammography,  $h_{int}$  representation. The discriminating factor between noise and calcifications is the appearance of the two types of structures. We will return to this issue later. The method is further developed in the following chapters of this thesis into a fully automated technique to detect microcalcifications.

### 3.2.1 Theory

One of the major characteristics that we have used in approaching microcalcification detection was the genuine difference that should be visible in the shape of microcalcification versus noise in mammographic imaging. While microcalcifications are anatomical structures with slightly blurred edges due to the scattering effect of X-ray beams passing through the breast [65], noise tends to have extremely sharp edges. In Figure 31, we show the plots of an intensity image and a  $h_{int}$  image showing the same characteristics for microcalcifications and noise. Noise does not represent an anatomical structure, therefore its shape is very well delimited and does not present the otherwise usual fluctuations in height or grey-level. Since the technique that leads to building  $h_{int}$  images from intensity images eliminates the side-effects produced by the X-ray imaging process [65], we will show in Chapter 4 how the elimination of undesired imaging artefacts helps the detection of microcalcifications on SMF images.



**Figure 31:** The shape differentiation between microcalcifications and noise; (a) the plot of a filtered (de-noised) intensity image sample containing a microcalcification; (b) the plot of a filtered intensity image sample containing noise; while the microcalcification has the appearance of a hill with less steep edges, the bit of noise is rather spiky and has a higher value of intensity; (c) the plot of a  $h_{int}$  image containing a microcalcification; (d) the plot of a  $h_{int}$  image containing noise. Each plot is taken from one line in an image.

Calcifications and noise differ significantly in their image characteristics or appearance. Shot-noise may drastically influence the local image characteristics and represents a main source of FP. The  $h_{int}$  representation can eradicate this type of noise, but since our method is designed to detect noisy structures as well, we did not remove shot noise in advance from the images on which we tested our algorithm. The appearance of  $h_{int}$  images would be extremely noisy mainly due to the removal of the glare effect, extra-focal and scattered radiation (which accounts for up to 40% of the total radiation exiting the breast [65]). This would lead to a difficult observation of the regions of interest, making it harder to distinguish small structures



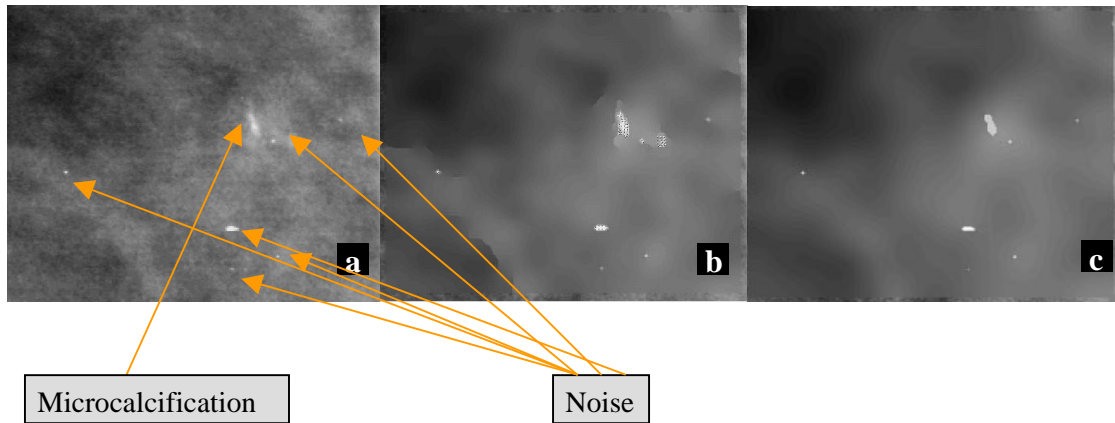
in mammograms. Since microcalcifications are tall in  $h_{int}$  images, only the most prominent spots of noise may lead towards FP, the smaller ones being easily removed by the diffusion process.

One should note that the term “contrast” in a  $h_{int}$  image does not correspond to the grey-level as in intensity images. It represents the height of interesting tissue in the region under observation (cf. Section 2.3). While calcifications are typically small and sparse structures, they appear in about 25% of mammograms. And even then they occupy a small fraction of the entire area of a mammogram. Hence, the percentage area of microcalcifications in a large set of typical mammograms is vanishingly small. For this reason, Highnam and Brady’s  $h_{int}$  generation algorithm [65] assumes only two types of tissue: fat and non-fat (i.e. parenchymal, tumour) and  $h_{calc}$  is omitted from the computation (12). Since the attenuation coefficient of calcium is typically 26 times higher than that of *interesting tissue*, microcalcifications are in effect an exception in the  $h_{int}$  representation. Therefore the  $h_{int}$  value of a region corresponding to calcifications does not represent the thickness of the corresponding area of the compressed breast. Those regions would appear much thicker than they really are and make the calcifications appear tall in the image. Yam *et al.*’s algorithm presented in Section 2.2.3 is based on exploiting this effect. The high values of calcifications in the  $h_{int}$  representation of the breast make them react differently to our filter than a background, an important assumption in our work.

### 3.2.2 Method

The filter model introduced here is a preliminary attempt towards the automatic detection of noise and calcifications in mammograms. It can operate on SMF images or grey-level pictures. The results we present are based on a set of samples of SMF images that show considerable variation in the size and visibility of calcifications. The filter implemented in this work is an anisotropic diffusion-based filter (cf. Section 3.1). It blurs the input mammographic image

while preserving significant intensity/thickness changes. The process relies on the use of a set of specific parameters, e.g. time, contrast, size, and it is critical to find the right choice of parameters that will lead to good repeatable results. Figure 32 shows different output images after using anisotropic diffusion on a grey-level digital mammogram containing both a calcification and noise.



**Figure 32:** (a) The original grey-level image containing a microcalcification in the centre-right of the image and a large spot of noise on the lower side of the image; (b) the diffused image with  $k=5$ ,  $\sigma=0.6$  and  $t=20$ , we notice that the edges of the important structures of the dense tissue are emphasised; (c) the diffused image with  $k=5$ ,  $\sigma=0.5$  and  $t=40$ , where only the important small structures are kept and their edges enhanced;

To choose the right diffusion tensor for our application, we initially tested some of the known tensors in literature. The diffusion tensors that we have tried for the anisotropic filtering were based on the following diffusivity-like functions:

$$g(|\nabla u|^2) = \frac{1}{1 + |\nabla u_\sigma|^2 / k^2} \quad (\text{Perona-Malik [132]}) \quad (48)$$

$$g(|\nabla u|^2) = \frac{1}{\sqrt{1 + |\nabla u_\sigma|^2 / k^2}} \quad (\text{Charbonier [171]}) \quad (49)$$

$$g(|\nabla u|^2) = 1 - \exp\left(\frac{-3.31488}{(|\nabla u_\sigma|/k)^8}\right) \quad (\text{Weickert [171]}) \quad (50)$$

$$g(|\nabla u|^2) = 1 - \exp\left(-\frac{|\nabla u_\sigma|^2}{2k^2}\right) \quad (51)$$

$$g(|\nabla u|^2) = 1 - \exp\left(\frac{-1}{(|\nabla u_\sigma|/k)^4}\right) \quad (52)$$

$$g(|\nabla u|^2) = 1 \quad (\text{linear diffusion}) \quad (53)$$

We found experimentally that Weickert's diffusion tensor is best suited to our application. We used a similar simplified tensor having the corresponding eigenvalues (54), (55) (since the shape of the diffusivity-like function can be easily altered by changing the values of its parameters):

$$\lambda_1 = \begin{cases} 1 & |\nabla u_\sigma| = 0 \\ 1 - \exp\left(\frac{-1}{(|\nabla u_\sigma|/k)^8}\right) & |\nabla u_\sigma| > 0 \end{cases} \quad (54)$$

$$\lambda_2 = 1 \quad (55)$$

Nonlinear anisotropic filtering proves to be highly flexible due to the variability of its parameters which help in covering a rather extensive set of possibilities in multi-scaling filtering with respect to the output one can get by filtering medical images, as Table 3 shows. Using the parameters of the process in multi-scale filtering lead us to make the following remarks:

- by increasing  $k$ , the contrast factor, one would increase the overall blurring of an image, would extinguish the anatomical structures that have smaller contrast to their

surroundings (gradient value) than  $k$  and would lose the edges that do not have very high contrast;

- by increasing  $\sigma$ , the scale factor, one would increase the overall blurring of an image, would extinguish the anatomical structures that are smaller than the kernel of the Gaussian (approximately  $10\sigma$ );
- by increasing  $t$ , the number of iterations, one would increase the overall blurring of an image, would strongly extinguish the anatomical structures, which become more and more diffused with an increasing  $t$ , but would preserve edges over a rather long number of iterations.


The flexibility of anisotropic diffusion occurs from the process parameters. More parameters may also mean less robust and the choice of the specific set of parameters becomes crucial. In general, we can only do this empirically.

The appearance of  $h_{int}$  images would still be extremely noisy, mainly due to the removal of the glare effect, extra-focal and scattered radiation [64, 65]. If glare is removed, facilitating the removal of shot-noise [66], the price to be paid is a massive decrease in the SNR in SMF images. Yam [178] attempts to overcome this substantial increase in noise by Wiener filtering the original images before generating the  $h_{int}$  surfaces, an approach that improves the SNR slightly. The  $h_{int}$  representation can eradicate shot-noise [65], an important source of FP in detecting microcalcifications. We prefer to work with glare de-convolved (no shot-noise removed) images and use anisotropic diffusion to differentiate edge sharpness of noise and microcalcifications. Figure 33 shows a phantom study on how the generation of SMF images and the glare de-convolution influence the appearance of a mammogram. The phantom was generated by adding noise to an image containing bright white blobs on a dark background (of values similar to the intensities of microcalcifications and fat in a mammogram); the noise was obtained by subtracting a filtered mammogram from the original one (high frequency structures in a mammogram).


**Table 3:** Variation of anisotropic diffusion parameters:  $k$  - the contrast factor,  $\sigma$  - the scaling factor and  $t$  - the number of iterations;  $\nearrow$  represents an increase in the associated feature, as opposite to decrease for  $\searrow$ .

	Blur	Anatomical features	Edges
$k \nearrow$	$\nearrow$	$\searrow$	$\searrow\searrow$
$\sigma \nearrow$	$\nearrow$	$\searrow$	$\searrow$
$t \nearrow$	$\nearrow$	$\searrow\searrow$	Well preserved over a long time

By varying the values of the three parameters of the process (cf. Table 3) different output images of the same input  $h_{int}$  would be obtained. In our application, small bright structures are salient; therefore an appropriate combination of contrast and scale is desirable. Having obtained the diffused image, we subtract it from the original. Some differences in the way microcalcifications, as opposed to noise, are diffused can be noticed in Figure 34. Microcalcifications tend to be smoothed faster than prominent noise spots, for an appropriate choice of parameters. After a certain number of iterations, the surface of the difference image contains significant changes for noise only. The subtraction of pre- and post-processed images just illustrates the difference in which microcalcifications and noise are smoothed, as an exemplification of the use of the principles of anisotropic diffusion for our specific purpose. The algorithm employs no such subtraction.



**Figure 33:** The changes in SNR during  $h_{int}$  generation: (a) the original phantom with simulated microcalcifications and noise; (b) the Wiener-filtered phantom; (c) the  $h_{int}$  image before glare deconvolution; (d) the  $h_{int}$  image after glare deconvolution; (e) the 3D plot of the original image in (a); (f) the plot of the smoothed image in (b) with improved SNR; (g) the plot of the  $h_{int}$  image in (c); (h) the noisier plot of the  $h_{int}$  image in (d).



**Figure 34** Image subtraction; (a) The original preprocessed SMF image containing a microcalcification on the left side, a large spot of noise on the lower right side and several other smaller noise structures; (b) the 3D plot of the difference image between the original image diffused with  $k=15$ ,  $\sigma=0.6$  and  $t=5$  and the same one diffused with  $k=15$ ,  $\sigma=0.6$  and  $t=10$ ; (c) the original image diffused with  $k=15$ ,  $\sigma=0.6$  and  $t=10$  and the same one diffused with  $k=15$ ,  $\sigma=0.6$  and  $t=15$ . We notice that after a few iterations the big changes appear at the location of noise only.


### 3.2.3 Results

This section starts by showing some results from applying nonlinear anisotropic diffusion filtering to samples of real mammograms containing microcalcifications. We de-noise  $h_{int}$  images while preserving only calcifications and significant bits of noise, Figure 35.

In order to reduce processing time and, more importantly in practice, remove the need for intervention of the operator in the filtering process, we initially chose a large value for the contrast factor  $k$ . We still chose a rather small value for the scaling factor  $\sigma$  for preserving tiny anatomical structures or noise over the first iterations in the process of diffusion. Due to the strong variability that exists in mammographic images (e.g. contrast, size of interesting tissue) a multi-scale approach is performed. Since the whole process should be robust and easy to use, we reduced the number of variable parameters to one, keeping constant the contrast and scale factors and varying only the number of iterations over a very small range. We found that the


time factor  $t$  gives optimal results for the filtering process over the whole set of SMF images when we used values between 3 and 7 iterations for that pair of values  $k$  and  $\sigma$ .

In demonstrating the efficiency of our method in increasing the number of true positives, we also considered images with very high likelihood of giving false positives. Such an example is presented in Figure 36.



**Figure 35:** Filtering example 1; (a) The original preprocessed SMF image containing a microcalcification on the left side and a large spot of noise on the lower right side and several other noise structures; (b) the diffused SMF image with  $k=15$ ,  $\sigma=0.6$  and  $t=5$ , we notice that the microcalcification has almost faded, while the noise is still preserved with high contrast; (c) the noisy 3D plot of the original SMF image in (a); (d) the surface of the diffused SMF image in (b), the microcalcification appears as a hill with smoother edges than those of the very sharp-edged noise structures in the same image.





**Figure 36** Filtering example 2; (a) The original preprocessed SMF image containing only noise structures, the largest piece of noise on the upper right side could be easily considered of being a microcalcification since it does not present very high contrast from the surrounding tissue; (b) the diffused SMF image with  $k=15$ ,  $\sigma=0.6$  and  $t=3$ ; (c) the 3D plot of the original SMF image in (a) with highly noisy appearance; (d) the 3D plot of the diffused SMF image in (b) where all structures have very sharp edges and are labelled as noise.

The detection method, both of calcifications and noise, was based initially on the association one can make between the original  $h_{int}$  mammograms containing the structures of interest and the surface we built from the filtered images after just a few iteration steps. Since radiologists may have doubts when searching the original image for microcalcifications, the surface we present would show either hill-shaped structures for microcalcifications or sharp-edged formations for noise in the locations corresponding to the structures of interest. Moreover, we found the simple visual comparison of the two  $h_{int}$  images - the original noisy one and the filtered one - to be quite reliable in differentiating between microcalcifications and

noise. While noisy structures tend to be better preserved by the filtering method applied with our specific choice of parameters, microcalcifications fade faster and look like imploding structures.


Our set of data images was obtained at the Breast Care Unit of the Churchill Hospital Oxford. Images were collected from the screening database and correspond to women aged between 50 and 64. There are a total of 102 images, 24 normal, while the rest contain microcalcifications. An experienced radiologist annotated a total of 98 microcalcification clusters, which have been previously proven by biopsy. The ground truth for the validation of our detection algorithm was the contour drawn by the radiologist on the film around each detected cluster of microcalcifications. The cluster positions were subsequently translated into x-y coordinates on the digital images. The films were subsequently digitised using the Lumisys scanner at the same location into 12-bit *.mit* images at a resolution of 50 $\mu$ m.

For the initial results on isolated calcifications that we present in this Chapter, we used a set of 33 samples from the same database. A clinician annotated each isolated calcification (10 macrocalcifications and 27 microcalcifications) in the dataset. We used the coordinates of calcifications as ground truth in the detection method.

### 3.2.3.1 Coarse Calcifications

The algorithm was tested initially on a set of 13 samples of average  $h_{int}$  mammograms containing 10 pre-labelled isolated coarse calcifications and several artefacts. The image samples were digitised at 50 $\mu$ m resolution. Samples were preferred, rather than whole mammograms, in order to reduce processing time, since, as noted earlier, the space occupied by microcalcifications tends to be small and mammograms digitised to 50 $\mu$ m are typically 4000 by 4000 pixels (which, at 12 bit resolution generates 32 Mbytes per image). Eight each samples contained one coarse calcification, one contained two calcifications, and four samples contained only noisy structures. The enhanced images were filtered by means of anisotropic diffusion and a surface representation was built for each filtered image. The values of the

diffusion parameters were  $k=15$ ,  $\sigma=0.6$ , while  $t$  varied between 3 and 7 iterations. The algorithm applied to the enhanced images gave a detection rate of 100%. It is hardly worth showing the free-response receiver operating characteristic (FROC) curve, but it is given in Figure 37.




**Figure 37:** The FROC curve of the detection method for the set of 13 samples with coarse calcifications.

### 3.2.3.2 Microcalcifications

While coarse calcifications are generally easy to detect and are typically benign, the real difficulty arises in the detection of microcalcifications, which are much smaller and variable in brightness. The algorithm was tested further on a set of 20 samples of  $h_{int}$  mammograms containing 27 pre-labelled isolated microcalcifications and various regions of noise at 50  $\mu\text{m}$ . The set was meant to offer an overview of possible clinical aspects related to microcalcification of different sizes, some of them clear while some others are faint, making sure that we will observe both the detection of microcalcifications and noise during the process of filtering. Sixteen of them contained one microcalcification each, two of them two microcalcifications, while there were two samples containing three and respectively four microcalcifications 25 of the microcalcifications were detected correctly, while 2 of them were labelled as noise. Also 2 FP occurred. The TP fraction was 92.6% for a number of 0.1 FP per image. In the initial

experiments, we tested the filtering method on isolated microcalcifications or small clusters to evaluate the method for every salt of calcium, rather than the general detection of a cluster. The detection of microcalcification clusters will be revisited in Chapter 4.

We further applied an implementation of Yam *et al.*'s algorithm (cf. Section 2.2.3) to the same set of microcalcifications. The process differed slightly in this case. The original  $h_{int}$  mammograms were not enhanced, in order to preserve a fixed scale for all mammogram samples. The diffusion parameters were:  $k=15$ ,  $\sigma=0.6$  while  $t$  varied between 2 and 5. The algorithm was applied to the filtered versions of the original SMF images. We obtained a 100% TP fraction with 0.3 FP per image. The FROC curve of the detection using the combination of the anisotropic diffusion filter and the algorithm implemented by Yam *et al.* is shown in Figure 38.




**Figure 38:** The FROC curve of the combined detection method for the set of 20 samples containing different types of microcalcifications.

### 3.2.4 Discussion

An important issue in the use of this new filtering method in X-ray mammography is the good preservation of tiny anatomical structures over the diffusion process. Unlike most filters that actually blur the whole image and blend small regions together, our method preserves the anatomical independence of many small structures encountered in an image. Figure 39 shows

the output of diffusing an image containing a cluster-like structure of tiny particles at a resolution of  $50 \mu\text{m}$ . A Wiener filter blurs the structure and its output would look like a single microcalcification, blurring together the individual tiny particles and possibly inducing some FP. In the case of anisotropic diffusion, the tiny bits of calcium stay independent of the rest, an essential feature in clustering techniques. The high variation between the diffusion parameters used in obtaining the two filtered versions proves the consistency of the filter.



**Figure 39:** Filtering example 3; the left column presents the original SMF image (a) and its two diffused versions for the sets of parameters  $k=5$ ,  $\sigma=0.5$ ,  $t=40$  (c) and  $k=15$ ,  $\sigma=0.6$ ,  $t=2$  (e); the right column shows the 3D surfaces of the three respective images.

A major source of FP in mammography corresponds to shot noise. The noise maps obtained after removing the glare-effect in the process of generating SMF images can be used as a further step to exclude this specific type of noise from mammograms and therefore reduce the number of FP. As Yam *et al.*'s algorithm is built to use a combination of a grey-level and an SMF image, using its original implementation on SMF only is expected to give poorer results. Detecting the small area changes over the height of shot noise (in a similar way to the detection of microcalcifications), would eliminate the imaging artefact and could generate a shot noise map similar to the one produced during the SMF generation.

Computational requirements are important in the development of real-time clinical applications and filtering algorithms are usually time-consuming because of the subsequent application of kernels over one image. In order to reduce the necessary time for the diffusion process, we used a higher value for the contrast factor  $k$ . A higher  $k$  leads to faster diffusion over the image and fewer iterations are requested. The consistence of our choice is based on the high  $h_{int}$  values corresponding to both shot noise and calcifications. Both structures preserve their characteristics for high contrasts over a few numbers of iterations.

### 3.3 Conclusion

We presented a filtering method based on anisotropic diffusion, a process known for its scale-space and edge detection properties. Our filter implements such nonlinear diffusion filtering for what seems to be the first time in digital mammography and aims to be an alternative to classic filters previously used in working with breast images and microcalcifications.

Our method uses the normalised representation of mammograms that the  $h_{int}$  generation provides only, a robust and consistent approach to digital mammography. The initial results are encouraging and further improvements to the method promise better rates of detection. The algorithm is also reliable in detecting both calcifications and noise in one go by taking into

account the “physical” appearance of different structures of interest. While the term noise refers only to shot noise, as a major source of FP, the term calcifications would include coarse calcifications as well as microcalcifications. Quantum mottle, an important source of errors in mammography, has little interference in our application as it is smoothed by our filter, with a good choice of the contrast and scaling factors. Furthermore, anisotropic diffusion blurs images making use of the edge enhancement property.

Having obtained the filtered mammograms (or as exemplified here, SMF images), we aim to develop the method towards a fully automatic real-time algorithm. Here are the steps to follow in extending the method to a full detection algorithm:

- a robust integration of the shot noise detection and removal, either further improving this filtering algorithm into a shot-noise removal or integrating Highnam and Brady’s artefact removal [65];
- the development of a non-parametric version of the present filtering method [101, 102, 103] - which enhances microcalcifications and smoothes the background into a more homogenous area – to improve its robustness;
- the development of a reliable method to depict microcalcifications from the set of possible candidates (as Yam *et al.*’s [179] implementation was used in the work presented in this Chapter).

In the following chapters we will show the expansion of the filter we have presented here into a novel non-parametric method to detect microcalcifications in X-ray mammography.

## CHAPTER 4

# 4 Adapting Characteristics of the Human Visual System to Digital Mammography

*Unlike the wine, science shouldn't  
be let to age*

Grigore Moisil

In this Chapter, we present the results of developing a method to identify microcalcifications in mammograms using a model of the human vision. It is critical that a programme that is designed to assist a radiologist detect microcalcification clusters misses few – if any – clinically important clusters, equally that it does not signal too many false positives (FP). However, no method is perfect and though some have reportedly reducing the numbers of missed calcification clusters by as much as 20% - they continue to return too many FP. Previous methods can be classified as primarily statistical [80, 81] or structural [179].

Our method relies upon using a quantitative representation of breast tissue, for example the Standard Mammogram Form (SMF) developed by Highnam and Brady [65]. The results presented below have been obtained using the SMF quantitative representation of breast tissue; however, the method *does not* depend upon the specific characteristics of the  $h_{int}$  and SMF representations. Our method will work interchangeably on any such representation, though, as we will demonstrate, there are significant advantages to applying it to the SMF representation.

The novelty of the technique derives from (a) the way in which the method analyses the statistical characteristics of the mammogram and (b) the particular combination of filters that



are applied in sequence (statistical analysis, image enhancement, adaptive segmentation). In particular, the algorithm adapts to digital mammography a model of how the human visual system functions for conventional visual images. We do this because the human eye proved, in combination with the mobilisation of specialist knowledge that radiologists use, to be generally more sensitive in depicting microcalcifications from a cluttered image than existing algorithms. For this reason, a novel measure for adaptively estimating image contrast is used, as explained further in this Chapter. The new measure adapts to the local brightness and adjusts a threshold. Critically, the approach does not depend on the careful setting of a number of parameters, despite the fact that it uses a number of methods (i.e. anisotropic diffusion), which depend crucially on the setting of numerical parameters. This is important for large-scale deployment in mammography, where it is simply unfeasible to re-set parameters for each mammogram. We do this by a novel method of adapting the filters to the characteristics of the particular image under consideration.


An overview of the algorithm is presented in Figure 40 with a clear separation of the stages we use to detect microcalcifications. While the top row in the diagram relates to the SMF generation (see Section 2.3) and Wiener filter (discussed in detail in Chapter 2), the bottom row underlines the steps that will be described in the following sections of this Chapter. Therefore, our input image is in SMF format and addressed as SMF-blurred-noGlare. In the pre-processing stage we remove shot noise and curvilinear structures (CLS) according to Section 4.1 and obtain SMF-noCLS. Then we compute the parameters of the anisotropic diffusion filter (addressed as statistical analysis and described in Section 4.2) and enhance the mammogram while removing the remaining noise; the result is SMF-AD. The map of microcalcifications is the output of the foveal segmentation explained in Section 4.3. Results and a comprehensive discussion about the parameter setting are included at the end of the chapter.

## 4.1 Pre-Processing

We have previously presented in Section 2.3 the new approach by Highnam and Brady for mammographic image normalisation offering a quantitative representation of the breast tissue, the  $h_{int}$  [65]. There is a drawback however, the extremely noisy appearance of these images that makes their analysis more difficult. Although the signal-to-noise ratio (SNR) of the mammogram is improved slightly by the process that generates the  $h_{int}$  representation prior to glare removal, the subsequent glare removal step, which has the advantage of reducing the number of FP, has the disadvantage of drastically decreasing the SNR due to the amplification of high frequency noise.

Yam *et al.* [178] introduce a de-noising algorithm, which attempts to remove radiographic mottle [65], a major source of FP in microcalcification detection. This is achieved by applying a Wiener filter, adapted to the characteristics of radiographic noise, to the original image *prior* to the generation of the  $h_{int}$  image. As a result, the SNR is increased and the overall appearance of the  $h_{int}$  image improved. Since microcalcifications are also subject to noise, the amplified high frequency noise in the original  $h_{int}$  tends to disrupt small structures, such as microcalcifications. The new smoother version would reduce the number of false negatives (FN) in our detection algorithm.

To detect microcalcifications, we aim to filter the image in a way that blurs the background, but enhances calcium. Prior to the image enhancement and detection of microcalcifications, the normalised mammogram is pre-processed for shot-noise and curvilinear structures (CLS) removal. This step, described in the following section, eliminates some major sources of FP, as will be seen in the detection results.



**Figure 40:** The diagram of the foveal algorithm to detect microcalcifications in Standard Mammogram Form images. The top row underlines the typical generation of an SMF image, including glare removal. The bottom row highlights the method described above: pre-processing, statistical analysis, image enhancement and adaptive segmentation.

The above flow diagram outlines the steps of the detection method including the foveal segmentation (c.f. Section 4.3). In the diagram, SMF is taken to mean any suitable quantitative representation of breast tissue, such as Standard Mammogram Form,  $I$  is the original grey-level image;  $k$  is the contrast parameter used in the anisotropic filter and the foveal segmentation:

#### 4.1.1 Shot Noise Removal

The main causes of FP that we aim to eliminate are shot-noise and curvilinear structures [65]. The SMF generation process can easily detect film-screen artefacts [3, 66], at the glare removal stage. Therefore, glare removal may reduce the SNR, but would also minimise the number of FP in the detection of microcalcifications. Shot-noise generally refers to a noise process in a sensor in the scanner. Dust, hair or scratches on the intensifying screen or on the surface of the film (when it is subsequently digitised) are light attenuating. They have sharp boundaries in mammograms and can be treated as shot-noise. It has visual properties that are similar to those of microcalcifications. Highnam and Brady [65] note the absence of blur in such structures of noise in mammograms. The 3D shape of noise is sharp, as seen in Chapter 3, Figure 31. Using

the point spread function of the intensifying screen; the points are identified and marked on a map (the energy imparted becomes negative after glare compensation where shot-noise is present, since it has been introduced to the image after the main blurring stage [66]). Using the binary information from the shot noise map, the artefacts can be removed from the SMF image by interpolating between their surrounding backgrounds. Figure 41 shows an example of shot noise removal.

#### **4.1.2 Curvilinear Structures Removal**

Curvilinear structures (CLS) in the mammogram arise due to anatomical features, such as milk ducts, ligaments, blood vessels and tumour spiculations. They appear relatively bright, though not necessarily with high contrast, in mammograms and are typically long thin lines crossing parts of the breast. They are locally linear, but may also curve on larger scales. Their detection is difficult due to the large variety of widths and lengths. Noise can easily disrupt the appearance of CLS in mammograms leaving isolated bright spots behind, which can easily be confused with microcalcifications. Also, at an intersection of two CLS in the image, corresponding to overlap of the CLS in the 3D compressed breast, the attenuations of the individual CLS add, producing a localised region of higher attenuation. Such points also appear similar to calcifications rather than noise.

The method we propose for CLS removal is based on the local energy model for feature detection of Kovesei [90] and is presented in [35, 36, 181]. For a more detailed review of local energy and phase congruency please refer to Appendix A. A recent development of CLS detection and removal is presented in [154] and we show results on initial experiments after removing CLS from mammograms using Schenk's multiresolution algorithm later in this Chapter.

The principles behind local energy and phase congruency methods are discussed next. From local energy we can search for features at points of maximum phase congruency, which

are the points where the local Fourier components are maximally in phase [90]. Phase congruency highlights those points with maximum local energy. It is invariant to changes in image brightness and contrast, but scale affects the detection of relevant features. Local energy ( $LE$ ) and phase congruency ( $PC$ ) are expressed by Evans *et al.* as in Equations ( 57) and ( 58) [36]. (We used their formulae for computation; but for Kovesi's approach refer to Appendix A.) For the spatial frequency  $j$ , the local Fourier coefficients of a one-dimensional signal are:

$$\sum_{j=1}^N (A_j + iB_j) \quad (56)$$

$$LE = \left( \sum_{j=1}^N A_j \right)^2 + \left( \sum_{j=1}^N B_j \right)^2 \quad (57)$$

$$PC = \frac{LE}{\sum_{j=1}^N (A_j^2 + B_j^2)} \quad (58)$$

When the Fourier components of the signal are in phase,  $PC$  becomes 1.

