Pattern Recognition ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at SciVerse ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

Detection and discrimination of disease-related abnormalities based on learning normal cases

K. Sai Deepak*, N.V. Kartheek Medathati, Jayanthi Sivaswamy

Center for Visual Information Technology, IIIT Hyderabad, Hyderabad 500 032, India

ARTICLE INFO

Article history:
Received 1 March 2011
Received in revised form
31 January 2012
Accepted 17 March 2012

Keywords:
Abnormality detection
Computer-aided diagnosis
Learning normal
Medical images
Shape descriptor
Texture descriptor

ABSTRACT

Detection of abnormalities from medical images is of key interest in developing computer-aided diagnostic tools. In this paper, we observe the key challenges for representation and feature extraction schemes to be met for detection of abnormalities by *learning normal cases*. We introduce an image representation, motivated by the effect of motion on perception of structures. This representation is based on a set of patterns called generalized moment patterns (GMP) generated via induced motion over regions of interest, for learning normal. The proposed GMP has been utilized to develop a scheme for addressing two well-known problems: lesion classification in mammograms and detection of macular edema in color fundus images. The strengths of this scheme are that it does not require any lesion-level segmentation and relies largely on normal images for training which is attractive for developing screening tools. The proposed scheme has been assessed on two public domain datasets, namely, MIAS and MESSIDOR. A comparison against the performance of state of the art methods indicates the proposed scheme to be superior.

1. Introduction

The use of computer-aided diagnosis (CAD) systems, to provide second opinion to human observers from images, is becoming an integral part of everyday clinical practice. These systems are intended to improve the accuracy and confidence of diagnosis performed by the user and also speed up the diagnosis process [1,2]. Given an image, the conventional CAD systems search for disease specific patterns, to make a decision on the presence/absence of an abnormality either in the entire image or within a region of interest (ROI). Therefore, detecting abnormality in an image requires solution to the challenging task of learning all the disease-specific visual patterns. Identifying and characterizing a single disease pattern is a difficult task in itself because of the high variability seen within medical images. High variability across and within images lead to loss of differentiation of disease-related abnormalities from background structures. Typical variations include illumination via bias fields and contrast which arises due to variations in imaging protocol and tissue types [3,4].

Towards addressing these problems, recently, an alternative approach to detection of disease-related abnormalities has gained interest where the visual patterns from normal cases are learnt and the presence of a disease is treated as a deviation from the normal cases [5–7]. There are some inherent advantages of using this approach. Firstly, the CAD system is required to learn only the known structures (along with their variabilities) observed in normal images. This effort is further simplified when pre-defined imaging protocols are used in acquisition. Multiple symptoms of a disease (which occur during different stages) can now be detected in one shot as a deviation from the normal. Secondly, the disease annotation requirement is reduced to image level rather than at the local, lesion-level. This results in a significant reduction of effort as obtaining annotations is a highly labor intensive task incurring huge costs and time. Further, it has been argued that a first step of disease/no disease classification of images is beneficial for the grading at the lesion level [8] as it helps to remove the true negatives or normal class of images. The reasoning behind this argument is that the main function of any grading system is to identify those patients who require further referral to experts. Hence, rejection of normal cases in screening scenario can significantly reduce the workload of a human grader.

In the dominant paradigm of designing disease-specific CAD systems, image representation and descriptors are used to describe and differentiate between disease symptoms. These are then used to detect/segment each manifestation separately. In the alternative approach based on learning of images from normal cases have similar requirements: derive representation and descriptors to capture the *normal* cases in a compact feature subspace in such a manner that the disease cases can be clearly

* Corresponding author. Tel.: +91 9886918469; fax: +91 40 6653 1413.
E-mail addresses: sai.deepak@research.iiit.ac.in (K.S. Deepak), nvkartheek@research.iiit.ac.in (N.V.K. Medathati), jsivaswamy@iiit.ac.in (J. Sivaswamy).

0031-3203/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.patcog.2012.03.020

discriminated as clear deviations from this subspace. In this paper, we examine this problem and propose a novel image representation based on patterns derived by inducing motion in static medical images. The rest of the paper is organized as follows: in Section 2, we draw the analogy between the existing background-learning problem in Image/Video Surveillance and learning normal cases in medical images and establish the challenges in directly using the former methods for the latter. In Sections 3 and 4, we derive the proposed representation and recommend features for learning normal cases. The representation is motivated by the effect of motion on human visual perception. In Sections 5 and 6, we experimentally validate the scheme on two popular screening problems using two public data sets, namely, classification of lesions in digital mammograms and detection of diabetic retinopathy from retinal images also known as color fundus images (CFI).

2. Background

The problem of learning normal in medical images shares similar motivation with the well researched problem of learning background in computer vision. Approaches based on learning background independently or in conjunction with foreground have been tried in the computer vision community for addressing problems like, object/motion detection on both dynamic and static scenes [9–13]. This background information is extracted either at each pixel or in local regions and a statistical model is imposed.

