A Biologically Inspired Algorithm for Microcalcification Cluster Detection

Marius George Linguraru, Kostas Marias, Ruth English, Michael Brady

A Biologically Inspired Algorithm
for Microcalcification Cluster Detection

Marius George Linguraru[1,2], Kostas Marias[1,3], Ruth English[4] and Michael Brady[1]

[1]University of Oxford, Medical Vision Laboratory, Oxford, UK
[2]Harvard University, Division of Engineering and Applied Sciences, Cambridge MA, USA
[3]Institute of Computer Science, Foundation for Research and Technology
- Hellas FORTH, Heraklion, Crete, Greece
[4]Churchill Hospital, Breast Care Unit, Oxford, UK

Abstract— The early detection of breast cancer greatly improves prognosis. One of the earliest signs of cancer is the formation of clusters of microcalcifications. We introduce a novel method for microcalcification detection based on a biologically inspired adaptive model of contrast detection. This model is used in conjunction with image filtering based on anisotropic diffusion and curvilinear structure removal using local energy and phase congruency. An important practical issue in automatic detection methods is the selection of parameters: we show that the parameter values for our algorithm can be estimated automatically from the image. This way, the method is made robust and essentially free of parameter tuning. We report results on mammograms from two databases and show that the detection performance can be improved by first including a normalisation scheme.

Index Terms— X-ray mammography; microcalcification clusters; image normalization; anisotropic diffusion; human visual system.

I. INTRODUCTION

Breast screening programs attempt to detect and eradicate cancer at the earliest possible stage to reduce the rate of mortality amongst women. However, detecting the early signs of breast cancer that appear in X-ray mammograms presents a significant challenge to radiologists, not least because such signs are often difficult to distinguish against the dense or highly textured breast anatomy. Studies prove that the performance of radiologists is increased when a good segmentation or classification system is used (Funovics et al., 1998; Giger et al., 2002; Petrick et al., 2002). However, such systems must have provably adequate performance and robustness to be used in a clinical setting.

Most ductal and lobular cancers have secretions that may form calcifications (Highnam and Brady, 1999). 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 microcalcifications are benign (Highnam and Brady, 1999; Veldkamp and Kerssemeijer, 1998), their shape and topology differentiating them from the malignant type.

In general, microcalcification detection algorithms consist of three stages: (i) a pre-processing step based on filtering the image, for example, subtracting a signal-enhanced image from a signal-suppressed image (Nishikawa et al., 1993); (ii) a segmentation step, for example based on global thresholding and morphological erosion (Nishikawa et al., 1993), on region growing (Shen et al., 1993) or adaptive thresholding (Betal et al., 1997); and (iii) a clustering step, typically using a fixed size kernel to eliminate noise points and isolated calcifications and to identify clusters. A major drawback of such methods is their poor sensitivity to small and faint microcalcifications. Other approaches use neural network classifiers (Edwards et al., 2000; Hara et al., 2000; Kim and Park, 1999; Rosen et al., 1996), several of which are trained on a single database and give poor results when tested on other databases. Wavelet filters have been used in (Aghdasi, 1996; McLeod et al., 1996) to highlight the high frequency signal from the background; unfortunately, noise also has significant power at high frequencies.

Yam et al. (2001) developed a structural approach to microcalcification detection, which exploits the Standard Mammogram Form (SMF) representation (therein called h_{int}), which is introduced in Section II of the paper. Yam et al.'s technique detects tall thin structures in the h_{int} image and tests with the height of such a structure to its footprint of the same structure in the image to detect possible calcifications. Noise is a major problem and some individual calcifications are overlooked and so statistical analysis of a cluster is compromised. The results appear promising, though it should be noted that they are from a considerably smaller database (Yam et al., 2001).

Perhaps the best-known published statistical method is that of Karssemeijer et al. (1993, Veldkamp and Kerssemeijer, 1998). The authors use adaptive noise equalization as part of a Bayesian approach, using Markov Random Fields to detect clusters of microcalcifications. Though it set a new standard for feature detection at the time it was published, the algorithm involves a large number of parameters, and there have been questions raised about its robustness or generalisability to other databases (Poissonnier and Brady, 2000).

R2 Technology's commercial product *ImageChecker*^{TM} (Roehrig, 2005) is a CAD system that is widely acknowledged as the state of the art of microcalcification detection (though, understandably, few details of the detailed operation of *ImageChecker*^{TM} have been published). Several large studies show the improvements from using *ImageChecker*^{TM} to detect breast cancer by reducing the false negative rate, while increasing the mean sensitivity of radiologists (Roehrig, 2005). Freer and Ulissey (2001) report that radiologists improve the detection of malignant microcalcification clusters from 68% to 100% after the use of CAD. In a more recent study, Birdwell et al. (2005) notify that the contribution of CAD increased by 7% the number of detected cancers, but without clear benefit of CAD systems in the detection of microcalcifications.

Herein we present a novel "foveal" segmentation algorithm that is based on a biological model of adaptive contrast-detection, complemented with complex pre-processing steps. The rationale for adopting a model inspired from the human visual system derives from the challenge posed by the highly variable mammogram image background that often confounds algorithms or human observers, especially when

'bright' microcalcifications are hidden within 'bright' dense tissue. The paper is structured as follows. Section II presents the pre-processing of SMF images (Highnam and Brady, 1999), in particular the optional, but helpful, normalisation step, the removal of shot-noise, curvilinear structures and image filtering. A novel statistical approach to derive automatically the parameters of the enhancing diffusion filter is also proposed. Section III introduces the "foveal" segmentation algorithm. In Section IV we perform a Free-Response Operating Characteristic (FROC) curve analysis as a step towards a validation of the microcalcification detection algorithm. Finally, we present a discussion and conclusions.

