DETECTION OF MASSES IN MAMMOGRAMS BY WATERSHED SEGMENTATION AND SPARSE REPRESENTATIONS USING LEARNED DICTIONARIES

Jostein Herredsvela, Kjersti Engan, Thor Ole Gulsrud, and Karl Skretting

University of Stavanger, NO-4036 Stavanger
E-mail: jostein.herredsvela@uis.no

ABSTRACT

In this paper a novel method for the detection of circumscribed masses in digital mammograms is presented. The proposed scheme uses morphological hierarchical watersheds with reconstructive morphological preprocessing in the segmentation process. The segmented regions represent suspicious regions in the mammogram that must be classified into two classes: Lesions and false detections.

The classification method is based on sparse representations of image blocks by learned dictionaries. The results are promising: 11 of 13 lesions are detected with 1.4 false positives per image.

1. DIGITAL MAMMOGRAPHY AND CAD

Breast cancer is a leading cause of cancer deaths among women. For women in the developed countries it is the most frequently diagnosed cancer. The most efficient way to reduce the mortality rate is early detection and removal of the cancerous tissue. To this end, large-scale mammographic screening programs are currently running in a large number of countries. In Norway, all women between 50 and 69 years of age are invited to participate in the national screening program. The probability of dying of breast cancer has been reduced by 30 % for women participating in the screening program [1].

Mammograms are x-ray projections of the breast tissue onto a detector array or a film plate, see Figure 1. Tumors often consist of dense tissue, and thus absorb most of the incident x-rays. They can therefore often be seen as bright regions in the mammograms.

The screening programs generate a vast number of mammograms which are to be carefully examined, usually by two radiologists. This is a costly and time consuming process. A major concern is the number of false negative errors, i.e. cases in which a mammogram containing a malignant tumor is classified as normal. It is seen that between 10 and 30 % of cancers are missed during routine screening [2].

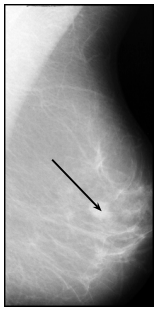

Figure 1: Example of a mammogram (*mdb010rm* from the MIAS public database). The location of the lesion is indicated by the arrow.

The last years have seen a large interest in Computer Aided Detection (CAD) of breast cancer. Most CAD systems are intended to give the radiologists a second opinion of the suspicious regions in mammograms. One problem with CAD systems is that "perfect" detection (i.e. detection of all tumors present) in practice leads to false positives (FPs).

The main factor that makes breast cancer detection in mammograms a very difficult task in image analysis is that there is a large variation in the appearance of both normal breast tissue and of cancerous tissue. Some breasts have very dense or glandular parenchymal tissue that is radiopaque, while other breasts are mostly fatty and radiolucent. There are several types of breast abnormalities that are visible in mammograms: Asymmetry between the breasts, calcifications, increase in breast tissue density, lesions, and architectural distortions. CAD performance for microcalcifications is on an acceptable level [3, 4] whilst the performance for lesions of various types is poorer [5]. The lesion class includes circumscribed lesions, which are compact and lobular or circular/oval shaped, and spiculated/stellate lesions which consist of a central mass (not always present) with radiating spicules in some or many directions. So far in our work we have focused on the class of lesions which is easiest to

detect: Circumscribed lesions.

2. METHOD

The proposed method consists of two steps. First the image is segmented using morphological watersheds [6]. The segmented regions are then classified using sparse representation of image blocks using learned dictionaries. Further removal of false positives can be done by using other standard feature extraction methods, but this is not considered in this paper. Each step in the method is explained below.

2.1. Image Segmentation by Morphological Watersheds

Morphological watershed segmentation is a simple yet powerful image segmentation method [7]. Picturing the image as a landscape with valleys (local minima) and peaks/ridges (local maxima), the watershed lines (i.e. dividing lines) between the valleys are found by flooding the valleys, using one source or *marker* in each valley, and building dams where the waterfronts meet. The dams together constitute the watershed lines. It is well known that the direct use of watersheds on a natural image almost always leads to oversegmentation, due to the large number of spurious extrema present in the image. Consequently the input image is filtered/simplified using a reconstructive alternating open/close sequential filter [7] prior to watershed segmentation. This filter flattens small details while keeping larger structures relatively unaltered. One important feature of the filter is that the edges of the remaining structures are preserved. Assuming that objects of interest have relatively sharp edges, the object edges can be found as the morphological gradient of the simplified image. This results in an image in which the valleys are the interior of the objects and the ridges are the object edges. It is this image that is flooded during the watershed segmentation.