2.1. Abnormality detection techniques in non-medical domains

Statistical models like Gaussian [13] and mixture of Gaussian [9,12] have been used to represent each background pixel given a static or dynamic scenes. A fundamental assumption here is that a mixture of Gaussian distributions can capture various complex background appearances [9]. The disadvantage with these kind of approaches is that foreground pixels with statistics similar to background are easily misclassified [10]. Principal component analysis (PCA) based approach for detection and recognition in the presence of cluttered background has also been tried for face recognition [11]. Here, the role of learning normal is limited just to add robustness to the task of face recognition in the presence of noise and face-like background structures.

2.2. Abnormality detection in medical images

In medical images, voxel intensity and contextual information have been used for segmentation of tumors in brain MRI images by [5]. Here, features corresponding to healthy tissues are learnt and used to find the deviation of a given voxel by identifying its membership within a class using an expectation–maximization based segmentation framework. This framework has been verified for segmentation of large tumors in brain MRI images which is relatively a simple problem. This method requires initialization of tumor boundary for segmentation which limits its use for automated detection and segmentation of pathology.

Automatic detection of Drusen from color fundus images has been attempted by suppressing normal background pixels in [14]. A wavelet feature for every pixel is computed and used with support vector data description (SVDD) to find the smallest n-dimensional sphere that fits all the normal pixels. All pixels lying outside this sphere are considered abnormal. This method has been verified on a very small dataset of six images. Higher order local auto-correlation (HLAC) features have been used in [6] to integrate local contextual information in histopathological images. Every image is thus represented with a feature vector of size 25 corresponding to all the HLAC patterns up to second order. These are used within a PCA framework to learn the normal subspace and find deviations to detect the degree of abnormality.

A similar technique has been attempted for white matter lesion detection in brain MRI [7]. Multiple subspaces at different scales, rather than a single subspace, is employed for learning distribution of normal intensities. Principal components are computed for each subspace and the original image is reconstructed to estimate the deviation from normal.

The past work on learning normal cases in medical images focuses on differentiating a disease manifestation from the normal background. While this is important in the abnormality detection task, no effort has gone into deriving appropriate representation and features that are suited for describing normal cases. Therefore there is a need for suitable image representation and feature that (a) can address the creation of a compact normal (cases) subspace which is simple to learn and (b) enhances the separability of the subspaces corresponding to normal and abnormal and cases. This approach of learning normal cases for abnormality detection is nascent which is indicted by the fact that most reported methods have been tested on relatively small datasets. We next present a method to generate a normal subspace and demonstrate how it can be used to learn normal cases.

3. Generalized moment patterns

The representation required for learning the normal has to capture a broad context of the appearance of normal. The following are some of the desired characteristics of such a representation:

1. Capacity to model the normal subspace without the need for segmentation of individual structures or tissues. This implies that the characteristic appearance of a given region of interest (ROI) as a whole has to be captured using a compact descriptor.
2. Robustness to deviations that occur in the appearance of the normal and in the presence of noise in the images. This is very important in medical images.
3. High enough sensitivity to reflect contextual disturbance in the presence of abnormalities.

In order to design a representation that meets the above criteria we take inspiration from human visual perception.

3.1. Impact of motion on imaging and human vision

It is well known that relative motion or long exposure of camera imaging a moving object results in motion blur which results in a streaky appearance of the moving object in the captured image. This is due to the smearing of intensities of the moving object at the receptor. This effect can be simulated by inducing motion in a static scene as in Fig. 1. The two test images show objects (two dots) on plain and textured backgrounds, respectively. The effect of induced translation is shown in the second column of Fig. 1, while the effect of induced rotation can be seen in the third column. These images were obtained by averaging a set of translated/rotated versions of the original image. It can be seen that the intensities of these objects are smeared along the path of motion thereby leaving a *motion signature*. This effect is much more interesting for the color image in Fig. 1. Here, the textured background attains certain homogeneity due to induced motion whereas the black and red dots appear as streaks. Thus, for a given type of induced motion, the signature produced by an object depends on the local contrast of the object and its spatial extent. The smearing for the low-contrast object is visibly less significant when compared to high contrast objects.

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020

Mathematically, the above observations can be expressed as follows. The response M of a sensor/receptor to an object $f(x,y)$ in motion at a constant velocity is a time integral of a transformation function G_t applied to the object:

$$M(x,y) = \int_0^T G_t\{f(x,y)\}\, dt \tag{1}$$

Here G_t represents the state of object $f(x,y)$ in motion at time t, $M(x,y)$ is the resultant effect of motion and T is the duration of interest. In the absence of motion, $M(x,y) = f(x,y)$.

We propose simulating this generation of patterns at the receptor level and consider the pattern of interest to be the responses of a fixed set of receptors which sample the intensities at a location from a moving image. The receptor's response is a result of information collation over the path traversed by the image. Given an image I, a moving image is generated by inducing motion to I. This collation of information should induce, in some sense, a wide sense stationarity in a given context and hence enable the use of the generated pattern for learning the normal class of static images.