II. PRE-PROCESSING

Due to the complexity of mammogram image formation, we pre-process images before applying contrast-detection. The aim is to remove structures that may give rise to false positives (e.g. shot noise or curvilinear structures) and to smooth out noise in the image without removing microcalcifications. Note that we also include in this section an image normalisation step SMF, which is optional to the method, but which we have found can improve detection results, since it models the image formation process and removes variability in imaging conditions (e.g. time of exposure, film speed). The pre-processing steps of our method are described in the following paragraphs.

A. The Standard Mammogram Form

A well known normalisation method that results from modelling the process of mammographic image formation is the Standard Mammogram Form (SMF) developed by Highnam and Brady (1999) and improved in (Yam et al., 1999). The SMF algorithm corrects for scatter and extra-focal radiation, glare and anode heel effect. We have also based our work on previous results by Highnam and Brady (1999) and Yam et al. (1999), which estimated radiographic mottle, and then used Wiener filtering to subtract its effect and estimate glare and local blur. Our paper started from the observation that sometimes Yam's algorithm gave very poor results, and that led us to use – and analyse - anisotropic diffusion.

Because of the relatively weak control over the image acquisition process, it is difficult to eliminate variability in image characteristics, such as contrast and brightness. It is well known that such differences in imaging conditions lead to a complicated functional relationship between the intensities of temporal pairs of mammograms that varies from patient to patient (Marias et al, 2005)0. This is a strong indication that the non-linearities introduced by varying imaging conditions may alter significantly the image-intensity profile and reduce the efficiency of automatic microcalcification detection algorithms. After developing a method for the calculation and correction for the scattered radiation, Highnam and Brady (1999) discovered a method for mammogram image normalisation that eliminates distracting variations. This was based on the following assumptions /observations:

  - The x-ray attenuation coefficients of normal healthy tissue and cancerous tissue are approximately equal over the range of tube voltages used in mammography; but are significantly different from that of fat;

  - Fat is clinically uninteresting, therefore normal healthy and cancerous tissues are collectively referred to as "interesting";

Highnam and Brady's method estimates – in millimetres – the amount of interesting tissue in each pixel column, as illustrated in Figure 1. This effectively provides objective quantitative information about the breast anatomy. If, for example, the separation between the Lucite plates is 6.5cm, the amount of interesting tissue at a location (x,y) might be 4.75cm, implying 1.75cm of fat.

This way, the interesting tissue representation refers only to (projected) anatomical structures – the SMF algorithm estimates, and then eliminates the effects of, the particular parameters that were used to form the mammographic image. In short, the image can be regarded as standardised (hence the more recent name Standard Mammogram Form or SMF) and quantitative since the 'normalised' pixel values are measurements in millimetres, not in arbitrary units of shading (contrast), and they convey anatomical information. The key point is that without normalisation, for example by the SMF process, there is inevitably the risk of confounding anatomically useful information and imaging parameter effects, whose choice significantly affect contrast, brightness and the level of noise. As we will show in the results

section, detection performance is improved if the normalisation step is included.


Figure 1: A schematic of h_{int} calculation. In the normalised image each pixel value reflects the corresponding amount of non-fatty tissue instead of acquisition-dependent attenuation measurements.

   The exception ignored in the generation of an SMF representation of a mammogram is the presence of microcalcifications. The attenuation coefficient of microcalcification is approximately 26 times that of normal/cancer tissue, and so, although microcalcifications are physically small, they appear in the SMF representation as non-fat tissue that is 26 times as high as it actually is, that is as very high, tall thin towers. Thus, although the breast parenchyma is normalised, calcifications remain small areas of higher contrast. Hence, the contrast of a feature (microcalcification in our case) is largely independent of the film contrast curve. It is its presence in a fat or dense neighbourhood that makes it easy or difficult to detect/visualise.

   We stress that the inclusion of a normalisation pre-processing step is not essential for our algorithm to work. However, our results indicate that the detection efficiency can be improved by including such a step, probably as a result of removing non-linearities and distracting variations, which in turn improves the contrast.

B. *Shot Noise and Curvilinear Structures Removal*

An important pre-processing step is the removal of noise, which is a major source of false positives (FP). In this paper, shot noise refers to an intrinsic noise process during scanning, due to dust and artefacts present on the film or in the scanner. Previous papers report the detection of film-screen artefacts associated with the SMF generation process (Ancelin et al., 2002; Highnam and Brady, 1999), for example at the glare removal stage. Although shot noise appears visually similar to microcalcifications, it is introduced at the image digitisation stage, that is, it is introduced to the image after the main blurring stage, and so the boundary shape of shot noise tends to be quite sharp. The point spread function of the intensifying screen (blur function) enables the identification of image locations that are locally bright but the energy imparted becomes negative where shot-noise is present after glare compensation. Using the binary information from the shot noise map, the artefacts can be removed from the SMF image by interpolating between their surrounding backgrounds.