Things become slightly more complicated in the case of two-dimensional images (intensity mammograms, SMF images).  $PC$  becomes a function of position in the image and filter orientation. Kovesi proposes the convolution of the logGabor function with the image using the Fast Fourier Transform over 6 orientations and 4 scales. The logGabor filter appears as a hill of approximative Gaussian shape on the positive side of frequency (see Figure 105). This is in the frequency domain a good approximation of the sum between an even symmetric and an odd symmetric filter (multiplied by  $i$ ) and the convolution with a logGabor can be seen as the sum of convolutions with both these symmetric filters. Performing the inverse FFT, we obtain the real part as the result of convolving with the even symmetric filter, while the imaginary part is the result of the convolution with the odd symmetric filter. For every scale and orientation, the magnitude and the phase of the convolution are computed. The size of the bandwidth is set as seen in Appendix A. The logGabor with a transfer function as in Equation ( 59) is used to

obtain oriented wavelet filter ( 61), where  $\omega_0$  is the centre spatial frequency of the filter (oriented wavelet filter),  $\theta_0$  is the orientation of the wavelet,  $\sigma_0$  is the angular spread of the wavelet,  $\log(k/\omega_0)$  is the wavelet's frequency spread and  $S(\theta)$  is the cross-section of the transfer function in the angular direction ( 60). Through the computation of  $LG(\omega)$ ,  $PC$  becomes  $PC_{\theta}(x)$ .

$$LG(\omega) = \exp\left(\frac{(\log(\omega/\omega_0))^2}{2(\log(k/\omega_0))^2}\right) \quad (59)$$

$$S(\theta) = \exp\left(-\frac{(\theta - \theta_0)^2}{2\sigma_\theta^2}\right) \quad (60)$$

$$G(\theta, \omega) = S(\theta) \times LG(\omega) \quad (61)$$

Evans, Yates and Brady [36, 181] make the following two assumptions in detecting CLS from  $PC$ :

- the intensity profile of a CLS perpendicular to its local orientation is a one-dimensional peak (e.g.  $\cos(\varphi)$ ), with a phase within  $-\pi/2 \leq \varphi \leq \pi/2$ ;
- $PC_{\theta+\pi/2} > PC_\theta$  at a CLS of orientation  $\theta$ , since CLS are long and thin.


The weighted mean local phase in the direction  $\theta$ ,  $\varphi_\theta$  is calculated using scaled wavelet filters and unless  $-\pi/2 \leq \varphi_\theta \leq \pi/2$  for a given pixel, then the pixel does not belong to a peak and is discarded. Alternatively,  $PC_\theta$  is calculated. The step is repeated at regularly spaced intervals of  $\theta$  and a pixel is labelled as CLS if  $|PC_{\theta+\pi/2} - PC_\theta| > 0$  for all orientations  $\theta$ . The CLS labelled pixels are now output in a binary map of CLS. The choice of scale and parameters is critical for a good CLS detection, as Evans and Yates remark [35, 181].

The CLS removal algorithm, in its original implementation [35, 36], gets the best estimation of CLS as in Figure 102. One major reason behind this over-estimation is the development of the method for mammograms of smaller resolutions than the one we use to find microcalcifications. Also, calcifications were not relevant for the validation of the original CLS

removal method and therefore their removal along the CLS not considered important. The implementation we propose uses a smaller wavelength, which makes it more sensitive to smaller scales. Furthermore, we reduced the number of scales over which features are detected by phase congruency to avoid marking as CLS the calcifications, as small structures that still have local orientation. The downfall is that there are discontinuities that appear in the true CLS and their removal leaves bright areas on the mammogram that do not correspond to real structure with that particular form and characteristics (similar to microcalcifications). To avoid having such discontinuities and major errors in image manipulations, we followed the original algorithm proposed by Evans *et al.* and dilated the features appearing in the binary CLS map. After dilation, a smooth interpolation is required.

To perform the smooth interpolation, we adopted the solution proposed by Evans, which uses the Matlab Image Toolbox function *roifill*. This function takes as input (in this particular application) the mammogram and the corresponding binary CLS map. The interpolation is done from the boundary of the non-zero contours in the CLS map inwards, using the pixel values from the mammogram. Each non-zero region is treated separately during the interpolation. The criterion used is to solve Laplace's equation with specified boundaries. This may be interpreted as finding the smoothest interpolation on the interior of the boundary by solving a heat equation. The Matlab implementation of Laplace's equation solves a sparse linear system using finite differences.

Figure 41 shows an example of CLS removal and a binary CLS map. The amount of estimated CLS is large and shows responses of the algorithm to very weak edges, which induces a smoothing effect in those areas when the removal/interpolation is performed. This is an undesired side effect of the algorithm. These are the results obtained using Evans *et al.*'s method, when a good compromise is achieved in removing CLS without removing microcalcifications. A more accurate CLS removal technique is proposed in Section 4.4.2 and will be employed in future work.



**Figure 41:** Removing artefacts; (a) the original image; (b) the binary shot-noise map (white dots are noise); (c) the binary CLS map; (d) the ‘clean’ image after shot-noise and CLS removal.

## 4.2 Statistical Analysis and Image Enhancement

There are various methods, which may be used to do image enhancement, for example anisotropic diffusion-based filtering (cf. Chapter 3) which aims to blur the input mammographic image while preserving certain intensity changes, such as small regions of interest. The process relies on the use of a set of numerical parameters, often referred to in the image processing literature as time, contrast, and size. It has been found in practice that it is essential to determine the right choice of parameters to obtain good reproducible results. Our method automates the parameter setting for each individual image, and in this way it eliminates the need to set parameters.



The method exploits Weickert's solution [172] for the diffusion tensor. It uses a similar simplified and stronger tensor having eigenvalues ( 62), where  $I$  is the initial image,  $I_\sigma$  the Gaussian smoothed image ( $\sigma$  is the standard deviation or scale),  $k$  a contrast measure and  $n$  a suitably high power, such as 8 or 12.

$$\lambda_1 = \begin{cases} 1 & |\nabla I_\sigma| = 0, \\ 1 - \exp\left(\frac{-1}{(|\nabla I_\sigma|/k)^n}\right) & |\nabla I_\sigma| > 0 \end{cases} \quad (62)$$

$$\lambda_2 = 1$$

As we have noted in the previous Chapter, the parametric format of anisotropic diffusion makes this process highly dependent on the fine-tuning of its input parameters. In practice, the more complex and variable the images in a dataset, the more problematical it is to choose a single set of values for these parameters that works well for the entire dataset. Critically, by any measure, mammograms are very complex images, whose appearance varies widely across a population (at a centre, hospital, region, country, or continent), yet any microcalcification detection algorithm must generate very few FP and have even fewer false negatives (FN). One possible approach is to provide for interactive setting of parameter values; but this is unacceptable in clinical practice.

Microcalcifications are typically extremely small, therefore a small scale  $\sigma$  is desirable. Since  $t$  is not an image characteristic, we choose to vary only  $k$ , which is image dependent. Now we can iterate the blurring filter for a constant number of iterations over the image, but this raises the problem of how to choose the contrast parameter to determine the values of the eigenvalues ( 62).

Our method uses an adaptive Gaussian derivative filter ( 63) - ( 66). The application of this filter to a de-noised, glare-removed  $h_{int}$  image results in a gradient-map that highlights suspicious regions as regards microcalcification detection. Furthermore, the same filter outputs the value of  $k$  for the subsequent diffusion. Figure 43.c.,d. shows two examples of such gradient maps. Having the gradient map and the resulting value of  $k$ , we apply the anisotropic diffusion

filter to the  $h_{int}$  image with the corresponding value of  $k$  for a fixed number of iterations and set scale. The new images will generally be blurred, except possibly for some suspicious regions that will have their edges preserved, as in Figure 43.e.,f. Note the emphasised outline of the microcalcifications, while most of the background is blurred. By using the adaptive filter to compute the value of  $k$ , the anisotropic diffusion filter becomes robust, easy to use and automatically adapted to the image characteristics. The value of  $k$  represents the threshold selecting only the top 4.4% outstanding features in an image, a statistical value related to the sum between the *mean* and two *standard deviations*. Since the mean of our matrix is null (see Figure 42), equation ( 66) takes a simpler form.


$$K_{\sigma}(I) = \frac{1}{2\pi\sigma^2} * \exp\left(-\frac{|I|^2}{2\sigma^2}\right) \quad (63)$$

$$M = |K_{\sigma}'(I)|; \quad (64)$$


$$g_i = M_i - \frac{1}{N} \sum_{j \in \delta_i} M_j \quad (65)$$

$$k = 2 * \text{std}(g) \quad , \quad (66)$$

where  $K_{\sigma}(I)$  is the Gaussian of image  $I$ ,  $M$  the Gaussian derivative,  $g_i$  a classical measure of local contrast in a neighbourhood of  $N$  pixels, while  $k$  is the computed contrast value to be subsequently used to filter the image.



**Figure 42:** An example of estimating  $k$ . The image shown on the left (after expanded display contrast) has the associated histogram of function  $g$  in the right. We note the zero mean value of  $g$ , as well as where the value of 0.47 of  $k$  falls.



**Figure 43:** Automated image enhancement; (a) an image with an isolated calcification; (b) an image with a microcalcification cluster; (c) the corresponding gradient map for image (a) depicting the microcalcification and some extra undesired regions; (d) the corresponding gradient map for image (b) with a good representation of the cluster, but some falsely suspicious areas as well; (e) the automatically diffused image (a); (f) the diffused image (b);

### 4.3 Foveal Segmentation

An important goal in present image processing and computer vision methods is to integrate spatial models that reveal the sensitivity of the human visual system (HVS) at various intensity

transitions and texture variations in an image. There is over one century of research behind trying to fully understand the way the HVS functions. A major contribution to this on-going research was brought by Holladay in 1926 [69] and later by Moon and Spencer in 1943 [118]. They introduced the fundamentals of the way the human processes visual scenes. Nevertheless, the process is very complex and still not entirely understood [6], which make it difficult to design comprehensive models of HVS.

The final step in our method to detect microcalcifications is an adaptive segmentation method. It is well known that the Human Visual System (HVS) is highly sensitive and can detect fine details in noisy or textured images. Image processing aims to reproduce the quality of these results. Recently, Heucke *et al.* [62] introduced a computational model for the HVS. In their method, foveal contrast depiction is adapted to the object surrounding and background.

The adaptation of the eye to light changes is a continuous process in HVS. We perceive objects differently if they are against a bright surface (for instance a window in a sunny day) or dark area (a dark wooden panel). The adaptation luminance proposed in [118] is the response of the eye in adding an average luminance within the central visual field or fovea ( $L_{fov}$ ) and an equivalent veiling luminance caused by the luminance of surfaces surrounding the peripheral field of view ( $L_{seq}$ ) as in ( 67). However, the foveal adaptation owes mainly to the luminance within the foveal field and only approximately ten percent to the luminance of the field of view outside of that of the fovea [118]. The literature proposes 7.7% of the adaptive luminance to be due to the background luminance [62], which gives a value of 0.923 to our weight  $w$  (see below). In practice, we studied the effect of varying  $w$  with 10% more or less than the proposed value and a comparative FROC curve is shown in Figure 61.

$$L_a = L_{fov} + L_{seq} \quad ( 67)$$

Furthermore, the visual perception of the eye is dependent on the spatial perception of the object we perceive. This effect is called lightness assimilation. The same object may appear lighter on a dark background and darker over a light surface (see Figure 44). The eye is still

perfectly capable of distinguishing the three central areas, but this is not obvious for a computer program.

In the mammographic environment, the main advantage of the human eye over CAD methods is the depiction of microcalcifications in dense areas of the breast. At the same time, the radiologist must not pick FP in the dark areas. This is the result of the adaptation to background, which is equivalent to lowering thresholds in bright areas and increasing them in dark areas, basically an inverse proportionality to the background. While in our synthetic images the background is constant, in a mammogram we must consider the neighbouring area of the foveal kernel, as shown in equation ( 70). The adaptive threshold model that we propose for the detection of microcalcification applies the above-mentioned concepts in the mammographic setting for the improvement of abnormality detection in the breast.



**Figure 44:** An illustration of the lightness assimilation. We show three synthetic images with dark (left), medium (middle) and bright (right) backgrounds. All have central objects of the same size and intensity, but are perceived differently by our eyes, due to the variance in background lightness.

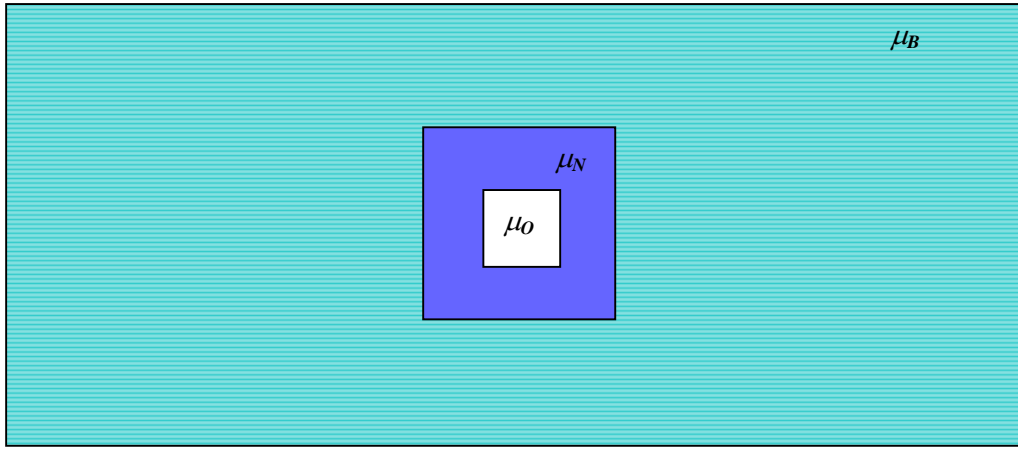
Previous contrast measures used in mammography seek to establish a minimal constant contrast difference all over the mammogram [80]. The contrast ( $C_{classic}$ ) is calculated at every pixel as the difference between that pixel value ( $p_0$ ) and a weighted sum of the pixel values in an immediate neighbourhood ( $N$ ) ( 68), where  $n$  is the number of pixels in  $N$ .  $C_{classic}$  is then compared with a fixed threshold,  $C_{thresh}$ , over the whole image and microcalcifications are marked. The variation in height in an SMF image or intensity in a typical mammogram makes

it far easier to detect microcalcifications against a fatty background but more difficult to detect correctly against a denser background [9]. In like manner, computer vision algorithms often find it more difficult to detect faint contrast changes against a bright background than against a dark background. The HVS, however, adapts to the local image contrast, and detects faint contrast changes in a manner essentially independent of the background. We have adopted a similar model of contrast detection in the HVS for mammography with the aim of improving the accuracy of the detection.

We initially remove the glare, shot-noise and CLS from the SMF image. Having the  $SMF_{noCLS}$  image, we compute a set of mean values using masks for the inner area (within the boundary of calcification), its neighbourhood (the local area around the calcification) and background (the rest of the breast tissue). The histogram of the inner surface will provide the mean of the object ( $\mu_o$ ), as the histogram of the whole image will give us the mean of the background ( $\mu_B$ ). The mean of the neighbourhood ( $\mu_N$ ) is defined as the weighted sum of intensities depending on the scale of the mask. A synthetic image designed to illustrate the kernels used for object ( $O$  with the mean  $\mu_o$ ), neighbourhood ( $N$  with the mean  $\mu_N$ ) and background ( $B$  with the mean  $\mu_B$ ) is shown in Figure 45. The perceivable contrast  $C$  is calculated according to equation ( 69):

$$C_{classic} = p_0 - \frac{1}{n} \sum_{i \in N} p_i \quad (68)$$

$$C = \begin{cases} \frac{\mu_o - \mu_N}{\mu_N}, & \text{if } \mu_o > \mu_N \\ 0, & \text{otherwise} \end{cases} \quad (69)$$




**Figure 45:** The foveal masks used for the computation of  $\mu_O$ ,  $\mu_N$  and  $\mu_B$ . The object  $O$  is the area of the *fovea centralis*,  $N$  its neighbourhood (twice the size of  $O$  in our applications) and  $B$  the background.

We then compute  $C_{min}$  (70), where  $\mu_A = w \mu_N + (1-w) \mu_B$  and  $w$  is a suitable weight between 0 and 1 affecting the amount of background implied in the computation of contrast. We have found that  $w=0.877$  to  $0.923$  gives good results. We will use the value of  $0.923$  for the default value in comparative studies, as proposed in literature [62]. The segmentation parameters are set-up automatically based on the image-adapted value of  $k$ , as computed from the statistical analysis, (discussed in Section 4.2). We found that  $c_w = \sqrt{k} / 200$  is appropriate for a conservative detection. Areas in the SMF image having  $C > C_{min}$  are marked as microcalcifications. In practice we have found that the value  $b=0.0808$  has given good values.

$$C_{min} = \begin{cases} \frac{c_w}{\mu_N} (b + \sqrt{\mu_A})^2, & \mu_A \geq \mu_N \\ \frac{c_w}{\mu_N} \left( b + \sqrt{\frac{\mu_N^2}{\mu_A}} \right)^2, & \mu_N > \mu_A \end{cases} \quad (70)$$

Using  $C_{min}$  instead of  $C_{thresh}$ , the contrast is adapted locally, not only globally, in a manner similar to that of the HVS. Figure 46 shows how the variation of perceivable contrast varies with the background in the HVS adaptation method versus classical methods. While  $C_{min}$  varies with the local image characteristics,  $C_{thresh}$  is constant over the whole image.


Microcalcifications can be correctly depicted in a fat background for both contrast measures, but in a dense background the detection is facilitated by the adaptability of  $C_{min}$ .




**Figure 46:** The simulation of a plot of a mammogram section containing microcalcifications over height/intensity variation. The variation of the perceivable contrast in the detection of microcalcifications is suited to the local characteristics for the adaptation of HVS using  $C_{min}$ . The classical minimal perceivable measure, (here called  $C_{thresh}$ ) is a global characteristic of the mammogram and less flexible in the elimination of FP in the detection of microcalcifications.

The function  $C_{min}$  (70) is a measure of contrast sensitivity. It sets the threshold from which objects in the image are visible for the observer, a measure of the eye's ability to perceive luminance gradients. Through the use of the minimal perceivable contrast ( $c_w$ ), it includes a measure of the image brightness. (Imagine varying the amount of objects we can distinguish with the naked eye by using a pair of sunglasses.) The variation of  $C_{min}$  over a set of synthetic images presenting variations between background (see Figure 47), neighbourhood and object is shown in Figure 48. Using the adaptive thresholding qualities of HVS we can depict all the high peaks in these images, which would be impossible with a simple threshold. For the parameter  $b$  in the  $C_{min}$  equation, we used the value 0.0808, as proposed in [118] with good segmentation results.






**Figure 47:** A set of five synthetic images with variations between object, neighbourhood and background and their associated cross-sections. These examples cover a wide aspect of contrasts in image processing: bright on dark, bright on bright, dark on dark. The corresponding values of  $\mu_O$ ,  $\mu_N$ ,  $\mu_B$  and  $C$  are shown in the table below the figure. The variation of the adaptive threshold is shown in Figure 48.




**Figure 48:** The variation of the adaptive threshold  $C_{min}$  for the synthetic images in Figure 47 in the central area of images, which corresponds to the object and neighbourhood surface. For this example, we used  $c_w=0.03$ . The value of  $C_{min}$  in the centre of fovea is shown in the table below the figure. We note that for the most delicate case (extreme right),  $C_{min}$  and  $C$  are in the same range of values. In such difficult cases, which approximate better the mammographic environment where transits between different intensities are much smoother, the adaptation of  $C_{min}$  become crucial. Moreover,  $C > C_{min}$  in all five cases and all peaks are detected.




**Figure 49:** The FROC curve of the microcalcification-detection method based on the adaptation of HVS in digital mammography.


Figure 49 shows the Free-response Receiver Operating Characteristic (FROC) curve of the tested microcalcification detection method based on the adaptation of HVS. We used a database of 102 samples of digital SMF images. 78 of them contain between 1 and 3 clusters per image, while 24 are normal mammogram samples. There are a total of 98 clusters of microcalcification annotated in the database. All images were digitised at a resolution of 50 $\mu$ m and have sizes under 1500x1500. We further show some examples of microcalcification clusters detection in Figure 50 -Figure 54. A cluster is detected if it contains at least three microcalcifications, where a distance of maximum 0.5 cm (approximate value) connects each calcification to the rest of cluster. Recursively, we noted that the distribution of detected FP/image is equal in samples with microcalcifications and in normal samples. Along with the original SMF sample and the microcalcification detection map, the figures show intermediate results from the following stages; we also show the gradient map resulting from the statistical analysis, the CLS map and the enhanced SMF sample after applying the automated anisotropic diffusion filter described previously. The detection process is non-parametric and fully automated, being adapted to the local and global image characteristics.




**Figure 50:** Detection example 1: (a) the original SMF images with a microcalcification cluster; (b) the gradient map from the statistical analysis depicting suspicious pixels; (c) the CLS map; (d) the enhanced image after diffusion; (e) the microcalcification detection map.




**Figure 51:** Detection example 2: (a) the original SMF images with a microcalcification cluster; (b) the gradient map; (c) the CLS map; (d) the enhanced image; (e) the microcalcification detection map.



**Figure 52:** Detection example 3: (a) the original SMF images with a microcalcification cluster; (b) the gradient map; (c) the CLS map; (d) the enhanced image; (e) the microcalcification detection map.




**Figure 53:** Detection example 4: (a) the original SMF images with a microcalcification cluster; (b) the gradient map; (c) the CLS map; (d) the enhanced image; (e) the microcalcification detection map.



**Figure 54:** Detection example 5: (a) the original SMF images with a very large microcalcification cluster; (b) the gradient map; (c) the CLS map; (d) the enhanced image; (e) the microcalcification detection map.

## 4.4 Discussion

We have presented a fully automated non-parametrical method to detect microcalcifications in digital mammography. The result of applying it to a mammogram (typically an SMF image) is a map of detection, which highlights the microcalcifications present in the image. As the FROC curve in Figure 49 shows, the detection rate on a 102 set of mammogram samples reaches 98% TP fraction at 0.1 FP/image. All the microcalcification clusters in the tested images are correctly detected at 1.1 FP/image. The most difficult case was a cluster in the breast margin (see Figure 55). Since the margin of the breast in mammogram has not been compensated, the characteristics of microcalcifications in that area are atypical and their detection becomes difficult.



**Figure 55:** A case of difficult detection with a faint microcalcification cluster in the breast margin. The contrast in the image has been enhanced for the reader to help in the visualisation of the cluster.

Adding adaptive contrast segmentation based on human foveal processing significantly enhances the detection of microcalcifications (see Figure 76). However, the variability of microcalcification appearance is very large and the algorithm may struggle with some difficult cases. The robustness of the method comes from the sequence of filters applied to the mammogram; the complex processing of images makes the algorithm slow when working with


entire mammograms (generally images of 4500x4500 pixels). Therefore, a faster implementation of it may be required.

#### 4.4.1 Comparative ROC Analysis

In this section we present comparative FROC curves to test the response of our method with variations in algorithm and input images. We will see the impact of CLS removal, the changes with the variation of the minimal perceivable contrast, the results on intensity images and on whole mammograms.


##### 4.4.1.1 CLS Removal

We will first test the influence of the CLS removal step in the detection of microcalcifications. Although the number of FP is lowered (see Figure 57 and Figure 56), the method is not perfect. To make sure that no microcalcifications are eliminated during this step, we preferred a conservative approach. This may leave some CLS residuals in the image, which may lead to FP in the detection. Also, we may still lose some small microcalcifications during the CLS removal, which does not seriously affect the detection of the cluster, as seen in the examples below.




**Figure 56;** Another example of CLS removal in detection: (a) the original mammogram; (b) the detection map using the CLS removal; (c) the detection map without CLS removal with a few extra FP detected.





**Figure 57:** CLS removal in detection: (a) the original mammogram; (b) the detection map using the CLS removal; (c) the detection map without CLS removal with a few extra FP detected.




**Figure 58:** The comparative FROC curve when CLS are removed or not prior to the microcalcification detection

In Figure 58 we show the comparative FROC curve of our method applied on original images, which are not pre-processed with the CLS removal algorithm versus CLS-free images. The response of the algorithm is vastly superior when CLS are eliminated prior to


microcalcification detection. The large number of FP caused by the presence of CLS delays the FROC response when CLS are not removed.

#### 4.4.1.2 Image Diffusion

The pre-processing of our database includes, as highlighted along the manuscript, the smoothing of images by anisotropic diffusion. The reason for employing smoothing along the algorithm has already been mentioned in previous sections. What is still left to be done is comparing the performance of the method with and without smoothing to get a clear understanding of the effect of diffusion on the detection results. The whole algorithm is replicated without the smoothing step and the comparative FROC curve is shown in **Figure 59**. An example is also shown in Figure 60, where we can compare the detection results with and without smoothing on a mammogram sample. Although the smoothing is expected to remove some of the very small or faint microcalcifications, the clusters are well preserved using the nonlinear qualities of anisotropic diffusion. Notably, the higher number of FP detected when no smoothing is present reduces considerably the performance of the algorithm.




**Figure 59:** The comparative FROC curve when smoothing by anisotropic diffusion is performed or not prior to the segmentation of microcalcifications.



**Figure 60:** Image smoothing in detection; (a) the original mammogram sample; (b) the detection map using anisotropic diffusion; (c) the detection map without using smoothing with FP marked.

#### 4.4.1.3 Perceptibility


In Section 4.3 we mentioned the significance of  $w$  in setting the minimal perceivable contrast for obtaining the best detection results when our algorithm is applied. We ran parallel tests to test the consistency of our conclusion to use the value 0.923 for  $w$ . Therefore, we varied the value of  $w$  over 5-10% and found that the detection peak is achieved when  $w$  is set between 0.877 and 0.923. However, a 5% difference in the value of  $w$  does not change significantly the detection results, although the efficiency of the algorithm is slightly lowered. Figure 61 shows the comparative detection results with the variation of  $w$ .



**Figure 61:** The comparative FROC curve when  $w$  is varied over a range of 5 to 10% of its default value of 0.923. The difference in detection results is quite small and all four algorithms converge smoothly to 100% TP ratio.

#### 4.4.1.4 Intensity Images

The obvious question we have to ask ourselves at this point is how well would the algorithm perform on intensity images and if the results are similar to those obtained on SMF images. We used the same 102 images, this time intensity images corresponding to the previously used SMF images, and tested our method over the same range of parameters as in the case of the SMF images that generated the above FROC curve. The results are shown in Figure 62. The detection algorithm performs similarly on intensity image, with a slower convergence, for the same set of parameters as for SMF images. The shape of the FROC curves makes a great difference to the number of microcalcification clusters detected at a particular number of FP/images, especially on the left side of the curves. This is the area of the FROC curve that corresponds to clinical results.



**Figure 62:** Intensity versus SMF comparative FROC curve. The detection algorithm converges slightly slower for intensity images, but reaches the same performance as for SMF images. One reason for the delay could be the use of same parameters when building the FROC curve, although the image characteristics (intensity versus SMF) are different.

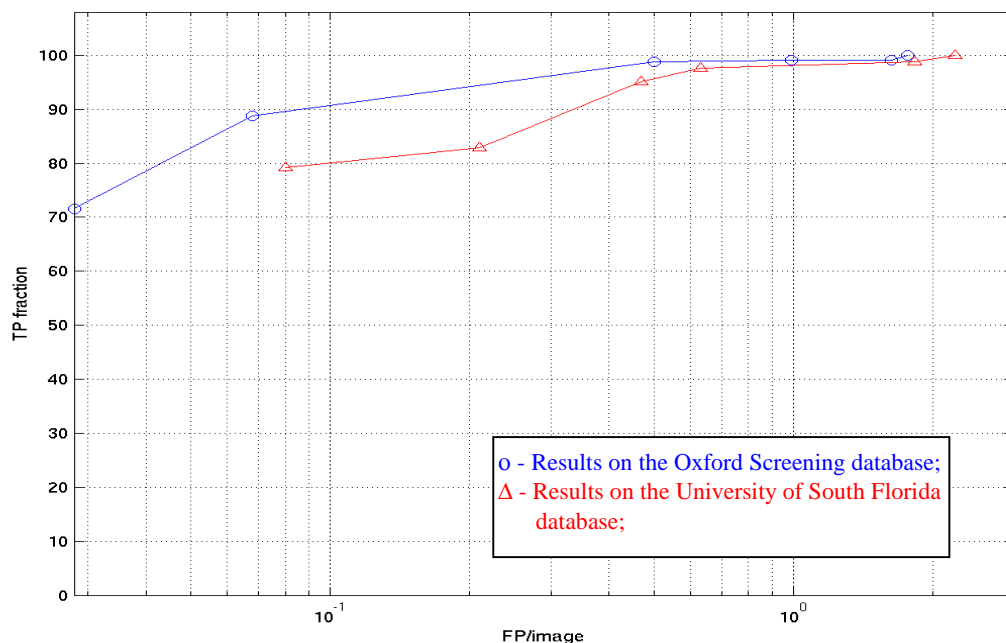
#### 4.4.1.5 University of South Florida Database

The ultimate goal of any CAD algorithm is to perform correctly on any given similar database, no matter where it comes from. As is acknowledged by many authors (not least those who constructed the University of South Florida database), without image normalisation this is hard to conceive of. The SMF generation algorithm is designed to help in this difficult situation, but excepting the Oxford database, no other image collections have mammograms in SMF format. Our detection algorithm, through its parametrical relation to the image attributes, facilitates the generalisation of detection standards, but without the use of a normalisation algorithm (a cornerstone in our reasoning), the results are not ideal.

We used for comparison a collection of images from the University of South Florida Digital Database for Screening Mammography (DDSM). More precisely, we applied our algorithm to mammogram samples digitised at  $43\mu\text{m}/\text{pixel}$  with annotated microcalcification clusters. They have similar sizes to the image samples from the Oxford Screening database

(c.f. Section 4.3). The major difference between these images and the ones from Oxford is the image resolution. To compensate for this difference, we converted the values of  $\sigma$  in the anisotropic diffusion step (see Section 4.5.2) to build a Gaussian kernel of approximately same size. In the Oxford database (at  $50\mu\text{m}/\text{pixel}$ ),  $\sigma$  was 0.6, which has the same physical size as a  $\sigma$  of 0.7 at the new resolution. Also, the kernel of fovea in the foveal segmentation (see Section 4.5.2) becomes 11 instead of 9 to be in accordance with the parameter setting reasoning. These conversions are done automatically at the launching of the application in agreement with the user specified resolution.