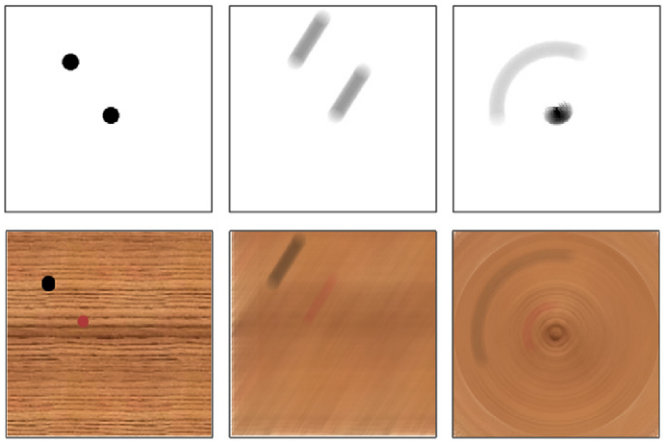

Fig. 1. Effect of inducing motion. Left to right: test images with uniform (top row) and textured background (bottom row); resulting images after translation and rotation. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

Next, we present a method to generate the desired motion signature at the receptor end.

3.2. Modeling receptor response to motion

Consider a set of static receptors whose response to a moving image is to be modeled as a spatio-temporal sampling. For a given image $I(\bar{r})$ where $\bar{r}$ is a position vector, after inducing motion, the intensity at location $\bar{r}$ at time t_i is denoted as $I_{t_i}(\bar{r})$.

We consider only planar motion, namely, translation and rotation around the center of the image. The given image, without any induced motion is defined to be the static image: $I_{t_0}(\bar{r}) = I(\bar{r}), \forall r$.

The number of samples N, received by a receptor $m_{\bar{r}_o}$ at location $\bar{r}_o$, over a duration T is $N = T/\tau$ where τ is the sampling interval. The response of the receptor at $\bar{r}_o$ depends on the sample set $\tilde{S_{\bar{r}_o}}$ which is given as

$$\tilde{S_{\bar{r}_o}} = \{I_{t_0}(\bar{r}_o), I_{t_1}(\bar{r}_o), I_{t_2}(\bar{r}_o), \ldots, I_{t_N}(\bar{r}_o)\} \tag{2}$$

The samples in this set are generated by a transformation function applied to I_{t_0}. If the induced motion is a translation at constant velocity, then the consecutive samples $I_{t_i}(\bar{r}_o)$ and $I_{t_{i+1}}(\bar{r}_o)$ are related as follows: $I_{t_{i+1}}(\bar{r}_o) = I_{t_i}(\bar{r}_o + \bar{v}\tau)$ where $\bar{v}$ is the velocity vector. The response of a receptor at location $\bar{r}_o$ is defined as

$$m_{\bar{r}_o} = \mathcal{F}\{\tilde{S_{\bar{r}_o}}\} = \mathcal{F}\{I_{t_0}(\bar{r}_o), I_{t_1}(\bar{r}_o), I_{t_2}(\bar{r}_o), \ldots, I_{t_N}(\bar{r}_o)\} \tag{3}$$

where $\mathcal{F}$ is the collation function of the receptor applied to the ensemble of motion generated samples. One option is to consider these samples as the outcome of a random experiment and the collation function as a moment generating function. In this case, the receptor response is an nth order moment of the ensemble $\tilde{S_{\bar{r}}}$. Thus, given an input image $I(\bar{r})$, we have a moment map $M(\bar{r})$

$$M(\bar{r}) = E_i\{\tilde{S_{\bar{r}}}\} \tag{4}$$

E_i indicates the i^{th} order moment.