Curvilinear structures (CLS) – ducts, blood vessels, ligaments or tumour spiculations - are locally linear, relatively thin structures, which often produce in a mammogram a characteristically textured appearance. Removing the CLS turns out to be important for the method specificity (Highnam and Brady, 1999), but their identification and removal remains a major challenge (Wai et al, 2004). The CLS appear relatively bright, though not necessarily with high contrast. They are locally linear, but typically curve at larger scales. Their detection is difficult due to the large variety of widths, contrasts, lengths, and intersections. Intersections are either genuine bifurcations or may correspond to the overlap of two projected CLS in the compressed breast. The attenuations of the individual CLS add, producing a localized 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 Kovesi (1999) and is presented in (Evans et al., 2002). From local energy (LE) we search points with maximum phase congruency (PC) to detect and then remove CLS from mammograms. For the spatial frequency j, if the local Fourier coefficients of a one-dimensional signal are as in equation (1), then LE and

PC can be calculated according to equations (2) and (3). In practice, logGabor filters were used to estimate the local energy and phase, with a spreading function and combination of responses from several discrete directions to extend the analysis to 2D (Kovesi, 1999).

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

$$\sum_{j=1}^{N}\left(A_j + iB_j\right); \tag{1}$$

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

$$PC = \sqrt{\frac{LE}{\sum_{j=1}^{N}\left(A_j^2 + B_j^2\right)}}; \tag{3}$$

Our implementation uses a small wavelength, which makes it more sensitive to smaller scales, and a reduced number of scales over which features are detected by phase congruency in order to avoid marking the CLS as calcifications. The limitation is that there are discontinuities that appear in the true CLS and their removal leaves bright areas on the mammogram similar to microcalcifications. We use a smooth interpolation process to remove CLS from mammograms. Figure 2 shows an example of removal of shot noise and CLS.

C. Image Enhancement

We have previously reported our work using anisotropic diffusion to filter and analyse mammographic images (Linguraru et al., 2001). Anisotropic diffusion has the desirable property of smoothing an image, that is reducing noise (hence increasing signal to noise, which is generally poor for mammographic images) while at the same time largely preserving image structure. In this paper we also address the issue of automatically defining the appropriate diffusion parameters.

Figure 2. Removing artefacts in the pre-processing step; (a) the original image; (b) the shot noise map (white dots); (c) the CLS map; (d) the "clean" image.

Previous methods used for noise reduction, such as the Wiener filter (originally embedded in the SMF generation), have limitations (Linguraru, 2002), primarily arising from deconvolution and inevitably simplified models of the complex physical processes involved in mammography. These possible sources of errors leave residual high-frequency noise in mammographic images. Digitiser noise, also of high frequency, adds to it. The anisotropic diffusion filter aims to smooth the remaining high-frequency noise that interferes with our specific application. We base our approach on Weickert's solution (Weickert, 1998) for the diffusion tensor for which we define the corresponding eigenvalues in equation (4), where I is the initial image, I_σ the Gaussian smoothed image (σ is the standard deviation or scale) and k is a contrast measure. I can be either an intensity or SMF image, depending if SMF normalization has been performed as a pre-processing step.

$$\lambda_1 = \begin{cases} 1 & |\nabla I_\sigma| = 0, \\ 1 - \exp\left(\dfrac{-1}{(|\nabla I_\sigma|/k)^{12}}\right) & |\nabla I_\sigma| > 0 \; ; \end{cases}$$
$$\lambda_2 = 1 \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad (4)$$

The parametric nature of anisotropic diffusion in general makes the 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. We address this problem by proposing an automatic scheme that defines the diffusion parameters using image characteristics. The parameters we deal with are: a) a contrast gradient k, b) the standard deviation σ and c) the number of iterations t.

Microcalcifications represent a tiny percentage of the mammogram surface. Typically, they are very small, indeed are only present on average in about a quarter of the total number of screening mammograms (Highnam and Brady, 1999). For these reasons, we generously estimate that at most 5% of the mammogram pixels is sufficient to accommodate the entire population of microcalcification pixels. Microcalcifications often appear relatively bright in an X-ray image, indeed were it not for surrounding dense tissue they would be the brightest/highest pixels of the mammogram. As illustrated in equation (5), we compute k as a measure of the gradient, where $K_\sigma(I)$ is the Gaussian of image I, M its derivative and N the number of pixels in the neighbourhood δ_i. k becomes a value with well-defined physical meaning that discriminates between microcalcifications (the brightest structures after noise removal) and background; more precisely, we chose the 4.4% structures with highest contrast, according to the statistical value in equation (5), given that mean(g) is 0. Figure 3 illustrates the estimation and meaning of the value of k.

$$k = 2*\text{std}(g), \text{ where } g_i = M_i - \frac{1}{N}\sum_{j\in\delta_i} M_j \text{ and } M = |K_\sigma'(I)|; \qquad\qquad (5)$$

Figure 3: An example of estimating k. a: From the image shown on the left (after expanded display contrast) the histogram of g is computed from equation (5), b: We note the zero mean value of g, as well as where the value of 0.47 of k (related to the 4.4% structures with highest contrast) falls.

The second parameter to be set is σ, the standard deviation of the Gaussian filter used to smooth the image. We need to choose a value for σ so that, on the one hand, it removes high-frequency noise (which is very localized, typically spread over a couple of pixels at, for example, 50μm/pixel), but, on the other hand, preserves microcalcifications (on average 0.5-1 mm in diameter). Digital images are convolved with a Gaussian kernel, whose width is given by σ ($n*\sigma$, for example $n=10$, where n gives the digital finite extent of the Gaussian). σ must be sufficiently small to clean noise and retain microcalcifications and is directly related to the image resolution. For example, the width of the kernel must be 300/50=6 pixels to preserve microcalcifications over 0.3mm (or 300μm) in diameter at an image resolution of 50μm/pixel.