The new database consists of 82 image samples, of which 58 show abnormalities in the form of microcalcification clusters and 24 are normal. The abnormal images contain between 1 and 4 clusters/image and the total number of clusters is 82. All images are intensity images, as termed before, therefore the FROC curve shown in Figure 63 compares the performance of the microcalcification detection algorithm between the Oxford Screening Database in intensity form and the DDSM collection.



**Figure 63:** The comparative FROC curve between the detection results on intensity images from the Oxford Screening Database and the University of South Florida Digital Database for Screening Mammography .


As expected, the algorithm performs better on the Oxford Screening database, on which the parameters were originally trained. Nevertheless, the detection results on the two databases converge at about 0.5 FP/image and they both achieve 100% TP fraction in the vicinity of 2 FP/image. A more appropriate test of the detection algorithm on the DDSM database will be done when images will be available in SMF form.

#### **4.4.1.6 Whole Mammograms**

The results used in building the previous FROC curves are based on processing cropped samples of mammograms from the Oxford Screening Database. All experiments were performed on a 1.2 GHz Pentium III class workstation with 1 GB RAM. Nevertheless, it was very computationally expensive to test the detection method on whole mammograms, which are sized 3549x4816. This is mainly caused by the implementation of our method in Matlab, which uses extensive memory and slows down the processing time due to the high-level programming. By doing some algorithm optimisation, the processing was made possible for the illustration of results on detecting microcalcifications in the full surface of breast.


The breast margin is detected in SMF, thus a threshold above 0 removes the background. Now we can compute the value of  $k$  for the inner area of the breast. The detection results are accurate and similar to those achieved on mammogram samples (see Figure 64). The size of the images processed does not influence substantially the number of FP/image in the left side of the FROC curve. On the right side of the FROC curve, where the algorithm aims to achieve a detection rate of 100%, the difference in the number of FP is more substantial. We used a total number of 83 mammograms in SMF format from the Oxford Screening Database. 59 of them contain between 1 and 5 clusters/image, adding the total number of clusters to 85, while 24 mammograms have no sign of abnormality. The clinicians of the Oxford Breast Care Unit of the Churchill Hospital set the ground truth.

The most challenging cluster to detect remains the one shown in Figure 55. The presence of CLS remains the main source of FP, or more precisely their imperfect removal. A few isolated calcifications were also depicted, but they were not labelled as microcalcification clusters (they were located in groups of less than 3 calcifications, the minimum number required in the detection method). In Figure 65 we present detection results on whole mammograms; we indicate with ellipses the TP microcalcification clusters and with arrows the locations of FP.



**Figure 64:** The comparative FROC curve of the detection of microcalcifications when mammogram samples are used versus full mammograms. The behaviour of the algorithm is similar and robust with the image size.





**Figure 65:** Detection results on whole mammograms; (a) and (c) are the MLO SMF images, while (b) and (d) the corresponding CC images. Ellipses indicate the locations and spread of the detected microcalcification clusters, while arrows indicate the positions of FP.

#### 4.4.2 An Alternative CLS Removal

This alternative method we present here to detect and remove CLS is a refined version of the phase congruency presented earlier in this Chapter (see Section 4.1.2). We have used local


energy (LE) and phase congruency (PC) to detect and then remove CLS from mammograms. The results are substantially improved (see Figure 58), but the CLS removal algorithm leaves artefacts in images that can be misinterpreted as microcalcifications. In this section we discuss an alternative method to detect and remove CLS based on multiresolution oriented local energy analysis [154]. We encounter the same major problem, namely the wide range of scales and orientations of CLS. Using multiscale analysis, CLS are detected as a collection of edge-ridge-edge lines with similar orientations at the correct scale. The condition that neighbourhoods must follow the same model removes a substantial number of the CLS candidates highlighted by LE. Through exact interpolation between the edges, no artefacts and undesired smoothing are introduced. The proposed algorithm is detailed below.

#### 4.4.2.1 Theory

The motivations behind detecting and removing CLS have already been discussed in this thesis, but to briefly summarise:

- CLS are thin bright structures corresponding mainly to healthy tissue in the breast;
- They complicate the textured appearance of a mammogram;
- Their scale and brightness resembles that of microcalcifications and CLS crossings correspond to bright blobs;
- CLS related to blood vessels, ducts or ligaments may be confused with tumour spiculations;

Schenk and Brady [154] make the following assumptions in their work on CLS detection: CLS are locally linear; they have well-defined orientation; CLS are high-frequency (small-scale) structures. Local energy (c.f. Appendix A) gives responses to all CLS feature (i.e. the centre and the edges of the CLS) and uses phase to distinguish between these as in Figure 66. An effective computation of orientation can be performed.




**Figure 66:** The response of local energy to a variety of input signals: (a) the input signal; (b) the local energy response; (c) the phase-angle response.

The local energy decomposition is done by polar-separable quadrature filters in the Fourier domain, which are implemented as steerable filters. The computation is performed using an even-symmetric filter (71) and a corresponding odd-symmetric filter for the radial part with a negative lobe on the negative frequency axis. For the angular component of the filters (spread at  $N$  orientations over a half-circle) the cube of the cosine is used (72) and its absolute value for the odd and respectively even-symmetric filters. Figure 67 shows the steerable quadrature filter pair. Empirically we found good practical results by setting  $N = 4$  over three scales.

$$F(\lambda) = \log\left(\frac{1}{\lambda} \cos(x)^2\right) \quad (71)$$

$$A(\phi_n) = \cos^3(\phi - \phi_n); \quad \phi_n = \frac{n\pi}{N} \quad (72)$$


The Fourier transform of the image is multiplied by each of the filters and the inverse transform calculated. The filter response is computed for each pixel at each scale by a vector sum weighted by the amplitude of the oriented filter coefficients.



**Figure 67:** The angular part: a steerable filter.

The method uses phase to find all locations corresponding to positive going lines (ridges) and locations corresponding to positive/negative going edges. A point on a CLS must consist of an edge-ridge-edge triplet of similar orientation and correct scale. Figure 68 shows an example of ridge-edge combinations. The immediate neighbours of the CLS must follow the same scale-dependent rule.

At each pixel of a CLS, an approximate width of the structure is computed by selecting the scale that has the largest magnitude coefficients. The width is used at each pixel to remove the CLS by interpolating between pixels neighbouring each edge. The algorithm uses histograms of immediate small neighbourhoods to randomly sample the points in-between, weighting the values by the distance to either step-edge.





**Figure 68:** A simplified example of edge-ridge-edge triplet. The background is a scale image; ridges are shown in green, while edges in red. The lengths of the vectors express scale.


#### 4.4.2.2 Initial Results

In Figure 69 we can see an example of CLS detection on a sample of real mammogram. The example is design to show the robustness of the detection and removal on an image containing CLS, microcalcifications and a mass. The CLS map over three scales is shown along with a close up exemplifying the edge-ridge-edge triplets detected from the filter response.

Having detected CLS, we removed these lines from the image and applied the microcalcification detection algorithm to both the original mammogram sample and the CLS-free image. The two sets of results are shown in Figure 70. A notable improvement can be seen from the incorporation of this multi-resolution CLS detection stage. The microcalcification clusters are preserved over the CLS removal, since CLS are now manually removed at different scales depending on the filter response. Furthermore, edges are not erroneously detected, as they do not respect the edge-ridge-edge rule, while microcalcifications are not affected, as they do not fit in the model at the right scale. At this initial testing stage, the method to remove CLS seems more performant than the phase congruency-based algorithm presented earlier. Note that the scale selection and CLS detection were done manually



**Figure 69:** An example of CLS detection: (a) the original mammogram sample containing a mass, two clusters of microcalcifications and CLS; (b) the CLS map over 3 scales; (d) a close-up of the upper right corner of image (a) showing the edge-ridge-edge triplets (edges in green, ridges in red). The CLS are detected using manual thresholding over each scale.



**Figure 70:** An example of microcalcification detection: (a) the original mammogram sample; (b) the CLS-free mammogram sample using Schenk and Brady's algorithm; (c) the CLS map before CLS removal with a large number of FP; (d) the CLS map after CLS removal with improved results and a significant reduction of the number of FP.

#### 4.4.2.3 Discussion

The new CLS detection method can reliably differentiate between CLS and other high-frequency components in a mammogram, which are congruent over scale. A better and more complete CLS removal can be performed without risking the removal of microcalcifications or excessive interpolation. However, when we performed experiments involving an automated selection of scales the results were different. We marked a CLS at the first scale that gives response to it. Unfortunately, most microcalcifications gave responses to one scale or another when no manual thresholding was used. While the number of FP was significantly lowered, most microcalcifications were lost and the features of the detected clustered seriously altered.

Thus, the CLS detection based on multiresolution oriented local energy analysis must be further developed before imbedding it in an automated detection algorithm for microcalcifications.

## 4.5 The Detection of Microcalcifications using SMF

### 4.5.1 Comparative Results

This section compares three algorithms that operate upon the  $h_{int}$  representation (or the SMF) to detect microcalcifications. Two of these algorithms have been described previously in this thesis; one is the method presented in the last sections (using image enhancement, CLS removal and foveal segmentation) and referred as the Foveal Approach, the second one is Yam *et al.*'s Physics Based Approach that was described in Chapter 3. The third is a variation of the statistical analysis introduced in Section 4.2. Using ROC analysis, we demonstrate the superiority of the first algorithm. First, however, we describe briefly the third algorithm that we compare here.

The third detection method, which we refer to as the Statistical Approach, differs from the foveal approach only in the final segmentation stage. In common with the foveal approach, the statistical approach applies an adaptive Gaussian derivative filter to the de-noised, glare-removed SMF images, which results in a gradient map and outputs the value of  $k$  for the subsequent diffusion stage. Next, the anisotropic diffusion filter is applied to the SMF image using a constant number of iterations  $t$ , with a pre-defined value of  $\sigma$  and the pre-computed value of  $k$ . With the diffused images ready, we need to ensure that the artefacts emphasised in the gradient maps will be eliminated.

A final segmentation stage uses statistical analysis to classify the content of an image into three main categories:

- background - where the variation of the image gradient is too low to correspond to microcalcifications;
- putative microcalcifications; they are conservative to avoid FN;



- shot noise - where the variation of the image gradient is too high to correspond to anatomical structures.

The filter is based on the same computation as the anisotropic diffusion, but also incorporates some adaptive thresholding suited to the image characteristics and the properties of microcalcifications. This last filtering process outputs a black-and-white map of detection (BWMD), where all pixels different than the black background correspond to calcifications in the breast.


The statistical analysis operates in the following manner. If  $I$  is the SMF image we process, the contrast of  $I$  is computed by comparing the value of the  $mean(I)$  with that of  $max(I)/2$ . A first threshold is adapted to the image contrast and  $I(SetPixels)=mean(I)$ , where  $SetPixels$  are the pixels in image  $I$  that have the absolute value of the Gaussian derivative smaller than a constant  $M$ . For images with low contrast, we found  $M=5$  to give good results, while for images with high contrast  $M$  has the value 20. The step is repeated over a few iterations in order to evaluate the steepness of the selected microcalcification candidates at each step. The final version of the image is contrast enhanced and the maximum value of  $M$  is used for thresholding. All the pixels that satisfy the thresholding criterion are marked as microcalcifications on the BWDM image.

In summary, the third method uses three filters in sequence:


- an adaptive Gaussian filter, which generates a gradient map and, more important, the value of  $k$ ;
- an anisotropic diffusion filter, which will enhance certain suspicious regions based on the previous computation of  $k$ ;
- some more statistical analysis built as an adaptive thresholding filter, which will discriminate between microcalcifications and the rest of the image.

Figure 71 to Figure 75 show the detection results on some mammogram samples containing microcalcification clusters. We present, along with the original contrast-enhanced SMF sample, the detection maps of the Physics-based Approach, Statistical Analysis and Foveal Approach.


We used a database of 102 samples of digital SMF images, 78 of them contain 1 to 3 clusters per image, while 24 are normal mammogram samples. There are a total of 98 clusters of microcalcification. All images were digitised at a resolution of  $50\mu\text{m}$  and have sizes under  $1500 \times 1500$ . Figure 76 shows the comparative FROC curves of the tested detection methods.




**Figure 71:** Comparative Results 1 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (c) the detection map of the Foveal Approach.




**Figure 72:** Comparative Results 2 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (c) the detection map of the Foveal Approach.




**Figure 73:** Comparative Results 3 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (d) the detection map of the Foveal Approach.



**Figure 74:** Comparative Results 4 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (d) the detection map of the Foveal Approach.



**Figure 75:** Comparative Results 5 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (d) the detection map of the Foveal Approach.




**Figure 76:** The FROC curves of the three microcalcification-detection methods, where we notice the better performance of the Foveal Approach.

#### 4.5.2 Setting the Parameters

The diffusion process is modelled by three parameters that determine its response to various image situations. Their influence on the smoothing of an input image has been described before in Section 3.2.2, their setting for a specific application, such as image smoothing prior to the detection of microcalcifications is the subject of the following paragraphs.

During the experiments performed and presented in Chapter 3, we aimed to characterise the effect of the choice of specific values for the three diffusion parameters,  $k$ ,  $\sigma$  and  $t$ , on glare-removed SMF images. Since anisotropic diffusion was originally meant to be used as an alternative smoothing technique to the Wiener filter (see Section 2.2.3), though it was subsequently used in combination with it, no such filtering was used on the trial images in Chapter 3 prior to the SMF generation. Hence, the contrast ( $k$ ) values used in these experiments, which were chosen empirically for the illustration of our first experiments, would be atypical for the validation of our algorithm on glare-removed Wiener-filtered SMF images, as in Chapter 4.


Microcalcifications represent a tiny percentage of a mammogram surface. Typically, they are very small and present in about a quarter of the total number of screening mammograms. Therefore, a percent of about five of the total number of mammogram pixels should be more than sufficient to accommodate the entire population of calcium salts. Furthermore, calcium is very bright in an X-ray image and it would be amongst the brightest/highest pixels of the mammogram. As illustrated in equation (66) we compute  $k$  as a value of the gradient (for tall steep microcalcifications) that discriminates between these brightest structures and the background; more precisely, we chose the 4.4% structures with highest contrast (a threshold at  $\text{mean}(g) + 2*\text{std}(g)$ ).  $k$  becomes a value with well-defined physical meaning (differentiating between high-pointed structures, referred to earlier as towers, on the SMF surface and the rest of the images), which is calculated automatically part of the detection algorithm. Figure 77 illustrates the effect of  $k$ .



**Figure 77:** In the left case, the  $k$  factor is bigger than the gradient and the diffusivity function  $g \rightarrow 1$ , which is equivalent to finding an edge and maximising diffusion; in the right case,  $g \rightarrow 0$  (for very big gradients) and diffusion is inhibited.

The second parameter to be set is  $\sigma$ , the standard deviation of the Gaussian filter used to smooth the image, which will give the size of the Gaussian kernel that removes noise by convolution with the mammogram. We need to choose a value for  $\sigma$  so that, on the one hand, it removes high-frequency noise (which are very small and spread over a couple of pixels at  $50\mu\text{m}/\text{pixel}$ ), but, on the other hand, preserves microcalcifications (on average 0.5-1 mm in diameter). To ensure that small calcifications (that is, those whose sizes are 300 microns or more) are preserved in the image, we set  $\sigma$  to 0.6. This will build a Gaussian kernel of 0.35mm

$(10\sigma + 1)$ , which is sufficiently small to clean noise and keep calcium salts. The principle is shown in Figure 78.



**Figure 78:** In the left case, the filter width is small compared to the structure, so it essentially detects a step edge; but in the second, it is not obvious that it will .

The experimental results show that  $k$  has small values for the inner area of the breast, which makes smoothing safe for faint microcalcifications, but that this process is also rather inefficient over the first couple of iterations. Weickert [172] notes that the number of iterations  $t$  is related to the spatial width of the Gaussian kernel. To blur features of the kernel order  $(10\sigma)$  requires  $t = (10\sigma)^2 / 2$ . Rounding this value for  $\sigma=0.6$ , yields  $t=5$ , which give excellent noise reduction results in the microcalcification preservation framework. We studied experimentally the effect of the variation of  $t$  and can offer the following comments: 1-3 iterations have too little influence and the output images are still noisy; for  $t$  greater than or equal to 7, not only is the process is time consuming, but most images are overly smoothed and valuable information is lost.

A reliable application must prove its robustness under different clinical imaging conditions. When parameters are involved, they must be suited for any type of input images. The danger that an algorithm is tuned to a particular data set is a major concern in the validation of clinical applications. Such a discussion is required when evaluating our method for the detection of microcalcifications.

One of the starting points in this method is the employment of a technique of mammographic image normalisation, the SMF. While various hospitals around the world use



different X-ray machines to acquire mammograms during screening trial, the SMF finds a common framework in which mammograms are presented as maps of dense tissue. The major advantage of SMF is the de-parameterisation of mammograms, the normalisation of their appearance. Ideally, the setting of parameters of our algorithm must be done once to make the algorithm work on any normalised image. A justification of the parameter tuning follows.

The first trials for the setting of parameters are done on small mammogram windows with isolated calcifications (c.f. Section 3.2.3). The samples are extracted from the Oxford Screening Database. They are generally of 500x500 pixels at 50 $\mu$ m/pixel, although some of them may be slightly smaller. These experiments are designed to evaluate the removal of noise associated with the differentiation in shape between shot-noise and background versus microcalcifications, as proposed by Yam [178, 179]. In Chapter 4, where the main evaluation of the algorithm is described, we also use mammogram windows (both intensity and SMF) from the Oxford Screening Database. They are samples from the same screening mammographic database, but different windows of different mammograms, since we search microcalcification clusters this time. The first tuning of parameters is therefore done on different images coming from the same screening centre, a different collection of images, but with similar imaging characteristics (as they are acquired and digitised using the same equipment).

The choice of the values for the parameters used in the anisotropic diffusion process has been explained above; but a few more comments are in order. First,  $k$  is computed directly from the image and its meaning (depict the most outstanding features in the mammogram) is independent of the data set.  $\sigma$  is strictly related to the size of microcalcifications and the image resolution; the condition imposed here is that databases must be digitised at 50 $\mu$ m/pixel.  $t$  is set according to the observed value of  $k$  and the proposed value of  $\sigma$ , which are consistent over the Oxford Screening, DDSM and MIAS databases (see Chapter 5). We perform tests with images from several databases to prove the robustness of the algorithm. Moreover, we obtain good results both on images with and without normalisation.

We presented a working example of tuning the parameters of anisotropic diffusion to an application. In a more general framework, anisotropic diffusion is a feature detector, namely an edge detector.  $k$ , being closely related to the gradient in the image, can be derived according to the percentage of features that we desire to enhance in an image.  $\sigma$  gives a measure of scale and must be set according to the size of searched features at the image resolution (multiscale analysis may be performed). The literature proposes stopping criteria for  $t$  as well [23, 171] as functions of  $\sigma$ , which can be well related to noise removing, but may be more difficult to combine with feature enhancement for some applications. With the automatic tuning of parameters that we propose, anisotropic diffusion can finally be used as a robust parameter-free process using little, but essential, a-priori knowledge. This should be a great relief in applications of diffusion, a controversial method for its parametrical dependency.

The other parameters employed by our method are the kernel of inner object ( $O$ ) used to compute  $\mu_O$  in Section 4.3 and the value of the minimal perceivable contrast  $c_w$ . The size of  $O$  is established according to the size of the microcalcifications for the resolution of the tested images. A kernel of 9x9 pixels has an area of  $0.4 \times 0.4 \text{mm}^2$ , which is just below the average surface of a microcalcification. It is desirable to have a slightly smaller kernel than the microcalcification diameter to assure the detection of small calcium salts, which are overlooked by larger kernels. Still, the size of  $O$  must not be extremely small to avoid overlapping  $O$  and  $N$  for slightly bigger microcalcifications. The last parameter,  $c_w$ , is related to the image contrast (c.f. Section 4.3) and empirically it has values typically between 0.002 and 0.005 for good detection results. We noted that adapting the value of  $c_w$  to the image characteristics gives better responses to microcalcifications than a constant  $c_w$ . Since  $k$  is a value of image contrast, we simply had to scale it to the suitable range of values for  $c_w$ .

### 4.5.3 Discussion

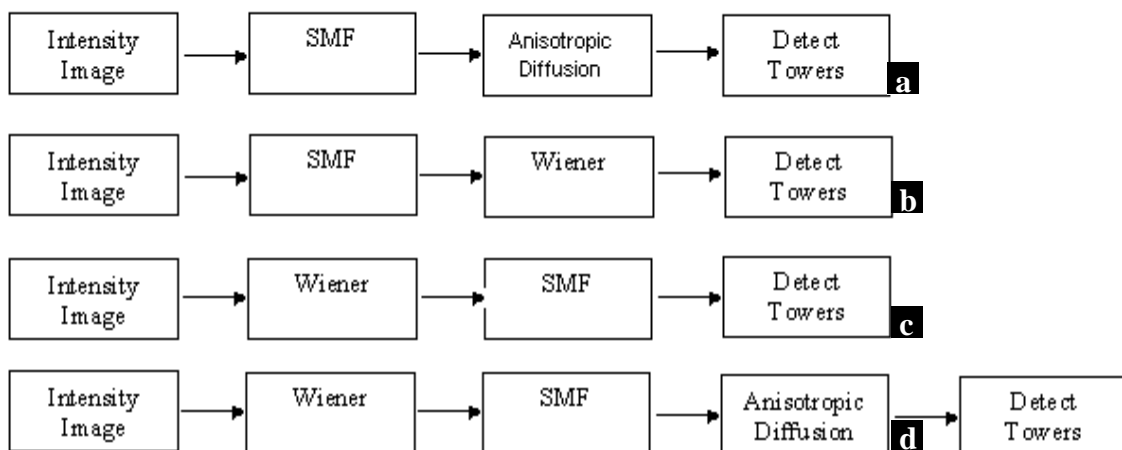
Though it is acknowledged that the state of the art of microcalcification detection is probably represented by R2's ImageChecker, unfortunately (though understandably) no details of the R2 detection algorithm are published. The R2 method was originally based on the work published by Nico Karssemeijer at the University of Nijmegen [80, 81], but it is known that the R2 implementation has been changed substantially over the years (some researchers claim, with no apparent justification, that the R2 implementation of ImageChecker takes the form of a neural network). All other published microcalcification detection algorithms have significant shortcomings, and sensitivities/specificities that fall short of R2's ImageChecker. This appears (to us) mainly because they don't work on normalised images. The key point is that without normalisation, for example by the  $h_{int}$ /SMF process described in the thesis, there is inevitably the risk of confounding anatomical information (of interest) and imaging parameter effects (not of interest), whose choice can affect contrast, brightness and level of noise, to name some of the classical limitations.

This is the framework where Yam's work (see Section 2.2.3) is the first to claim results that appear to rival those published by R2, though it must be acknowledge that her results are from a smaller database. Nevertheless, her work is based on SMF, which was the point of departure of the thesis. The advantages of her technique are straightforward and the incentive that we should build on Yam's work followed, but it is important to understand what we considered to be the problems with Yam's algorithm, hence why it needed improvement.

In very few words (more details can be found in Section 2.2.3) Yam's technique works in two steps: (i) detect thin towers in the SMF (towers appear because the SMF generation process is based on the fundamental assumption that the breast (more precisely, the vast majority of projected pixels) contain only fat and "interesting tissue", not calcium. A consequence is that a calcification appears approximately 26 times higher than it should (this is the factor by which the attenuation of calcium exceeds "interesting tissue"); and (ii) then test a 26 fold reduction

with footprint in image as a plausible calcification. The major problem arises in step (i) since the SMF is extremely noisy. Yam's solution is to introduce a Wiener filter; but we found that in many cases, which we considered during the first year of our research that this blurred the edges of the microcalcifications, therefore hampering both steps (i) and (ii) above. Faint microcalcifications become "too small" and "less sharp" to be detected, while the remaining noise imposes strong constraints on removing false positives and the algorithm further misses some microcalcifications (see Figure 76 for comparative results). While clusters of microcalcification are generally still well detected as a whole, some individual calcifications are overlooked and so statistical analysis of a cluster is compromised.

Our suggestion was to develop an alternative way to smooth the image/SMF, namely anisotropic diffusion, a process well known for its quality to smooth while preserving edges. Our initial idea was proposed as in Figure 79.a, as an alternative to Figure 79.b, but in reality, what is done corresponds to Figure 79.c. This is the point where the work done in this thesis intervenes.



**Figure 79:** The original idea for the detection of microcalcifications: (a) what we initially proposed; (b) what we thought it was done; (c) what was done in reality; (d) our solution.

The final diagram that we propose is represented in Figure 79.d, adding the anisotropic diffusion filter after the SMF generation and before the detection of towers. Of course, this immediately poses the question: why use two different smoothing filters, namely the Wiener filter and anisotropic diffusion. The noise removal is carried out in two separate steps; the

Wiener filter models a noise process (quantum mottle and film granularity), but not glare. Glare, along with scatter and extra-focal radiation and the anode-heel effect, is removed within the SMF generation, which massively reduces SNR by amplifying high frequency noise (see Table 4). Thus, glare removal solves the problem associated with the presence of this type of low frequency noise and sharpens the image, but also creates a new difficulty, the lower SNR. The newly amplified high-frequency noise is a major source of false positives and governs the difference in SNR between columns 2 and 3 in Table 4.

At this point of the SMF generation, most imaging artefacts are dealt with either by the  $h_{int}$  model (scatter, glare, anode heel, extra-focal) or the Wiener filter (quantum mottle, film grain). However, as emphasised in Section 2.3, the SMF generation has its imperfections. Yam comments in [178] about the drawbacks of her noise deconvolution arising from the simplification existing in the estimating theory of quantum mottle and the available physics parameters of film-screens. These possible sources of errors leave residual high-frequency noise in mammographic images. Digitiser noise, also of high frequency, and errors of SMF generation add to it. Nevertheless, we can consider the Wiener filter together with the  $h_{int}$  generation as one main stage designed to remove imaging specific parameters/errors, basically the image normalisation. The anisotropic diffusion filter aims to smooth the remaining high-frequency noise that interferes with our specific application, the process of detecting microcalcifications.

**Table 4:** The variation of SNR in generating SMF.

	SMF with Glare no Wiener	SMF no Glare no Wiener	SMF no Glare with Wiener
SNR	39.86	5.37	29.24

More specifically, would replacing the Wiener filter with anisotropic diffusion have eliminated the need for Wiener and subsequent anisotropic diffusion? Glare would amplify any residual noise after SMF generation, hence would have continued to interfere with the detection

of microcalcifications. It follows that smoothing after SMF is needed for any application dealing with very small anatomical features, in particular microcalcifications. On the other hand, not smoothing before computing the SMF would amplify noise to very high levels during the image normalisation process (see Table 4), which would make it extremely difficult to detect small structures and the subsequent filtering would be inefficient in properly removing noise while preserving microcalcifications. The Wiener filter has the advantage, stressed by Yam, of modelling particular kinds of mammographic noise (see Section 2.2.3).

We implemented the diagram in Figure 79.d for the detection of microcalcifications, as illustrated in Section 3.2 and further in Chapter 4, where a simulation of the detectability properties of the human visual system replaces the detection of towers. Thus, besides the SMF generation and its normalising action, we use two non-linear blurring filters: anisotropic diffusion (see Section 3.1) and Wiener (fundamentally a probabilistic assignment to signal and to noise). Through them, we address two different problems: the Wiener filter models quantum mottle and granularity of film, whereas the anisotropic diffusion filter smoothes the SMF image, but retains the signal (in the form of towers).

The subsequent filters (Wiener, SMF related, diffusion) model and correct for specific image analysis problems, rather than trying to amalgamate into a single (linear or nonlinear) filter that attempts to do everything. Separating them should make things clearer for the developer of such a filter, even if, for the end user, it is all reduced to a “black-box” that detects microcalcifications. Hence, we have a collection of blurring/low-pass and deblurring/high-pass filters. Is there any danger that they simply counteract/undo each other? As explained above, the Wiener filter and the deconvolutions within the SMF generation address the physics and the condition of the imaging process its controlling parameters. They separately attack different problems on the way to a normalised image. So does the anisotropic diffusion filter for the purpose of detecting salts of calcium, using data estimated parameters. The Wiener filter is a deconvolution tuned for quantum mottle and film grain noise; the glare unblurs and sharpens the image revealing anatomical details (calcium for example), but noise as well; the diffusion filter blurs again, but this time just the high-frequency noise, without losing the anatomical

details emphasised in the previous step. The last filter may have to undo the noise amplification, but will preserve the valuable anatomical information that would not have been available without the glare removal. Furthermore, we have a normalised image to work on.

## 4.6 Conclusion

In this Chapter we presented an algorithm for the detection of microcalcifications in X-ray mammography. The robustness of the algorithm has been demonstrated by the ROC analysis performed over a range of parameters. The method converged in each case to 100% TP ratio. Similar results were obtained on intensity images, although for the lower scale of FP/image there is a more significant difference in results.

Our method was tested on mammographic samples for faster processing and simpler validation, but on whole mammograms too. We also compared the performance of our algorithm on data from different databases with good detection results. Adding adaptive contrast segmentation based on characteristics of the human visual system significantly enhances the detection of microcalcifications. The parameters are set according to the image attributes and the method is fully automated. In future work, we aim to develop the algorithm by incorporating additional knowledge of X-ray attenuation.

## CHAPTER 5

# 5 Temporal Comparison of Feature Enhanced Mammograms For Mass Detection

*What hath night to do with sleep?*

John Milton – “Comus”

The development of reliable CAD systems for mass detection in digital mammograms remains an important problem that is still only partially solved. In this Chapter, we present a new method for prompting the clinician to “suspicious” dense regions in temporal mammogram sequences that combines feature detection and temporal comparison. The particular context that we envisage is post-screening assessment, where a radiologist recalls a woman after a suspicious screening mammogram for more in-depth analysis. At this point, the clinician compares the most recent mammogram to previous ones in order to detect significant changes. An important problem in automatic mass detection is the large number of candidate masses. The method presented in this Chapter uses anisotropic filtering (cf. Chapter 3) as a pre-processing step in order to significantly reduce the number of candidate masses, while preserving the important anatomical information about each mass. The method has been tested on the 15 temporal pairs currently available from the Oxford Database, where pathology has been diagnosed in the most recent image. Though we detect 100% of the masses, the number of false positives remains significant, necessitating further work.