Since most objects of interest (i.e. tumors) have relatively high intensity, the markers (i.e. sources) used in the flooding are the regional maxima in the image. Only maxima with contrast larger than a certain *dynamics threshold* are used. In this manner only regions with relatively large contrast with respect to their background are segmented.

A more complete description of watershed segmentation of mammograms is found in [6].

2.2. Sparse Representation using Learned Dictionaries

The watershed segmentation results in a number of binary regions indicating suspicious regions in the mammogram, see Figure 5(b). In general, several of these are false positives. The next step is to reduce the number of false positives while keeping the true lesions. Sparse representations of image blocks by learned dictionaries are used to classify the suspicious regions. Such an approach has been successfully used in classification of various textures in [8]. Below we give a summary of the theory of learned dictionaries.

Any N-dimensional vector can be written as a linear combination of $K \geq N$ vectors that span the space. A good approximation to an N-dimensional signal vector $\mathbf{x}$ can often be obtained by linearly combining only a few of these K vectors. Mathematically, $\mathbf{x} = \mathbf{F}\mathbf{w}+\mathbf{n}$, where $\mathbf{F}$ is a learned frame/dictionary [9] in the form of an $N \times K$ ($K \geq N$) matrix, and where $\mathbf{w}$ is a sparse coefficient vector. $\mathbf{n}$ represents the approximation error. The columns of $\mathbf{F}$ are *dictionary vectors*. We emphasize that $\mathbf{F}$ must be *learned*, contrary to the case of orthogonal expansions.

Given a collection of L vectors that are to be approximated using $\mathbf{F}$ an $N \times L$ data matrix $\mathbf{X}$ can be formed, of which the columns are the collection of vectors. Then the representation can be written $\mathbf{X} = \mathbf{F}\mathbf{W} + \mathbf{N}$.

The *training matrix* from which $\mathbf{F}$ is to be learned is denoted $\mathbf{Y}$. In order to be able to make a good approximation to $\mathbf{X}$ it is important that $\mathbf{Y}$ has more or less the same "qualities" as $\mathbf{X}$. Learning $\mathbf{F}$ implies minimizing the representation error $||\mathbf{Y} - \mathbf{F}\mathbf{W}||$ subject to $\mathbf{W}$ being sparse [9]. The sparsity constraint is that only s dictionary vectors may be used in the representation of each of the columns in $\mathbf{X}$. In this work the Method of Optimal Directions (MOD) is used for learning [9]. In short MOD works as follows: First an initial dictionary $\mathbf{F}$ is formed e.g. by selecting every L/K'th of the labelled data vectors, as done in this work. The columns of $\mathbf{F}$ are then normalized. Starting from the initial dictionary, an iterate of $\mathbf{W}$ is found using the Order Recursive Matching Pursuit (ORMP) algorithm [10]. Then a new dictionary $\mathbf{F}$ is found from this $\mathbf{W}$ as

$$\mathbf{F} = \mathbf{Y}\mathbf{W}^{\dagger}, \tag{1}$$

where $\mathbf{W}^{\dagger}$ is the pseudoinverse of $\mathbf{W}$. The procedure is repeated until the error has converged. The algorithm does not guarantee convergence to an optimum solution, but usually provides a good suboptimal solution. The dictionary $\mathbf{F}$ is now *learned*, and it may now be used to approximate or represent any column vector in $\mathbf{X}$ using a linear combination of s of the dictionary vectors.

2.3. Classification of Segmented Suspicious Regions

We now return to the question of how to use learned dictionaries in the classification of the watershed segmented regions.

An $n \times n$ image block consisting of a pixel and its neighborhood pixels can be reshaped into a (column) vector of dimension $N = n \cdot n$. Each such vector in the segmented regions that are to be classified are represented/approximated using two different learned dictionaries. The regions are considered one at a time. An example of lesion tissue and the corresponding segmented binary region is shown in Figure 4(a) and 4(b).