Weickert (1998) notes that the number of iterations t is related to the spatial width of the Gaussian kernel. To blur features of the kernel order it is required that $t = (n*\sigma)^2/2$, which gives excellent noise reduction results in the microcalcification preservation framework.

The reasoning behind the parameter setting in the anisotropic diffusion process has been explained above. As a concluding remark, we note that our method computes all parameters automatically: k is computed directly from the image and its meaning (it depicts the most outstanding features in the mammogram) and is independent of the data set; σ is strictly related to the size of microcalcifications and the image resolution; and t is set according to the proposed value of σ. We show a typical example of

automated image enhancement for microcalcification detection in Figure 4.

With the automatic tuning of parameters that we propose, anisotropic diffusion becomes, for microcalcification detection, a robust process using little, but essential, a-priori knowledge. This is important since anisotropic diffusion has often attracted criticism because of its strong dependency on a number of parameters.


Figure 4. Automated image enhancement; (a) the original noisy image with a microcalcification cluster; (b) the map of features selected by the value of k; (c) the enhanced image.

III. FOVEAL SEGMENTATION

An important part in image processing and computer vision is to understand and implement models of the human visual system (HVS), partly because understanding the HVS is an important scientific goal in its own right, and partly because expert radiologists have generally outperformed computer systems, especially for the detection of subtle microcalcifications and/or masses in dense textured backgrounds. There is over a century of research behind trying to fully understand the way HVS functions, particularly how it responds to intensity transitions and texture variations. A major contribution to this on-going research was made by Holladay (1926) and later by Moon and Spencer (1943). They developed several of the fundamentals of the way the human processes visual scenes, at least in the early stages. Nevertheless, the process is very complex and still not fully understood (Ashdown, 1996).

The adaptation of the eye to light changes is a continuous process in the HVS. We perceive objects differently if they are against a bright surface (for instance a window on a sunny day) or dark area (a dark wooden panel). The adaptation luminance proposed in (Moon and Spencer, 1943) 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 surfaces surrounding the peripheral field of view (L_{seq}) as in equation (6).

$$L_a = L_{fov} + L_{seq};$$ (6)

However, foveal adaptation stems mainly from the luminance within the foveal field and relates to approximately only ten percent to the luminance of the field of view outside that of the fovea (Moon and Spencer, 1943). Furthermore, the visual perception of the eye is dependent on the spatial perception of the object we visualise. This effect is called lightness assimilation. The same object may appear lighter on a dark background and darker over a light surface (see Figure 5).

In mammography, the main advantage of human eye over CAD methods is the detection of microcalcifications in dense areas of the breast, usually the clusters that make the True Positive (TP) ratio inferior to the 100% of the ground truth. At the same time, neither the radiologist nor a computer program must pick FP in the dark areas. This is the result of the adaptation to background, which is equivalent to lowering thresholds in bright areas (to increase the TP ratio by detecting clusters in dense breast areas) and increasing them in dark areas (to decrease the FP number), basically an inverse proportionality to the background. In a mammogram we must consider the neighbouring area of the foveal kernel, as shown in equation (6). The adaptive thresholding 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 5. 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.

The classical 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*), as in equation (7), where *n* is the number of pixels in *N* (Gonzalez and Woods, 1993). In mammography, such contrast measures were used in (Karssemeijer, 1993) combined with noise/contrast equalisation and (Kegelmeyer and Allmen, 1994) together with local thresholding. In a simpler scenario, $C_{classic}$ is 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 (or contrast changes) against a fatty (dark) background but more difficult to detect correctly against a denser (bright) background (Highnam and Brady, 1999). The HVS, however, adapts to the local image contrast, and detects faint contrast changes in a manner essentially independent of the background. We adopted a model of contrast detection similar to HVS for mammography with the aim of improving the accuracy of the detection.

$$C_{classic} = p_0 - \frac{1}{n}\sum_{i \in N} p_i \; ; \tag{7}$$

We initially pre-process the images as described in the previous section (shot-noise and CLS removal, etc.). Then, we compute a set of mean intensity values of the inner area of the object to visualise (i.e.

within the boundary of calcification), its neighbourhood (the local area around the object) and background (the rest of the breast tissue). The histogram of the inner foveal surface provides the mean of the object (μ_O), as the histogram of the whole image will give us the mean of the background (μ_B) and a measure of the density of breast. The mean of the neighbourhood (μ_N) is found from the intensities of pixels within the neighbourhood and excluding the object pixels according to Figure 6. μ_N determines whether the visualised object is on a dense (bright) or a fat (dark) area of breast. The size of the kernel of the inner object (O) used to compute μ_O is established according to the average size of microcalcifications (Highnam and Brady, 1999) and the resolution of the tested images, in a similar manner to the kernel of the Gaussian in Section II. 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 too small to avoid overlapping O and N for slightly bigger microcalcifications. In our application N is set twice the size of O. Then the perceivable contrast C is calculated in agreement with equation (8).

$$C = \begin{cases} \dfrac{\mu_o - \mu_N}{\mu_N}, & if \quad \mu_o > \mu_N \; ; \\ 0, & otherwise \end{cases} \qquad (8)$$
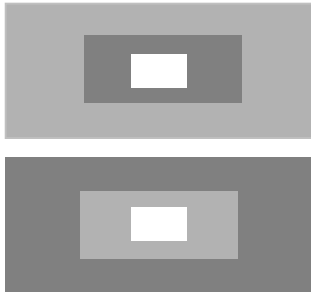

Figure 6. The foveal masks used for the computation of μ_O, μ_N and μ_B. The object O is the area of the *fovea centralis*, N its neighbourhood (twice the size of O in our applications) and B the background.

