Image Processing in Scale-Orientation Space

Reyer Zwiggelaar
School of Information Systems
University of East Anglia
Norwich, NR4 7TJ, UK
reyer.zwiggelaar@sys.uea.ac.uk

Abstract

Scale-orientation signature space represents a non-linear transformation from normal grey-level image space. From this point of view there are similarities with other representations such as radial basis functions, Fourier, Mellin and wavelet transforms. Morphological related techniques are used to generate the scale-orientation space. A series of filtering techniques, based on selectively removing values in scale-orientation space, have been developed. Transportation across scales, but not orientations, in scale-orientation space is discussed as a way to preserve the base of image structures. Evaluation of the described techniques is based on a set of standard images which contain linear structures superimposed on a complex texture background. Both quantitative and qualitative results are presented.

Keywords. Image features, scale-orientation space, thresholding, filtering.

1 Introduction

Morphology [7] and sieves [1] have been used for the segmentation of image features. In general, when these techniques are applied to images the results contain enhanced structures, although these results depend on the structuring element associated with the techniques. However, recently these techniques have been used to obtain a scale-orientation representation at a pixel level [11]. Such a representation can be used, in combination with statistical modelling, for the detection of image structures. The main reason for incorporating statistical modelling is to generalise and clean the signatures. It has been shown that this mainly affects the small values in the scale-orientation representation with a bias towards smaller scales in the scale-orientation space [11]. It is known that the scale-orientation representation has desirable [2] and less desirable [3] properties.

Here we investigate a number of basic image processing techniques by using characteristics of scale-orientation space. In the first instance this concentrates on simple techniques such as thresholding, filtering and thinning.

The developed techniques can be used for the enhancement of image structures, and here we concentrate on the detection of linear structures. Because a standard image set is used the results can be compared directly with previous evaluations [12, 9]. We present both quantitative and qualitative results on the standard image set and mammographic data. The quantitative results are presented as receiver operating characteristic (ROC) curves [5], showing the trade-off between true positive and false positive detection of the linear structures.

BMVC 2002 doi:10.5244/C.16.18

2 Scale-Orientation Signatures

Two approaches to the extraction of scale-orientation signatures exist, one based on recursive median filtering [8, 1] and another on morphology [7]. Detailed accounts of the usage of both approaches can be found in [11] and [3].

The Recursive Median Filter (RMF) is one of a class of filters, known as sieves, that remove image peaks or troughs of less than a chosen size [1]. An efficient implementation of such an approach uses run-length coding to process the signal. By applying sieves of increasing size to an image, then taking the difference between the output image from adjacent size sieves, it is possible to isolate image features of a specific size. Sieves have been shown to have desirable properties when compared to other methods [4] of constructing a scale space [2]. In particular the results at different positions on the same structure are similar (local stationarity) and the interaction between adjacent structures is minimised (i.e. there is no leakage between structures).

For 2-D images, a 1-D RMF can be applied at any chosen angle, by covering the image with lines at this angle, ensuring that every pixel belongs only to one line, i.e. Bresenham algorithm [6]. By performing 1-D Directional Recursive Median Filtering (DRMF) at several orientations, a scale-orientation signature ($\Psi_{x,y}(\sigma,\theta)$) can be built for each pixel (x, y). The signature is a 2-D array in which the columns represent measurements for the same orientation (θ), the rows represent measurements for the same scale (σ), and the values in the array represent the change in grey-level at the pixel, resulting from applying a filter at the scale and orientation corresponding to the position in the array. The grey-level changes are measured with respect to the image filtered at the next smaller scale at the same orientation. A locally brighter structure is represented by positive values in the scale-orientation signatures, whilst locally darker structures result in negative values (small values are more indicative of noise). An example of such a signature can be found in Figure 1 (for display purposes all signature values equal to zero have been left blank). The columns represent measurements for the same orientation and the rows represent measurements for the same scale. It should be noted that a residual value is not incorporated in the signature. The residual value is the grey-level value left after the final filter has been applied. In general, when the final scale is equivalent to the image size, the residual value is similar across the image and equal to the mean intensity, but can be orientation dependent.

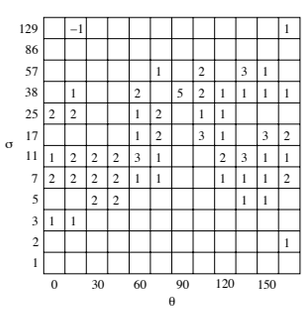

Figure 1: Typical scale-orientation signature from a linear structure.