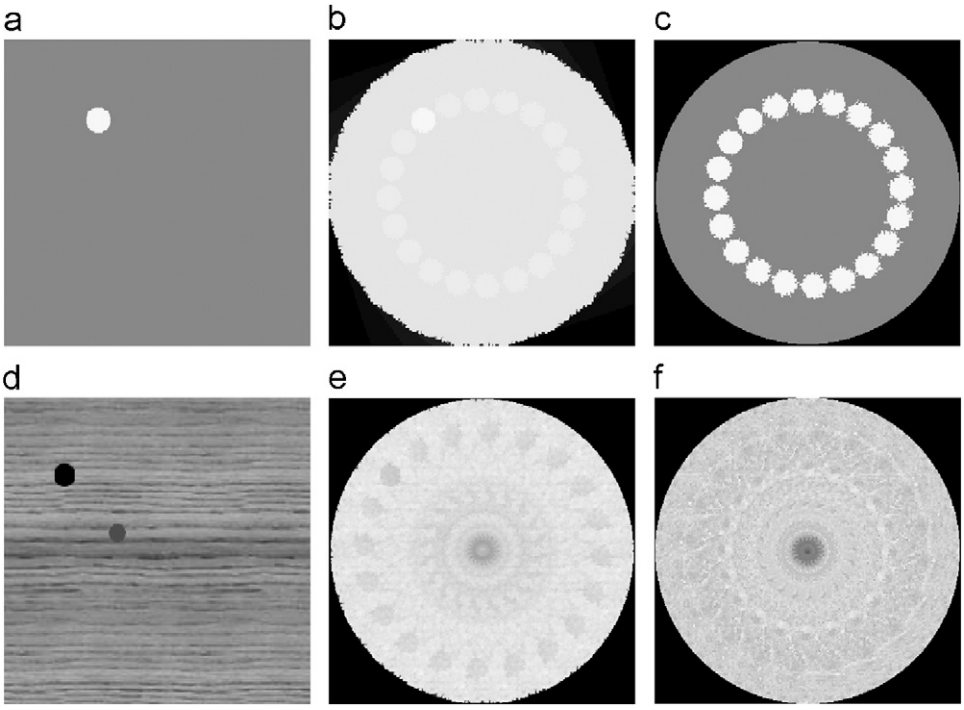

Fig. 2. Effect of using different collating functions on sample moment patterns. All patterns are obtained with induced rotation. (a)–(d) Original, (e) mean and (f) maxima.

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020

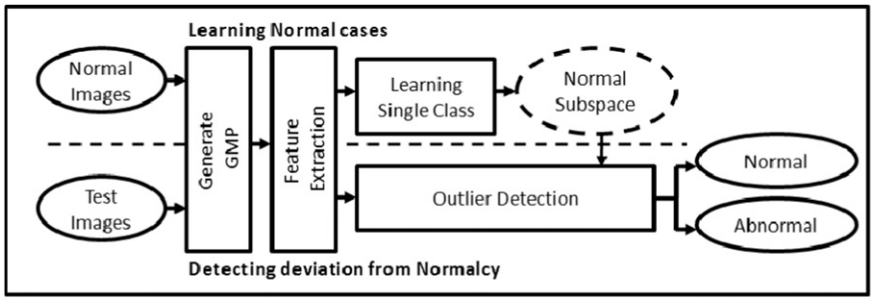

Fig. 3. Workflow for detection of abnormal cases as deviation from normal.

Alternately, other kinds of collation functions can also be considered, such as extrema, L-norm and characteristic function. When the collation function is an average, the moment map for the moving image is (from Eq. (4))

$$M(\bar{r}) = \sum_{t_i} I_{t_i}(\bar{r}) \tag{5}$$

where $\bar{r}$ is a position vector.

Generally, the derived moment pattern $M(\bar{r})$ will be dependent on the choice of velocity vector and the collation function. We call the moment maps generated by inducing different kinds of motion on a given static image, as generalized moment patterns (GMP).

Given a region of interest, different GMPs can be generated based on the choice of the collation function. Fig. 2 shows GMPs with induced rotation for two different collation functions, namely, the mean and maximum for images with uniform and textured backgrounds. It can be noted that the generated GMPs are consistent with the effect of motion that was observed earlier. Thus, it follows that an image region capturing a particular context, for example, a particular region of interest (ROI) from healthy tissue samples, will have a characteristic GMP. Following an introduction of a high contrast structure, there will be a disturbance in this GMP map. Based on this observation, the problem of learning a normal ROI can be approached as follows. A set of feature descriptors derived on GMPs generated using the normal ROIs can be used to represent the normal subspace. The normal subspace can be learnt using a single class classifier. An image belonging to the test set will be considered normal if it falls within the normal subspace. Any image deviating from this normal behavior is deemed abnormal. Fig. 3 shows this general workflow. Next, we look at the feature descriptors that can be used to describe the GMP response for images.

4. Learning the normal subspace

4.1. Features used for learning normal subspace

We hypothesize that any contextual disturbance in a given ROI can be encoded as shape/texture deviation in the moment pattern of the ROI. Feature descriptors can be. We test this hypothesis using different descriptors for capturing texture or shape. The local binary patterns (LBP) [15] and radon transform (RT) [16] are well-established descriptors of texture and shape information and hence, we consider some variants of these descriptors for our purpose. Thus, three feature descriptors were considered for learning the normal subspace: (i) a descriptor based on the radon transform (RT) constructed by vectorizing the RT of the ROI by column ordering. The resulting descriptor is referred to as SRT and it provides a compact shape representation; (ii) a descriptor derived from the texture of projections which was shown in [17] to be effective for representation of shape information in gray scale images. In that work, the image matching problem was considered. A descriptor consisting of a combination of LBP computed on the intensity *and* the projection spaces was shown to improve the performance over LBP computed in the intensity space. Given the GMP representation for an ROI, we wished to investigate if the texture of projection of the GMP is sufficient to capture shape information. Hence, the desired TOP descriptor was constructed as follows: the RT of the GMP is computed first, following which the LBP of this RT is computed; (iii) the last descriptor considered was the standard LBP computed over the intensity space.