We then compute C_{min} (9), a measure of contrast sensitivity, where $\mu_A = w\,\mu_N + (1\text{-}w)\,\mu_B$ and w is a suitable weight between 0 and 1 affecting the amount of background implied in the computation of contrast. C_{min} 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. The literature proposes 7.7% of the adaptive luminance to be due to the background luminance (Heucke et al., 2000), which gives a value of 0.923 to our weight w. In practice, we studied the effect of varying w with 10% more or less than the proposed value. We use w=0.923 as the default value in comparative studies. For parameter b in equation (9), we have found that the literature proposed value b=0.0808 (Heucke et al., 2000) gives good results. Areas in the image having $C>C_{min}$ are marked as microcalcifications.

$$C_{min} = \begin{cases} \dfrac{c_w}{\mu_N}\left(b + \sqrt{\mu_A}\right)^2, & \mu_A \geq \mu_N \\[3mm] \dfrac{c_w}{\mu_N}\left(b + \sqrt{\dfrac{\mu_N^{\,2}}{\mu_A}}\right)^2, & \mu_N > \mu_A \end{cases}$$

(9)

The segmentation parameters are established automatically based on the image-adapted value of k, as computed from the statistical analysis, (discussed in Section II). Through the use of the minimal perceivable contrast c_w, we include a measure of the image contrast (previously denoted as k) - imagine varying the amount of objects we can distinguish with the naked eye by using a pair of sunglasses. c_w is a crucial variable for assessing our method and is consequently used for producing the FROC curves in the results session. From the FROC analysis, we noted that adapting the value of c_w to the image characteristics gives a better balance between TP and FP than a constant c_w. We performed simple preliminary tests on 13 random cropped mammograms digitised at 50µm resolution and containing 10 pre-labelled isolated coarse calcifications. These images originate from Oxford, but are not part of the Oxford database on which the algorithm is evaluated in the next section. We found that $c_w = \sqrt{k}/200$ is appropriate

for conservative detection of microcalcifications.

Using C_{min} instead of C_{thresh}, the contrast is adapted locally, not just globally, in a manner similar to that of HVS. Figure 7 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 easily depicted in a fat background for both contrast measures, but C_{min} reduces the number of FP. The main advantage of using C_{min} becomes clear in a dense background, where its adaptability facilitates the detection of microcalcifications.

IV. RESULTS

In this section we present comparative Free-Response Receiver Operating Characteristic (FROC) curves to test the outcome of our method with variations in algorithm and input images. We used a database of 102 samples of digital SMF images from 83 mammograms. 78 of the samples contain between 1 and 3 clusters per image, while 24 are normal mammogram samples. An experienced radiologist annotated a total of 98 microcalcification clusters, which had previously been 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 films were subsequently digitised using the Lumisys scanner at a resolution of 50μm. The cluster positions were then translated into x-y coordinates on the digital images. The cropped image sizes are of maximum 1500x1500 pixels.

Figure 7. 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.

Figure 8 shows a schematic of the method flow described in the previous sections. As noted earlier, the normalisation step in the flow is optional and results are reported for both original and normalised (SMF) images. A cluster is detected if it contains at least three microcalcifications, where a distance of maximum 5 mm connects each calcification to the rest of cluster. When a cluster is identified at the position indicated by the radiologist, we count a TP. If a cluster is detected outside the marked areas or more than one cluster are found at the approximate location of a single contour indicated by the radiologist, we count a FP. Recursively, we noted that the distribution of detected FP/image is equal in samples with microcalcifications and in normal samples. In Figure 9 we note the impact of CLS removal, image smoothing and image normalisation (SMF images) on detecting microcalcification clusters. The results are superior when the proposed algorithm (Figure 8) is employed (the top line in the FROC curve); when the algorithm is applied to intensity (not normalised) images the results are also very satisfactory.
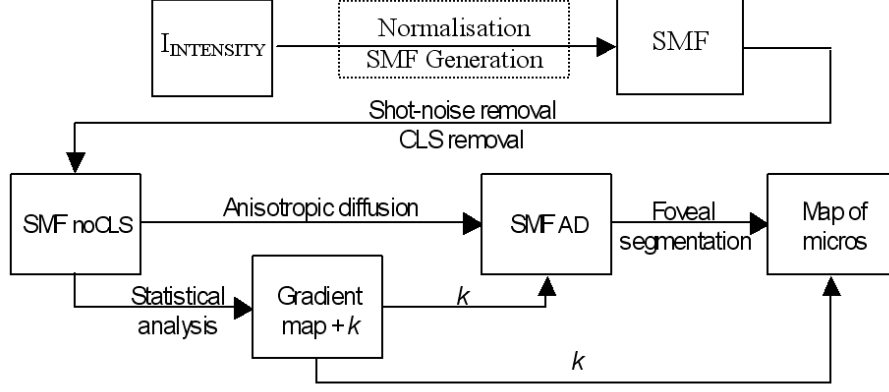
Figure 8. The diagram of the biologically driven algorithm to detect microcalcifications. The top row underlines the optional image normalisation that generates SMF images. The rest of the diagram highlights the steps introduced by our algorithm, which are described in this paper.

In Section III we noted the significance of the value 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, by varying the value of w over 5-10%. Figure 10 shows the comparative detection results with the variation of w. A small variation of w around the proposed value of 0.923 does not influence the evolution of our algorithm."