The training matrices are extracted from the segmented regions in a small set of mammograms. All true (i.e. lesion) regions are used, but only a subset of the false regions (the latter regions are dominant in number). One dictionary has been learned using blocks from false regions and the other dictionary has been learned using blocks from lesion regions as training matrices, $\mathbf{Y}$. For the application in this work it seen that subtraction of the regional grayscale average is crucial in order to obtain good results. This must be done during both training and testing. Since the breast tissue as seen in mammograms often have a dominant direction all image blocks are rotated 90°, 180°, and 270°, prior to reshaping them into training vectors, to avoid ending up with dictionaries with directional qualities. Naturally the unrotated blocks are used as well. We make sure that none of the block pixels are located outside the segmented region. The image blocks used overlap each other in the original images. The learned dictionary blocks/vectors obtained in one case with $s = 3$ and $N = K = 121$ is shown in Figure 2(a) for the dictionary learned using lesion data, and in 2(b) for the dictionary learned using normal tissue data. The most frequently used (when representing the training set) block is shown at the top left in each figure. The least frequently used block is in the bottom right. The ordering is top-down, left to right.

Figure 3 shows the learned dictionary blocks/vectors obtained with $s = 1$ and $N = K = 121$. The most frequently used blocks are slightly more uniform for the lesion class, indicating that the interiors of the lesions are smoother than the falsely segmented normal tissue.

Each of the vectors generated from the region pixels are represented using both dictionaries, and the corresponding representation error is calculated. We now have two representation error images for the region. These error images are now filtered using a gaussian low-pass filter, see Figures 4(c) and 4(d). Note that there is almost no visible difference between the two images, but still a correct classification is obtained.

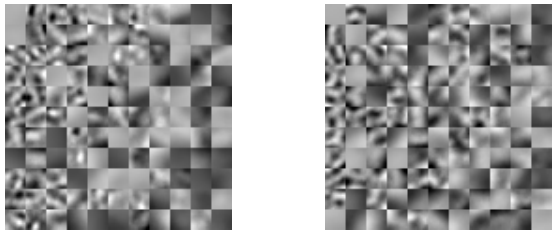

(a) (b)

Figure 2: An example of learned dictionary blocks. $s = 3$. $N = K = 121$. Figure (a): Ranked learned dictionary blocks for the lesion class. Figure (b):Ranked learned dictionary blocks for the lesion class.

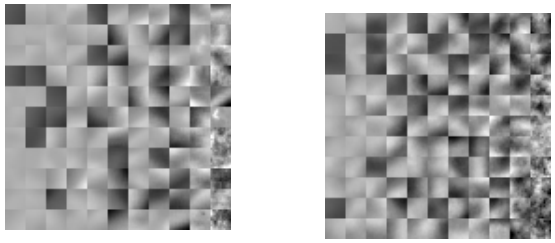

(a) (b)

Figure 3: An example of learned dictionary blocks. $s = 1$. $N = K = 121$. Figure (a): Ranked learned dictionary blocks for the lesion class. Figure (b):Ranked learned dictionary blocks for the lesion class.

This is common in texture classification. In this work it is justified because we are not interested in locating very small texture regions. A pixel-wise classification of the interior of the region is now performed. If error image no. i (= 1 or 2) has the smallest value for the pixel considered, this pixels is assigned class

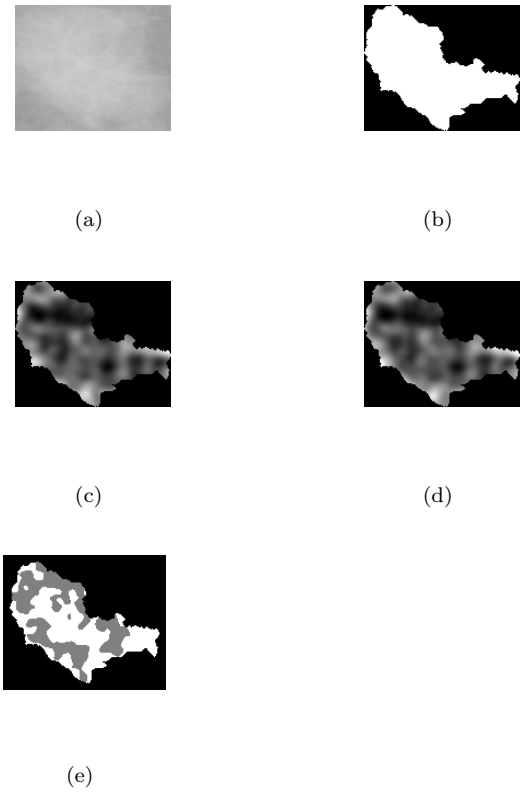

Figure 4: Illustration of the regional classification. $N = 121, K = 218$. Figure (a): A subimage containing a lesion. Figure (b): The corresponding detected binary region. Figure (c): The smoothed representation error obtained using the "normal" dictionary. $\sigma = 3$ was used in the smoothing. Figure (d): The smoothed representation error obtained using the "lesion" dictionary. $\sigma = 3$ was used in the smoothing. Figure (e): The classified pixels.

i, see Figure 4(e). Normal/false regions are assigned the value 1 and lesion pixels are assigned the value 2. Finally the entire region is assigned class j (= 1 or 2) if the relative number of pixels of this class in the region surpasses a certain percentage threshold p. The specificity of the method can be varied by adjusting this percentage threshold.