4.2. K-nearest neighbor-based classifier for abnormality detection

The features derived over the GMP are used for learning the subspace corresponding to normal cases. A single-class classifier, *K*-nearest neighbor data description (KNN DD) that uses the natural distribution of data is used for this purpose. The ROIs corresponding to normal cases are divided into training and test set. Since the abnormal cases are not used for training the KNN DD classifier, the corresponding ROIs are considered to be part of the test set. In the validation stage, the number of nearest neighbors (K) within a specified distance (d) from a test ROI is found in the normal subspace. This test ROI is deemed to be normal if the number of nearest neighbors is greater than a specified *threshold* (see Eq. (6)). The distance to nearest neighbors is determined using *Euclidean* distance

$$Class(ROI) = \begin{cases} normal & \text{if } K \geq \text{threshold} \\ abnormal & \text{otherwise} \end{cases} \tag{6}$$

The KNN DD is known to be one of the simplest one-class classifiers and requires the features to be distinctive enough for classification. Unlike linear classifiers, KNN DD can perform well when the feature subspace corresponding to normal cases form several localized clusters. It is also known to perform well in the high-dimensional spaces.

5. Implementation details

Given an ROI a GMP for our experiments were generated by inducing motion to the given ROI. As explained earlier, this results in a vector of samples at each pixel location which are combined as per the chosen collation function. The vector length depends on the rate and type of motion. Both rotation and translation motions were considered. Rotations of $\theta_k = k \cdot \theta_0$; where $k = \{0, 1, 2, \ldots, K\}$ were used with different choices of θ_0, ranging from 5° to 135°. Translation of $x_l = l \cdot x_0$; where $l = \{0, 1, 2, \ldots, L\}$ was induced by applying shift in a particular direction ϕ with $x_0 = \{1, 2, \ldots, M\}$. Once again, a range of values for x_0 starting from 10% to 40% of ROI width were considered for different sampling rates at the receptor. Collation functions that were considered were *mean* and *maximum*. A challenge in generating a moment map is to evaluate the best

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020

motion parameter values for a given classification problem. Therefore, the motion parameter values that achieve highest discriminability between a small set of sample normal and abnormal cases are found first and then used for generating this moment map.

6. Experimental results

Two case studies were taken up for demonstrating the utility of the proposed method. The first study involves a classification problem, namely classification of lesions in mammograms as benign or malignant. This is a challenging task as both types of lesions share many similar characteristics (see top row of Fig. 4) and the research activity in this area spans two decades. The ROI is the output of a candidate selection module in a mammogram analysis system [18]. The traditional approach to classification of a candidate ROI has been to detect and segment lesions within the ROI and then classify them based on derived shape and texture features. In our approach, the aim will be to learn the normal (benign) subspace and use it to verify the malignancy or benign nature of the lesions of a given ROI.

The second case study is an example of abnormality detection, namely macular edema detection. Sample color fundus images are shown in the bottom row of Fig. 4. The ROI is a circular region of fixed diameter from the center of macula, a substructure of the retina. Macular edema is the swelling of this region and the presence of diabetic retinopathy symptoms within this ROI is known to be one of its indicators. This is traditionally detected by identifying (locating) the presence of subtler lesions namely hard exudates, which appear as tiny bright lesions and red irregular structures which are hemorrhages [19]. In our approach, the subspace of normal ROIs will be learnt and this should enable detection of the presence of lesions, regardless of their locations within the macula.

In the mammogram study, the decision is on whether a detected lesion is malignant or not. If several lesions are present in a given mammogram then each of them has to classified before deciding if the given mammogram is of a normal or a disease-free breast. In contrast, in the edema case, the decision of normal versus abnormal macula directly leads to a decision of whether a given image is of an eye with or without macula edema. Thus, the two case studies not only provide two different types of challenges, they also illustrate the applicability of the proposed method in different scenarios.

The processing pipeline (see Fig. 3) in both case studies is as follows: (i) given an ROI a GMP is computed for the ROI and a feature descriptor is derived for the GMP and (ii) the normal subspace is learnt by training a naive *K*-nearest neighbor (KNN DD) based classifier on a set of normal ROI images. Receiver operating characteristics (ROC) plots are generated by varying the normalized *distance* (d) from 0 to 1 for evaluating the algorithm's performance.

Two publicly available datasets that are commonly used in benchmarking algorithms for the detection and classification tasks were selected: the mini-MIAS dataset for lesion classification [20] and the Messidor dataset for macular edema detection [21].