The results used in building the FROC curves in Figure 9 and Figure 10 are based on processing cropped samples of mammograms. To illustrate results on whole mammograms, 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 pooled opinions of clinicians at the Oxford Breast Care Unit established ground truth. The breast margin is detected in SMF, thus a threshold above 0 removes the background. Now we can compute the value of k (see Section II) for the inner area of the breast. The detection results are accurate and similar to those achieved on mammogram samples (see Figure 11). 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. The method is more robust to detecting

microcalcification in samples because the probability of generating FP is increased when applied to the whole image (i.e. the search space for microcalcification detection is larger).
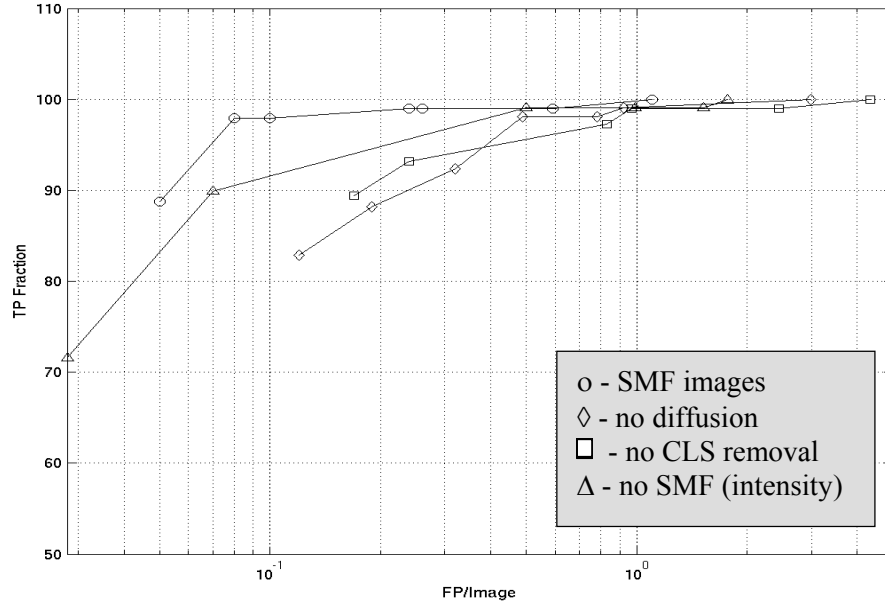

Figure 9. The comparative FROC curves for the detection of microcalcifications. o - represents the detection on SMF images; ◊ - are the results without noise removal by anisotropic diffusion; □ - are the results without CLS removal; ∇ - shows results on intensity images, when no SMF generation is present. All algorithms reach 100% TP fraction with a clear better performance on SMF images, which were de-noised and CLS-removed.

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, without image normalisation this is hard to achieve. 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 extraction of parameters from image attributes, facilitates the generalization of detection standards.

Figure 10. 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.


Figure 11. 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.

We used for comparison a collection of images from the University of South Florida Digital Database

for Screening Mammography (DDSM). 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 annotated clusters is 82. All images are intensity images, as termed before, digitised at 43μm/pixel.

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 σ in the image enhancement step and that of the kernel of fovea in the foveal segmentation change accordingly. These conversions are done automatically in agreement with the user specified image resolution. The FROC curve shown in Figure 12 compares the performance of the microcalcification detection algorithm between the Oxford Screening Database in intensity form and the DDSM database on original images. The algorithm performs better on the Oxford Screening database, but 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. The DDSM database was chosen for its availability; no similarity between databases was considered and no additional tuning of parameters performed. It remains desirable to test the detection algorithm on the DDSM database when images will be available in SMF form, as normalisation would improve the flexibility of the results of the algorithm on other databases.

Finally, this section compares three algorithms to detect microcalcification clusters that operate upon the SMF representation of mammograms from the Oxford Database. The first has been described previously in this paper and is addressed as the "foveal algorithm"; the second one is Yam *et al.*'s physics based approach that was described in (Yam et al., 2001). The third is a variation of the statistical analysis introduced in Section II, here addressed as the "statistical approach", and is presented in (Linguraru and Brady, 2002). Using FROC analysis, we demonstrate the superiority of the foveal algorithm in Figure 13. In Figure 14 we present examples of microcalcification detection.
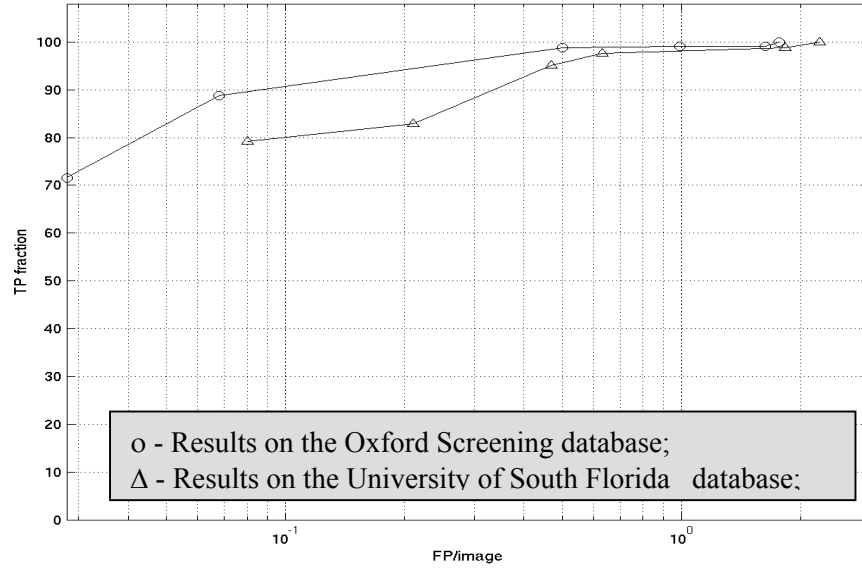
Figure 12. The comparative FROC curve between the detection results on intensity images from the Oxford Screening Database and the University of South Florida Digital Database for Screening Mammography. Note that the results are for mammogram samples without SMF normalisation.