3. RESULTS

The detection method was tested on mammograms from the MIAS database provided by the Mammographic Image Analysis Society (MIAS) in the UK [11]. There were 4 training mammograms and 12 test mammograms. 11 of the mammograms contains one lesion each while one mammogram contains two lesions. Prior to use the mammograms were downsampled from $50\mu m$ resolution to $100\mu m$ resolution. The training matrix for the lesion class contained approximately $L = 15000$ vectors and the training matrix for the normal class contained approximately $L = 6600$ vectors. The initial dictionaries were generated by selecing every L/K'th training vector, as stated previously. Other initial dictionaries may be chosen, and the resulting learned dictionary will be different. However, as long as normalized vectors from the training matrix (randomly or uniformly selected) are used as initial dictionary, the performance of the resulting dictionaries will be approximately the same.

Based on experience, $s = 3$ vectors were used in the representation. Intuitively one might expect the performance to become better with increasing s beyond 3. However when s becomes large the dictionaries may represent blocks from lesions and normal regions equally well, since a linear combination of s blocks is used.

The parameters of the segmentation step were fixed in all the experiments: A dynamics/contrast threshold of 4 gray levels was used, and the reconstructive open/close preprocessing operation was repeated 8 times using increasingly large structuring elements. This yielded a good segmentation quality and a relatively large number of detections. In general using large image blocks performed better than smaller ones. In this work $n = 11$, i.e. a block size of 11×11, was used.

A representative mammogram with corresponding segmentation and classification results, as well as ground truth information, is shown in Figure 5.

The classification results obtained are summarized in Table 1 for different low-pass filter widths and percentage thresholds. A detection is counted as true if 20% of the truth circle (i.e. a circle indicating the location and size of the lesion) is covered by it. The

$\sigma = 1$	$\sigma = 5$	$\sigma = 5$
$p = 50\%$	$p = 50\%$	$p = 44\%$
$TPr = 69\%$	$TPr = 77\%$	$TPr = 85\%$
$FP = 1.3$	$FP = 1.0$	$FP = 1.5$

Table 1: Detection performance in terms of True Positive rate (TPr) and False Positives per image (FP) with varying low-pass filter widths σ and classification thresholds p

detection performance of the method also varies with the number of dictionary vectors K. See Table 2. In general the best results are obtained using a relatively large K.

The lesions in the mammograms *mdb080rm* and *mdb091lm* are not classified correctly using any reasonable parameter settings. For the former, this may be due to the fact that the lesion differs much in appearance from the lesions used for training. Normal tissue from other regions in the breast has been projected onto the lesion. For *mdb091lm* the classifica-

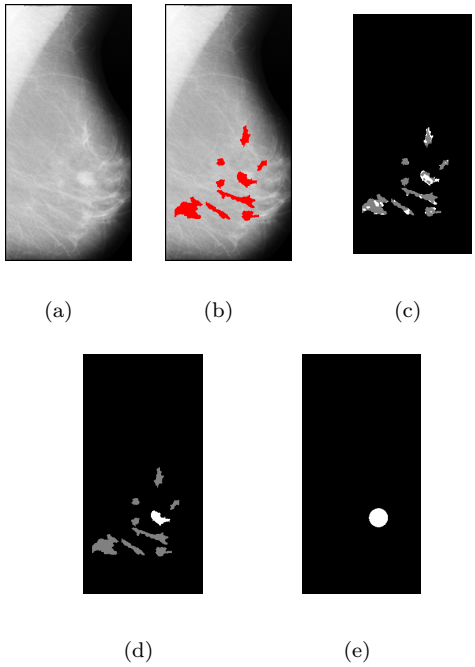

(a) (b) (c) (d) (e)

Figure 5: Detection results for mammogram *mdb010rm*. $N = K = 121$, $s = 3$, $\sigma = 7$, $p = 50\%$, and 150 iterations were used. Figure (a): Mammogram. Figure (b): Detected regions. Figure (c): Pixel-wise classified regions. The white regions are classified as lesions. Figure (d): Further regional classification. Figure (e): Truth image.