The organization of the rest of this section is as follows. We begin with a description for two publicly available datasets. Next, we describe the procedure for determining the optimal motion parameters that maximize the discriminability (based on Shannon's entropy) between the normal and abnormal cases. Finally, we report the results of validation (with the optimal motion parameters) on public datasets: ROC curves, sensitivity, specificity and AUC are presented first for different features descriptors followed by an analysis of performance using scatter plots. Since the type of classifier is fixed (KNN D) for all validation, the scatter plots help to analyze the influence of a descriptor on classification performance.

6.1. Classification of lesions in mammograms

6.1.1. Dataset

mini-MIAS: This dataset comprises of 322 digital mammograms provided by Mammographic Image Analysis Society (MIAS) and annotated by expert radiologists. The annotations provide information on individual locations, extent of lesions and their tissue characteristics. Additionally, information on the benign (normal) or malignant nature of the lesions is also available.

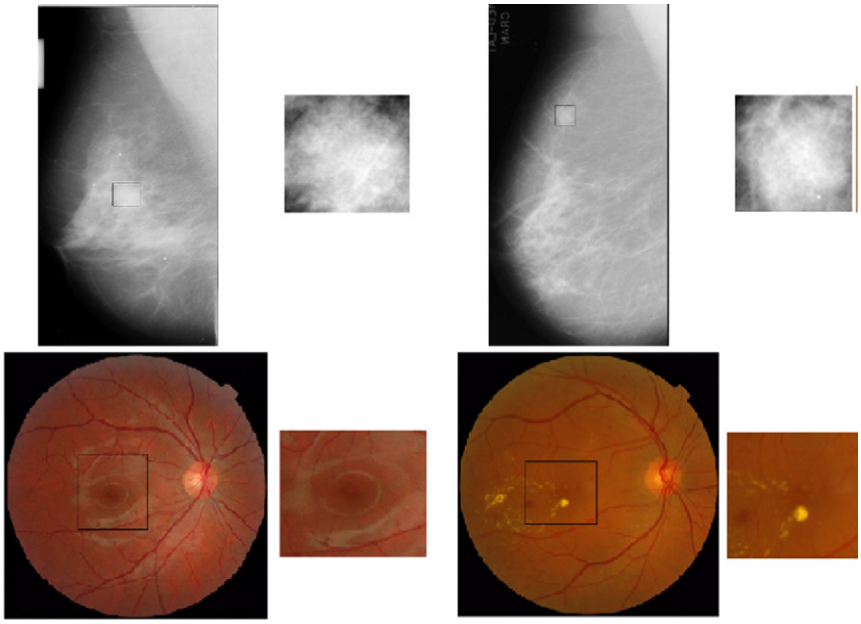

Fig. 4. Sample images and the regions of interest (enclosed in a black rectangle). First row (left to right): mammograms with benign and malignant lesions respectively. Second row (left to right): sample normal and abnormal color fundus images. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020

The number of benign regions are 68 and malignant are 51. Sample benign and malignant ROI's can be seen in Fig. 5. It should be noted that apart from high intra class variability there are very subtle intensity variations in ROIs across classes as well. 20 (of 68) benign cases were randomly selected and used for learning the normal subspace. The remaining ROIs were used in testing if a given ROI is benign or malignant.

6.1.2. Determining motion patterns

In order to determine the best motion parameters for the task of lesion classification, three representative lesions from each class (benign and malignant) were selected. These lesions were resized to 50×50 pixels. Two kinds of motion, namely, rotation and translation were induced to generate the respective GMPs, GMP_r and GMP_t by varying the motion parameters. For GMP_r, the value of θ_0 was varied from 5° to 135° and the value of $\theta_k = 360^\circ$ was selected. Consequently, the resultant sampling at the receptor occurs for one complete cycle of rotation. For generating GMP_t, the value of x_0 was varied from 10% to 40% of ROI width and the direction of motion was set to $\phi = 0^\circ$. The value of x_l was selected to be the width of the ROI.

As discussed earlier, inducing motion over the ROI results in some contextual disturbance. As the extent of motion is increased over the ROI, so does the extent of disturbance in the surroundings. This local contextual disturbance, observed in the ROI can be characterized with the Shannon's entropy. Starting with a GMP M, its gradient magnitude M_g is computed. For every location X in M_g, the local entropy (in a neighborhood $\Omega(X)$ chosen as 9×9 in our work) is computed to derive an entropy map M_e. A total entropy M is then derived by summing all these values of M_e. The difference δ_{GMP}, between this total value for the benign and malignant case, is used as a metric to determine the separability between the two classes.

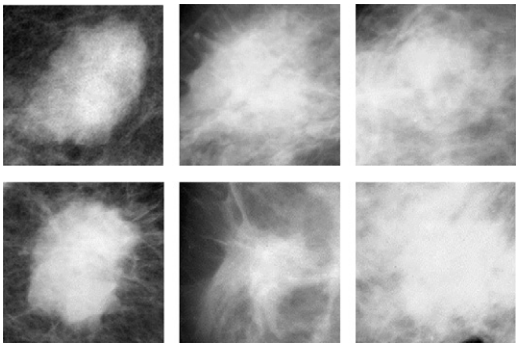

Fig. 5. Sample benign and malignant lesion examples. Top row: benign lesions; bottom row: malignant lesions.