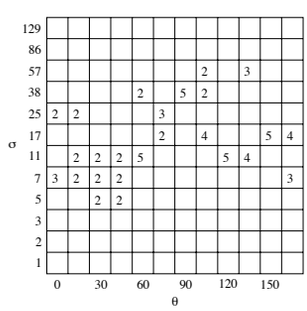

Figure 2: Illustration of the basic thresholding approach (with transportation).

This DRMF approach results in a large set of images, each containing one particular

part of the scale-orientation range, all extracted from the original image, i.e.

$$I(x,y) \Rightarrow \bigcup_{(\sigma,\theta)\in\Gamma\times\Theta} \{I(x,y,\sigma,\theta)\}\,, \tag{1}$$

where Γ is the set of scales and Θ is the set of orientations. Every image, $I(x,y,\sigma,\theta)$, has structures up to scale σ at orientation θ. All the scale-orientation signature information can be derived from this set of images. Every signature entry is determined by $\Psi_{x,y}(\sigma,\theta) = I(x,y,\sigma,\theta) - I(x,y,\sigma+\delta\sigma,\theta)$, where $\sigma+\delta\sigma$ indicates the scale after σ.

2.1 Reconstruction

When using the approach described in the previous sections a reconstruction of the original image is trivial, as the information at one particular orientation is sufficient for this process. This can be achieved by selecting any orientation (θ_0) and performing a summation over all the scales, i.e.

$$I(x,y) = \sum_{\sigma\in\Gamma} \Psi_{x,y}(\sigma,\theta), \tag{2}$$

which is equivalent to a column summation for every scale-orientation signature. The choice of orientation is arbitrary as each gives the same result (where it is assumed that the residual value is used in the reconstruction process).

However, as soon as an approximation, either to remove noise, filter out particular aspects or for generalisation purposes (e.g. dimensionality reduction by using principal component analysis [11]), is used the signatures reconstruction becomes less trivial and aspects of this will be discussed in the remainder of this work.

From a reconstruction point of view this now means that when using Eq. 2 the results can be orientation dependent. This can partially be circumvented by determining the reconstruction value for all orientations and obtaining the maximum (or median, or mean) value of these results. When considering the maximum Eq. 2 can be rewritten as

$$I(x,y) = \arg\max_{\theta} \left(\int_0^{\infty} \Psi_{x,y}(\sigma,\theta)\partial\sigma \right). \tag{3}$$

2.2 Thresholding

Thresholding in scale-orientation space is achieved by setting small values of $\Psi_{x,y}(\sigma,\theta)$ equal to zero. For well behaved image structures this means that the integral in Eq. 3 reduces to $\int_{\sigma_{low}}^{\sigma_{high}} \partial\sigma$, which indicates that both the top as well as the base of the structure is removed and only the large scale-orientation signature values in between σ_{low} and σ_{high} are preserved (see also Figure 5 in Section 3.5).

2.3 Transportation Aspects

The approach described in the previous section has the undesirable effect that the base of the image structures are removed as these tend to be represented by small values. Our second approach partially addresses this problem by allowing transportation of low values (below the threshold) at large scales to high values (above the threshold) at lower scales.

Effectively this means that at a particular orientation the extend of an image structure is determined by the signature values above the threshold and that the base of that structure is preserved. In effect this should result in an improved signal to noise ratio. To avoid discontinuities in the image structures transportation is only allowed in the direction of decreasing scale values. For the reconstruction this means that the integral in Eq. 3 can be written as $\int_{\sigma_{low}}^{\infty} \partial\sigma$ and it should be pointed out that this is only the case when $\int_{\sigma_{low}}^{\sigma_{high}} \partial\sigma$ is finite else the base $\int_{\sigma_{high}}^{\infty} \partial\sigma$ is not added.

3 Image Processing in Scale-Orientation Space

In general thresholding/filtering is based on discarding values that are below the threshold and results in binary images, or low values are discarded whilst keeping all the values above the threshold [8].

However, thresholding or filtering in scale-orientation signature space is less trivial. There are a number of issues that will need consideration. The scale-orientation signatures can contain both positive and negative values, so a decision has to be taken if these should be treated equally or that only positive values should be taken into account. A second aspect that will need to be addressed is the fact that the spacing in the scale dimension in the scale-orientation signatures does influence the results. This should be clear when an extreme case is considered where the scale dimension covers all the integer values between zero and the maximum scale. In this case it is very unlikely that the values in the scale-orientation signatures are large and the thresholding will result in very sparse data. This aspect is addressed by using a set of scales with a logarithmic distribution.