Figure 13. The FROC curves of the three microcalcification-detection methods on SMF images; the Foveal Approach yields the best results.

Figure 14. Two examples of microcalcification detection. On the left we present the input images in SMF format: on top an image with two large microcalcification clusters, on bottom an image with faint microcalcifications in varying background, while on the right we note the results of our detection method.

V. Discussion and conclusion

We have presented a novel method for detecting microcalcification clusters based on a biologically inspired contrast detection algorithm in combination with a pre-processing step (shot noise and CLS removal, image enhancement), as well as an optional image normalisation step. Many methods to detect microcalcifications attempt to tune the variety of parameters used in the implementation of the algorithm to best suit the studied cases. The consistency and reproducibility of results becomes highly dependant on the

operators and their capability to find the best parametrical configuration for the detection. We proposed a fully automated method to detect microcalcifications and we reported results on both SMF and intensity images.

The first contribution of the presented work is the automatic tuning of the parameters of anisotropic diffusion. The contrast k is found from its direct relation to the image gradient and the percentage of features that we desire to enhance. The second parameter σ reflects the scale of features of interest and must be set according to their size and the given the image resolution (multiscale analysis may be performed). In our application, both σ and k become measures related to the anatomy of breast. The last parameter, the number of iterations $t,$ is derived as a function of σ to remove noise in the context of preserving microcalcifications. In a general application for feature enhancement, the automatic tuning of t may prove to be more difficult than in our particular framework. With the automatic tuning of parameters that we propose, anisotropic diffusion, a method with a solid theoretical basis, but controversial for its parametrical dependency, can be used as a robust process using little, but essential, a-priori knowledge.

However, the main contribution of the paper is the method for adaptively thresholding the filtered results in order to segment microcalcifications. Based on the biologically driven idea that human eye is a better feature detector than a computer algorithm and that both local and global adaptability are necessary for subtle contrast detection, our detection model mimics the human eye's function for improved segmentation results. Therefore, the algorithm combines filters, analysis of image characteristics and adaptive thresholding based on a human visual system paradigm.

The robustness of the algorithm has been demonstrated by the FROC analysis performed over a range of parameters. It is important to note that the parameters are set according to the image attributes and the method is fully automated. 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. We also compared the performance of our algorithm on data from different databases with good detection results. The results indicate that adding adaptive contrast segmentation based on

characteristics of the human visual system significantly enhances the detection of microcalcifications.

REFERENCES

F. Aghdasi, Application of neural network adaptive wavelets for signal representation and classification in digital mammography, in: K. Doi, M.L. Giger, R.M. Nishikawa, R.A. Schmidt (Eds.), Digital Mammography, Elsevier, Amsterdam, 1996, pp. 307-310.

B. Ancelin, M. Yam, R.P. Highnam, J.M. Brady, Automated detection of artifacts on film-screen mammograms using a novel iterative detection and correction scheme, in: Digital Mammography, Lecture Notes in Computer Science, Springer, New York, 2002, pp. 174-178.

I. Ashdown, Luminance gradients: photometric analysis and perceptual reproduction, Journal of Illuminating Engineering Society 25(1) (1996) 69-82.

D. Betal, N. Roberts, G.H. Whitehouse, Segmentation and numerical analysis of microcalcifications on mammograms using mathematical morphology, British Journal of Radiology 70 (1997) 903-917.

R.L. Birdwell, P. Bandodkar, D.M. Ikeda, Computer-aided detection with screening in a university hospital setting, Radiology 236 (2005) 451-457.

D.C. Edwards, M.A. Kupinski, R. Nagel, R.M. Nishikawa, J. Papaioannou, Using Bayesian neural network to optimally eliminate false-positive microcalcification detections in a CAD scheme, in: M.J. Yaffe (Ed.), Digital Mammography, Medical Physics Publishing, Madison, 2000, pp. 168-173.

C.J. Evans, K. Yates, J.M. Brady, Statistical characterisation of normal curvilinear structures in mammograms, in: Digital Mammography, Lecture Notes in Computer Science, Springer, New York, 2002, pp. 285-291.

T.W. Freer, M.J. Ulissey, Screening mammography with computer-aided detection: prospective study of 12860 patients in a community breast center, Radiology 220 (2001) 781-786.

M. Funovics, S. Schamp, B. Lackner, P. Wunderbaldinger, G. Lechner, G. Wolf, Computer-assisted diagnosis in mammography: the R2 ImageChecker system in detection of spiculated lesions, Wien Med Wochenschr 148(14) (1998) 321-324.

M.L. Giger, Z. Huo, C.J. Vyborny, L. Lan, R.M. Nishikawa, I. Rosenbourgh, Results of an observer study with an intelligent mammographic workstation for CAD, in: Digital Mammography, Lecture Notes in Computer Science, Springer, New York, 2002, pp. 297-303.

R.C. Gonzales, R.E. Woods, Digital Image Processing, Addison-Wesley Publishing Company, New York, 1993.