## 5.1 Introduction

Computer Assisted Diagnosis (CAD) systems for breast cancer aim to assist the clinician in interpreting images and establishing reliable and early diagnosis of pathology. In particular, mass detection algorithms aim to automatically detect masses and/or characterise cancerous tissue against normal parenchymal tissue and benign diseases. Their purpose is to assist clinicians in the early detection of cancers. In digital mammography, CAD systems can be directly integrated in a soft-copy environment [12, 37, 142, 177]. To date, research in mass detection has given moderate results. For instance, using the R2 system, the accuracy in detection rate was reported to be less than 81.6% true positives (TP) for the detection of masses [41].

In this Chapter we present a new approach to prompt breast tumours with a central mass. The method aims to identify and segment such mammographic anatomical structures. Given the presence of central mass, we assume in Section 5.3.3 that masses are to be found in the dense and very dense areas of the breast. We do not expect architectural distortions, for example, to be prompted by our mass detection algorithm. We discuss our preliminary results on mammograms containing masses along with some initial features we propose towards mass classification.

Mass detection is a rapidly developing field; the main trends in approaching the subject are reviewed in Chapter 2. From the radiologist's point-of-view, detection is performed on single mammograms, bilateral mammogram matching and temporal mammogram matching. We aim to use a number of detection features that prompt masses in single mammograms and then use temporal matching of mammogram pairs to reduce the number of FP. The reasoning in the algorithm design was inspired by the mammographic screening programmes, where radiologists compare temporal mammograms to depict changes and discard FP.

Our method is similar to the work of Li *et al.* [97] for the segmentation step, where each mammogram is decomposed into several tissue classes according to their statistical properties,

but we use temporal comparison for improved results. However, our approach makes use of a more sophisticated pre-processing step (adaptive anisotropic diffusion-based filter) that enhances dense regions in mammograms and a simpler but more intuitive mammogram segmentation that defines dense regions that prompt the mass-candidates. Moreover, we investigate the possible role of prompting in mammogram sequences (after registering them) for reducing the number of false positives via temporal comparison.

The method is an extension of previous work on image registration [109, 111] and the mammogram filtering for microcalcification detection presented in Chapter 3. Next we discuss how anisotropic diffusion can assist in general mass detection and feature analysis. The following sections focus more on the method designed for temporal mass detection; they illustrate the concepts behind the image registration, image diffusion and texture analysis that are used for prompting the mass candidates. Section 5.4 presents the results of our method applied to a set of clinically assessed mammograms. The chapter concludes with final remarks and improvements that will be the basis for future work.

## 5.2 Diffusion and Masses


In the previous Chapter, we have described an automated anisotropic diffusion filter to enhance microcalcifications in digital mammography. The promising results encouraged us to perform more tests on images that contained masses to see if a similar filter can emphasise the relevant mass characteristics and help in distinguishing them from parenchymal tissue.

While microcalcifications have distinct characteristics in both intensity and SMF images (thin tower shape, see Figure 31), their enhancement can be very well tuned for their shape and intensity/interesting height values. Most important, a microcalcification filter must preserve calcium and clean the small noise. When referring to masses, filter design must address regions with different characteristics:

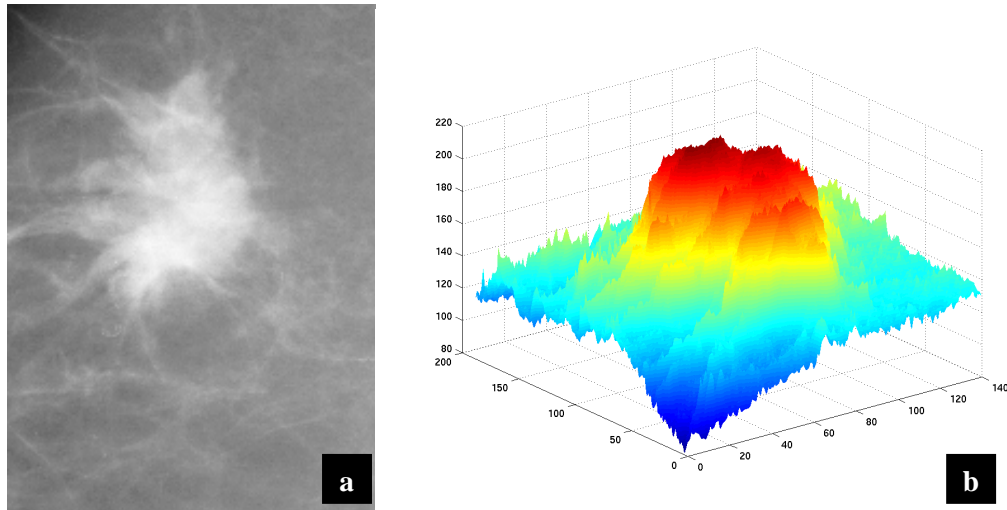
- masses are larger;

- masses have similar intensity/interesting height to normal mammographic, i.e. parenchymal tissue, the main source of FP in mass detection;
- noise (typically small structures) does not influence significantly the number of FP;


Hence, using the anisotropic diffusion filter presented in Chapter 3 with the same characteristics would be expected to be unsuitable for mass detection. The scale of the filter,  $\sigma$ , would need to be enlarged when searching for larger structures (missing microcalcifications is irrelevant in the search for central masses, although their eventual association is important for characterisation); the number of iterations must be increased, since a mild diffusion would enhance similarly the edges of normal tissue and we search only for well-defined self-contained dense regions. The parameters of the new filter force it to highlight larger regions of dense content with strong boundaries (masses have better defined boundaries than normal tissue). The anisotropic diffusion filter targets benign and malignant tumours with a central mass. Figure 80 and Figure 81 show a couple of examples of mammogram samples where masses are enhanced using the adapted automated anisotropic diffusion filter.




**Figure 80:** The plot of a benign mass; (a) a mammogram sample containing a benign tumour with well-defined margins; (b) the 3D plot of image (a) where the tumour appears as a high hill surrounded by several smaller structures of normal tissue and noise.




**Figure 81:** The plot of a malignant mass; (a) a mammogram sample containing a malignant spiculated tumour; (b) the 3D plot of image (a) where the tumour appears as a structures of high ridges descending along the spicules and surrounded by several smaller structures of normal tissue and noise.




**Figure 82:** Diffusing a benign mass 1; (a) the mammogram sample containing a benign tumour in Figure 80 after diffusion with  $t=10$ ,  $k=15$ ,  $\sigma=0.8$  (small number of iterations and high contrast); (b) the 3D plot of the diffused image in (a); (c) the SSD image between the original not-blurred imaged and the diffused one – the latter image is cleaned by its high-frequency component, but the inner surface of the mass is also diffuse because of its iso-density.



**Figure 83:** Diffusing a benign mass 2; (a) the mammogram sample containing a benign tumour in Figure 80 after diffusion with  $t=400$ ,  $k=5$ ,  $\sigma=0.8$  (large number of iterations and small contrast); (b) the 3D plot of the diffused image in (a) with flat background; (c) the SSD image between the original not-blurred imaged and the diffused one – the latter image has a ‘clean’ background, since almost everything else, but large dense regions has been removed; at this high number of iterations, the inner surface of the benign lesion (which is roughly uniform) is still diffused.



**Figure 84:** Diffusing a malignant mass 1; (a) the mammogram sample containing a malignant spiculated tumour in Figure 81 after diffusion with  $t=8$ ,  $k=8$ ,  $\sigma=0.8$  (small number of iterations and high contrast); (b) the 3D plot of the diffused image in (a); (c) the SSD image between the original not-blurred imaged and the diffused one – the latter image is cleaned by its high-frequency component, while the complex geometry of the malignant mass is seen as a combination of edges/ridges and is not diffused.



**Figure 85:** Diffusing a malignant mass 2; (a) the mammogram sample containing a malignant spiculated tumour in Figure 81 after diffusion with  $t=40$ ,  $k=5$ ,  $\sigma=0.8$  (large number of iterations and small contrast); (b) the 3D plot of the diffused image in (a) with mainly flat background; (c) the SSD image between the original not-blurred imaged and the diffused one – the high-frequency map is similar to the one in Figure 84, since the inner surface of the malignant lesion is irregular and the ridges are perceived as edges.

### 5.2.1 Defining Diffusion Parameters for Mass Detection


From Figure 82 to Figure 85 we can deduce that low contrast and a large number of iterations (as one choice in the set of parameters) yield well-preserved and clearly enhanced edges, an important characteristic in mass classification, for which the roughness of edges is an important feature. The SSD images refer to the difference image between the original (un-blurred) mammogram and the diffused version. We used the minimisation of the sum-of-square-differences (SSD) to compute a weighting factor used to scale the diffused image before subtraction. The background becomes flatter as well, but the process is highly time consuming. Furthermore, the 3D plot of the diffused image doesn't necessarily show a great improvement from the process using a second choice of diffusion parameters, since the evolution of diffusion is slow with the number of iteration for a small  $k$ . Hence, in our further experiments we prefer to use the same automated setting of parameter  $k$  as described in Chapter 4, increasing the number of iterations  $t$  and the scale  $\sigma$ .

The difference in the shape of the enhanced margins of the benign versus malignant masses becomes more pronounced: benign masses have regular shaped margins, while malignant tumours have irregular jagged boundaries. From the 3D plot, also visible to a smaller extent in the 2D images, one can notice a relative iso-density inside the benign tumour, while the malignant lesion has several peaks or ridges within its boundary. One observation is that benign masses diffuse differently than malignant masses, due to their internal smoothness versus the relative complexity in the case of malignancy. This characteristic will be further exploited in Section 5.6. Also note the resemblance between the SSD images in Figure 82 to Figure 85 and the CLS maps shown in the previous chapters (if we ignore the surface of the benign mass). Since microcalcifications are not the subject of this Chapter, high-frequency components can be eliminated prior to mass detection.

We ran the diffusion algorithm on a set of 20 mammogram samples from the MIAS database. Half of the mammograms were benign, the other half malignant. The database annotation was considered to be the ground truth in looking for features that differentiate cancerous masses from benign tumours. Figure 86 and Figure 87 show a couple of examples of benign and malignant masses presenting the same diffusion in the SSD image characteristics as described above. We used the statistical analysis presented in Chapter 4 to compute the contrast factor  $k$ , but we used the value  $3*k$  for diffusing the mass images. While the value  $k$  was sufficient in the previous application (to clean images for microcalcification detection), we need a higher contrast value to diffuse masses (to enhance their boundaries). The number of iterations  $t$  was also increased from 5 to 20, for a stronger effect. There were also exceptions and Figure 88 shows a benign mass and a malignant lesion that have associated SSD images which do not follow the same rule. The observation is based on visually inspecting the images. From the total of 20 tested mammograms, 3 did not follow the diffusion rule; 1 contained a malignant tumour, while the other 2 showed benign lesions. Therefore, the error of the proposed classification feature went up to 15%. Yet, the *diffusion feature* (a measure of the ‘amount’ of diffusion within the mass boundaries), with classification TP ratio of 85%, can be


combined with a set of other features (as commonly done in literature) for a more reliable classification.

The following sections will refer to the method we propose for mass prompting in temporal mammograms. After the presentation of prompting results, Section 5.6 will expand on feature detection and mass classification.




**Figure 86:** The *diffusion feature* on benign masses; (a) and (c) the original mammogram samples showing benign lesions; (b) and (d) the SSD images corresponding to (a) and (c), respectively; both show that diffusion is allowed within the mass area.





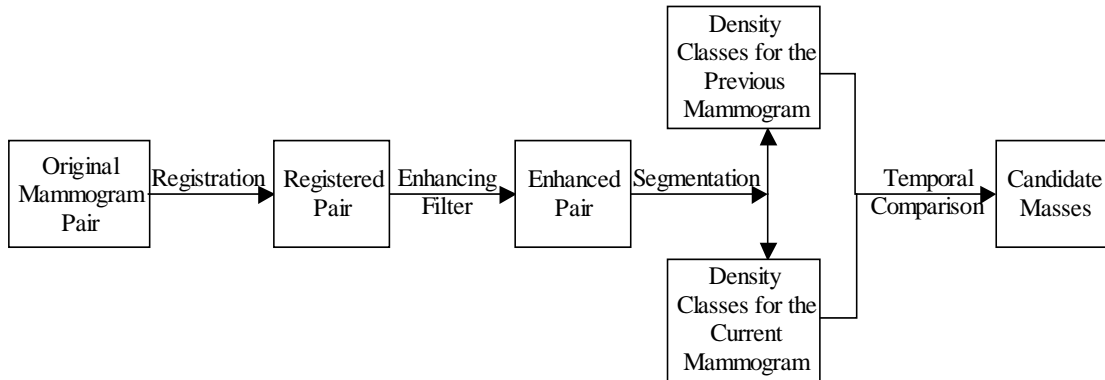
**Figure 87:** The *diffusion feature* on malignant masses; (a) and (c) the original mammogram samples showing malignant lesions; (b) and (d) the SSD images corresponding to (a) and (c), respectively; only high frequency structures are diffused within the mass area.



**Figure 88:** The *diffusion feature* on exceptions; (a) an original mammogram samples showing a benign lesions; (b) the SSD images corresponding to (a) where the diffusion was mainly inhibited within the mass area; (c) an original mammogram samples showing a malignant lesions; (d) the SSD image corresponding to (b) where diffusion is allowed within the mass area.

### 5.3 Method

The method we propose comprises two steps that pre-process the original mammogram prior to the detection of dense regions: mammogram registration and anisotropic diffusion of the registered mammograms. The basic assumption in our work is that masses appear as dense regions in mammograms. We believe that temporal comparison of automatically generated “prompts” in mammogram sequences can reduce the number of false positives. Figure 89 shows the algorithm flow we propose for mass prompting.



**Figure 89:** The diagram of the algorithm proposed for temporal mass prompting.


### 5.3.1 Mammogram Registration

First, the temporal pair is registered using a mammogram registration method developed by Marias *et al.* [109, 111]. In summary, this is a three-stage registration algorithm:

- Initially, the images are aligned based on the boundary. An algorithm that automatically detects 3 anatomically significant points in the outline of both mammograms does this. A thin-plate spline interpolation is used to calculate the image transformation that aligns the boundaries of the two mammograms.
- Using a wavelet-analysis segmentation algorithm [111] we define internal regions of dense tissue in each mammogram. The boundary transformation, together with scale and area information, is used to match the segmented internal structures.
- Subsequently, a regularised approximation scheme is used to calculate the refined transformation. This accounts for possible inaccuracies in the selection of the internal landmarks.

Registration is performed in order to facilitate the comparison between any temporal pair of mammograms emanating from successive screening visits. Figure 90 shows an example of mammogram registration where a large deformation is required (notice the displacement of the film edge) in order to geometrically align the images. The main aspect of registration for this specific work is to aid mass detection by comparing “suspicious” regions in the registered

mammogram sequence, where false positives can be reduced by visually inspecting the correspondence of “temporal prompts”. This is further discussed in the result section.



**Figure 90:** Mammogram registration is performed as a pre-processing step in order to facilitate the comparison between the temporal mammogram pair for mass detection. (a) and (b) comprise the original mammogram pair, while (c) is the registered mammogram (b) to (a).

Appendix B offers more detailed insight into the mammogram registration method used by our mass-prompting algorithm.

### 5.3.2 Anisotropic Diffusion of the Temporal Mammograms

In order to detect only the most important features of the mammogram, the images are processed using an adaptive anisotropic diffusion-based filter, which enhances the suspicious features in mammograms [100, 102, 103]. The parameters of the filter are computed from a statistical analysis of the image gradient (cf. Chapter 4) and the mammogram is blurred anisotropically. The new version of the image will generally be blurred, while the prominent areas will be enhanced.

### 5.3.3 Tissue Classification Based on Texture Analysis

Finally, a texture-based classifier is used to segment the image into different tissue types and the “denser” classes are used to automatically prompt to “suspicious” regions in each mammogram of the sequence.

The image is divided in small patches (10x10 pixels for 300 micron images). Since small masses are between 3 and 15 mm diameter [97], they are likely to occupy over 10 pixels in a 300-micron digital image. For this reason the choice of a 10x10 pixel window is a reasonable compromise between speed and mass detectability.

In each patch, normalised second order statistics are calculated. For example the correlation measure:

$$f = \frac{\sum_i \sum_j ij p(i, j) - \mu_x \mu_y}{\sigma_x \sigma_y}, \quad (73)$$

where

$$\mu_x = \sum_i i p_x(i), \quad p_x(i) = \sum_j p(i, j), \quad \sigma_x^2 = \sum_i (i - \mu_x)^2 p_x(i), \quad (74)$$

$$\mu_y = \sum_j j p_y(j), \quad p_y(j) = \sum_i p(i, j), \quad \sigma_y^2 = \sum_j (j - \mu_y)^2 p_y(j), \quad (75)$$

and  $p(i, j) = P(i, j) / R$  is the normalised joint probability of the pixels  $i$  and  $j$ ;  $R$  is the number of co-occurrences (pixel intensities transitions) [57, 112]. For each image patch  $i$ , a texture vector  $T_i$  is calculated from (73) and all the vectors are classified in a desired number of classes using hierarchical clustering. The texture classification is extended to temporal pairs, since we are interesting in detection only. In our work we detect 4 classes corresponding to:

- Very dense tissue. This class usually includes the pectoral muscle and very dense regions of the breast parenchyma. The separation of masses versus pectoral muscle can


be done either by visual inspection or on a geographical basis since the pectoral muscle consistently appears on the top left (right) part of a left (right) mammogram.

- Dense tissue. This class includes all the remaining parts of the dense parenchymal cone (fibrotic stromal tissue and glandular tissue).
- Fatty tissue. This class effectively represents the fatty background of the mammogram or according to Wolfe classification [174], “normal” involuted breast patterns.
- Fatty breast edge. This is the last segmented class and it’s a homogenous, low-intensity region near the breast edge.


A candidate mass is expected to appear either as a “very dense” or “dense” tissue region according to the above classification depending on the presence or not of the pectoral muscle and on the local density variations. Density variations around an iso-dense contour in the denser classes indicate the presence of abnormalities, namely tumours. These classes are explained in more detail in the result session.

## 5.4 Results


Figure 91 and Figure 92 show typical results on both the original and the diffused MLO mammogram pair. The main effect of diffusing the images is the significant reduction of the variations in the denser regions, which in our work are considered the “suspicious” regions for mass detection. While Figure 91 prompts numerous density variations in the image and therefore a large number of candidate masses, Figure 92 shows more homogeneous regions and reduces dramatically the number of possible tumours. For that reason, as illustrated in Figure 93, temporal mammogram comparison can be further facilitated by prompting the clinician in the denser regions of the mammogram pair. A second example is shown in Figure 94, Figure 95 and Figure 96 on a CC mammogram pair.




**Figure 91:** Texture classification 1 of the registered MLO mammogram pair into the classes A, B, C and D described in the previous section. The top row shows the firstly taken mammogram, while the bottom row shows the most recent mammogram. Both mammograms are registered, but not enhanced using anisotropic diffusion before the texture classification is applied.



**Figure 92:** Texture classifications 1 of the diffused and registered MLO mammogram pair into the classes A, B, C and D. The top row shows the firstly taken mammogram, while the bottom row shows the most recent mammogram. Both mammograms are registered and enhanced using anisotropic diffusion before the texture classification is applied.




**Figure 93:** The detected “suspicious” regions 1 (from the diffused pair) superimposed in the original MLO mammogram. The current mammogram (b) prompts a real tumour with no correspondent in the earlier mammogram (a).




**Figure 94:** Texture classification 2 of the registered mammogram pair into the classes A, B, C and D in a CC mammogram pair. The top row shows the firstly taken mammogram, while the bottom row shows the most recent mammogram. Both mammograms are registered, but not enhanced using anisotropic diffusion before the texture classification is applied.





**Figure 95:** Texture classifications 2 of the diffused and registered CC mammogram pair into the classes A, B, C and D. The top row shows the firstly taken mammogram, while the bottom row shows the most recent mammogram. Both mammograms are registered and enhanced using anisotropic diffusion before the texture classification is applied.



**Figure 96:** The detected “suspicious” regions 2 (from the diffused pair) superimposed in the original CC mammogram. The current mammogram (b) prompts a real tumour with no correspondent in the earlier mammogram (a).

Table 5 shows the preliminary results we have obtained in mass detection (true positives and false positives for the “suspicious” regions detected) in 15 temporal pairs. An experienced radiologist annotated the masses, which we used as ground truth. Note the improvement in detection by including anisotropic diffusion as a pre-processing step. In addition, the same table shows the improvement in the mass detection rate by visually inspecting the generated prompts in the mammogram sequence. Only candidate masses that have been detected in the current mammogram and not in the previous one, or which have significantly evolved between the screening sessions could prompt a tumour.

The detection of mass candidates is based only on the visual inspection of the temporal prompts using a light-box approach. The implementation of the temporal prompts comparison as an automatic step of the mass detection algorithm must take into account several factors describing the evolution of dense areas of the breast over time. Amongst them, the quantification of area and intensity changes in and around denser areas of the breast (to assess the appearance or evolution of a tumour and remove FP in involuted areas) and the distribution of very dense and dense areas over less dense areas of the breast (to remove FP in dense breasts and increase sensitivity towards the detection of faint tumours). We aim to incorporate this comparison as an automatic step in future work.

The number of FP is excessively high if no enhancement or temporal comparison is involved. It decreases dramatically when the enhancement pre-processing filter is used on single mammograms without temporal comparison (from 3.93 to 0.86 FP/image). Using the typical clinical procedure of comparing temporal mammograms (without image enhancement), the FP number is also lowered by over 71%. Using both image enhancement and temporal comparison of registered mammograms we obtain 0.4 FP/image.

**Table 5:** True positives and false positives in 15 pairs of mammograms (a mass has previously been diagnosed in each pair).

Temporal comparison of “prompts” (visual)	Anisotropic diffusion preprocessing	True positives	False positives
No	No	15	59
No	Yes	15	13
Yes	No	15	17
Yes	Yes	15	6

Note that only mammograms where an abnormality is present have been considered in validating the method. Finally, the potential to facilitate temporal comparison for mass detection has to be tested within the screening environment.

## 5.5 Conclusion


We have presented an automatic method to detect masses in digital mammography, which prompts the clinician to candidate masses in temporal mammograms. The first pre-processing step is registration of temporal mammograms, which aligns the main landmarks in both images of the pair for a better comparison of the temporal changes in mammograms. Subsequently, we blur the mammograms with an adaptive non-parametric anisotropic diffusion-based filter. This scale-space process enhances the “suspicious” regions in the breast, while halving the number of false positives. A texture-based classifier performs the segmentation of the image into classes of different densities. The number of candidate masses drops significantly after diffusing the images. Still, the FP fraction is unacceptably high, since no association between the candidates in the two images has been implemented yet. However, the visual inspection of prompts in mammogram sequences (Table 5) indicates that a further improvement can be achieved. Future work will concentrate in developing an automatic method that would enable the temporal comparison of prompts detected in any mammogram of a particular registered patient-sequence. This could reduce the FP problem that appears in most of the suggested mass-detection algorithms.

One other future application of the method presented here is the automatic identification and characterisation of the important tissue groups in mammograms and subsequent classification of the dense tissue according to the BI-RADS criteria [127]. The American College of Radiologists suggests that breast composition should be reported in all patients using the BI-RADS classification: I describes an almost entirely fat breast, II breasts with scattered fibroglandular densities, III heterogeneously dense breasts and IV extremely dense breasts. Little work has been done on the identification and characterization of significant tissue categories prior to feature detection. Such an analysis framework could have a significant impact on mammographic examination and breast cancer epidemiology since tissue content has been related to the risk of developing breast cancer.

More mammogram sequences are required in order to test the method in a sufficiently large database of temporal mammograms. It is also necessary to consider methods to validate the value of this approach for temporal mass detection; parallel prompting of the current and previous mammograms in the screening environment. We will further investigate the results of this method on images in Standard Mammogram Form [65].


## 5.6 Future Work on Mass Characterisation Features

Further to the *diffusion feature*, which aims to detect the ‘amount’ of anisotropic diffusion ‘permitted’ within the mass area (a mathematical model of the feature must still be developed), this section introduces work on what we call the *uniformity feature*. This new feature searches the uniformity versus roughness of the boundary of masses. We have observed that benign masses, as described in literature [127], have smoother surfaces, while malignant tumours have jagged or more complex 3D plots. In the following examples, we use the vector flow of the image Gaussian derivative (63),(64). Figure 97 and Figure 98 illustrate the concept behind the *uniformity feature*.




**Figure 97:** The vector flow of a benign mass; (a) the vector flow of the original mammogram sample; excepting the boundaries of the mass, the vectors flow chaotically; (b) the vector flow of the diffused mammogram; note the radial pattern of vectors pointing towards the centre of the mass.

We used the same 20 case database, as in Section 6.2 to perform tests on the *uniformity feature*. Figure 99 and Figure 100 illustrate more examples of vector flows on malignant and benign lesions. Only 50% of the tested cases showed consistent results regarding smoother flows along the boundaries or a certain radial/chaotic pattern within the mass area, when the diffusion was performed as described in Section 6.2. The large variety of benign masses and the presence of CLS create an outburst of vector flows within the mass and interfere with the general radial pattern. Repeating the experiments with a more carefully chosen set of parameters (set manually for each individual mass), we observed promising results in 75% of the tested mammogram samples.




**Figure 98:** The vector flow of a malignant mass; (a) the vector flow of the original mammogram sample; excepting the boundaries of the mass, the vectors flow chaotically; (b) the vector flow of the diffused mammogram; although the vectors point towards the inside of the lesion, there is no definite centre of the mass; this ‘multi-focality’ is due to the rough surface of the cancerous mass.



**Figure 99:** The *uniformity feature* on benign masses; (a) the vector flow of the original mammogram sample showing a benign lesion; (b) the vector flow of the diffused images corresponding to (a) showing a smooth transition along the edge of the mass;





**Figure 100:** The *uniformity feature* on malignant masses; (a) the vector flow of the original mammogram sample showing a malignant lesions; (b) the vector flow of the diffused images corresponding to (a) where the transitions along the edge are still rapid;

We are sufficiently encouraged by the initial results to have begun the implementation of a mathematical model of the *uniformity feature* through vector flow analysis. In future work it can be combined with the *diffusion feature* and other measures of mass roughness and contrast in building a discriminating mass classification criterion. A second order Gaussian derivative should offer extra information about the smoothness of the inner surface of a mass. We should expect that removing the CLS would, as in the case of microcalcifications, again improve the results. Although we would lose some of the spiculations associated with malignant lesions, the main shape of the mass would not be altered. The number and variety of tested cases must be increased.

## CHAPTER 6

### 6 Discussion And Future Work

*Any man who reads too much and  
uses his own brain too little falls into  
lazy habits of thinking.*

Albert Einstein

We presented in this thesis a method for feature detection in digital mammography. The detection of microcalcification clusters and breast masses have been investigated on both Standard Mammogram Form (SMF) images (c.f. Chapter 2) and original intensity images. The following section summarises the work presented in the previous chapters, along with the encountered limitations and some concluding remarks. Then we will present suggestions for future work and, where possible, some initial results on improving feature detection in mammographic image analysis.

#### 6.1 Summary and Discussion

##### 6.1.1 Mammographic Image Filtering

The thesis commences with a filter model for SMF images to ‘clean’ the noisy appearance of  $h_{int}$  images and enhance the structures of interest. Initially, the filter was designed for the



detection of microcalcification (c.f. Chapter 3), but later in Chapter 5 the same filter is used to enhance mammographic masses as well as in a temporal detection method.

The need for a pre-processing filter comes from the poor signal-to-noise ratio (SNR) that SMF images have and the difficulty that arises from this in spotting small salts of calcium in a mass of textured ambiguity. Furthermore, it is desirable to have enhanced structures of interest prior to segmentation for a simpler differentiation between microcalcifications and mammographic background.

The use of partial differential equations (PDE) and in particular anisotropic diffusion led to the following observations:

- The parametric format of the proposed filter allows the user to obtain very reliable results given the right choice of parameters. The appearance of the image is improved, but a non-parametric approach is desirable to make the algorithm robust;
- The filter is stable in time and offers a good alternative to filters previously described in the literature;
- Quantum mottle is smoothed by the filter and errors are minimised, without disrupting microcalcifications;
- Microcalcifications can be depicted according to their physical appearance as thin towers in a smoothed background;
- The main source of false positives (FP) are shot-noise, caused by dust or hair on the film, and curvilinear structures (CLS), long thin bright structures in a mammogram corresponding to blood and lymph vessels, ligaments, ducts or tumour spiculations.

### **6.1.2 Complex Pre-processing**

To eliminate some of the difficulties outlined above, a complex pre-processing step was introduced prior to image segmentation to reduce the number of FP in the detection algorithm. Our aim has been to eliminate shot-noise and CLS from the mammographic image, as well as

to filter the mammogram with an adaptive filter suited to the detection of microcalcifications. This step was designed to improve the robustness of the algorithm and to obtain consistent results independent of the user's experience and dexterity.

By removing glare in the  $h_{int}$  generation process, shot-noise can be marked on a binary map that can subsequently be used for its removal (c.f. Chapter 4). Information gathered from local energy (LE) and phase congruency (PC) is used to detect and remove CLS from mammograms (c.f. Chapter 4). Using the filter design described in Chapter 3 and the statistical analysis of the local image characteristics, the parameters of the diffusion process can be computed automatically and the pre-processing step becomes fully automated. A few remarks are necessary here:

- The number of FP is significantly lower after pre-processing;
- CLS removal, being scale dependent, tends to eliminate some microcalcification candidates. To avoid this, we only use a subset of the scales at which we detect CLS. However, as a result, some disruptions appear in the CLS map and they can leave small bright dots on the processed image, which may be interpreted as FP by the segmentation algorithm. A better solution could be provided by an alternative CLS removal algorithm;
- The computation of the diffusion parameters by statistical analysis is robust and consistent over the database. The same computation was used both for intensity and SMF images with similar results, although the image characteristics are different. The value calculated here is further used to automate the segmentation step. We do not alter the 'natural' contrast of the images we process. Images with very high contrast will be smoothed more and therefore a prior contrast enhancement of the analysed images (as found in the literature) will lead to deteriorating the values of the parameters. The results will be similar, since the method is adapted to the image characteristics, but a different computation of the parameters used to set the minimal perceivable contrast (70) may give optimal results.