3.1 Thresholding - Manual Approach

In our most basic approach to the thresholding in scale-orientation signature space all the values smaller than a threshold are set to zero. An example of this, based on the scale-orientation signature shown in Figure 1, can be found in Figure 2, where the threshold used is equal to two and transportation has been used. As both higher (bright) and lower (dark) than average intensity image structures are important the thresholding is applied in an absolute sense.

3.2 Thresholding - Automatic Approach

Instead of using a fixed threshold value for all the signatures we have investigated a more local/adaptive approach, where the threshold value is determined at a single signature (and hence pixel) level. The threshold value is determined by two basic assumptions. The first assumption is that for a signature to be part of an image structure it should have signature values at all orientations (i.e. none of the columns of the signature contains only zeros). This assumption is not true for image structures which are larger than the largest scale used in the signature extraction stage, e.g. any straight linear structure which extends between two image boundaries. If any of the columns only contain zeros all signature values are set equal to zero, i.e. making the assumption that it is of no interest. For the signatures that do have values at all orientations a threshold needs to be determined. For such a signature to *remain* a structure signature the *all orientations* rule set out above

needs to be preserved. To do so the highest value for each orientation is determined and of these the lowest is chosen to be the threshold value. For the signature shown in Figure 1 the resulting thresholded signature can be found in Figure 2 (which in this case is equal to using a fixed threshold value equal to two).

3.3 Filtering - Single Scale

Filtering in scale-space becomes almost a trivial exercise with scale-orientation signatures. After selecting a particular scale all the signature values above that scale are checked (those at a smaller scale are set to zero). Transportation at a particular orientation is only allowed if there is a non-zero (or positive) value at the selected filtering scale. This means that structures at the selected scale are preserved whilst smaller and larger scale structures are suppressed. An example of such a filtering approach, at scales six to seven, applied to the signature shown in Figure 1 can be found in Figure 3 (it should be noted that this process not necessarily preserves the *all orientations* rule as described in the previous section).

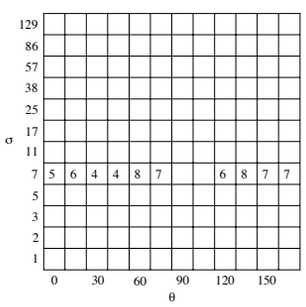

Figure 3: Illustration of filtering.

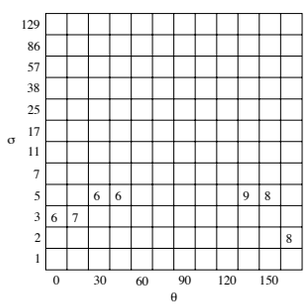

Figure 4: Illustration of thinning.

3.4 Thinning

It is possible to extend the scale-space filtering to include a range of scales. Especially when this is done for a band at the lower scales the resulting reconstruction only contains image structures over a limited set of scales which is closely related to thinning in normal image space. In this case a maximum scale is selected above which all the signature values are transported to the largest scale signature value available (i.e. if no signature values at or below the selected scale exist no transportation occurs). An example of a resulting signature, using the scale range from one to five, based on the signature shown in Figure 1 can be found in Figure 4.

3.5 One Dimensional Example

To visualise some of the aspects described so far a one dimensional Gaussian function will be used which is shown in Figure 5a. The effect of thresholding (see Section 2.2) (without transportation) in scale-orientation space (in essence this is just scale space as there is no orientation dependence for this one dimensional example) can be found in Figure 5b, which indicates that the top and the base of the Gaussian function have been

removed. When transportation (see Section 2.3) is added to the thresholding the base of the Gaussian function is preserved where the basic thresholding had a non-zero response. This is shown in Figure 5c. The effect of single-scale filtering (and transportation) (see Section 3.3) is shown in Figure 5d, which shows that the top of the Gaussian function has been removed but the base at the selected scale has been preserved. For this simple example this results in a binary response. Finally, the top of the Gaussian that was removed in Figure 5d can be preserved by using the thinning as described in Section 3.4 and the resulting one dimensional function can be found in Figure 5e.

