Foveal Algorithm for the Detection of Microcalcification Clusters: a FROC Analysis

Marius George Linguraru [a,b], Michael Brady [a] and Ruth English [c]

[a] Medical Vision Laboratory, University of Oxford,
Ewert House, Ewert Place, Summertown, Oxford OX2 7BZ, UK
[b] EPIDAURE Research Project, INRIA,
2004 route des Lucioles, B.P. 93, 06902 Sophia Antipolis Cedex, France
[c] Breast Care Unit, the Churchill Hospital, Oxford OX3 9DU, UK
Marius.Linguraru@sophia.inria.fr

Abstract. Clusters of microcalcifications are often the earliest signs of breast cancer and their early detection is a primary consideration of screening programmes. We have previously presented a method to detect microcalcifications based on normalised images in standard mammogram form (SMF) using a foveal segmentation algorithm. In this paper, we discuss the selection and computation of parameters, which is a key issue in automatic detection methods. Deriving the parameters of our algorithm from image characteristics makes the method robust and essentially removes its dependence on parameters. We carry out a FROC analysis to study the behaviour of the algorithm on images prior to normalisation, as well as the contribution of the stages employed by our method. We report results from two different image databases.

1 Methodology

Our method of microcalcification detection relies on using a normalised representation of breast tissue, for example the standard mammogram form (SMF) [3] Having an SMF image as input, the first objective is the removal of curvilinear structures (CLS), which turn out to be an important consideration in the specificity but whose computation is itself a major challenge. This is performed using the local energy model for feature detection of Kovesi [4] and is presented in [1].

The SMF model accounts for the majority of imaging artefacts (scatter, glare, anode heel, extra-focal, quantum mottle, film grain). Inevitably, however, there are imperfections arising from deconvolution and model simplifications that are intrinsic to the SMF (or any similar) generation process [3,9]. These possible sources of errors leave residual high-frequency noise in mammographic images. Digitiser noise, also of high frequency, adds to it. We employ an anisotropic diffusion filter to smooth the remaining high-frequency noise [4]**.** A priori, this is 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 preserves image structure.

Unfortunately, the substantial number of parameters required for anisotropic diffusion makes the results of this process highly dependent on 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 the parameters that works well for the entire dataset. To address this problem, we propose using image characteristics to set the diffusion parameters for each individual image in the dataset.

It should be noted from the outset that although they may be important early indicators of breast cancer, microcalcifications generally occupy only a tiny percentage of the pixels of a mammogram image. Typically, they are very small and present in about a quarter of the total number of screening mammograms.

For these reasons, we consider that at most five percent of the total number of mammogram pixels suffices to account for the entire population of calcium salts. Since the x-ray attenuation of calcium is far larger than that of normal tissue, microcalcifications are expected to appear bright in a mammogram, indeed to be amongst the brightest/highest pixels of the mammogram, though the brightness can be reduced by scattering of x-ray photons. As made explicit in equation (1), we compute the contrast k as a measure of the gradient, where $K_s(I)$ is a Gaussian blurring of the image I. k becomes a value with well-defined physical meaning that discriminates between these brightest structures and the background; more precisely, we select the 4.4% structures with highest contrast.

The second parameter to be set is s, the standard deviation of the Gaussian filter used to smooth the image. We need to choose a value for s such that, on the one hand, it removes high-frequency noise; but, on the other, preserves microcalcifications. We compute s according to (2), where S is the maximal size of features to be smoothed and R the image resolution.

The number of iterations t of the anisotropic diffusion process is related to the spatial width of the Gaussian kernel [8]. To blur features of the kernel order ($n*s$, n=ct.) t is computed according to (3), which gives excellent noise reduction results while preserving microcalcifications.

$$M = |K_s'(I)|, \quad g_i = M_i - \frac{1}{N}\sum_{j \in d_i} M_j, \quad k = 2*\mathrm{std}(g) \tag{1}$$

$$s = S / R \tag{2}$$

$$t = (n*s)^2 / 2 \tag{3}$$

The process is completed by applying what we call a foveal algorithm to segment microcalcifications; the name derives from a model of the adaptability of the human eye to detect features in textured images [6]. The eye's ability to perceive luminance gradients is controlled by the w factor in (4), where μ_N is the weighted mean of the neighbourhood of the object to segment and μ_B the mean value of the image background. w is a suitable weight (between 0 and 1) which controls the amount of background used in the computation of contrast, which gives a global measure to the locally adapted contrast. Once more, the parameters are computed from the image characteristics, as seen in (5), where $C_{\min}$ is the adapted threshold used to segment microcalcifications and b a constant.