### 6.1.3 Human Vision Based Segmentation

The final step is the segmentation of microcalcification clusters. The aim of this stage is to produce a reliable detection map that indicates the location of microcalcification clusters by maximising the number of individual calcium salts detected and minimising the impact of FP. A contrast measure is introduced that is derived from a model of the human vision detection mechanism (c.f. Chapter 4). This new measure is locally adaptive in each processed image and uses the results of the pre-processing step (statistical analysis) to become non-parametric. Combining the pre-processing and the segmentation steps in building a microcalcification detection method, we can conclude the following:

- Our method is non-parametric and fully automated and gives similar good results on both SMF and original intensity images;
- The variation of parameters showed consistent results with small differences in the detected ratios when parameters are varied around the optimal value;
- The receiver operating characteristic (ROC) curve analysis demonstrated that adding adaptive contrast segmentation based on human foveal processing significantly enhanced the detection of microcalcifications;
- The method was tested on mammographic samples for faster processing and simpler validation; a faster implementation would be necessary for better clinical manipulation;
- The detection rate reaches 98% TP fraction at 0.1 FP/image and converges to 100% for each ROC curve when the number of FP/image reaches 1.1.

### 6.1.4 Temporal Mass Prompting

The final application presented was mass prompting by temporal comparison of feature enhanced mammogram pairs (c.f. Chapter 5). It is an expansion to breast masses of the

mammographic image filter described above, which was originally designed for microcalcifications. The algorithm enhances structures of interest and blurs the background. Alteration of the diffusion parameters lead to good results on mammograms presenting masses. The goal of this application was to assist the clinician in interpreting images and establishing reliable and early diagnosis of pathology. The prompting of masses comes as a natural completion of the microcalcification detection in the early detection of cancers.

The algorithm consists of three steps that conclude with mass candidates being prompted in the pair of temporal mammograms: image registration for accurate temporal comparison; image filtering for feature enhancement (an adaptive anisotropic diffusion-based filter); and mammogram segmentation that defines dense regions that prompt the mass-candidates. Note the assumption that masses appear denser in mammograms than other tissue. From our experimental results, we were able to draw the following observations:

- The adaptive anisotropic-diffusion filter is robust when used for mass enhancement in digital mammography; some of the diffusion parameters must be stronger than for microcalcification enhancement (size and time), but the preceding statistical analysis used to remove the user-set parameters from the algorithm is used to automatically compute the contrast;
- The algorithm correctly detected all the masses in the tested images; the number of candidate masses drops significantly after diffusing the images, but the number of FP remains high;
- Registration facilitates further reduction of FP in the temporal comparison process through a more accurate matching between the mass candidates depicted after texture analysis;
- A more refined texture analysis would give a better segmentation of the candidate masses; the algorithm should also be further developed to incorporate automated mass matching and FP discarding in the temporal sequence.

Using a similar computation as in the image pre-processing, we can depict some breast tumour related features. Once the tumour is detected, a number of features can be signalled (c.f. Chapter 5). We have noted the following characteristics of benign and malignant masses:

- Benign masses diffuse differently than malignant masses. This is because benign masses are smoother within their boundaries, while malignant tumours have a more complex internal geometry;
- The vector flow of the image Gaussian derivative prompts a measure of the uniformity versus roughness of boundaries in benign and respectively malignant masses through radial or chaotic patterns within the masses and along their boundaries;

It still remains to incorporate a mathematical model of these two features into a feature analysis algorithm.

### 6.1.5 Conclusion

This thesis presented the results of research to develop better detection tools to assist radiologists in their task of detecting breast cancer at the earliest possible stage. The need for efficiency of screening programmes has become obvious [4, 29, 164] and the necessity of computer aided diagnosis (CAD) systems to assist clinicians in evaluating such large amounts of information is increasingly important in contemporary society. The improvement of CAD systems is expected to improve feature detection and classification in digital mammography, especially with the introduction of soft-copy environments in hospitals around the world. At the present time, there are very few systems ready for clinical use. This is a primary motivation for the work presented in this thesis.

We have presented a method to detect microcalcifications in X-ray mammography using the scale-space properties of anisotropic diffusion and a model of the characteristics of the human visual system. The method consists of three stages: image pre-processing for eliminating some of the sources of FP and computing the filter parameters; image enhancement

for emphasising the structures of interest and blurring the background; and image segmentation for the detection of microcalcifications. The algorithm was designed for SMF images, normalised representations of the breast, but similar results are obtained on intensity images. It is fully automated and non-parametric.

We have further investigated the prompting of masses in pairs of temporal mammograms. This is an extension of the automated filter designed for the detection of microcalcifications and previous work done in mammographic image registration [109, 111]. The characterisation of masses is also considered at this stage.

The method shows promising results both for microcalcifications and masses. While the detection rates are very high, minimising the number of FP is the main concern for future work. The next section will expand on the main ideas to be further explored to this end. First, a few more general remarks on the actual state and limitations of the detection algorithm will be discussed.

The large variability of microcalcification size and brightness combined with poor control in imaging conditions when dealing with such small structures makes the detection of microcalcifications a difficult task. Some measure of the calcium X-ray attenuation combined with image normalisation would be useful in the detection process. The CLS removal step should also be refined for more accurate results. A faster implementation of the detection algorithm would be necessary for better clinical manipulation of full mammograms. A cluster classification step is desirable for a complex detection and classification of microcalcification clusters method. Possible features to be used by such a detection-classification system include size, shape, distribution, density and orientation.

As for prompting masses, more experiments must be performed on larger databases. An automated implementation of the image-matching algorithm must be developed, since visual inspection is performed at present. A more sophisticated segmentation should depict the boundary of the detected mass more accurately for the subsequent classification. Only preliminary work on such a boundary detection system has been investigated in this thesis.

In the next section we will concentrate on some possibilities for future work for the immediate improvement of the detection algorithm for microcalcification clusters in X-ray mammography.

## 6.2 Future Work

The most obvious improvement in the detection of microcalcifications is the elimination of FP. The main observed source of FP remains the presence of CLS. These have high attenuation and their thin appearance may make their disruptions or overlapping look like microcalcifications. The first idea presented here for future research is the use of an expectation maximisation algorithm to differentiate between the detected salts of calcium and the FP corresponding to CLS. An alternative method for the detection and removal of CLS is also discussed in Section 4.4. In Chapter 3 we presented an edge-enhancing anisotropic diffusion approach. Coherence enhancing is a modification of the same model smoothing along flow-like structures using a coherence-enhancing anisotropic diffusion filter [172]. CLS can be enhanced before detection and removal; the enhanced CLS are expected to be contiguous and their elimination more accurate. Finally, both CLS and parenchymal tissue may produce FP, but they are laid on a basis that must be much larger or longer than that of a microcalcification. An iso-level segmentation algorithm is proposed for differentiating between calcifications and other breast tissue for improved detection results. These algorithms are presented below with some initial tests on mammograms.

### 6.2.1 Expectation Maximisation

Many of the FP prompted by the microcalcification detection algorithm result from artefacts of the CLS removal step that leaves some disrupted bright blobs in the mammogram. Therefore a simple differentiation between the candidates would aim to split them into two categories:

- Putative microcalcifications (including FP that do not occur from CLS);
- FP caused by the inadequate removal of CLS.

An estimation maximisation (EM) [10, 26, 122] algorithm was implemented to try to solve this problem. Here are some theoretical foundations of the algorithm followed by some first results.

### 6.2.1.1 Theory

The EM algorithm is one of the best-known methods to estimate maximum likelihood for problems of incomplete data with large applications from image and signal processing to statistics and applied science. The maximum likelihood parameters are computed iteratively using an initial estimation. At each iteration, there is:

- An expectation step: starting from the observed variables and the current parameters, the unobserved variables are estimated;
- A maximisation step: assuming that the expectation is correct, the parameters are re-evaluated to maximise likelihood.

The algorithm converges to a steady state once a local maximum is reached. For the current problem of microcalcification detection we can use the EM algorithm to solve a mixture estimation problem. We will express the distribution function as a sum of  $k$  Gaussians:

$$f(x) = \sum_{n=1}^k \pi_n \cdot \frac{1}{\sqrt{2\pi}\sigma_n} \cdot \exp\left(-\frac{(x - \mu_n)^2}{2\sigma_n^2}\right) = \sum_{n=1}^k \pi_n \cdot G(x; \mu_n; \sigma_n), \quad (76)$$

where

$$\sum_{n=1}^k \pi_n = 1 \quad (77)$$

$$\mu_n = \frac{1}{N} \sum_{i=1}^N x_i \quad (78)$$

$$\sigma_n^2 = \frac{1}{N} \sum_{i=1}^N (x_i - \mu_n)^2 \quad (79)$$



The parameters of the distributions are computed in the expectation step. Let  $X = (x_1, x_2, \dots, x_N)$  be the sequence of observations from a mixture of  $k$  Gaussians (as above) and  $\theta = \{\pi_1, \pi_2, \dots, \pi_N, \sigma_1, \sigma_2, \dots, \sigma_n, \mu_N, \mu_2, \dots, \mu_N\}$  the parameters that must be estimated from  $X$ . The values of the parameters  $\pi, \sigma$  and  $\mu$  will be updated with each iteration until the algorithm converges. The likelihood maximisation becomes that of maximising:

$$\ln P(X | \theta) = \sum_{i=1}^N \ln P(x_i | \theta) \quad (80)$$

Note that we are trying to determine the ‘belongingness’ of a pixel to one of the  $n$  classes considered and denominated by  $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ . In the present problem we have 3 classes: microcalcification, CLS and ‘uncertainty’. Using the parameters  $\theta$  we can compute  $P(x_i \in \Gamma_n)$ .

Using Bayes’s law:  $P(A | B) \cdot P(B) = P(B / A) \cdot P(A)$ , it results that:

$$P(x_i \in \Gamma_n | x_i) = \frac{P(x_i | x_i \in \Gamma_n) \cdot P(x_i \in \Gamma_n)}{P(x_i)}, \quad (81)$$

where

$$P(x_i) = f(x_i) \quad (82)$$

$$P(x_i \in \Gamma_n) = \pi_n \quad (83)$$

$$P(x_i | x_i \in \Gamma_n) = G(x_i; \mu_n; \sigma_n) \quad (84)$$

Therefore:

$$P(x_i \in \Gamma_n) = \frac{\pi_n \cdot G(x_i; \mu_n; \sigma_n)}{\sum_j \pi_j \cdot G(x_i; \mu_j; \sigma_j)} \quad (85)$$

The maximisation step will update the values of the variables as follows:

$$\mu_n = \frac{\sum_i P(x_i \in \Gamma_n) \cdot x_i}{\sum_i P(x_i \in \Gamma_n)} \quad (86)$$

$$\sigma_n^2 = \frac{\sum_i P(x_i \in \Gamma_n) \cdot (x_i - \mu_n)^2}{\sum_i P(x_i \in \Gamma_n)} \quad (87)$$

$$\pi_n = \frac{\sum_i P(x_i \in \Gamma_n)}{\sum_{i,n} P(x_i \in \Gamma_n)} \quad (88)$$

The EM algorithm is a simple accessible tool to solve maximum likelihood problems, when a good estimate of the initial parameters is available. It offers a numerically stable solution that is easy to implement. However, it may converge very slowly and additional constraints or assumptions may need to be introduced to speed up the process.

### 6.2.1.2 Initial Results

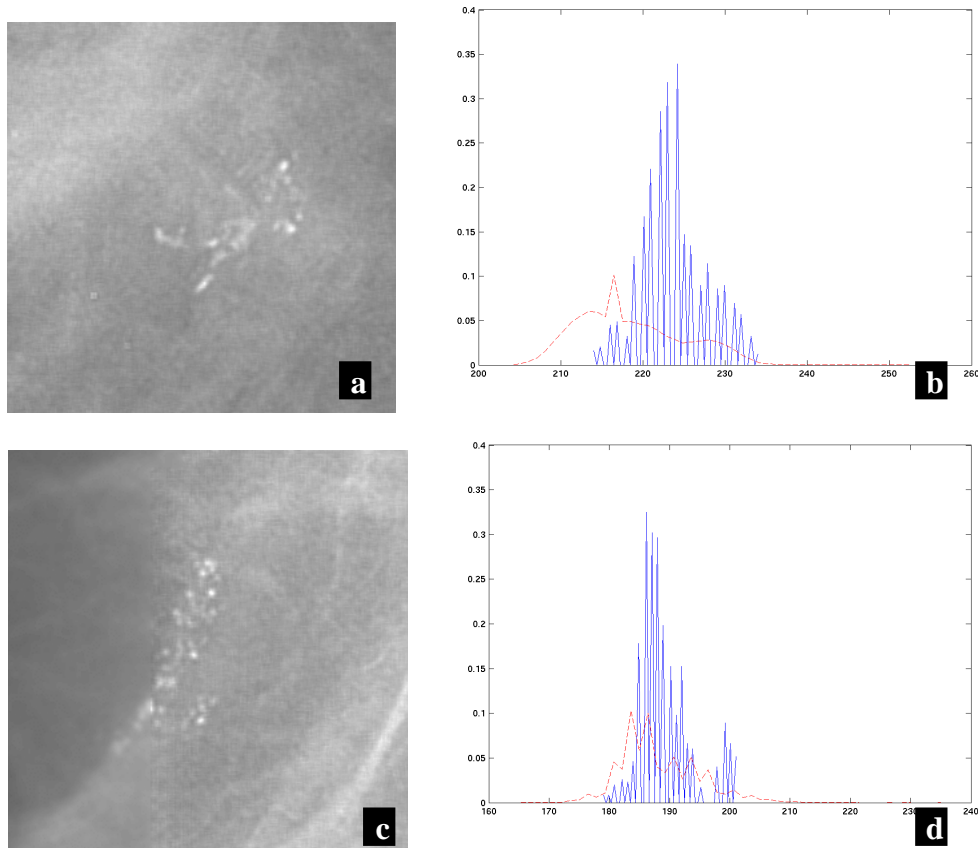
For testing the result of the EM implementation on the set of processed mammograms, we first produced a microcalcification and a CLS map. In both cases the detected microcalcifications and CLS pixels are marked on a black background with the value of the original SMF/intensity image corresponding to the putative pixel. The values of variance and mean are computed from the approximation with a Gaussian of the histogram of the non-zero elements of the two images (microcalcification map and CLS map). The histograms of the two structures overlap, as seen in

Figure 101. Three classes for the mixture estimation are defined:

- Calcifications: where there is a non-zero pixel in the microcalcification map and a zero pixel in the CLS map at the same location;
- CLS: where there is a zero pixel in the microcalcification map and a non-zero pixel in the CLS map;
- Uncertainties: where there is a non-zero pixel both in the microcalcification map and the CLS map;

The final goal of the algorithm is to split the uncertainty class between microcalcifications and CLS. For a good estimate of the CLS, most of the putative microcalcifications will overlap with some CLS and therefore be labelled as uncertainties. Figure 102 shows the best estimate of CLS versus the estimate used in the CLS removal. The initial estimates of the maximum likelihood parameters incline to give priority to CLS, both because of the much larger number


of CLS versus calcifications and because most microcalcifications are labelled as uncertainties. For that reason, we forced the uncertainties to be part of the calcification class in the expectation step and then re-evaluated all the probabilities in the maximisation.



**Figure 101:** Histogram comparison: (a) and (c) are two mammogram samples with microcalcification clusters and CLS; (b) and (d) are their respective histograms, where the blue continuous plot corresponds to microcalcifications, while the red dotted plot is related to CLS;

We applied the EM algorithm to reduce the number of FP in the detection of microcalcifications. While we found the number of FP to be lowered, the number of correctly detected microcalcifications also decreased. Generally, most or all calcifications are labelled as CLS when the algorithm converges. In a few cases, when the histograms were minimally overlapping, the calcification clusters were preserved or even some of the CLS were labelled as calcium. Some *a priori* information about the percentage of microcalcifications versus CLS should be embedded in the algorithm, as well as different convergence criteria restricting the algorithm from eradicating too many TP. This is subject for further work.

The expectation maximisation algorithm that we propose is an iterative Bayesian classifier, which could be related to the work proposed in [81]. The authors propose a Bayesian approach using Markov Random Fields to detect clusters of microcalcifications. This well-known algorithm has been closely studied and intensively tested and is briefly described in Section 2.2.2.



**Figure 102:** Errors in CLS estimation: (a) the original contrast-enhanced mammogram sample with a microcalcification cluster; (b) the best estimation of CLS, which erroneously includes the microcalcification cluster; (c) the CLS map used in our algorithm, which may cause the disruption of CLS in the removal step


## 6.2.2 Iso-Level Segmentation

The algorithm presented here is based on work done by Hong and Brady [70] for segmentation of mammograms using a topographic approach.

### 6.2.2.1 Theory

From a topographic perspective, images are seen as surfaces that rise higher with the level of intensity (grey-level) in the original mammogram or interesting height in SMF images (Figure 103). By thresholding the image over a large number of intensity levels, a set of iso-contours is obtained to form a topographic representation. The intensity quantisation is similar to that used

by Yam *et al.* [179] in their microcalcification detection algorithm and is equivalent to slicing the 3D surface along its height.



**Figure 103:** 3D image representation and quantisation: (a) a 2D phantom of Gaussian intensity variation used here as the original image; (b) the 3D surface of image (a); (c) the grey-level threshold levels ( $l_1, l_2, l_3, \dots$ ) on a surface model; (b) the iso-contours super-imposed on image (a).

An iso-level contour  $C$  for the level  $l$  of the image  $f$  is represented by a closed curve that does not cross itself at any point along the curve (89).  $\Omega$  is a domain of an image in  $\mathbb{R}^2$ . The continuity of the curve is tested in its 8 adjacent neighbours. A bilinear interpolation is used to resample the image for a continuous approximation of the digital image. Gaussian smoothing is then applied to reduce noise and smooth curves prior to segmentation.

$$C(l) = \{(x, y) \mid f(x, y) = l\}, \quad \forall (x, y) \in \Omega \quad (89)$$

The algorithm generates quasi-concentric iso-contours. A nested relationship is defined to provide a hierarchical representation between iso-contours to examine the topological characteristics of an image. Saliency is seen as a significant gradient value across the boundary



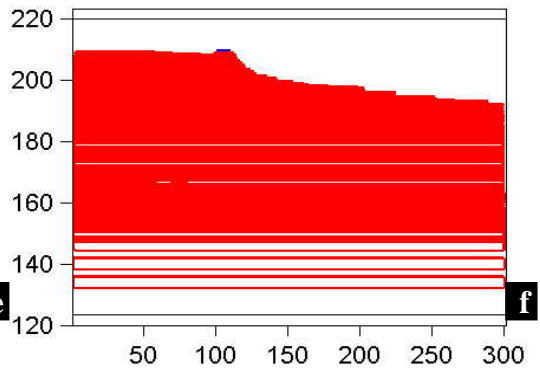




of a region. By considering the nesting depth from the innermost contours of each nesting structure within a given contour, a minimum nesting depth is set and used as saliency estimator. An appropriate choice for the minimum nesting depth is critical for achieving good results.


### **6.2.2.2 Initial Results**

Topographic representation provides extra information about a mammogram; shape, size and location are key issues in the detection and classification algorithms. We are interested in the size and shape of the lowest intensity contour above the noise level. If the contour corresponds to a microcalcification, it should be approximately round and small, while FP should have thin elongated contours (for CLS) or large surfaces (for parenchymal tissue). For an approximation of the noise level we can use the minimisation of the sum-of-square-differences (SSD) and estimate noise as the high frequency component (c.f. Chapter 5).

For our initial tests, we selected mammograms which generated microcalcification maps, including both correctly labelled microcalcifications and FP. Our motivation is to use the iso-level segmentation method to remove FP. A labelling algorithm identified the weighted-centre of each putative microcalcification to be used as a seed pixel in the iso-contour algorithm. The seed pixel is used as the innermost contour on the original mammogram and concentric contours are generated out of this seed. Using the minimal nesting depth, we can test the shape and size of the obtained iso-contours and identify if the seed pixel corresponds to a microcalcification or an FP. Figure 104 shows some examples of iso-contours obtained from a mammogram sample containing a microcalcification cluster correctly detected and a number of FP.

The initial results have proved to be highly accurate when a small slicing step is used. Future work will be to apply the algorithm to a mammographic image database in order to better assess the robustness of this approach through ROC analysis. A natural extension to this work would be to incorporate an automatic method of estimating the noise level.





**Figure 104:** Iso-contours and their 3D plots used to reduce the number of FP in the detection of microcalcifications: (a) the original mammogram sample with a microcalcification cluster; (b) the microcalcification map with a conservative selection of parameters; (c), (e) the iso-contours around two selected FP as seed pixels (the FP are marked in blue); (d), (f) the projections of the 3D plots of (c) and (e) with large elongated contours below the seed point; (g), (i) the iso-contours around two selected TP as seed pixels (marked in blue); (h), (j) the projections of the 3D plots of (g) and (i) with thin round contours below the seed point and above the noise level.

### 6.3 Concluding Remark

In this thesis we have presented a method to detect features in mammographic images. The primary application has been the detection of microcalcifications. We have developed a three-step automated algorithm designed to aid in the early detection of breast cancer. Moreover, we have investigated the detection of mammographic masses as part of a more complex detection algorithm for breast anomalies. The method has been motivated by known characteristics of the human visual system and facts from the manner clinical examinations are conducted in the screening programmes.



## APPENDIX A

### Phase Congruency and the Structure Multivector

What is the status of feature detection prior to phase congruency? If one would summarise, then the following remarks should be drawn to attention:

- Anisotropic diffusion, implemented properly, gives good responses to isolated step changes in intensity between locally constant regions (c.f. Chapter 3). Unfortunately, such isolated edges are rare in medical image analysis. Anisotropic diffusion remains a good tool in smoothing images with local enhancements, but does not suffice for feature detection.
- The many kinds of intensity change (e.g. lines, ramps) that occur in practice are not detected reliably by anisotropic diffusion.
- There are a number of techniques dubbed “corner detectors” for estimating locally two-dimensional changes. Of these, the algorithms developed by Harris and Stephens [58] and by Smith and Brady [155] are perhaps the best known.

Ignoring point 3 for the moment, we may restrict attention to locally one-dimensional intensity changes. It is a remarkable fact that since 1987 there has been a single, unifying theory of feature detection which (a) works well, and (b) has been largely ignored by researchers in computer vision. It is called *phase congruency* or *local energy*. We will follow the Morrone-Owens-Kovesi [119, 90] development, leading to recent work by Felsberg [38]. The following account is based on [13].

Morrone and Owens [119] observed that all commonly occurring one-dimensional signal changes - steps, ramps, thin bars, dots - correspond to signal locations at which the local Fourier components are all in phase. Furthermore, they found that even complex intensity changes often correspond to points of maximum phase “congruence”. They established the link between a signal features to the point of maximum phase congruency. They also proved that their definition of phase congruency (PC) is a normalised measure of a well-known computation, the *local energy*. This provides an important explanation to the concept of the local energy: peaks in the energy function correspond to feature points where phase congruency is a maximum. Being a measure of phase ‘congruence’, PC (or local energy) has the advantage of not being sensitive to the image contrast and brightness. According to this approach, a feature is defined as a location in a signal where the PC is high.

## A.1 Local Energy and Phase Congruency (PC)

Let  $f(x)$  be a one-dimensional signal. It can be reconstructed from its Fourier spectrum by

$$f(x) = \int_{-\infty}^{\infty} a_{\omega} \cos(\omega x + \phi_{\omega}) d\omega \quad (90)$$

where  $a_{\omega}$  is the amplitude and  $\omega x + \phi_{\omega}$  is the phase offset. Since steps, roofs, etc. all correspond to points in a signal where the components of the spectrum are in phase, the *Phase Congruency* PC(x) at each point  $x$  in the signal is defined as:

$$PC(x) = \underset{\theta \in (0, 2\pi)}{\text{Max}} \frac{\int_{-\infty}^{\infty} a_{\omega} \cos(\omega x + \phi_{\omega} - \theta) d\omega}{\int_{-\infty}^{\infty} a_{\omega} d\omega} \quad (91)$$

The  $\theta$  that maximises this expression for PC represents the amplitude-weighted mean phase angle. By definition, PC is a dimensionless value between (0, 1).

It is inconvenient to compute PC directly from its definition. It is normally obtained from the local energy, computed as below, using the relationship:

$$PC(x) = \frac{E(x)}{\int_{-\infty}^{\infty} a_{\omega} d\omega} \quad (92)$$

The local energy can be obtained from the analytic wavelet transform, which is equivalent to convolving the signal with a pair of quadrature filters. It turns out that the choice of quadrature filters plays an important role in determining the quality of results obtained.

An analytic wavelet  $\Psi$  is a function whose Fourier transform is zero for non-positive frequencies (hence the issue about DC correction).  $\tilde{\psi}(\omega) = 0$  if  $\omega \leq 0$ , ( $\tilde{\cdot}$  denotes the Fourier transform). Let  $f$  be the signal, then the result  $W$  of transforming the signal with  $\Psi$  is the inner product  $W = \langle f, \Psi \rangle$ . In general,  $W$  is complex. The local energy  $E$  is the amplitude of the transform.


$$E = |W| = |\langle f, \Psi \rangle|, \quad (93)$$

Similarly,

$$p = Arg(W), \quad (94)$$

gives the phase angle at which the phase congruency occurs, and will be used later to specify the feature type. The analytic function has zero phase. It does not change the signal's phase. Instead, it separates the signal's amplitude and phase. This is so for any analytic wavelet, so the question arises which one to use.

Gabor functions are widely used and are considered approximately analytic, if the non-positive frequencies are small enough. This requirement restricts the use of Gabor filter: one cannot construct Gabor functions of arbitrarily wide bandwidth and still maintain a reasonably small non-positive frequency component.



**Figure 105:** Log-Gabor functions of different bandwidths and centre frequencies.


The log-Gabor function is found to have better properties. It is analytic by definition. It has a long high frequency tail, which is useful for detecting fine features. The log Gabor function is defined in the frequency domain as

$$G(\omega) = e^{\frac{-\log(\omega/\omega_0)^2}{2(\log(\kappa/\omega_0))^2}} \text{ if } \omega > 0, \text{ and all zero otherwise,} \tag{95}$$

where  $\omega_0$  is the filter's centre frequency. The term  $\kappa/\omega_0$  is held constant. A  $\kappa/\omega_0$  value of 0.74 will result in a filter bandwidth of approximately one octave, 0.55 will result in two octaves, and 0.41 will produce three octaves. Figure 105 is a plot of log-Gabor functions of

different bandwidths and centre frequencies. The function's centre frequency and bandwidth are to be chosen according to the application.


Figure 106 shows the Phase Congruency of a sample signal calculated using log-Gabor functions and the Hilbert Transform. The solid line of PC obtained from the log-Gabor function has peaks for every feature point, which are marked with '\*'. Note, however, that the dotted line of the Hilbert transform has only two peaks, instead of four in the second half of the signal. The reason is that the Hilbert transform of 1D signal is a special case of the analytic wavelet transform where the whole positive frequency spectrum is evenly covered, unlike the band-pass filter of the log-Gabor functions. Therefore its time-space coverage is too narrow to produce the all feature peaks. We next discuss a method to recover all the feature points using phase angles.



**Figure 106:** Phase congruency and phase of a sample signal

## A.2 Detecting Feature Type from Phase Angle

Consider the third sub-plot of Figure 106, which is the phase angle  $p(x)$  obtained from the Hilbert transform of the sample signal. Each feature point corresponds to a specific phase angle. Figure 107 shows the feature type and the phase angle.



**Figure 107:** Phase angles of different feature type

Combining the phase information and with the PC (or local energy), it is possible to distinguish the peaks in the local energy (or phase congruency) function. More precisely, PC and phase can be combined in the following way, which gives a new measure of the features, denoted as  $PC_f$ ,


$$PC_f(x, \phi_i) = PC(x) \cdot \sum_i [\cos(p(x) - \phi_i)]^2 \quad , \quad (96)$$

where  $p(x)$  is the phase of the signal as defined in Section A.1, and  $\phi_i$  is the phase at the following feature types

$$\phi_i = \begin{cases} 0 & \text{peak} \\ \pi/2 & \text{up step} \\ \pi & \text{valley} \\ -\pi/2 & \text{down step} \end{cases} \quad ,$$

,  $i$  is one of the feature types listed above, and  $[ \ ]$  denotes that the enclosed quantity is equal to itself when its value is positive and zero otherwise.  $PC_f$  can be used to detect any of the feature types, or, more importantly, any combination of them.

The usefulness of the above combination is illustrated in Figure 108, where different features were detected from the sample signal of Figure 106. The two feature points in the second half of the signal are clearly identified in the up-step feature. Even more features show up; for example, the peak of the first half the signal and the long down slope in the second half of the signal.



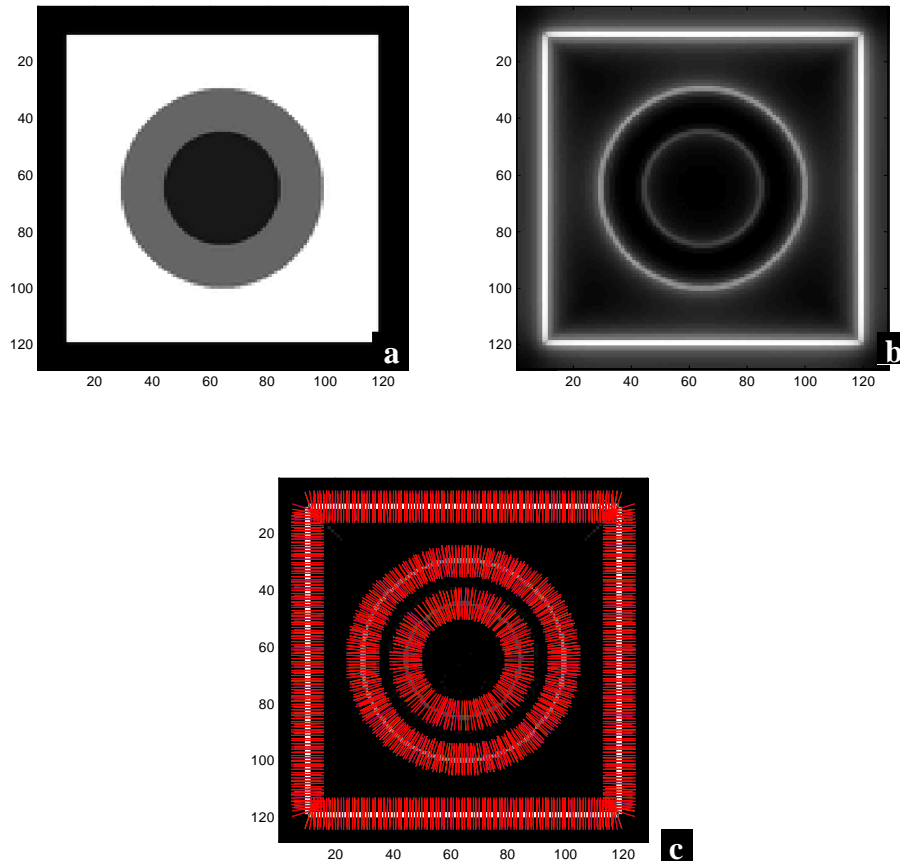
**Figure 108:** Detecting features using the phase angle

### A.3 Extending PC to 2D

The local energy and Phase Congruency computation can be extended to 2D if it is assumed that the signals of interest have simple neighbourhoods, that is, varies locally only in one direction. With this kind of image data, it is possible to interpolate the local energy and estimate the orientation from a minimum of three energy outputs obtained from three symmetrical distributed directions:  $30^\circ$ ,  $90^\circ$  and  $150^\circ$ . The energy was computed in each direction with an analytic wavelet function constructed at this direction and extended with a spread function  $\cos^2(\varphi)$ . In practice, six directions ( $0^\circ$ ,  $30^\circ$ ,  $60^\circ$ ,  $90^\circ$ ,  $120^\circ$  and  $150^\circ$ ) were used to accommodate the complexity of the images.