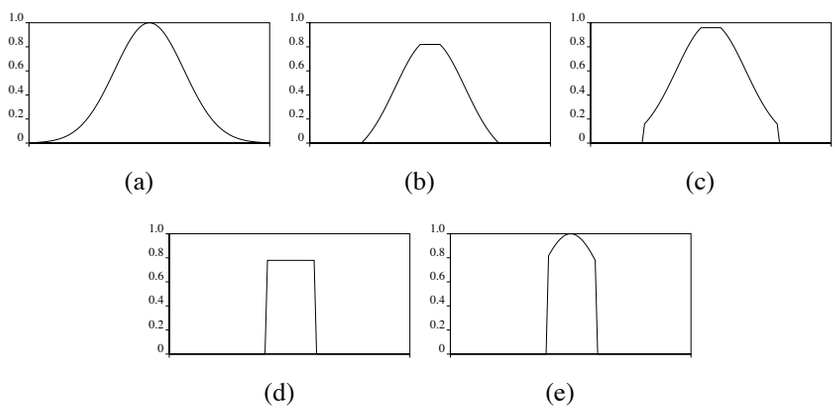

Figure 5: (a) A one dimensional Gaussian function. Processing of the Gaussian function using: (b) basic thresholding, (c) thresholding in combination with transportation, (d) single-scale filtering, and (e) thinning.

4 Evaluation

The evaluation of the developed approach is based on a standard set of images that have been used before in the evaluation of the detection of linear structures [12, 9]. This will cover both quantitative and qualitative evaluation aspects. In addition, qualitative results based on mammographic data are presented.

4.1 Standard Image Set

The performance of the methods described above is evaluated on a set of synthetic images designed to test the application to digitised mammograms, an example of which can be found in Figure 6. Linear structures were generated at known positions and orientations. The profiles of these linear structures were generated by a statistical model trained on spicules (the linear structures which radiate from spiculated lesions) and blood vessels, both of which are linear structures found in mammograms [10]. The line patterns were superimposed on real mammographic backgrounds. Figure 6 shows the patterns superimposed on a dense mammographic backgrounds. This set of synthetic images is representative of real mammographic linear structures as these are modelled on real mammographic linear structures and real mammographic background texture.

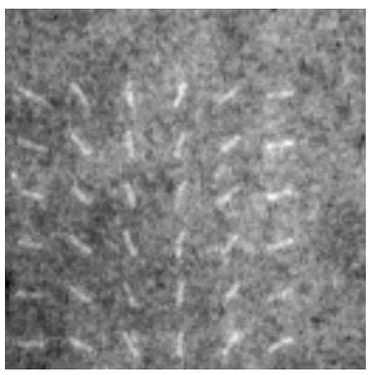

Figure 6: Artificial lines (for display purposes the image has been intensity stretched).

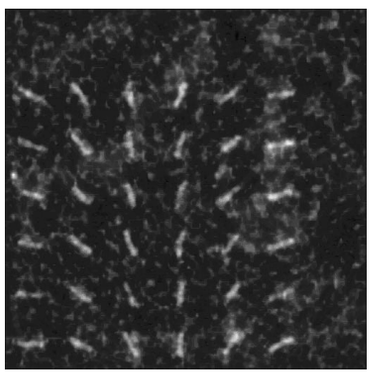

Figure 7: Reconstruction based on absolute thresholding in combination with the transportation approach.

The results of thresholding in scale-orientation space (see Section 3.1) applied to the image shown in Figure 6 can be found in Figure 7. Here we have used a constant threshold value equal to two. It should be clear that the linear structures have been enhanced, but it should also be mentioned that some of the background noise has not been suppressed. The remaining background noise is mainly made up of small blob-like structures.

Based on the image shown in Figure 6 the result of using scale based filtering (see Section 3.3) can be found in Figure 8. This shows only structures with a width of respectively three (Figure 8a) and six to seven (Figure 8b) pixels have been reconstructed. It also shows that, especially at the lower scales, this means that the image structures have been broken up which indicates discontinuities in scale space (which are caused by the way scale-space is sampled).

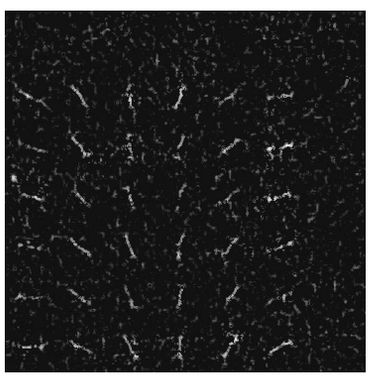

(a)

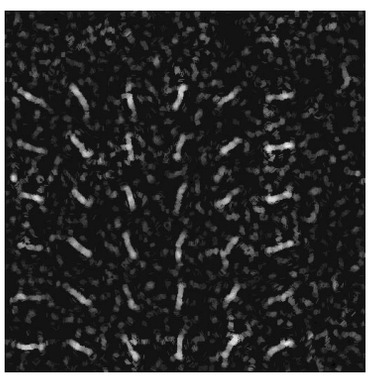

(b)

Figure 8: Reconstruction based on scale filtering where (a) scale equal to three and (b) scale equal to seven.