$$\mu_A = w + (1\text{-}w)\,\mu_B \tag{4}$$

$$C_{\min} = \frac{c_w}{m_N}\left(b + \sqrt{\frac{m_N^{\,2}}{m_A}}\right)^2,\ c_w = \sqrt{k}\,/\,200 \tag{5}$$

The detection method is described in greater detail in [6].

2 FROC Analysis

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 plot the true positive (TP) ratio against the number of false positives (FP) per image. 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 microcalcification clusters annotated in the database. All images were digitised at a resolution of 50 μm and have sizes at most 1500x1500 pixels. A cluster is detected if it contains at least three microcalcifications, where a distance of maximum 0.5 cm connects each calcification to the rest of cluster. In Figure 1 we note the impact of CLS removal, image smoothing and image normalisation (SMF images) on detecting microcalcification clusters. The results are clearly superior when the detection algorithm uses all of the stages in the algorithm, though they are not very different when the algorithm is applied to intensity (not normalised) images.

In the previous section, we noted the significance of w in setting the minimal perceivable contrast for obtaining the best detection results when our algorithm is applied. The literature proposes 7.7% of the adaptive luminance to be due to the background luminance 1, which gives a value of 0.923 to our weight w. 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 2 shows the comparative detection results with the variation of w.

The results used in building the FROC curves in Figures 1 and 2 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 the clinicians at the Oxford Breast Care Unit, Churchill Hospital were used as 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 Methodology) for the inner area of the breast. The detection results are accurate and similar to those achieved on the mammogram samples (see Figure 3). The most challenging cluster to detect is located in the breast margin. 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 ultimate goal of any CAD algorithm is to work reliably 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 precisely to cater for this 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 corner stone 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). 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 (using the earlier terminology) so the FROC curve shown in Figure 4 compares the performance of the microcalcification detection algorithm between the Oxford Screening Database in intensity form and the DDSM collection.

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.

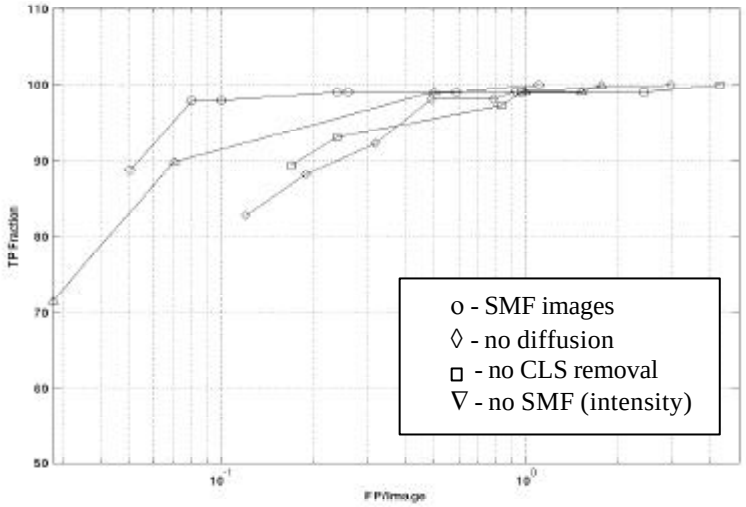

Fig. 1. The comparative FROC curves for the detection of microcalcifications. o - represents the detection on SMF images; ◊ - are the results without noise removal by anisotropic diffusion; □ - are the results without CLS removal; ∇ - shows results on intensity images, when no SMF generation is present. All algorithms reach 100% TP fraction with a clear better performance on SMF images that were de-noised and CLS-removed