$K = 81$	$K = 81$	$K = 218$	$K = 218$
$p = 50\%$	$p = 40\%$	$p = 50\%$	$p = 40\%$
$TPr = 62\%$	$TPr = 77\%$	$TPr = 85\%$	$TPr = 77\%$
$FP = 0.66$	$FP = 1.3$	$FP = 1.4$	$FP = 1.1$

Table 2: Detection performance in terms of True Positive rate (TPr) and False Positives per image (FP) with varying K and classification thresholds p. A low-pass filter width of $\sigma = 7$ was used.

tion problems may due to a poor segmentation. The lesion is detected by the watershed algorithm, but the detection is too large. See Figure 6. In this way a quite large percentage of the detection may represent a texture atypical of lesions. Note that the area classified as "lesion", i.e. the white area, fits well with the truth circle. This case indicates that a good segmentation is important in order to get a correct classification. In this respect the proposed method does not differ from other known detection methods.

Note that in previous experiments, classification of these lesions using conventional feature extraction has also been problematic.

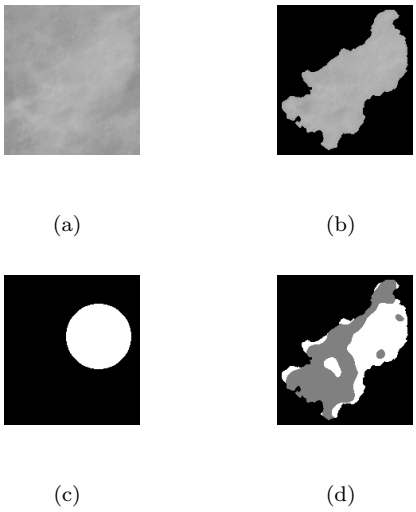

(a) (b) (c) (d)

Figure 6: The lesion in the mammogram *mdb091lm*. $N = 121, K = 121$. Figure (a): A subimage containing a lesion. Figure (b): The segmented region. Figure (c): The truth image. Figure (d): The pixel-wise classified detected region. $\sigma = 5$ was used in the smoothing.

4. CONCLUSION

A novel method for mammogram segmentation and classification has been presented. The segmentation step is based on morphological watersheds and the classification is based on sparse representation by learned dictionaries. The results obtained are promising. The use of learned dictionaries in the context of mammography has great potential and should be further explored.

5. REFERENCES

[1] *www.kreftregisteret.no*.

[2] R. E. Bird, T. W. Wallace, and B. C. Yankaskas. Analysis of cancers missed at screening mammography. 184(3):613–617, 1992.

[3] X. Zhang, T. Hara, H. Fujita, N. Shinohara, Y. Ooe, T. Iwase, and T. Endo. Artificial neural network method for detecting clustered microcalcifications in masses on mammograms. In *Proc. 6th International Workshop on Digital Mammography*, pages 391–393, 2002.

[4] T. O. Gulsrud. Optimal feature extraction in digital mammograms. In *Proc. 6th International Workshop on Digital Mammography*, pages 402–404, 2002.

[5] G. M. te Brake and N. Karssemeijer. Segmentation of suspicious densities in digital mammograms. *Medical Physics*, 28(2):259–266, February 2001.

[6] J. Herredsvela, T. O. Gulsrud, and K. Engan. Detection of circumscribed masses in mammograms using morphological segmentation. In *Proc. SPIE*, volume 5747, pages 902–913, April 2005.

[7] Edward R. Dougherty and Roberto A. Lotufo. *Hands-on Morphological Image Processing*. 2003.

[8] K. Skretting and J. H. Husøy. Texture classification using sparse frame based representations. *Eurasip JASP*. Accepted for publication.

[9] K. Engan, S. O. Aase, and J. H. Husøy. Multi-frame compression: Theory and design. *Signal Processing*, (80):2121–2140, October 2000.

[10] M. Gharavi-Alkhansari and T. S. Huang. A fast orthogonal matching pursuit algorithm. In *Int. Conf. on Acoust. Speech and Signal Proc.*, pages 1389–1392, Seattle, U.S.A, May 1998.

[11] *http://www.wiau.man.ac.uk/services/MIAS/MIAScom.html*.