Applying thinning (as described in Section 3.4) to the original image gives the reconstructed images shown in Figure 9. Here the maximum scale is set to three and seven. This

results in respectively image structures covering the scales one to three and one to seven. Allowing the larger scale range results in better connected image structures. But it should be clear that even for the lower scale range most of the image structures are connected and this shows a substantial improvement over the single scale results (see Figure 8a).

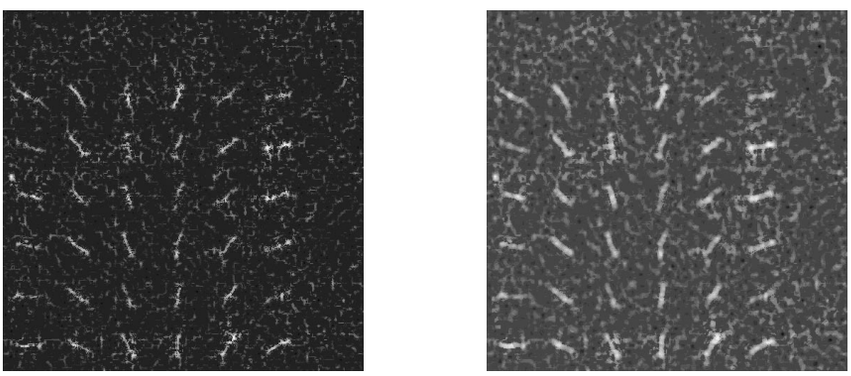

(a) (b)

Figure 9: Reconstruction based on thinning where (a) scale equal to three and (b) scale equal to seven.

ROC curves were obtained by varying a (reconstructed) pixel-intensity threshold and comparing the resulting pixel labels (line versus non-line) with ground truth. The closer the curve approaches the top left-hand corner (detection of all true positives at the cost of zero false positives) the better the detection [5]. In general, ROC curves will have a false positive range from zero to one. However, for display purposes we have reduced this and used a logarithmic scale.

The ROC curves associated with the images shown in Figures 7 - 9 can be found in Figure 10. These results show that the simple processing in scale-orientation space provides a significant improvement when compared with results based on non-processed scale-orientation signatures. However, the various processing techniques do not give a substantial difference in the ROC results, which can be explained by the fact that in all cases non-structures are suppressed and only small threshold values are used.

It should be mentioned that this processing in scale-orientation signature space produces results that are a significant improvement on previously published methods which were evaluated on the same standard image set [12]. To illustrate this we have included the ROC curve based on a non-linear line detection approach [12]. The ROC results based on filtering in scale-orientation space are similar to those based on the statistical modelling of the non-processed scale-orientation signatures [9]. This indicates the likelihood that all the statistical modelling was doing was to remove non-structure (sparse) signatures and small signatures values (see also the results in [11, 9]).

4.2 Mammographic Data

In Figure 11a a region from a mammographic image, containing a cluster of micro-calcifications and a number of linear structures, is shown. The resulting of thinning using

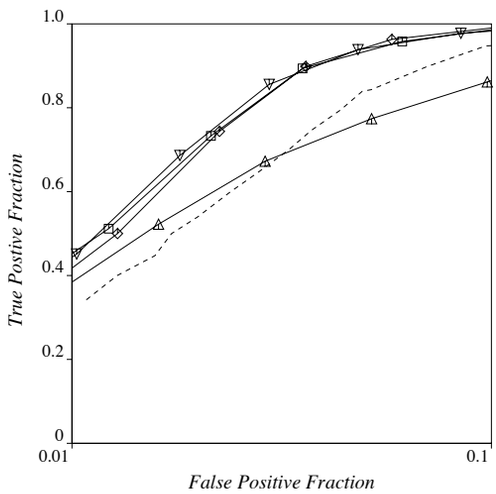

Figure 10: ROC classification curves, where the results are based on: the original signatures ($\triangle$), and reconstruction with transportation based thresholding ($\square$), filtering at a single scale ($\diamond$), and thinning using a range of scales ($\bigtriangledown$). The dashed line represents detection based on a non-linear detection approach [12].

scales one to seven results in the image shown in Figure 11b. This shows a clear enhancement of the linear structures.