For illustration, this method of extending PC to two dimensions is applied to the idealised image shown in Figure 109.a, with the result shown in Figure 109.b. The method also gives the approximate local signal orientation as shown in Figure 109.c.






**Figure 109:** Test (a) Idealised test image; (b) Phase Congruency; (c) Orientation vectors

### A.4 Multiscale Analysis

Phase Congruency can be applied to an image at multiple scales and at different filter bandwidths. Figure 110 shows the scalogram of applying PC on a section of an incised tablet image. The first and second sub-images of Figure 110 show the tablet image and a signal from the image. The next three scalograms were obtained with three log Gabor filters of different bandwidths. The horizontal axes of the scalograms correspond directly with the signal's horizontal axes. The vertical axes of the scalograms are scales ranging from 1-20, and correspond to the filter wavelength from 3 to 20. The scaling factor between the filters is 1.1. The scalogram shows the lifeline of a feature across scale. We can choose the most suitable scale and bandwidth to do a single scale analysis of PC or local energy, or multi-scale analysis using Kovesei's method [90]. However we have found that the performance of Kovesei's method

depends on noise estimation. In the case of tablet images, we find that single scale analysis achieve better results.



**Figure 110:** Phase Congruency at different scales.

A recent development has been the Structure Multivector developed by Felsberg and Sommer [39]. The initial observation is purely technical: it has frequently been asserted that the Hilbert Transform – fundamental to developing quadrature filters – only exists in one-dimension, so that the PC/Local Energy model inevitably involves one-dimensional signals that are then extended to two- and three-dimensions by steering. Felsberg and Sommer point out that this is only true for scalar valued functions; there does indeed exist a 2- and 3-dimensional Hilbert Transform for vector valued functions. More precisely, they develop a 2- or higher dimensional generalisation of the analytic function called the *monogenic function*. The monogenic function is given by:

$$f_M(x_1, x_2) = (f, h_1 * f, h_2 * f)(x_1, x_2) \quad (97)$$

,where  $h_i$  are the inverse FT of two filters  $H_i(\omega_1, \omega_2)$  defined in the 2D Fourier plane. The local energy turns out to be  $f_M / |f_M|$  and this leads to a representation of the image in terms of spherical polars that contain the local phase and the orientation of image features. In this way, the structure multivector/monogenic signal embodies all of the information that has been found useful for detecting features of *all* kinds even in textured images. This is subject to considerable ongoing effort.

## APPENDIX B

# Registration Framework for Mammography

## B.1 Introduction

The method described in this Appendix was developed by Kostantinos Marias [109, 110, 111] in order to overcome the main problems in temporal mammogram registration which are due to a combination of a non-rigid tissue motion due to different compression between acquisitions, differences in the imaging parameters and the temporal changes in tissue composition and structure of the breast. Since the intensity distribution and the structural morphology can significantly differ in temporal sequences, we need a photometrically invariant method that can exploit the variable similarity of temporal pairs of mammograms. An application is illustrated in Chapter 5 for the prompting of masses in pairs of temporal mammograms.

The method relies on two stages:

- Initially the images are aligned based on the boundary. This is done by developing an algorithm that automatically detects 3 points with characteristic curvature in the outline of both mammograms. A thin plate spline interpolation is used to calculate the image transformation that aligns the boundaries of the two mammograms.
- Using a wavelet-analysis segmentation algorithm internal regions of dense tissue are defined that have good spatial characteristics in each mammogram. The boundary transformation together with scale and area information of the segmented regions is

used to match internal structures and refine the registration. In this second optimisation stage, a regularised approximation scheme is used to account for possible inaccuracies in the selection of the internal landmarks, especially because the centre of mass of each matched region pair is used to calculate the image transformation.

Figure 111 is a flow chart of the registration algorithm. Though the technique could be fully automated, the acceptance or not of the internal landmarks should in practice be decided or confirmed by the user. If the suggested internal landmarks do not meet the clinician's satisfaction (e.g. possibly in involuted breast pairs) the boundary-based registration is the final result, otherwise an approximation scheme (including internal and boundary landmarks) is employed in order to better approximate the deformation necessary to align the mammograms.




Figure 111: The basic steps of our breast registration algorithm (reproduced from [110]).

## B.2 Partial Registration Using the Boundary

The breast boundary is the most useful feature of the mammogram in terms of temporal consistency. It provides information about the difference in compression between two acquisitions and enables the calculation of landmarks that allow the approximation of the transformation that relates the geometry of a temporal pair of mammograms. Still, the registration using the boundary is not sufficient as internal structures move to different extents under different compressions, because of differences in shape and tissue density. However, accurate detection of the breast outline and calculation of temporally invariant geometrical landmarks is a key first step for mammogram registration. The steps that comprise the boundary registration method are:

- Boundary outline detection;
- Curvature analysis of the outline(s) and detection of consistent landmarks;
- Anatomical significance of detected boundary landmarks;
- Thin-plate spline interpolation to align the boundaries.

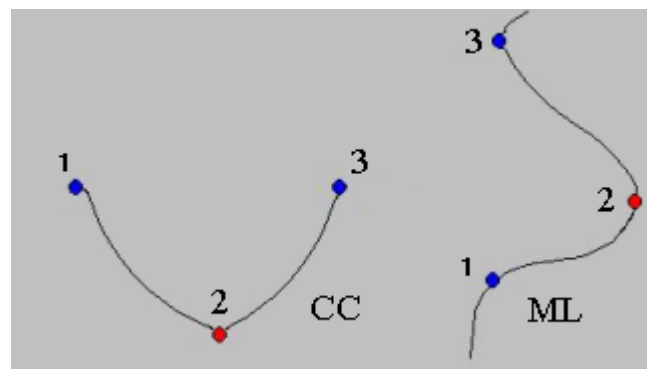
### B.2.1 Breast Outline Detection

To generate the breast outline, the image is thresholded in the first “valley” between two peaks (inside the breast and outside the breast) of the intensity histogram resulting in a binary image. Subsequently, an 8-connected component outline is obtained using mathematical morphology (closing followed by dilation and then subtraction).

This method yields an approximation to the boundary. This tracks the points along the boundary since the breast outline can have an irregular shape (e.g. very often 2 points in the  $x$  axis can correspond to 1 in the  $y$ ). However, the segmented curve is “jagged” and this poses a problem for the robust detection of boundary landmarks, as is discussed in the next section.

## B.2.2 Curvature Analysis on the Breast Outline-consistent Landmarks

In order to be able to register mammogram pairs one needs to be able to establish correspondences between the breast outline of each mammogram, since the “beginning” and “end” of each outline highly depend on the segmentation result and the image acquisition (e.g. in some medio-lateral mammograms a larger part of the rib is visible than in others). The aim is to “translate” the geometrical consistencies in mammograms into an automatic algorithm for the detection of *consistent* boundary landmarks. Three points are considered, as in Figure 112: in the cranio-caudal case (CC), the points 1 and 3 can be assumed to be near in the chest wall (and thus invariant) and are approximated by the first and last points of the breast-outline respectively; point 2, is the maximum curvature point (negative curvature by convention). The medio-lateral oblique mammograms (MLO) represent the most difficult case. Geometrically, these points can be described as two maxima of positive (by convention) curvature (points 1 and 3) and 1 point of maximum negative curvature (point 2).



**Figure 112:** Consistent landmarks in the CC and ML “idealised” outlines.

In order to build a robust detection algorithm for the three points discussed above, the curvature profile of the breast outline is calculated. The steps are as follow:

- An algorithm based on the separation of positive and negative curvature automatically detects the three suggested points (Figure 112);



- Define an optimum sampling rate ( $S_{opt}$ ) of points along the segmented breast-outline and run a spline to approximate the breast boundary. Different sampling rates preserve different amount of detail at a trade-off with overall smoothness. This optimum would depend on the pixel dimension (300 microns in the images we used).

### **B.2.3 Anatomical Significance of Boundary Landmarks**

The boundary registration technique is based on the robust detection of 3 points of characteristic curvature. This makes the boundary alignment more robust as consistent boundary points are calculated for a mammogram pair, instead of attempting to match the whole segmented breast outline. As is shown in Figure 112, the detected landmarks very often correspond to the anatomical location (in MLO) of the rib (point 1 in Figure 112), the nipple (point 2), and the axilla (point 3). However, this anatomical correspondence is *not* a requirement for the boundary registration to work. A good example is the case where the nipple is not visible but there is always a global maximum of negative curvature.

The user can refine the location of the boundary landmarks by shifting the calculated points along the breast-edge. Making the registration process completely “blind” can have an effect in the robustness and adaptability of the method. The most important correspondence is the nipple, since the glandular tissue converges to it.

### **B.2.4 Partial Registration from the Breast Boundary**

Sampling between these three points (that from now on will be referred as rib, nipple and axilla) and more specifically, between the axilla-nipple and nipple-rib segments, any temporal or bilateral mammogram pair can be aligned based on the boundary.

For temporal mammogram registration, a good initial alignment can be achieved using at least five points along the breast boundary. However, for greater accuracy in aligning the

boundaries, seven points uniformly sampled between the “axilla” and “nipple” boundary landmarks, and another seven between the “nipple” and the “rib” landmarks (total of 17 points) are used. Using these points, the images are registered using thin-plate spline interpolation.

Radial basis functions (RBF) are used for the elastic image deformation in this registration scheme. RBF are used in two contexts – firstly, for aligning only the boundary. Later, when internal landmarks are included, information about the spatial characteristics of the deformation points is used to implement a more sophisticated regularisation that is based on an approximation method.

In RBF interpolation, a set of  $n$  landmarks  $(p_i, q_i)$  is used to define a transformation function  $u: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ , where  $p_i = (x_{i1}, y_{i1})$  are the landmarks in the first image,  $q_i = (x_{i2}, y_{i2})$  are the landmarks in the second, and the interpolated transformation function  $u(x)$  must satisfy the interpolation constraint:

$$u(p_i) = q_i, i = 1, \dots, n \quad (98)$$

Marias uses the thin-plate spline radial basis function  $R_{TPS}$ , since it is a stable method to recover deformations (including local deformations due to breast tissue motion). By weakening the interpolation constraint, the smoothness of the transformation can be controlled and the uncertainty in the localisation of landmarks can be taken into consideration.

Once the interpolating function has been calculated, “warped” images are produced by forcing every point  $(x, y)$  in a mammogram to take the intensity value of the point where the interpolating function maps the  $(x, y)$  point of the previous mammogram. After image warping, difference (subtraction) images can be generated and used to search for regions of large intensity differences. These regions can be either new growths (e.g. a cancer), changes due to involution, or they can be due to local inaccuracy in registration.

### B.3 Multi-scale Landmark Selection for Improved Registration

It is demonstrated that a usually small, but significant number of internal correspondences greatly improves registration and better approximates the complex internal tissue deformation due mainly to differences in compression.

The main steps of the multi-scale segmentation algorithm to detect regions of dense tissue are:

- The mammogram pair is decomposed using the Coiflet wavelet packets. This particular wavelet was chosen because it yields good spatial localisation (e.g. it is edge preserving) and has compact support.
- After each mammogram is decomposed into a set of high-frequency and low-frequency images (with good spatial localisation of features), these are ranked by information content using an entropy measure. This construction is used to track significant features through scale space and forms the basis of the feature segmentation.
- A region growing is performed from the lowest scale towards the highest. A merging operator tracks the feature to the highest scale so that each feature can be represented with more “detailed” information. In a mammogram pair, the  $n$  most important regions (usually  $n \leq 5$ ) are tracked and subsequently matched.

### B.4 Landmark Matching and Registration Refinement

Based on the regions that are detected using the scale-space segmentation approach described above, a set of internal landmarks is defined by a matching algorithm that includes the partial transformation (induced by the boundary alignment) in conjunction with scale, size and area information of the candidate matches. In the registration process the segmented regions are represented by their centroids.

The initial search-space for a match in the first image is defined as a window in the second image whose size is proportional to the amount of displacement of the transformed internal landmarks using the boundary transformation. This is used to limit the possible matches to a “window” or neighbourhood. All the feature parameters (size, scale, relative motion) are used to drive a simple spatial searching. Essentially these criteria are used as the basis of a “match-rejection” filter. A distance measure is evaluated between landmarks to ensure that landmarks classified as a “match” have similar spatial properties and have demonstrated a change in geometrical correspondence as a result of the boundary deformation. On average, depending on the degree to which the breast is involuted, 3 to 5 internal landmarks are defined at the centres of mass of the corresponding wavelet-defined regions.

The last step in the registration process is to include both the boundary (curvature-based) landmarks and the internal landmarks. However, at this stage of the registration process, an approximation (rather than an interpolation) scheme is used to compute the elastic deformation. This is to account for possible inaccuracies in landmark representation, as well as to produce a smooth deformation that takes into consideration the relative importance of the matched regions (represented by their centre of mass).

The boundary points and the internal landmarks (computed by the wavelet analysis) together control a thin-plate spline approximation technique, which gives the final registration.

## **B.5 Results an Discussion**

Marias [110] reports the following validation results:

**Table 6:** Comments on registration results in 50 mammogram pairs: The viewer classified the results in three categories according to the alignment of image features.

<b>Boundary alignment</b>	<b>Internal correspondence</b>
Good: 100%	Good: 70%
Average: 0%	Average: 25%
Poor: 0%	Poor: 5%

**Table 7:** Clinical assessment of the improvement in registration using internal landmarks in 25 mammogram pairs and comparison with the reduction in the standard deviation of the difference image after geometrical alignment for the same cases.

<b>Improvement of the registration result when including internal landmarks</b>	<b>Reduction in the standard deviation of the difference image</b>
Significant: in 20 mammogram pairs	10%-35%
Not significant: in 5 mammogram pairs	0%-15%

Multi-scale analysis provides a reliable framework for establishing correspondences between significant regions inside the breast. Using internal landmarks, the registration result is improved, as was asserted by the clinician and the difference images after registration as well as the joint histograms of the aligned images. Even though matching points inside the breast is difficult due to temporal changes and depends upon the extent to which the architecture (or topology of the surface) is preserved, the multi-scale segmentation method used, reliably locates regions of dense tissue that appear in both temporal mammograms. Additionally, using the thin-plate approximation scheme, the internal landmarks can be weighted according to their size and scale and therefore compensate for landmark localisation errors.

## APPENDIX C

### **Receiver Operating Characteristic (ROC) Curve Analysis**

Clinical information from patients is routinely collected in all hospital across the country. However, in order to understand the correlation of the different clinical signs, symptoms and diagnostic tests to the likelihood of a disease, this information has to be accurately analysed and interpreted through objective measures and definitions.

In this Appendix we shall be concerned with issues related to the evaluation of detection and classification algorithms. The ROC curve analysis is the most common measure of diagnostic accuracy and an essential tool in assessing the performance of a developed method.

The questions addressed are:

- How well can we detect microcalcifications or masses?
- How do a few selected detection algorithms compare with each other, and how should we assess the performance of an algorithm?
- Given a mass or microcalcification cluster, can we classify it as benign or malignant?

A detection or classification algorithm in digital mammography is traditionally assessed by applying the algorithm to a representative set of mammograms, for which ground truth information is known. It is then possible to compare the results with the ground truth to determine if the microcalcification cluster or mass regions have been correctly detected or classified. In this work we assess the performance of microcalcification detection algorithms using a set of 102 normal and abnormal mammographic cases taken from the Oxford database.

Each abnormal mammogram has an associated truth file containing the outline of the microcalcification cluster present in the image. We believe that the tested cases are a reasonable representation of different types of microcalcification clusters a clinician is likely to see in women of breast screening age. Although we will further refer to microcalcification clusters, the same assessing principle is used in evaluating mass detection and classification algorithms.

The most common assessment criterion in the radiological literature is the ROC curve. Since detection algorithms will output a 'positive' or a 'negative' flag for each processed image, a ROC analysis must evaluate their ability over a mammographic database, over a range of parameters. Consequently, a database of processed mammographic images can be split into four disjoint subsets, given the ground truth:

- TP - True Positives, correctly identified as containing a microcalcification cluster;
- FP - False Positives, incorrectly identified as containing a microcalcification cluster;
- TN - True Negatives, correctly identified as being clear;
- FN - False Negatives, incorrectly identified as being clear.

Denoting the total number of malignant (or detected) outcomes as  $M$ , we have  $M=TP+FN$ , similarly the total number of benign (or non-detected) outcomes is  $N=TN+FP$ . The following definitions are related to clinical diagnostic tests and the main interest from the image analysis standpoint is to improve the understanding of clinical articles as well as the reported results on CAD systems:

- Sensitivity:  $TP/M$ , also referred to as the true positive fraction (TPF), is the probability of a positive test among patients with a disease;
- Specificity:  $TN/N$  is the probability of a test being negative among healthy patients;
- The false positive fraction (FPF): is defined as 1-specificity, which is equivalent to  $FP/N$ ; typical ROC curves plot the TPF as a function of the FPF;
- Diagnostic accuracy:  $(TP+TN)/(M+N)$  is the ratio of correct detections over the total number of assessments;

- Incidence: is the probability that a healthy patient develops the disease during an interval (e.g. in a year, or in the screening interval for the case of breast cancer); this definition is important when determining the incidence of breast cancer, since a significant number of cancers are missed;
- Prevalence: is the probability of a disease in the entire population. Using the definitions of sensitivity, specificity and prevalence, the probability of disease given a positive test can be defined according to Bayes' equation:

$$P(\text{disease} | \text{positive test}) = \frac{\text{prevalence} \cdot \text{sensitivity}}{[(\text{prevalence} \cdot \text{sensitivity}) + ((1 - \text{prevalence}) \cdot (1 - \text{sensitivity}))]} \quad (99)$$

Although sensitivity and specificity are important in describing diagnostic tests, they do not always offer sufficient information to interpret the results of a test. For this reason, the predictive values are more useful to the clinicians:

- Positive predictive value:  $TP/(TP+FP)$  is the probability of disease among patients that had a positive test;
- Negative predictive value:  $TN/(TN+FN)$  is the probability of no disease among patients that has a negative test.

It can be the case that although the sensitivity of the test is very high, the probability of disease among patients with a positive test to be lower, which is the probability that a patient has breast cancer if the CAD test for cancer is positive.

If now, we can redefine what is 'positive' or 'negative' according to the CAD (e.g. change the threshold at which a clustered is detected), the values for the specificity and sensitivity will change. By repeating the procedure for several thresholds, we can plot the values of sensitivity versus (1-specificity), the Receiver Operating Characteristic curve. Examining different thresholds represents the trade-off between sensitivity and specificity, or between false positives and negatives.


In traditional ROC analysis, the image is viewed as a single entity and the binary value P, (Positive: cluster detected) or N (Negative: no cluster detected) is assigned. As we change our



parameters, e.g. threshold value or contrast value in Chapter 4, the ROC curve is drawn out. Figure 113 shows a comparative analysis between two representative ROC curves. The straight line connecting the points (0,0) to (1,1) is referred to as the chance line. The name arises because an algorithm that just assigns the value P to an image with a probability 1/2, i.e. it guesses, lies on the chance line. An ROC curve should always lie above the chance line.

What features do we expect in a ROC curve? The ROC curve of a good algorithm will show a sharp monotonic rise as we move away from the origin; it detects the majority of microcalcification clusters whilst generating very few false results. Towards its end, the curve should ideally become horizontal. If a ROC curve's gradient falls below 1 then it is starting to label more things wrong than right, and it is of no real use beyond that point.

Referring again to Figure 113 we note that curve A always lies above B, indicating that algorithm A outperforms B. Typically, two ROC curves C and D cross so that for part of the plot  $D < C$  and then  $D > C$ . One would have to qualify the statement D outperforms C by stating the TPF range over which the statement held. Choosing the right trade-off between FPF and TPF and comparing the curves at that point could be used as a comparison. This is generally insufficient, since the general behaviour of the algorithm might be more important than just one point on the ROC curve. Another generally accepted comparison criterion between two algorithms is the integral below each curve, an exact and exhaustive measure in discriminating between healthy and non-healthy patients or benign and malignant. Therefore, the ROC measure gives a clear indicator of an algorithms performance.



**Figure 113:** An example of ROC analysis.

A number of variations on ROC curves are used in clinical evaluation. The free-response receiver operating characteristic (FROC) curve is largely used in literature. Unlike the classical ROC curve plotting TPF versus FPF, the FROC curve plots TPF versus FP/image. The ROC analysis performed in this thesis is based on plotting FROC curves, as exemplified in Chapters 3 and 4.

## Bibliography

- [1] Aghdasi, F.: Application of Neural Network Adaptive Wavelets for Signal Representation and Classification in Digital Mammography. In Doi, K. Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): Digital Mammography, Elsevier, Amsterdam (1996) 307-310
- [2] Alberdi, E. Lee, R. Taylor, P.: Radiologists' Description and Interpretation of Mammographic Calcifications: Knowledge Elicitation for Computerized Decision Support. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 135-141
- [3] Ancelin, B. Yam, M. Highnam, R.P. Brady, J.M.: Automated Detection of Artefacts on Film-Screen Mammograms Using a Novel Iterative Detection and Correction Scheme. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 174-178
- [4] Andersson, I.: Breast Cancer Screening Results. In Karssemeijer, N. Thijssen, M. Hendricks, J. van Erning, L. (eds.): Digital Mammography, Kluwer Academic Publishers, Dordrecht Boston London (1998) 427-433
- [5] Andolina, V.F. Lillé, S.L. Willison, K.M.: Mammographic Imaging: A Practical Guide. J.B. Lippincott Company (1992)
- [6] Ashdown, I.: Luminance Gradients: Photometric Analysis and Perceptual reproduction. In: Journal of Illuminating Engineering Society, Vol. 25(1) (1996) 69-82.
- [7] Beral V., Banks E., Reeves G., Appleby P: Use of HRT and the Subsequent Risk of Cancer. In: Journal Epidemiology Biostatistics, Vol. 4(3) (1999) 191-210
- [8] Betal, D. Roberts, N. Whitehouse, G.H.: Segmentation and Numerical Analysis of Microcalcifications on Mammograms using Mathematical Morphology. In: British Journal of Radiology, Vol. 70 (1997) 903-917
- [9] Birdwell, R.L. Ikeda, D.M. O'Shaughnessy, K.F. Sickles, E.A.: Mammographic Characteristics of 115 Missed Cancers Later Detected with Screening Mammography and the Potential Utility of Computer-aided Diagnosis. In: Radiology, Vol. 219 (2001) 192-202
- [10] Bishop, C. M.: Neural Networks for Pattern Recognition. Oxford University Press, Oxford (1995)
- [11] Blanks, R.G. Wallis, M.G. Moss S.M.: A Comparison of Cancer Detection Rates Achieved by Breast Cancer Screening Programmes by Number of Readers, for One and Two-View Mammography: Results from the UK National Health Breast Screening Programme. In: J Med Screen, Vol. 5 (1998) 195-201

- [12] Boedicker, A. Evertsz, C.J.G, Brady, J.M., Hendriks, J.H.C.L. Juergens, H. Karssemeijer, N. Peitgen, H.O.: SCREEN-TRIAL: Softcopy Reading Environment in European Screening Mammography – SCREEN. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 423-427
- [13] Brady, J.M. Pan, X.B. Schenk, V.U.B. Terras, M. Molton, N.: Shadow Stereo, Image Filtering and Constraint Propagation. In: Proc. Roy. Soc. (2002) in press
- [14] te Brake, G.M. Karssemeijer, N. Hendriks. J.H.C.L.: An Automatic Method to Discriminate Malignant Masses from Normal Tissue in Digital Mammograms. In: Phys. Med. Biol. (2000) 2843-2857
- [15] te Brake, G.M.: Computer Aided Detection of Masses in Digital Mammograms, PhD Thesis. The Catholic University of Nijmegen (2000)
- [16] te Brake, G.M. Karssemeijer, N.: Single and Multi-Scale Detection of Masses in Digital Mammography. In: IEEE Transactions on Medical Imaging, Vol. 18(7) (1999) 628-639
- [17] Brown, M. Eccles, C. Wallis, M.G.: Geographical Distribution of Breast Cancers on the Mammogram: An Interval Cancer Database. In: The British Journal of Radiology, Vol. 74 (2001) 317-322
- [18] Burcher, M. Noble, J.A Han, L.: Deformation Correction in Ultrasound Imaging Using Contact Force Information. In: Proc. IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (2001) 63-70
- [19] Buscombe, J. Hill, J. Parbhoo, S.: Scintimammography: A Guide to Good Practice. Gibbs Associates Limited (1998)
- [20] Byng, J.W., Critten, J.P. Yaffe, M.J.: Thickness-equalization Processing for Mammographic Images. In: Radiology, Vol. 203 (1997) 564-568
- [21] Cancer Research Campaign: Breast cancer fact sheet (6.1). Cancer Research Campaign, Oxford (1996)
- [22] Caseldine, J. Blamey, R. Roebuck, E.J. Elston, C.: Breast Disease for Radiographers. Wright (1998)
- [23] Catté, F. Lions, P.L. Morel, J.M. Coll, T.: Image Selective Smoothing and Edge Detection by Nonlinear Diffusion. In: SIAM Journal on Numerical Analysis, Vol. 29 (1990) 629-639
- [24] Chan, H.P. Sahiner, B. Helvie, M.A. Petrick, N. Roubidoux, M.A. Wilson, T.E. Adler, D.D. Paramagul, C. Newman, J.S. Sanjay-Gopal, S.: Improvements of Radiologists' Characterization of Mammographic Masses by Using Computer-aided Diagnosis: A ROC Study. In: Radiology (1999) 817-827
- [25] Chan, H.P. Doi, K. Galhorta, S. Vyborny, C.J. MacMahon, H. Jokich, P.M.: Image Features Analysis and Computer-Aided Diagnosis in Digital Radiography: 1. Automated Detection of Microcalcifications in Mammograms. In: Medical Physics, Vol. 14 (1987) 538-548

- [26] Couvreur, C.: The EM Algorithm: A Guided Tour. In Proceedings IEEE European Workshop on Computationally Intensive Methods in Control and Signal Processing. Prague (1996) 115-120
- [27] Dash, N.: Magnetic Resonance Imaging in the Diagnosis of Breast Diseases. In: American Journal of Roentgenology, Vol. 146 (1986) 119-125
- [28] Davis, R.R. Dance, D.R.: Automatic Detection of Clustered Calcifications in Digital Mammograms. In: Physics in Medicine and Biology, Vol. 35 (1990) 1111-1118
- [29] Dean, P.B.: Overview of Breast Cancer Screening. In Doi, K. Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): Digital Mammography, Elsevier, Amsterdam (1996) 19-26
- [30] Dixon, J.M.: ABC of Breast Disease. BMJ Publishing Group, London (1995)
- [31] Doyley, M.M. Bamber, J.C. Fuechsel, F. Bush, N.L.: A Freehand Elastographic Imaging Approach for Clinical Breast Imaging: System Development and Performance Evaluation. In: Ultrasound in Med. and Biol., Vol. 27(10) (2001) 1347-1353
- [32] Edwards, D.C. Kupinski, M.A. Nagel, R. Nishikawa, R.M. Papaioannou, J.: Using Bayesian Neural Network to Optimally Eliminate False-Positive Microcalcification Detections in a CAD Scheme. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 168-173
- [33] Van Engeland, S. Karssemeijer, N. Hendriks, J.H.C.L.: Using Information from Two Mammographic Views to Improve Computer-Aided Detection of Mass Lesions. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 377-381
- [34] EstEve, J. Krickler, A. Ferlay, J. Parkin, D.M. (eds.): Facts and Figures of Cancers in the European Community. IARCPress, Lyon (1993)
- [35] Evans, C.J. Yates, K. Brady, J.M.: Statistical Characterisation of Normal Curvilinear Structures in Mammograms. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 285-291
- [36] Evans, C.J.: Detecting and Removing Curvilinear Structures from Mammograms. Internal Report, Department of Engineering Science, University of Oxford (2001)
- [37] Evertsz, C.J.G. Bodicker, A. Bohnenkamp, S. Dechow, D. Beck, C. Peitgen, H.O., Berger, L. Weber, U. Jurgens, H. Hendricks, J.H.C.L. Karssemeijer, N. Brady, J.M.: Soft-Copy Reading Environment for Screening Mammography – SCREEN. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 566-572
- [38] Fatemi, M. Wold, L.E. Alizad, A. Greenleaf, J.F.: Vibro-Accoustic Tissue Mammography. In: IEEE Transactions on Medical Imaging, Vol. 21(1) (2002) 1-8
- [39] Felsberg, M. Sommer, G: The Structure Multivector. In: AGACSE 2001, Cambridge, to appear at Birkhäuser
- [40] Ferlay, J. Bray, F. Pisani, P. Parkin, D.M.: GLOBOCAN 2000: Cancer Incidence, Mortality and Prevalence Worldwide. IARCPress, Lyon (2001)