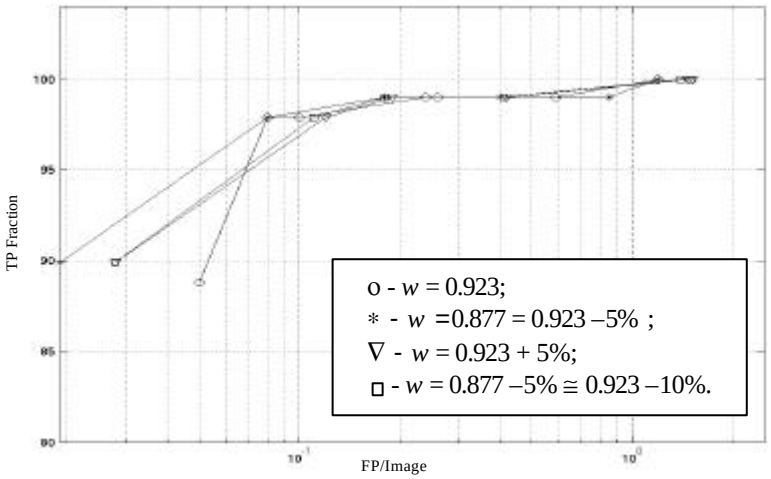

Fig. 2. The comparative FROC curve when *w* is varied over a range of 5 to 10% of its default value of 0.923. The difference in detection results is quite small and all four algorithms converge smoothly to 100% TP ratio.

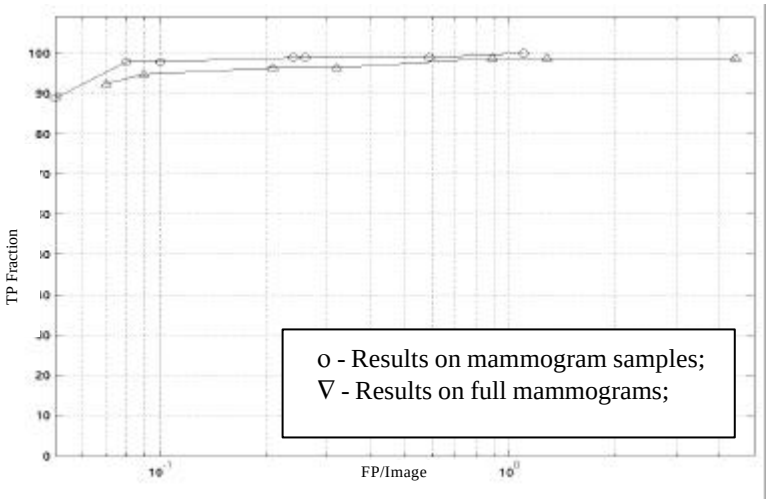

Fig. 3. The comparative FROC curve of the detection of microcalcifications when mammogram samples are used versus full mammograms. The behaviour of the algorithm is similar and robust with the image size.

Finally, this section compares three algorithms to detect microcalcification clusters that operate upon the SMF representation of mammograms. 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 9. The third is a variation of the statistical analysis introduced in Methodology, here addressed as the "statistical approach", and is presented in [6]. Using FROC analysis, we demonstrate the superiority of the foveal algorithm in Figure 5.

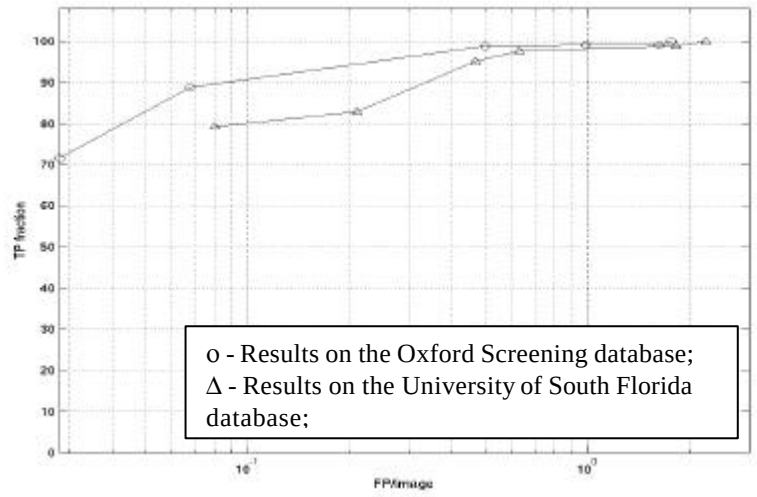

Fig. 4. The comparative FROC curve between the detection results on intensity images from the Oxford Screening Database and the University of South Florida Digital Database for Screening Mammography.