5 Discussion and Conclusions

We have developed filtering techniques in scale-orientation signature space. Three distinct approaches were investigated. The first approach was based on the basic thresholding of the values in the scale-orientation signatures and allowed for the transportation of signature values below the threshold at large scales to signature values above the threshold at smaller scales. This results in a preservation of the base of image structures. The remaining approaches use a selected scale range for the filtering process. Again, transportation as described above was allowed to improve the enhancement process.

All three methods show a distinct improvement in the enhancement of image structures when compared to the original image set and also showed a substantial improvement when compared with other detection approaches [12]. In addition, based on the similarity in ROC results and other publications [11, 9] it seems likely that the statistical modelling of the original signatures [9] basically has the same effect as the presented scale-orientation space filtering, but this will need to be confirmed by future research.

In addition, for mammographic data visual evaluation showed a realistic enhancement of linear structures.

References

[1] J.A. Bangham, P. Ling, and R. Young. Multiscale recursive medians, scale-space and transforms with applications to image processing. *IEEE Transactions on Image Processing*,

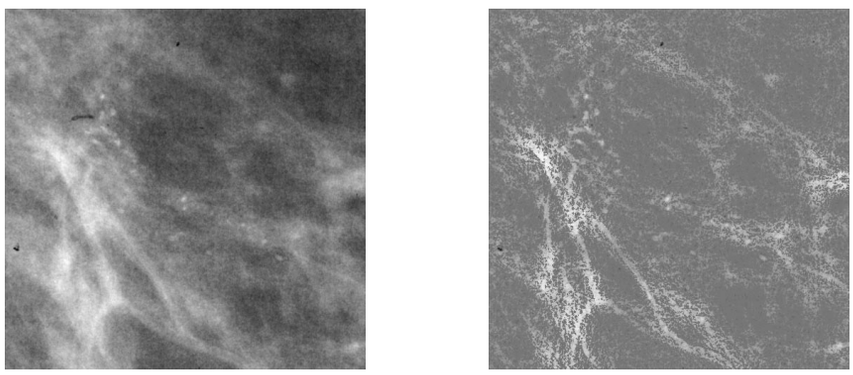

(a) (b)

Figure 11: (a) Mammographic region and (b) reconstruction based on thinning covering scale up to seven.

5(6):1043–1048, 1996.

[2] R. Harvey, A. Bosson, and J.A. Bangham. The robustness of some scale-spaces. In *Proceedings of the* 8^{th} *British Machine Vision Conference*, pages 11–20, Colchester, UK, 1997.

[3] A.S. Holmes and C.J. Taylor. Developing a measure of similarity between pixel signatures. 10^{th} *British Machine Vison Conference*, Nottingham, England:614–622, 1999.

[4] S.N. Kalitzin, B.M. ter Haar Romeny, and M.A. Viergever. Invertible orientation bundles on 2d scalar images. In 1^{st} *International Conference on Scale-Space Theory in Computer Vision*, number 1252 in Lecture Notes in Computer Science, pages 77–88, 1997.

[5] C.E. Metz. Evaluation of digital mammography by roc analysis. *Excerpta Medica*, 1119:61–68, 1996.

[6] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1992.

[7] P. Soille, E.J. Breen, and R. Jones. Recursive implementation of erosions and dilations along discrete lines at arbitrary angles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(5):562–567, 1996.

[8] M. Sonka, V. Hlavac, and R. Boyle. *Image Processing, Analysis and Machine Vision*. Chapman and Hall Publishing, 1993.

[9] R. Zwiggelaar, R. Marti, and C.R.M. Boggis. Detection of linear structures in mammographic images. 4^{th} *Conference on Medical Image Understanding and Analysis*, London, UK:97–100, 2000.

[10] R. Zwiggelaar, T.C. Parr, C.R.M. Boggis, S.M. Astley, and C.J. Taylor. Statistical modelling of lines and structures in mammograms. *Proceedings of SPIE*, 3034:510–521, 1997.

[11] R. Zwiggelaar, T.C. Parr, J.E. Schumm, I.W. Hutt, S.M. Astley, C.J. Taylor, and C.R.M. Boggis. Model-based detection of spiculated lesions in mammograms. *Medical Image Analysis*, 3(1):39–62, 1999.

[12] R. Zwiggelaar, T.C. Parr, and C.J. Taylor. Finding orientated line patterns in digital mammographic images. In *Proceedings of the* 7^{th} *British Machine Vision Conference*, pages 715–724, Edinburgh, UK, 1996.