- [41] Freer, T.W. Ulissey, M.J. Screening Mammography with Computer-aided Detection: Prospective Study of 12,860 Patients in a Community Breast Cancer. In: *Radiology*, Vol. 220 (2001) 781-786
- [42] Funovics, M. Schamp, S. Lackner, B. Wunderbaldinger, P. Lechner, G. Wolf, G.: Computer-assisted Diagnosis in Mammography: the R2 ImageChecker System in Detection of Spiculated Lesions. In: *Wien Med Wochenschr*, Vol. 148(14) (1998) 321-324
- [43] Le Gal, M. Ollivier, L. Asselain, B. Meunier, M. Laurent, M. Vielh, P. Neuenschwander, S.: Mammographic Features of 455 Invasive Carcinomas. In: *Radiology*, Vol. 185 (1992) 705-708
- [44] Le Gal, M. Durand, J.C. Laurent, M. Pellier, D.: Conduite a Tenir Devant une Mammographie Revelatrice de Microcalcifications Groupees sans Tumeur Palpable. In: *La Nouvelle Presse Medicale*, Vol. (5) (1973) 1621-1627
- [45] Gale, A.G. Cowley, H.C.: Analysis of Breast Cancer Screening Results. In Karssemeijer, N. Thijssen, M. Hendricks, J. van Erning, L. (eds.): *Digital Mammography*, Kluwer Academic Publishers, Dordrecht Boston London (1998) 19-26
- [46] Garra, B.S. Cespedes, E.I. Ophir, J. Spratt, S.R. Zuurbier, R.A. Magnant, C.M. Pennanen, M.F.: Elastography of Breast Lesions: Initial Clinical Results. In: *Radiology*, Vol. 202(1) (1997) 79-86
- [47] Geman, S. Geman, D.: Stochastic Relaxation, Gibbs Distribution and the Bayesian restoration of Images. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 6(6) (1984) 721-741
- [48] Giger, M.L. Huo, Z. Vyborny, C.J. Lan, L. Nishikawa, R.M. Rosenbourgh, I.: Results of an Observer Study with an Intelligent Mammographic Workstation for CAD. In: *Digital Mammography, Lecture Notes in Computer Science*, Springer-Verlag, Berlin Heidelberg New York (2002) 297-303
- [49] Giger, M.L. Maloney, M. Huo, Z. Vyborny, C.J. Kupinski, M. Venta, L.: Computerized Classification of Lesions on Digital Mammography. In Yaffe, M.J. (ed.): *Digital Mammography*, Medical Physics Publishing, Madison (2000) 189-192
- [50] Grady, L. Schwartz, E.: Anisotropic Interpolation on Graphs: The Combinatorial Dirichlet Problem. Technical Report CAS/CNS-TR-03-014, Department of Cognitive and Neural Systems, Boston University, Boston, MA (2003). Submitted to *IEEE Transactions on Pattern Analysis and Machine Intelligence*
- [51] Gonzales, R.C. Woods, R.E.: *Digital Image Processing*. Addison-Wesley Publishing Company (1993)
- [52] Groshong, B.R. Kegelmeyer, W.P.: Evaluation of a Hough Transform Method for Circumscribed Lesion detection. In Doi, K. Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): *Digital Mammography*, Elsevier, Amsterdam (1996) 361-366
- [53] Haagensen, C.C.: *Disease of the Breast*. W.B. Saunders Company (1986)
- [54] ter Haar Romeny, B.M.(ed.): *Geometry-Driven Diffusion in Computer Vision*. Kluwer Academic Publisher, Dordrecht (1994)

- [55] Hadjiiski, L. Sahiner, B. Chan, H.P. Petrick, N: Classification of Malignant and Benign Masses Based on Hybrid ART2LDA Approach. In: IEEE Trans. Med. Imaging, Vol. 18(12) (1999) 1178-1187
- [56] Hara, T. Yamada, A. Fujita, H. Iwase, T. Endo, T.: Automated Classification Method of Mammographic Microcalcifications by Using Artificial Neural Network and ACR BI-RADS Criteria for Microcalcification Distribution. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 198-204
- [57] Haralick, R.M. Shapiro, L.G.: Computer and Robot Vision, Vol. I. Addison-Wesley Publishing Company, Reading MA (1992)
- [58] Harris, C.G. Stephens, M.: A Combined Corner and Edge Detector. In: Proceedings of the Fourth Alvey Vision Conference, The University of Sheffield Printing Unit (1988) 147-152
- [59] Hayton, P. Brady, J.M. Tarassenko, L. Moore, Niall: Analysis of Dynamic MR Breast Images using a Model of Contrast Enhancement. In: Medical Image Analysis, Vol. 1(3) (1997) 207-224
- [60] Hendrinck, R.E. Lewin, J.M. Isaacs, P. Kuni, C.K. D'Orsi, C.J. Moss, L. Karellas, A. Sisney, G.A. Cutter, G.R.: Interim Clinical Evaluation of FFDM in a Screening Cohort: Comparison with Screen-Film Mammography in 4,965 Exams. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 468-475
- [61] Hendrinck, R.E. Lewin, J.M. D'Orsi, C.J. Kopans, D.M. Conant, E. Cutter, G.R. Sitzler, A.: Non-Inferiority Study of FFDM in an Enriched Diagnostic Cohort: Comparison with Screen-Film Mammography in 625 Women. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 475-481
- [62] Heucke, L. Knaak, M. Orglmeister, R.: A New Image Segmentation Method Based on Human Brightness Perception and Foveal Adaptation. In: IEEE Signal Processing Letters, Vol. 7, No. 6 (2000) 129-131
- [63] Heywang-Koebrunner, S.H. Wolf, A. Pruss, E. *et al.*: MR Imaging of the Breast with Gd-DTPA: Use and Limitations. In: Radiology, Vol. 171 (1989) 95-103
- [64] Highnam, R.P. Brady, J.M. Shepstone, B.J.: Computing the Scatter Component of Mammographic Images. In: IEEE Transactions in Medical Imaging, Vol. 13 (1994) 301-313
- [65] Highnam, R.P. Brady, J.M.: Mammographic Image Analysis. Kluwer Academic Publishers, Dordrecht Boston London (1999)
- [66] Highnam, R.P. Brady, J.M. English, R.: Detecting Film-Screen Artifacts in Mammography using a Model-Based Approach. In: IEEE Transactions in Medical Imaging, Vol. 18 (1999) 1016-1024
- [67] Highnam, R.P. Brady, J.M. Shepstone, B.J.: A Quantitative Feature to Aid Diagnosis in Mammography. In Doi, K. Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): Digital Mammography, Elsevier, Amsterdam (1996) 201-206
- [68] Hodge, G.: The Poison on Your Plate. In The Independent on Sunday, LifeEtc, No.649 (14 July 2002), 4-5

- [69] Holladay, L.L.: The Fundamentals of Glare and Visibility. In: J.Opt. Soc. Amer, Vol. 12 (1926) 271-320
- [70] Hong, B.W.: Segmentation of Mammograms in Topographic Approach. First Year Report. University of Oxford (2002)
- [71] Horsch, K. Giger, M.L. Huo, Z. Vyborny, C.J. Venta, L.A.: Computer-Aided Diagnosis of Breast Lesions on Ultrasound. . In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 233-326
- [72] Hume, A. Thamisch, P. Hartswood, M. Procter, R.: On the Evaluation of Microcalcification Detection Algorithms. In Doi, K. Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): Digital Mammography, Elsevier, Amsterdam (1996) 273-276
- [73] Huo, Z. Giger, M.L. Vyborny, C.J. Bick, U. Lu, P. Wolverton, D.E. Schmidt, R.A.: Analysis of Spiculations in the Computerized Classification of Mammographic Masses. In: Med. Phys. Vol. 22(10) (1995) 1569-1579
- [74] Jiang, Y. Nishikawa, R.M. Maloney, M.M. Giger, M.L. Venta, L.L.: Computer Classification of Malignant and Benign Microcalcifications in Small-Field Digital Mammograms. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 237-242
- [75] Kallergi, M.: Digital Mammography: From Theory to Practice. In: Cancer Control Vol. 5(1) (1998) 72-79
- [76] Karssemeijer, N.: State of the Art in Computer Aided Detection. In: 6<sup>th</sup> International Workshop on Digital Mammography, Bremen (2002), Invited Talk
- [77] Karssemeijer, N.: Computer Aided Detection and Interpretation in Mammography. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 243-252
- [78] Karssemeijer, N. te Brake, G.M.: Combining Single View Features and Asymmetry for Detection of Mass Lesions. In Karssemeijer, N. Thijssen, M. Hendricks, J. van Erning, L. (eds.): Digital Mammography, Kluwer Academic Publishers, Dordrecht Boston London (1998) 95-102
- [79] Karssemeijer, N. te Brake, G.M.: Detection of Stellate Distortions in Mammograms. In: IEEE Transaction on Medical Imaging, Vol 15(10) (1996) 611-619
- [80] Karssemeijer, N.: Adaptive Noise Equalisation and Recognition of Microcalcification Clusters in Mammograms. In: Int. J. of Pattern Recognition and Artificial Intelligence, Vol. 7(6) (1993) 1357-1372
- [81] Karssemeijer, N.: Stochastic Model for Automated Detection of Calcifications in Mammograms. In: Image and Vision Computing, Vol. 10 (1992) 369-375
- [82] Kedar, R.P. Cosgrove, D. McCready, V.R. Bamber, J.C. Carter, E.R.: Microbubble Contrast Agents for Colour Doppler US: Effect on Breast Masses. In: Radiology, Vol. 198 (1996) 679-686



- [83] Kegelmeyer, W.P. Pruneda, J.M. Bourland, P.D. Hillis, A. Riggs, M.W. Nipper, M.L.: Computer-Aided Mammographic Screening for Spiculated Lesions. In: *Radiology*, Vol. 191 (1994) 331-337
- [84] Kichenassamy, S: The Perona-Malik Paradox. In: *SIAM Journal of Applied Mathematics*, Vol. 57 (1997) 1328-1342
- [85] Kim, J.K. Park, H.W.: Statistical Textural Features for Detection of Microcalcifications in Digitized Mammograms. In: *IEEE Transactions on Medical Imaging*, Vol. (18) (1999) 231-238
- [86] Kita, Y. Highnam, R.P. Brady, J.M.: Correspondence between Different View Breast X-Rays Using a Simulation of Breast Deformation. In: *Computer Vision and Pattern Recognition* (1998) 700-707
- [87] Kita, Y. Tohno, E. Highnam, R.P. Brady, J.M.: A CAD System for 3D Locating of Lesions in Mammogram. In Niessen, W. Viergever, M. (eds.): *Medical Image Computing and Computer-Assisted Intervention 2001*, Lecture Notes in Computer Science, Vol. 2208. Springer-Verlag, Berlin Heidelberg New York (2001) 94-102
- [88] Kobatake, H. Murakami, M. Takeo, H. Nawano, S: Computerized Detection of Malignant Tumours on digital Mammograms. In: *IEEE Transactions on Medical Imaging*, Vol. 18(5) (1999) 369-378
- [89] Kok-Wiles, S.L. Brady, J.M. Highnam, R.P.: Comparing Mammogram Pairs for the Detection of Lesions. In Karssemeijer, N. Thijssen, M. Hendricks, J. van Erning, L. (eds.): *Digital Mammography*, Kluwer Academic Publishers, Dordrecht Boston London (1998) 103-110
- [90] Kovese, P.: Image Features from Phase Congruency. In: *Videre: Journal of Computer Vision Research*, Vol. 1 (1999) 1-26
- [91] Kupinski, M.A. Giger, M.L.: Automated Seeded Segmentation on Digital Mammograms. In: *IEEE Transactions on Medical Imaging*, Vol. 17 (1998) 510-517
- [92] Krupinski, E.A. Nodine C.F.: Gaze Duration Predicts the Locations of Missed Lesions in Mammography. In Gale, A.G. Astley, S.M. Dance, D.R. Cairns, A.Y. (eds.): *Digital Mammography*, Elsevier, Amsterdam (1994) 399-405
- [93] Lado, M.J. Tahoces, P.G. Souto, M. Vidal, J.J.: Orthogonal Wavelets and Decision-Making Techniques in Digital Mammography: Computerized Detection of Microcalcifications. In: *IEEE Transactions on Medical Imaging*, (2003) to be published.
- [94] Lai, S.M. Li, X. Bischof, W.F.: On Techniques for Detecting Circumscribed Masses in Mammograms. In: *IEEE Transactions on Medical Imaging*, Vol. 8 (1989) 377-386
- [95] Lee, L. Stickland, V. Wilson, P. Roebuck, E.: *Fundamentals of Mammography*. W.B. Saunders Company (1995)
- [96] Lefebvre, F. Benali, H. Gilles, R. Di Paola, R.: A Simulation of Clustered Breast Microcalcifications. In *Medical Physics*, Vol. 21(12) (1994) 1865-1874
- [97] Li, H. Wang, Y. Liu, K.J.R. Lo, S.C.B. Freedman, M.T.: Computerized Radiographic Mass Detection – Part I: Lesion Site Selection by Morphological Enhancement and Contextual Segmentation. *IEEE Transaction on Medical Images* 20 (2001) 289-301

- [98] Li, H.D. Kallergi, M. Clarke, L.P. Jain, V.K. Clark, R.A.: Markov Random Field for Tumour Detection in Digital Mammography. In: IEEE Transactions on Medical Imaging, Vol. 14 (1995) 565-576
- [99] Linguraru, M.G. Brady, J.M. Kadir, T. Yam, M.: A Novel Method to Detect Microcalcifications for Early Signs of Breast Cancer. In: Tenth World Congress on Health and Medical Informatics Medinfo2001. IOS Press, Amsterdam (2001) 946
- [100] Linguraru, M.G. Brady, J.M. Yam, M.: Filtering  $h_{int}$  Images for the Detection of Microcalcifications. In Niessen, W. Viergever, M. (eds.): Medical Image Computing and Computer-Assisted Intervention 2001, Lecture Notes in Computer Science, Vol. 2208. Springer-Verlag, Berlin Heidelberg New York (2001) 629-636
- [101] Linguraru, M.G. Brady, J.M.: An Anisotropic Diffusion Approach for Early Detection of Breast Cancer. In: Acta Universitatis Cibiniensis, Vol. XLIII. University of Sibiu (2001) 49-60
- [102] Linguraru, M.G. Brady, J.M.: A Non-Parametric Approach to Detecting Microcalcifications. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 339-341
- [103] Linguraru, M.G. Brady, J.M. Yam, M.: Detection of Microcalcifications using SMF. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 342-346
- [104] Linguraru, M.G. Marias, K. Brady, J.M.: Temporal Mass Detection. In: I Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 347-349
- [105] Logan-Young, W. Yanes Hoffman, N.: Breast Cancer: A practical Guide to Diagnosis, Vol. 1- Procedures. Mt. Hope Publishing (1994)
- [106] Loscocco, M.F. Thomas, J.A. Chakrabarti, K. Kaczmarek, R.V.: Preference Studies of Workstation Monitor Performance. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 612-616
- [107] Mallat, S.: A Wavelet Tour of Signal Processing. Academic Press (1999)
- [108] Marchese, M.: An Eye on CAD. In Medical Imaging, February (2002) 54-59
- [109] Marias, K., M., Brady, J.M. Highnam, R.P., Parbhoo, S., Seifalian, A.M.: Registration of Temporal Mammograms. In Hawkes, D. Hill, D. Gaston, R. (eds.): Medical Image Understanding and Analysis 99. BMVA, London (1999) 97-100
- [110] Marias, K.: Registration and Quantitative Comparison of Temporal Mammograms (with Application to HRT Data). Ph.D Thesis, University College London (2001)
- [111] Marias, K. Behrenbruch, C.P. Brady, M. Parbhoo, S. Seifalian, A.: Multi-Scale Landmark Selection for Improved Registration of Temporal Mammograms. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 580-586
- [112] Marias, K.: Texture Analysis for Mammography. MSc Dissertation, Imperial College of Science Technology and Medicine London (1997)

- [113] McLeod, G. Parkin, G.J.S. Cowen, A.R.: Automatic Detection of Clustered Microcalcifications in Mammograms. In Doi, K., Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): Digital Mammography. Elsevier, Amsterdam (1996) 311-316
- [114] McLoughlin, K.J. Bones, P.J. Kovesi, P.D.: Detection of Microcalcification Clusters in Digital Mammogram Images. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 353-357
- [115] Miller, L. Ramsay, N.: The Detection of Malignant Masses by Non-linear Multiscale Analysis. In Doi, K., Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): Digital Mammography. Elsevier, Amsterdam (1996) 335-340
- [116] The Million Women Study Collaborative Group: The Million Women Study: design and characteristics of the study population. In: Breast Cancer Research, Vol. 1 (1999) 73-80
- [117] Montagnat, J. Sermesant, M. Delingette, H. Malandain, G. Ayache, N.: Anisotropic Filtering for Model Based Segmentation of 4D Cylindrical Echocardiographic Images. In: Pattern Recognition Letters (2001)
- [118] Moon, P. Spencer, D.E.: The Specification of Foveal Adaptation. In: J. Opt. Soc. Amer., Vol. 33 (1943) 233-248
- [119] Morrone, M.C. Owens, R.A.: Feature Detection from Local Energy. In: Pattern Recognition Letters, Vol. 6 (1987) 303-313
- [120] Mudigonda, N.R. Rangayyan, R.M. Desautels, J.E.L.: Gradient and Texture Analysis for the Classification of Mammographic Masses. In: IEEE Trans. Med. Imaging, Vol. 19(10) (2000) 1032-1043
- [121] Nakahara, H. Namba, K. Fukami, A. Watanabe, R. Mizutani, M. Matsu, T. Nishimura, S. Jinnouchi, S. Nagamachi, S. Ohnishi, T. Futami, S. Flores, I.I.L.G. Nakahara, M. Tamura, S.: Computer-Aided Diagnosis (CAD) for Mammography: Preliminary Results. In: Breast Cancer, Vol. 5(4) (1998) 401-405
- [122] Neal, R.M. Hinton, G.E.: A New View of the EM Algorithm that Justifies Incremental and Other Variants In Jordan, M. (ed.): Learning in Graphical Models. Kluwer Academic Publishers, Dordrecht Boston London (1998) 355-368
- [123] Ng, S.L. Bischof, W.F.: Automated Detection and Classification of Breast Tumours. In: Comput.Biomed.Res, Vol. 25 (1992) 218-237
- [124] Nishikawa, R.M. Giger, M.L. Vyborny, C.J. Schmidt, R.A.: Computer-Aided Detection of Clustered Microcalcifications: An Improved Method for Grouping Detected Signals. In: Medical Physics, Vol. 10 (1993) 1661-1666
- [125] Noble, J.A.: Descriptions of Image Surfaces. DPhil thesis, University of Oxford (1989)
- [126] Orel, S.G. Mendonca, M.H. Reynolds, C. Schnall, M.D. Solin, L.J. Sullivan, D.C.: MR Imaging of Ductal Carcinoma in Situ. In: Radiology, Vol. 202(2) (1997) 413-420
- [127] D'Orsi, C.J., Basset, L.W. Feig, S.A. Jackson, V.P. Kopans, D.B. Linver, M.N. Sickles, E.A. Stelling, C.B.: BI\_RADS: Illustrated Breast Imaging Reporting and Data System. American College of Radiology, Third Edition (1998)

- [128] Ozdemir, A. Oznur, I.I. Vural, G. Atasever, T. Karabacak, N.I. Gokcora, N. Isik, S. Unlu, M.: TI-201 Scintigraphy, Mammography and Ultrasonography in the Evaluation of Palpable and Nonpalpable Breast Lesions: A Correlative Study. In: *Eur J Radiol* Vol. 24(2) (1997) 145-154
- [129] Parker, J. Dance, D.R., Davies, D.H.: Classification of Ductal Carcinoma In-Situ by Image-Analysis of Calcifications from Digital Mammograms. In: *British J. of Radiology*, Vol. 68 (1990) 241-246
- [130] Parr, T.C. Astley, S.M. Taylor, C.J. Boggis, C.R.M.: Model Based Classification of Linear Structures in Digital Mammograms. In Doi, K. Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): *Digital Mammography*, Elsevier, Amsterdam (1996) 351-356
- [131] Patrick, E.A. Moskowitz, M. Mansukhani, V.T. Gruenstein, E.I.: Expert Learning System Network for Diagnosis of Breast Calcifications. In: *Investigative Radiology*, Vol. 26 (1991) 534-539
- [132] Pearson, H.: Study Refines Breast Cancer Risks. In: *Nature News Service*, Macmillan Magazines Ltd (2002)
- [133] Perona, P. Malik, J.: Scale-space and Edge Detection using Anisotropic Diffusion. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol. 12 (1990) 629-639
- [134] Petrick, N. Chan, H.P. Sahiner, B. Helvie, M.A. Paquerault, S.: Preclinical Evaluation of a CAD Algorithm for Early Detection of Breast Cancer. In Yaffe, M.J. (ed.): *Digital Mammography*, Medical Physics Publishing, Madison (2000) 328-333
- [135] Petrick, N. Sahiner, B. Chan, H.P. Helvie, M.A. Paquerault, S. Hadjiiski, L.M.: Breast Cancer Detection: Evaluation of a Mass-Detection Algorithm for Computer-aided Diagnosis. In: *Radiology*, Vol. 224 (2002) 217-224
- [136] Pisano, E.T. Cole, E. Hemminger, B.M. Muller, K. Shumak, R. Yaffe, M. Chakraborty, D. Conant, E. Fajardo, L.L. Feig, S. Jong, R. Kopans, D. Maidment, A. Staiger, M. Williams, M.: Accuracy of Digital Mammography vs. Screen-Film Mammography in a Diagnostic Mammography Population. In Yaffe, M.J. (ed.): *Digital Mammography*, Medical Physics Publishing, Madison (2000) 504-511
- [137] Pohlman, S. Powell, K.A. Obuchowski, N.A. Chilcote, W.A. Grundfest-Broniatowski, S.: Quantitative Classification of Breast Tumours in Digitized Mammograms. In: *Med. Phys.* Vol. 23 (1996) 1337-1345
- [138] Poissonnier, M. Brady, J.M.: 'Noise Equalisation' for Microcalcification Detection? In Yaffe, M.J. (ed.): *Digital Mammography*, Medical Physics Publishing, Madison (2000) 334-341
- [139] Rangayyan, R.M. El-Faramawy, N.M. Desautels, J.E.L. Alim, O.A.: Measures of Acutance and Shape for Classification of Breast Tumours. In: *IEEE Trans. Med. Imaging*, Vol. 16(6) (1997) 799-810
- [140] Rodriguez-Florido, M.A. Krissian, K. Ruiz-Alzola, J. Westin, C.F.: Comparison of Two Restoration Techniques in the Context of 3D Medical Imaging. In Niessen, W. Viergever, M. (eds.): *Medical Image Computing and Computer-Assisted Intervention 2001*, Lecture Notes in Computer Science, Vol. 2208. Springer-Verlag, Berlin Heidelberg New York (2001) 1031-1039

- [141] Roehrig, J. Doi, T. Hasegawa, A. Hunt, B. Marshall, J. Romsdahl, H. Schneider, A. Sharbaugh, R. Zhang, W.: Clinical Results with R2 Imagechecker System. In Karssemeijer, N. Thijssen, M. Hendricks, J. van Erning, L. (eds.): Digital Mammography, Kluwer Academic Publishers, Dordrecht Boston London (1998) 395-400
- [142] Roelof, T. van Woundenberg, S. Hendriks, J.H.C.L. Boedicker, A. Evertsz, C.J.G. Karssemeijer, N.: Performance Evaluation of a Digital Reading Station for Screening Mammography. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 455-459
- [143] Rosen, D. Martin, b. Monheit, M. Wolff, G. Stanton, M.: A Bayesian Neural Network to Detect Microcalcifications in Digitized Mammograms. In Doi, K. Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): Digital Mammography, Elsevier, Amsterdam (1996) 277-286
- [144] Sahiner, B. Hadjiiski, L.M. Chan, H.P. Petrick, N. Helvie, M.A.: The Use of Joint Two-View Information for Improving Computerized Microcalcification Detection Accuracy. In: Digital Mammography, Lecture Notes in Computer Science, Springer-Verlag, Berlin Heidelberg New York (2002) 358-362
- [145] Sahiner, B. Chan, H.P. Petrick, N. Hadjiiski, L.M. Helvie, M.A. Paquerault, S: Active Contour Models for Segmentation and Characterisation of Mammographic Masses. In Yaffe, M.J. (ed.): Digital Mammography, Medical Physics Publishing, Madison (2000) 357-362
- [146] Sahiner, B. Chan, H.P. Petrick, N. Helvie, M.A.: Computerized Characterisation of Masses on Mammograms: The Rubber Band Straightening Transform and Texture Analysis. In: Medical Physics, Vol. 25 (1998) 516-526
- [147] Sahiner, B. Chan, H.P. Petrick, N. Wei, D. Helvie, M.A. Adler, D.D. Goodsitt, M.M.: Classification of Mass and Normal Breast Tissue: A Convolution Neural Network Classifier with Spatial Domain and Texture Images. In: IEEE Transactions on Medical Images, Vol. 15(10) (1996) 598-610
- [148] Samani, A. Bishop, J. Yaffe, M.J. Plewes, D.B.: Biomechanical 3-D Finite Element Modelling of the Human Breast Using MRI Data. In: IEEE Transactions on Medical Imaging, Vol. 20(4) (2001) 271-279
- [149] Sanchez-Ortiz, G.I. Rueckert, D. Burger, P.: Knowledge-Based Tensor Anisotropic Diffusion of Cardiac MR Images. In: Medical Image Analysis, Vol. 3(1) (1999) 77-101
- [150] Savage, C.J. Gale, A.G. Pawley, E.F. Wilson A.R.M.: To Err is Human; to Compute Divine? In Gale, A.G. Astley, S.M. Dance, D.R. Cairns, A.Y. (eds.): Digital Mammography, Elsevier, Amsterdam (1994) 405-414
- [151] Schlieff, R.: Ultrasound Contrast Agents. In: Radiology, Vol. 3 (1991) 198-207
- [152] Schmidt, R.A. Newstead, G.M. Linver, M.N. O'Shaughnessy, K.: Estimated Expected benefit of CAD in Mammographic Clinical Practice. In: 6<sup>th</sup> International Workshop on Digital Mammography, Bremen (2002), Invited Talk
- [153] Shen, L. Rangayyan, R.M. Desautels, J.E.L: Detection and Classification of Mammographic Calcifications. In: Int. Journal of Pattern Recognition and Artificial Intelligence, Vol. 71 (1993) 1403-1416

- [154] Schenk, V.U.B. Brady, J.M.: Finding CLS Using a Multiresolution Oriented Local energy Feature Detection Framework. In: *Digital Mammography, Lecture Notes in Computer Science*, Springer-Verlag, Berlin Heidelberg New York (2002) 64-68
- [155] Smith, S. M. Brady, J. M.: SUSAN - A New Approach to Low-level Image Processing. In: *International Journal of Computer Vision*, Vol. 23(1) (1997) 45-78
- [156] Spiegel, M. R.: *Probability and Statistics*. McGraw-Hill Book Company, New York (1975)
- [157] Stack, J.P. Redmond, O.M. Codd, M.B. Dervan, P.A. Ennis, J.T: Breast Disease: Tissue Characterisation with Gd-DTPA Enhancement Profiles. In: *Radiology*, Vol. 174 (1990) 491-494.
- [158] Stafford, G.: *Digital Mammography*.  
<http://members.ozemail.com.au/~glensan/digmam2.htm>
- [159] Swann, C.A. Kopans, D.B. Koerner, F.C. McCarthy, K.A. White, G. Hall, D.A.: The Halo Sign and Malignant Breast Lesions. In: *American Journal of Roentgenology*, Vol. 149 (1987) 1145-1147
- [160] Tanner, C. Schnabel, J.A. Chung, D. Clarkson, M.J. Rueckert, D. Hill, D.L.G. Hawkes, D.J.: Volume and Shape Preservation of Enhancing Lesions when Applying Non-rigid Registration to a Time Series of Contrast Enhancing MR Breast Images. In *Proc. MICCAI 2000, Lecture Notes in Computer Science*, Springer-Verlag, Berlin Heidelberg New York (2000) 462-471
- [161] Taylor, P. Fox, J. Todd-Pokropek, A.: Evaluation of a Decision Aid for the Classification of Microcalcifications. In Karssemeijer, N. Thijssen, M. Hendricks, J. van Erning, L. (eds.): *Digital Mammography*, Kluwer Academic Publishers, Dordrecht Boston London (1998) 237-244
- [162] Taylor, P. Lee, R. Alberdi, E.: Incorporating Explanations into a CAD Tool for the Classification of Calcifications. In: *Digital Mammography, Lecture Notes in Computer Science*, Springer-Verlag, Berlin Heidelberg New York (2002) 372-376
- [163] Thurfjell, E.L. Lindgren, J.A.: Breast Cancer Survival Rates with Mammographic Screening: Similar Favorable Survival Rates for Women Younger and those Older than 50 years. In: *Radiology*, Vol. 201(2) (1996) 421-426
- [154] Valverde, F.L. Munoz, J. Nishikawa, R. Doi, K: Elimination of Calcified False Positives in detection of Microcalcifications in Mammograms Using Hough Transform. In Yaffe, M.J. (ed.): *Digital Mammography*, Medical Physics Publishing, Madison (2000) 383-389
- [165] Veldkamp, W.J.H. Karssemeijer, N. Hendricks, J.H.C.L: Automated Classification of Clustered Microcalcifications into Malignant and Benign Types. In: *Medical Physics*, Vol. 27 (2000) 2600-2608
- [166] Veldkamp, W.J.H. Karssemeijer, N. Hendrick, J.H.C.L.: Fully Automated Classification of Microcalcification Cases Referred from a Nationwide Screening Program. In Yaffe, M.J. (ed.): *Digital Mammography*, Medical Physics Publishing, Madison (2000) 390-394
- [167] Veldkamp, W.J.H. Karssemeijer, N.: Accurate Segmentation and Contrast Measurement of Microcalcifications in Mammograms: A Phantom Study. In: *Medical Physics*, Vol. 25(7) (1998) 1102-1110