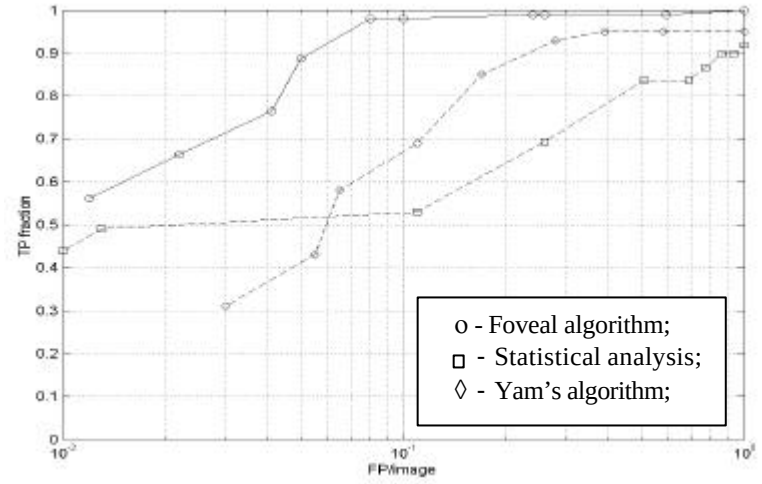

Fig. 5. The FROC curves of the three microcalcification-detection methods on SMF images, where we notice the better performance of the Foveal Approach.

3 Discussion

From a combination of solutions to partial differential equations (PDE), wavelet methods and statistics, the developed technique presents the user with a map of

detected microcalcifications. In the first original step of our method, 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. The contrast *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. *s* 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 number of iterations *t* can be expressed as a function of *s*, 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 may be used in a way that has minimal dependence on preset parameters and which uses some limited, but essential, a-priori knowledge. Potentially, this is more widely applicable in diffusion, a method that has attracted criticism for its parametrical dependency.

The second original step is the development of a method for adaptively thresholding the filtered results in order to segment microcalcifications. The combination of filters, statistical analysis and adaptive thresholding adds to the novelty of our technique. We compare detection methods on SMF images as well as the outcome of our algorithm on both SMF and intensity images. An example of microcalcification detection is shown in Figure 6.

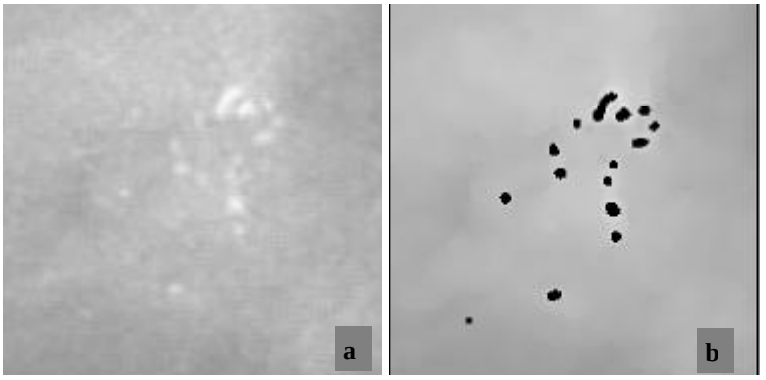

Fig. 6. An example of microcalcification detection. On the left (a) we present the input image in SMF format with a microcalcification cluster, while on the right (b) we note the result of our detection method.

The subsequent filters (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.

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 propose a fully automated non-parametric method to detect microcalcifications using the SMF normalised representation of the breast.

4 Conclusion

In this paper we presented a FROC analysis of an algorithm for the detection of microcalcification clusters mammography. The robustness of the algorithm has been demonstrated by the FROC 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. 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.

References

1. 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
2. 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
3. Highnam, R.P. Brady, J.M.: Mammographic Image Analysis. Kluwer Academic Publishers, Dordrecht Boston London (1999)
4. Kovesi, P.: Image Features from Phase Congruency. In: Videre: Journal of Computer Vision Research, Vol. 1 (1999) 1-26
5. 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
6. Linguraru, M.: Feature Detection in Mammographic Image Analysis. Ph.D Thesis, University of Oxford (2002)
7. 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
8. Weickert, J.: Anisotropic Diffusion in Image Processing. B.G. Teubner, Stuttgart (1998)
9. 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
10. 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