The notation for this metric is as follows: the superscripts θ and Δ refer to the type of motion namely, rotation and translation, respectively. The subscript max and mean indicate the type of collation function used in generating the GMPs. The entropy plots are shown in Fig. 6. For the translation motion, the experiment was conducted by fixing the direction of motion ($\phi = 0^\circ$). From these plots it can be observed that the discriminability between the normal and abnormal is maximum around $\theta_0 = 60^\circ$ in the case of rotation and around $x_0 = 20\%$ of ROI width in the case of translation. Furthermore, the achievable discriminability is better with translation than with rotation. Next, we studied the influence of direction of translation on the class separability by varying the direction of translation while keeping the translation rate fixed at $x_0 = 20\%$. Fig. 7 shows the effect of changing the direction of translation. It should be noted that the magnitude of variation (or amplitude) in this plot is actually insignificant (only 15%). Hence, we conclude that for translation, the choice of direction is not important.

6.1.3. Results

Empirically, a combination of two translations was found to be effective for GMP generation. Hence, two GMPs were generated with linear motions in two different directions (l=15; $x_0 = 15\%$; $\phi = 0^\circ$ and 45°) using the collation function of maximum in each case. The final GMP was taken as an average of these two GMPs. This set of parameters is then used to generate GMPs from 20 benign samples to train a KNN DD classifier. The value of K was

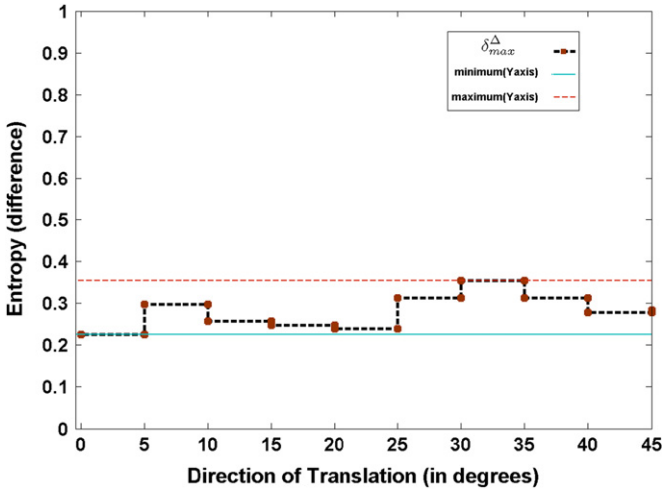

Fig. 7. Effect of change in direction of translation motion on class separability of lesions.

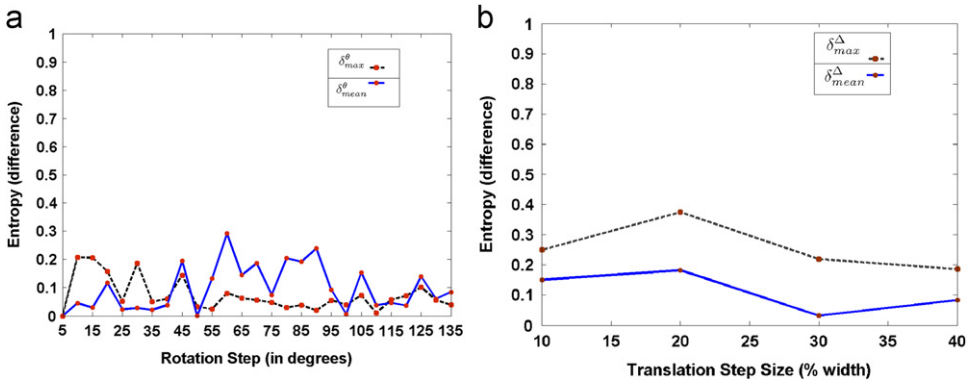

Fig. 6. Effect of increasing motion step on the class separability of lesions. (a) δ^{θ}: class separability and (b) δ^{Δ}: class separability.

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020

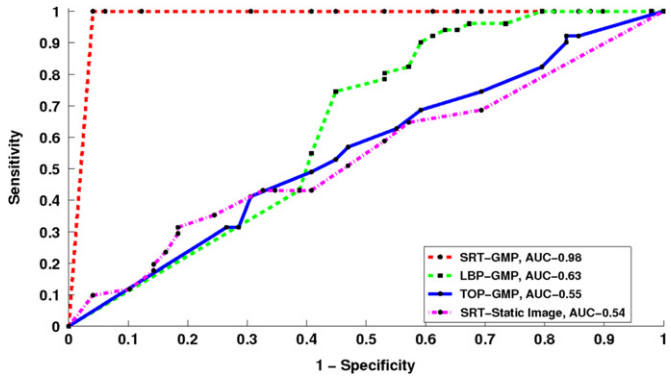

Fig. 8. ROC curves: classification performance for lesion classification.