T. Hara, A. Yamada, H. Fujita, T. Iwase, T. Endo, Automated classification method of mammographic microcalcifications by using artificial neural network and ACR BI-RADS criteria for microcalcification distribution, in: M.J. Yaffe, (Ed.), Digital Mammography, Medical Physics Publishing, Madison, 2000, pp. 198-204.

L. Heucke, M. Knaak, R. Orglmeister, A new image segmentation method based on human brightness perception and foveal adaptation, IEEE Signal Processing Letters 7(6) (2000) 129-131.

R.P. Highnam, J.M. Brady, Mammographic Image Analysis, Kluwer Academic Publishers, London, 1999.

L.L. Holladay, The fundamentals of glare and visibility, J.Opt. Soc. Amer. 12 (1926) 271-320.

N. Karssemeijer, Adaptive noise equalisation and recognition of microcalcification clusters in mammograms, Int. J. of Pattern Recognition and Artificial Intelligence 7(6) (1993) 1357-1372.

W.P. Kegelmeyer, M.C. Allmen, Dense feature maps for the detection of microcalcifications, in: A.G. Gale, S.M. Astley, D.R. Dance, A.Y. Cairns (Eds.), Digital Mammography, Elsevier, Amsterdam, 1994, pp. 3-13.

J.K. Kim, H.W. Park, Statistical textural features for detection of microcalcifications in digitized mammograms, IEEE Transactions on Medical Imaging 18 (1999) 231-238.

P. Kovesi, Image features from phase congruency, Videre: Journal of Computer Vision Research 1 (1999) 1-26.

M.G. Linguraru, J.M. Brady, M. Yam, Filtering h_{int} images for the detection of microcalcifications, in: W. Niessen, M. Viergever (Eds.), Medical Image Computing and Computer-Assisted Intervention 2001, Lecture Notes in Computer Science 2208. Springer, New York, 2001, pp. 629-636.

M. Linguraru, Feature Detection in Mammographic Image Analysis. PhD Thesis, University of Oxford, 2002.

M.G. Linguraru, J.M. Brady, A non-parametric approach to detecting microcalcifications, in: H.O. Peitgen (Ed.), Digital Mammography, Springer, New York, 2002, pp. 339-341.

K. Marias, C. Behrenbruch, S. Parbhoo, S. Seifalian, M. Brady, A registration framework for the comparison of mammogram sequences, IEEE Transactions on Medical Imaging 24(6) (2005) 782-790.

G. McLeod, G.J.S. Parkin, A.R. Cowen, Automatic detection of clustered microcalcifications in mammograms, in: K. Doi, M.L. Giger, R.M. Nishikawa, R.A. Schmidt (Eds.), Digital Mammography. Elsevier, Amsterdam, 1996, pp. 311-316.

P. Moon, D.E. Spencer, The specification of foveal adaptation, J. Opt. Soc. Amer. 33 (1943) 233-248.

R.M. Nishikawa, M.L. Giger, C.J. Vyborny, R.A. Schmidt, Computer-aided detection of clustered microcalcifications: an improved method for grouping detected signals, Medical Physics 10 (1993) 1661-1666.

N. Petrick, B. Sahiner, H.P. Chan, M.A. Helvie, S. Paquerault, L.M. Hadjiiski, Breast cancer detection: evaluation of a mass-detection algorithm for computer-aided diagnosis, Radiology 224 (2002) 217-224.

M. Poissonnier, J.M. Brady, 'Noise equalisation' for microcalcification detection?, in: M.J. Yaffe (Ed.), Digital Mammography, Medical Physics Publishing, Madison, 2000, pp. 334-341.

J. Roehrig, The manufacturer's perspective, The British Journal of Radiology 78 (2005) 41-45.

D. Rosen, B. Martin, M. Monheit, G. Wolff, M. Stanton, A Bayesian neural network to detect microcalcifications in digitized mammograms, in: K. Doi, M.L. Giger, R.M. Nishikawa, R.A. Schmidt (Eds.), Digital Mammography, Elsevier, Amsterdam, 1996, pp. 277-286.

L. Shen, R.M. Rangayyan, J.E.L Desautels, Detection and classification of mammographic calcifications, Int. Journal of Pattern Recognition and Artificial Intelligence 71 (1993) 1403-1416.

W.J.H. Veldkamp, N. Karssemeijer, Accurate segmentation and contrast measurement of microcalcifications in mammograms: a phantom study, Medical Physics 25(7) (1998) 1102-1110.

L.C.C. Wai, M. Mellor, M. Brady, A multi-resolution CLS detection algorithm for mammographic image analysis, in: C. Barillot, D.R. Haynor, P. Hellier (Eds.), Medical Image Computing and Computer-Assisted Intervention (2), Lecture Notes in Computer Science 3217, Springer, New York, 2004, pp. 865-872.

J. Weickert, Anisotropic Diffusion in Image Processing, Teubner, Stuttgart, 1998.

M. Yam, J.M. Brady, R.P. Highnam, R. English, Denoising h_{int} surfaces: a physics-based approach, in: C. Taylor, A. Colchester (Eds.), Medical Image Computing and Computer-Assisted Intervention, Springer, New York, 1999, pp. 227-234.

M. Yam, J.M. Brady, R.P. Highnam, C.P. Behrenbruch, R. English, Y.:Kita, Three-dimensional reconstruction of microcalcification clusters from two mammographic views, IEEE Trans. Med. Imaging 20(6) (2001) 479-489.