- [168] Viehweg, P. Lampe, D. Buchmann, D. Heywang-Koebrunner, S.H.: In Situ and Minimally Invasive Breast Cancer: Morphologic and Kinetic Features on Contrast-Enhanced MR Imaging. In: MAGMA, Vol. 11(3) (2000) 129-137
- [169] Warren Burhenne, L.J. Wood, S.A. D'Orsi, C.J. Feig, S.A. Kopans, D.B. O'Shaughnessy, K.F. Sickles, E.A. Tabar, L. Vyborny, C.J. Castellino, R.A.: The Potential Contribution of Computer Aided Detection to the Sensitivity of Screening Mammography. In: Radiology Vol. 215 (2000) 554-562
- [170] Webb, S.(ed.): The Physics of Medical Imaging. Institute of Physics Publishing, Bristol Philadelphia (1995)
- [171] Weickert, J.: A Review of Nonlinear Diffusion Filtering. In ter Haar Romeny, B. Florack, L. Koenderink, J. Viergever, M. (eds.): Scale-Space Theory in Computer Vision, Lecture Notes in Computer Science. Springer-Verlag, Berlin Heidelberg New York (1997) 3-28.
- [172] Weickert, J.: Anisotropic Diffusion in Image Processing. B.G. Teubner, Stuttgart (1998)
- [173] Westin, C.F. Maier, S.E. Khidhir, B. Everett, C. Jolesz, F.A. Kikinis, R.: Image Processing for Diffusion Tensor Magnetic Resonance Imaging. In MICCAI (1999) 441-452
- [174] Wolfe, J.N.: Breast Patterns as an Index of Risk for Developing Breast Cancer. In: American Journal of Roentgenology, Vol. 126 (1976) 1130-1139
- [175] Woods, K. Bowyer, K.: A General View of Detection Algorithms. In Doi, K. Giger, M.L. Nishikawa, R.M. Schmidt, R.A. (eds.): Digital Mammography, Elsevier, Amsterdam (1996) 385-390
- [176] Xiao, G.: 3D Free-hand Ultrasound Imaging of the Breast, Ph.D Thesis. Department of Engineering Science, The University of Oxford (2001)
- [177] Yaffe, M.: Development of Full Field Digital Mammography. In Karssemeijer, N. Thijssen, M. Hendricks, J. van Erning, L. (eds.): Digital Mammography, Kluwer Academic Publishers, Dordrecht Boston London (1998) 3-10
- [178] Yam, M. Brady, J.M. Highnam, R.P. English, R.: Denoising  $h_{int}$  Surfaces: a Physics-based Approach. In: Taylor, C. Colchester, A. (eds.): Medical Image Computing and Computer-Assisted Intervention, Springer-Verlag, Berlin Heidelberg New York (1999) 227-234
- [179] Yam, M. Brady, J.M. Highnam, R.P. English, R.: Detecting Calcifications in Mammograms using the  $h_{int}$  Representation. In: Lemke, H.U. Vannier, M.W. Inamura, K. Farman, A.G. Computer Assisted Radiology and Surgery, Elsevier (1999) 373-377
- [180] Yam, M., Brady, J.M. Highnam, R.P. Behrenbruch, C.P. English, R. Kita, Y.: Three-dimensional Reconstruction of Microcalcification Clusters from Two Mammographic Views. In: IEEE Trans Med Imaging, Vol. 20 No. 6 (2001) 479-489
- [181] Yates, K. Evans, C. Brady, J.M: Improving teBrake's Mammographic Mass-Detection Algorithm Using Phase Congruency. In: Proc. Digital Image Computing Techniques and Applications (DICTA2002), Melbourne (2002)

- [182] Yin, F.F. Giger, M.L. Doi, K. Metz, C.E. Vyborny, C.J.: Computerized Detection of Masses in Digital Mammograms: Analysis of Bilateral Subtraction Images. In: Medical Physics, Vol. 18 (1991) 955-963
- [183] Zhang, Y. Brady, J.M. Smith, S.: Segmentation of Brain MR Images Through a Hidden Markov Random Field Model and the Expectation Maximization Algorithm. In: IEEE Transactions on Medical Imaging, Vol. 20(1) (2001) 45-57
- [184] Zwiggelaar, R., Parr, T.C. Schumm, J.E. Hutt, I.W. Taylor, C.J. Astley, S.M. Boggis, C.R.M.: Model-Based Detection of Spiculated Lesions in Mammography. In: Medical Image Analysis, Vol. 3(1) (1999) 39-62
- [185] Zwiggelaar, R., Taylor, C.J. Rubin, C.M.E.: Detection of the Central Mass of Spiculated Lesions – Signature Normalisation and Model Data Aspects. In: IPMI'99, Springer (1999) 406-411
- [186] Altruis Biomedical Network - <http://www.lung-cancers.com/>
- [187] Altruis Biomedical Network - <http://www.breast-cancers.com/>
- [188] The Center for Orthopaedic Oncology and Musculoskeletal Research - <http://www.sarcoma.org>
- [189] The Breast Clinic - <http://www.thebreastclinic.com/>
- [190] American College of Radiology – Mammography Quality Control Manual (1999)



## List of Figures

- Figure 1 The appearance of young versus menopausal breasts in mammography: (a) A pair of left and right cranio-caudal mammograms of a pre-menopausal woman with very dense appearance due to the presence of milk-producing tissue; (b) a pair of left and right cranio-caudal mammograms of a post-menopausal woman where there is a larger amount of fat-tissue, which makes the depiction of dense areas simpler. .... 4
- Figure 2 A brief anatomy of the breast showing the branching internal structure of ducts and lobules. .... 5
- Figure 3: The FROC curve of the microcalcification-detection method. ....16
- Figure 4: Detection example 1; (a) the original contrast-enhanced SMF sample with a very large microcalcification cluster in a dense area of the breast; (b) the detection map of the detection method presented in this thesis depicting correctly the cluster. .... 17
- Figure 5: Detection example 3; (a) the original contrast-enhanced SMF sample with a subtle microcalcification clusters in a dense breast area; (b) the detection map..... 17
- Figure 6: Detection example2; (a) The original contrast-enhanced SMF sample with a widespread microcalcification cluster; (b) the detection map .....18
- Figure 7 Fig Example4. (a) The original contrast-enhanced SMF sample with a cluster of very small microcalcifications in an area with several curvilinear structures; (b) the detection map.....18
- Figure 8: The clinical mammographic film-screen machine. The X-rays pass through the compressed breast from the X-ray source towards the film cassette..... 24
- Figure 9: A representation of mammographic image formation and scattered radiation. .... 24
- Figure 10: The combination of four images used by the clinician in diagnosing the breast; (a) a pair of MLO images showing the breast, the pectoral muscle and axilla; (b) a pair of CC images from the same patient. .... 26
- Figure 11: A soft copy environment produces digital mammograms, which are stored directly into a computer. This figure shows the Opdima System from Siemens, which was designed for near real-time computer guided biopsy..... 28
- Figure 12: The MRI scanner with a patient in the right position for an MRI of the breast. .... 29
- Figure 13: An axial T1-weighted MRI slice of the breast (using a gradient echo sequence), the closest view to the X-ray CC image. In this image, fatty tissue appears brighter, as its T1 value (around 200 ms) is considerably less either that that of normal healthy tissue (typically 700-1000 ms) and of cancerous tissue (typically 1500 ms). .... 30
- Figure 14: A sequence of contrast-enhanced MR images of the breast; (a) is the MR slice before the injection of contrast agent; (b) is the MR slice after contrast enhancement, where the tumour is highlighted due to its higher vascularity.....31
- Figure 15: An ultrasound machine; as the probe is swept over the patient's body, the clinician can visualise in real-time the US images on the machine monitor. .... 32

- Figure 16: Two US images of the breast; (a) the image of a cyst, a compact dark area squeezed by the probe; (b) the image of a tumour, an elongated dense area with less well-defined margins. .... 33
- Figure 17: A PET scanner..... 35
- Figure 18: An example of scintimammograms showing the sensitivity of SPECT images to multi-focal tumours. .... 35
- Figure 19: Some examples of masses in X-ray mammography; (a) a benign cyst will well-defined boundaries and an ellipsoidal self-contained shape; (b) a stellate mass with spiculations radiating from the central mass into the surrounding tissue; (c) an ill-defined mass with low contrast at the boundary, which can be easily confused with the neighbouring tissue. .... 39
- Figure 20: Some examples of microcalcifications in X-ray mammography; (a) an isolated large calcification; (b) a group of two isolated macrocalcifications; (c) a subtle cluster of microcalcifications following the shape of a duct (ductal carcinoma *in-situ*); (d) a compact malignant microcalcification cluster..... 46
- Figure 21: The extraction of candidate microcalcification regions in Yam's algorithm.....51
- Figure 22: Computing the interesting tissue volume (removing the background) in Yam's algorithm..... 52
- Figure 23: The  $h_{int}$  surface; (a) a mammogram presenting a lump; (b) the SMF that is generated from the mammogram where the ducts become ridges, and the mass a mountainous area..... 54
- Figure 24:** The glare process..... 58
- Figure 25:** The cylinder used in scatter estimation..... 59
- Figure 26:** Modelling the breast for the estimation of extra-focal radiation component.....61
- Figure 27: The Standard Mammogram Form image of a breast; (a) an MLO digitised intensity mammogram; (b) the correspondent SMF image. .... 63
- Figure 28: The plot of the 2D Gaussian, where the central point (the top of the hill-like shape) will be gradually smoothed into the background. .... 67
- Figure 29: The diffusion tensor eigenvectors;  $v_1$  is parallel with the edge gradient and the smoothing is inhibited across the edge;  $v_2$  is orthonormal to  $v_1$  and the diffusion is permitted along the edge. .... 70
- Figure 30: Some comparative diffusion results. (a) the original image; (b), (c), (d) the smoothed image with nonlinear isotropic diffusion (Perona-Malik) after 20, 40 and 100 iterations; (e), (f), (g) the smoothed image with nonlinear anisotropic diffusion (Weickert) after 20, 40 and 100 iterations; (h), (i), (j) the smoothed image with linear diffusion after 20, 40 and 100 iterations. .... 72
- Figure 31: The shape differentiation between microcalcifications and noise; (a) the plot of a filtered (de-noised) intensity image sample containing a microcalcification; (b) the plot of a filtered intensity image sample containing noise; while the microcalcification has the appearance of a hill with less steep edges, the bit of noise is rather spiky and has a higher value of intensity; (c) the plot of a  $h_{int}$  image containing a microcalcification; (d) the plot of a  $h_{int}$  image containing noise. Each plot is taken from one line in an image. .... 76
- Figure 32: (a) The original grey-level image containing a microcalcification in the centre-right of the image and a large spot of noise on the lower side of the image; (b) the diffused image with  $k=5$ ,  $\sigma=0.6$  and  $t=20$ , we notice that the edges of the important structures of the dense tissue are emphasised; (c) the diffused image with  $k=5$ ,  $\sigma =0.5$  and  $t=40$ , where only the important small structures are kept and their edges enhanced; ..... 78
- Figure 33: The changes in SNR during  $h_{int}$  generation: (a) the original phantom with simulated microcalcifications and noise; (b) the Wiener-filtered phantom; (c) the  $h_{int}$  image before glare deconvolution; (d) the  $h_{int}$  image after glare deconvolution; (e) the 3D plot of the

- original image in (a); (f) the plot of the smoothed image in (b) with improved SNR; (g) the plot of the  $h_{int}$  image in (c); (h) the noisier plot of the  $h_{int}$  image in (d). ..... 82
- Figure 34 Image subtraction; (a) The original preprocessed SMF image containing a microcalcification on the left side, a large spot of noise on the lower right side and several other smaller noise structures; (b) the 3D plot of the difference image between the original image diffused with  $k=15$ ,  $\sigma=0.6$  and  $t=5$  and the same one diffused with  $k=15$ ,  $\sigma=0.6$  and  $t=10$ ; (c) the original image diffused with  $k=15$ ,  $\sigma=0.6$  and  $t=10$  and the same one diffused with  $k=15$ ,  $\sigma=0.6$  and  $t=15$ . We notice that after a few iterations the big changes appear at the location of noise only. .... 83
- Figure 35: Filtering example 1; (a) The original preprocessed SMF image containing a microcalcification on the left side and a large spot of noise on the lower right side and several other noise structures; (b) the diffused SMF image with  $k=15$ ,  $\sigma=0.6$  and  $t=5$ , we notice that the microcalcification has almost faded, while the noise is still preserved with high contrast; (c) the noisy 3D plot of the original SMF image in (a); (d) the surface of the diffused SMF image in (b), the microcalcification appears as a hill with smoother edges than those of the very sharp-edged noise structures in the same image..... 84
- Figure 36 Filtering example 2; (a) The original preprocessed SMF image containing only noise structures, the largest piece of noise on the upper right side could be easily considered of being a microcalcification since it does not present very high contrast from the surrounding tissue; (b) the diffused SMF image with  $k=15$ ,  $\sigma=0.6$  and  $t=3$ ; (c) the 3D plot of the original SMF image in (a) with highly noisy appearance; (d) the 3D plot of the diffused SMF image in (b) where all structures have very sharp edges and are labelled as noise. .... 85
- Figure 37: The FROC curve of the detection method for the set of 13 samples with coarse calcifications..... 87
- Figure 38: The FROC curve of the combined detection method for the set of 20 samples containing different types of microcalcifications. .... 88
- Figure 39: Filtering example 3; the left column presents the original SMF image (a) and its two diffused versions for the sets of parameters  $k=5$ ,  $\sigma=0.5$ ,  $t=40$  (c) and  $k=15$ ,  $\sigma=0.6$ ,  $t=2$  (e); the right column shows the 3D surfaces of the three respective images. .... 89
- Figure 40: The diagram of the foveal algorithm to detect microcalcifications in Standard Mammogram Form images. The top row underlines the typical generation of an SMF image, including glare removal. The bottom row highlights the method described above: pre-processing, statistical analysis, image enhancement and adaptive segmentation..... 95
- Figure 41: Removing artefacts; (a) the original image; (b) the binary shot-noise map (white dots are noise); (c) the binary CLS map; (d) the 'clean' image after shot-noise and CLS removal..... 100
- Figure 42: An example of estimating  $k$ . The image shown on the left (after expanded display contrast) has the associated histogram of function  $g$  in the right. We note the zero mean value of  $g$ , as well as where the value of 0.47 of  $k$  falls..... 102
- Figure 43: Automated image enhancement; (a) an image with an isolated calcification; (b) an image with a microcalcification cluster; (c) the corresponding gradient map for image (a) depicting the microcalcification and some extra undesired regions; (d) the corresponding gradient map for image (b) with a good representation of the cluster, but some falsely suspicious areas as well; (e) the automatically diffused image (a); (f) the diffused image (b);..... 103
- Figure 44: An illustration of the lightness assimilation. We show three synthetic images with dark (left), medium (middle) and bright (right) backgrounds. All have central objects of the same size and intensity, but are perceived differently by our eyes, due to the variance in background lightness..... 105
- Figure 45: The foveal masks used for the computation of  $\mu_O$ ,  $\mu_N$  and  $\mu_B$ . The object  $O$  is the area of the *fovea centralis*,  $N$  its neighbourhood (twice the size of  $O$  in our applications) and  $B$  the background.....107

- Figure 46: The simulation of a plot of a mammogram section containing microcalcifications over height/intensity variation. The variation of the perceivable contrast in the detection of microcalcifications is suited to the local characteristics for the adaptation of HVS using  $C_{min}$ . The classical minimal perceivable measure, (here called  $C_{thresh}$ ) is a global characteristic of the mammogram and less flexible in the elimination of FP in the detection of microcalcifications. .... 108
- Figure 47: A set of five synthetic images with variations between object, neighbourhood and background and their associated cross-sections. These examples cover a wide aspect of contrasts in image processing: bright on dark, bright on bright, dark on dark. The corresponding values of  $\mu_O$ ,  $\mu_N$ ,  $\mu_B$  and  $C$  are shown in the table below the figure. The variation of the adaptive threshold is shown in Figure 48. .... 109
- Figure 48: The variation of the adaptive threshold  $C_{min}$  for the synthetic images in Figure 47 in the central area of images, which corresponds to the object and neighbourhood surface. For this example, we used  $c_w=0.03$ . The value of  $C_{min}$  in the centre of fovea is shown in the table below the figure. We note that for the most delicate case (extreme right),  $C_{min}$  and  $C$  are in the same range of values. In such difficult cases, which approximate better the mammographic environment where transits between different intensities are much smoother, the adaptation of  $C_{min}$  become crucial. Moreover,  $C > C_{min}$  in all five cases and all peaks are detected. .... 110
- Figure 49: The FROC curve of the microcalcification-detection method based on the adaptation of HVS in digital mammography. .... 111
- Figure 50: Detection example 1: (a) the original SMF images with a microcalcification cluster; (b) the gradient map from the statistical analysis depicting suspicious pixels; (c) the CLS map; (d) the enhanced image after diffusion; (e) the microcalcification detection map. ... 112
- Figure 51: Detection example 2: (a) the original SMF images with a microcalcification cluster; (b) the gradient map; (c) the CLS map; (d) the enhanced image; (e) the microcalcification detection map. .... 112
- Figure 52: Detection example 3: (a) the original SMF images with a microcalcification cluster; (b) the gradient map; (c) the CLS map; (d) the enhanced image; (e) the microcalcification detection map. .... 113
- Figure 53: Detection example 4: (a) the original SMF images with a microcalcification cluster; (b) the gradient map; (c) the CLS map; (d) the enhanced image; (e) the microcalcification detection map. .... 113
- Figure 54: Detection example 5: (a) the original SMF images with a very large microcalcification cluster; (b) the gradient map; (c) the CLS map; (d) the enhanced image; (e) the microcalcification detection map. .... 114
- Figure 55: A case of difficult detection with a faint microcalcification cluster in the breast margin. The contrast in the image has been enhanced for the reader to help in the visualisation of the cluster. .... 115
- Figure 56; Another example of CLS removal in detection: (a) the original mammogram; (b) the detection map using the CLS removal; (c) the detection map without CLS removal with a few extra FP detected. .... 116
- Figure 57: CLS removal in detection: (a) the original mammogram; (b) the detection map using the CLS removal; (c) the detection map without CLS removal with a few extra FP detected. .... 117
- Figure 58: The comparative FROC curve when CLS are removed or not prior to the microcalcification detection. .... 117
- Figure 59: The comparative FROC curve when smoothing by anisotropic diffusion is performed or not prior to the segmentation of microcalcifications. .... 118

- Figure 60: Image smoothing in detection; (a) the original mammogram sample; (b) the detection map using anisotropic diffusion; (c) the detection map without using smoothing with FP marked..... 119
- Figure 61: The comparative FROC curve when  $w$  is varied over a range of 5 to 10% of its default value of 0.923. The difference in detection results is quite small and all four algorithms converge smoothly to 100% TP ratio..... 120
- Figure 62: Intensity versus SMF comparative FROC curve. The detection algorithm converges slightly slower for intensity images, but reaches the same performance as for SMF images. One reason for the delay could be the use of same parameters when building the FROC curve, although the image characteristics (intensity versus SMF) are different.... 121
- Figure 63: The comparative FROC curve between the detection results on intensity images from the Oxford Screening Database and the University of South Florida Digital Database for Screening Mammography . .....122
- Figure 64: The comparative FROC curve of the detection of microcalcifications when mammogram samples are used versus full mammograms. The behaviour of the algorithm is similar and robust with the image size. ....124
- Figure 65: Detection results on whole mammograms; (a) and (c) are the MLO SMF images, while (b) and (d) the corresponding CC images. Ellipses indicate the locations and spread of the detected microcalcification clusters, while arrows indicate the positions of FP. ....125
- Figure 66: The response of local energy to a variety of input signals: (a) the input signal; (b) the local energy response; (c) the phase-angle response.....127
- Figure 67: The angular part: a steerable filter. .... 128
- Figure 68: A simplified example of edge-ridge-edge triplet. The background is a scale image; ridges are shown in green, while edges in red. The lengths of the vectors express scale. ....129
- Figure 69: An example of CLS detection: (a) the original mammogram sample containing a mass, two clusters of microcalcifications and CLS; (b) the CLS map over 3 scales; (d) a close-up of the upper right corner of image (a) showing the edge-ridge-edge triplets (edges in green, ridges in red). The CLS are detected using manual thresholding over each scale. .... 130
- Figure 70: An example of microcalcification detection: (a) the original mammogram sample; (b) the CLS-free mammogram sample using Schenk and Brady's algorithm; (c) the CLS map before CLS removal with a large number of FP; (d) the CLS map after CLS removal with improved results and a significant reduction of the number of FP..... 131
- Figure 71: Comparative Results 1 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (c) the detection map of the Foveal Approach. ....134
- Figure 72: Comparative Results 2 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (c) the detection map of the Foveal Approach. ....134
- Figure 73: Comparative Results 3 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (c) the detection map of the Foveal Approach. ....135
- Figure 74: Comparative Results 4 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (c) the detection map of the Foveal Approach. ....136
- Figure 75: Comparative Results 5 for the Detection of Microcalcifications in SMF; (a) the original SMF image sample; (b) the detection map of the Physics-based Approach; (c) the BWMD of the Statistical Analysis; (c) the detection map of the Foveal Approach. ....137
- Figure 76: The FROC curves of the three microcalcification-detection methods, where we notice the better performance of the Foveal Approach. .... 138

- Figure 77:** In the left case, the  $k$  factor is bigger than the gradient and the diffusivity function  $g \rightarrow 1$ , which is equivalent to finding an edge and maximising diffusion; in the right case,  $g \rightarrow 0$  (for very big gradients) and diffusion is inhibited.....139
- Figure 78:** In the left case, the filter width is small compared to the structure, so it essentially detects a step edge; but in the second, it is not obvious that it will ..... 140
- Figure 79: The original idea for the detection of microcalcifications: (a) what we initially proposed; (b) what we thought it was done; (c) what was done in reality; (d) our solution.....144
- Figure 80: The plot of a benign mass; (a) a mammogram sample containing a benign tumour with well-defined margins; (b) the 3D plot of image (a) where the tumour appears as a high hill surrounded by several smaller structures of normal tissue and noise..... 151
- Figure 81: The plot of a malignant mass; (a) a mammogram sample containing a malignant spiculated tumour; (b) the 3D plot of image (a) where the tumour appears as a structures of high ridges descending along the spicules and surrounded by several smaller structures of normal tissue and noise. ....152
- Figure 82: Diffusing a benign mass 1; (a) the mammogram sample containing a benign tumour in Figure 80 after diffusion with  $t=10$ ,  $k=15$ ,  $\sigma=0.8$  (small number of iterations and high contrast); (b) the 3D plot of the diffused image in (a); (c) the SSD image between the original not-blurred imaged and the diffused one – the latter image is cleaned by its high-frequency component, but the inner surface of the mass is also diffuse because of its iso-density. ....152
- Figure 83: Diffusing a benign mass 2; (a) the mammogram sample containing a benign tumour in Figure 80 after diffusion with  $t=400$ ,  $k=5$ ,  $\sigma=0.8$  (large number of iterations and small contrast); (b) the 3D plot of the diffused image in (a) with flat background; (c) the SSD image between the original not-blurred imaged and the diffused one – the latter image has a ‘clean’ background, since almost everything else, but large dense regions has been removed; at this high number of iterations, the inner surface of the benign lesion (which is roughly uniform) is still diffused. ....153
- Figure 84: Diffusing a malignant mass 1; (a) the mammogram sample containing a malignant spiculated tumour in Figure 81 after diffusion with  $t=8$ ,  $k=8$ ,  $\sigma=0.8$  (small number of iterations and high contrast); (b) the 3D plot of the diffused image in (a); (c) the SSD image between the original not-blurred imaged and the diffused one – the latter image is cleaned by its high-frequency component, while the complex geometry of the malignant mass is seen as a combination of edges/ridges and is not diffused.....153
- Figure 85: Diffusing a malignant mass 2; (a) the mammogram sample containing a malignant spiculated tumour in Figure 81 after diffusion with  $t=40$ ,  $k=5$ ,  $\sigma=0.8$  (large number of iterations and small contrast); (b) the 3D plot of the diffused image in (a) with mainly flat background; (c) the SSD image between the original not-blurred imaged and the diffused one – the high-frequency map is similar to the one in Figure 84, since the inner surface of the malignant lesion is irregular and the ridges are perceived as edges. ....154
- Figure 86: The *diffusion feature* on benign masses; (a) and (c) the original mammogram samples showing benign lesions; (b) and (d) the SSD images corresponding to (a) and (c), respectively; both show that diffusion is allowed within the mass area. ....156
- Figure 87: The *diffusion feature* on malignant masses; (a) and (c) the original mammogram samples showing malignant lesions; (b) and (d) the SSD images corresponding to (a) and (c), respectively; only high frequency structures are diffused within the mass area. .... 157
- Figure 88: The *diffusion feature* on exceptions; (a) an original mammogram samples showing a benign lesions; (b) the SSD images corresponding to (a) where the diffusion was mainly inhibited within the mass area; (c) an original mammogram samples showing a malignant lesions; (d) the SSD image corresponding to (b) where diffusion is allowed within the mass area.....158
- Figure 89: The diagram of the algorithm proposed for temporal mass prompting.....159

- Figure 90: Mammogram registration is performed as a pre-processing step in order to facilitate the comparison between the temporal mammogram pair for mass detection. (a) and (b) comprise the original mammogram pair, while (c) is the registered mammogram (b) to (a). ..... 160
- Figure 91: Texture classification 1 of the registered MLO mammogram pair into the classes A, B, C and D described in the previous section. The top row shows the firstly taken mammogram, while the bottom row shows the most recent mammogram. Both mammograms are registered, but not enhanced using anisotropic diffusion before the texture classification is applied. ....163
- Figure 92: Texture classifications 1 of the diffused and registered MLO mammogram pair into the classes A, B, C and D. The top row shows the firstly taken mammogram, while the bottom row shows the most recent mammogram. Both mammograms are registered and enhanced using anisotropic diffusion before the texture classification is applied.....163
- Figure 93: The detected “suspicious” regions 1 (from the diffused pair) superimposed in the original MLO mammogram. The current mammogram (b) prompts a real tumour with no correspondent in the earlier mammogram (a). .....164
- Figure 94: Texture classification 2 of the registered mammogram pair into the classes A, B, C and D in a CC mammogram pair. The top row shows the firstly taken mammogram, while the bottom row shows the most recent mammogram. Both mammograms are registered, but not enhanced using anisotropic diffusion before the texture classification is applied. ....164
- Figure 95: Texture classifications 2 of the diffused and registered CC mammogram pair into the classes A, B, C and D. The top row shows the firstly taken mammogram, while the bottom row shows the most recent mammogram. Both mammograms are registered and enhanced using anisotropic diffusion before the texture classification is applied.....165
- Figure 96: The detected “suspicious” regions 2 (from the diffused pair) superimposed in the original CC mammogram. The current mammogram (b) prompts a real tumour with no correspondent in the earlier mammogram (a). .....165
- Figure 97: The vector flow of a benign mass; (a) the vector flow of the original mammogram sample; excepting the boundaries of the mass, the vectors flow chaotically; (b) the vector flow of the diffused mammogram; note the radial pattern of vectors pointing towards the centre of the mass. ....169
- Figure 98: The vector flow of a malignant mass; (a) the vector flow of the original mammogram sample; excepting the boundaries of the mass, the vectors flow chaotically; (b) the vector flow of the diffused mammogram; although the vectors point towards the inside of the lesion, there is no definite centre of the mass; this ‘multi-focality’ is due to the rough surface of the cancerous mass.....170
- Figure 99: The *uniformity feature* on benign masses; (a) the vector flow of the original mammogram sample showing a benign lesion; (b) the vector flow of the diffused images corresponding to (a) showing a smooth transition along the edge of the mass;.....170
- Figure 100: The *uniformity feature* on malignant masses; (a) the vector flow of the original mammogram sample showing a malignant lesions; (b) the vector flow of the diffused images corresponding to (a) where the transitions along the edge are still rapid; ..... 171
- Figure 101: Histogram comparison: (a) and (c) are two mammogram samples with microcalcification clusters and CLS; (b) and (d) are their respective histograms, where the blue continuous plot corresponds to microcalcifications, while the red dotted plot is related to CLS;..... 183
- Figure 102: Errors in CLS estimation: (a) the original contrast-enhanced mammogram sample with a microcalcification cluster; (b) the best estimation of CLS, which erroneously includes the microcalcification cluster; (c) the CLS map used in our algorithm, which may cause the disruption of CLS in the removal step..... 184
- Figure 103: 3D image representation and quantisation: (a) a 2D phantom of Gaussian intensity variation used here as the original image; (b) the 3D surface of image (a); (c) the grey-

level threshold levels ( $l_1, l_2, l_3 \dots$ ) on a surface model; (b) the iso-contours super-imposed on image (a). .....185

Figure 104: Iso-contours and their 3D plots used to reduce the number of FP in the detection of microcalcifications: (a) the original mammogram sample with a microcalcification cluster; (b) the microcalcification map with a conservative selection of parameters; (c), (e) the iso-contours around two selected FP as seed pixels (the FP are marked in blue); (d), (f) the projections of the 3D plots of (c) and (e) with large elongated contours below the seed point; (g), (i) the iso-contours around two selected TP as seed pixels (marked in blue); (h), (j) the projections of the 3D plots of (g) and (i) with thin round contours below the seed point and above the noise level. .... 188

Figure 105: Log-Gabor functions of different bandwidths and centre frequencies. ....192

Figure 106: Phase congruency and phase of a sample signal .....193

Figure 107: Phase angles of different feature type .....194

Figure 108: Detecting features using the phase angle .....195

Figure 109: Test (a) Idealised test image; (b) Phase Congruency; (c) Orientation vectors .....197

Figure 110: Phase Congruency at different scales. .... 198

Figure 111: The basic steps of our breast registration algorithm (reproduced from [106]). .....202

Figure 112: Consistent landmarks in the CC and ML "idealised" outlines.....204

Figure 113: An example of ROC analysis.....214



## List of Tables

Table 1: The strengths and weaknesses of breast imaging procedures .....	37
Table 2: The linear coefficients for various tissue types reported by Highnam and Brady [65] after Johns and Yaffe. The coefficients of fibrous tissue and tumour overlap, while that of fat is clearly distinctive. Microcalcifications also have different attenuation coefficients, much higher than that of fibrous tissue. ....	55
Table 3: Variation of anisotropic diffusion parameters: $k$ - the contrast factor, $\sigma$ - the scaling factor and $t$ - the number of iterations; $\nearrow$ represents an increase in the associated feature, as opposite to decrease for $\searrow$ .....	81
Table 4: The variation of SNR in generating SMF. ....	145
Table 5: True positives and false positives in 15 pairs of mammograms (a mass has previously been diagnosed in each pair). ....	167
Table 6: Comments on registration results in 50 mammogram pairs: The viewer classified the results in three categories according to the alignment of image features. ....	209
Table 7: Clinical assessment of the improvement in registration using internal landmarks in 25 mammogram pairs and comparison with the reduction in the standard deviation of the difference image after geometrical alignment for the same cases. ....	209