Table 1
Performance comparison of lesion classification on MIAS database—AUC.

Method	Sensitivity	Specificity	AUC
Mudigonda et al. [22]	~0.85	~0.7	**0.79**
Buciu and Gacsadi [23]	0.84	0.8	**0.78**
Martins et al. [24]	0.9	0.71	NA
Proposed method	1	0.97	**0.98**

Table 2
Performance comparison of lesion classification on MIAS database—accuracy.

Method	Mass			Microcalcification		
	Benign (%)	Malignant (%)	Average (%)	Benign (%)	Malignant (%)	Average (%)
Mousa et al. [25]	100	83.3	**93.7**	75	100	**87.5**
Proposed method	100	100	**100**	97	100	**98.9**

considered as 3 for computing the sensitivity and specificity values in the classification tasks.

Receiver operating characteristics: The ROC plots for lesion classification is shown in Fig. 8 for a set of chosen feature descriptors: SRT, LBP, TOP, SRT-Static. The last descriptor refers to the case where SRT was computed over a static image: no motion is induced and hence the GMP is the original image. It can be observed that SRT-GMP performs best and that too significantly better than the rest. The computed AUC is 0.98 for the best performing descriptor GMP-SRT.

Lesion classification (as benign or malignant) in mammograms is a well-researched problem. We present a comparison of the proposed scheme in Table 1 for abnormality detection against the methods reported in the literature. However, it should be noted that while our figures are for 68 benign and 51 malignant lesions the existing methods report on different (smaller) number of ROIs. Overall, we can observe that the proposed method achieves high specificity (0.97) for the sensitivity of 1.0 on the test set. There is a significant increase compared to other methods, in both the sensitivity and the specificity. Table 2 shows the results in terms of *percentage* accuracy of classification for the masses and calcification. Ref. [25] reports an average accuracy of 93.7% for masses and 87.5% for microcalcification, whereas our method has 100% and 98.9% accuracies, respectively.

Scatter plots: The ROC plots indicate that SRT-GMP outperforms the rest. In order to understand the reason behind this behavior, the feature scatter plots for the descriptors were analyzed in a lower dimensional (2d) space.

In order to generate a scatter plot, the feature vectors (for each of the descriptors) were computed and then a PCA was performed on the vectors. The distribution of features along their two principal components are shown in Fig. 9 for the mini-MIAS dataset. The axes in the scatter plots are the normalized representation of descriptors computed corresponding to each ROI for first two directions in principal component space. In these figures, blue and red data points indicate elements of the normal (benign) and abnormal (malignant) subclasses, respectively. From the feature scatter plots, it can be observed that among the three feature descriptors, the GMP-SRT is the most effective for lesion

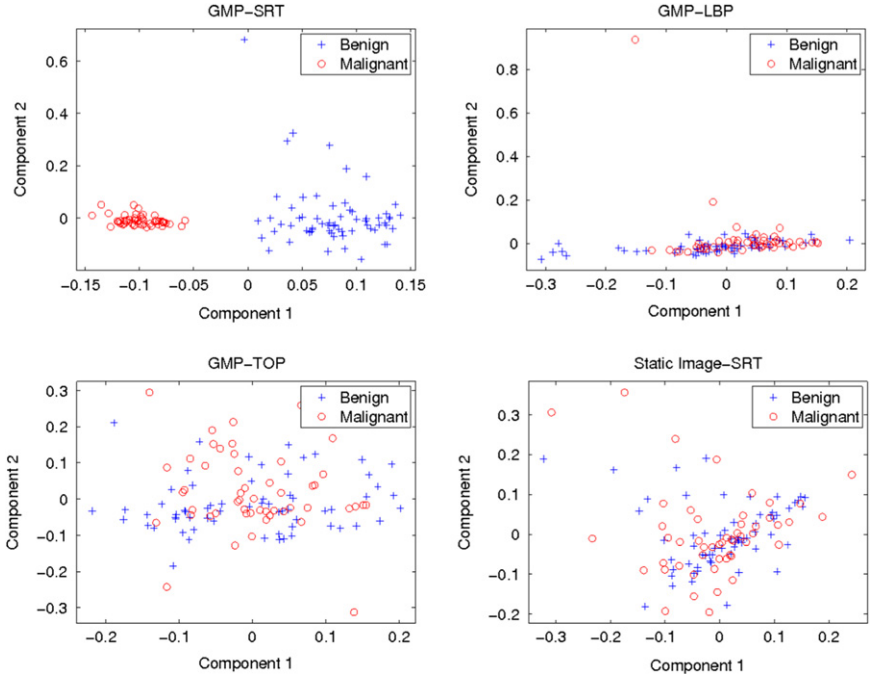

Fig. 9. Feature scatter plot for lesion classification in mammograms. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020

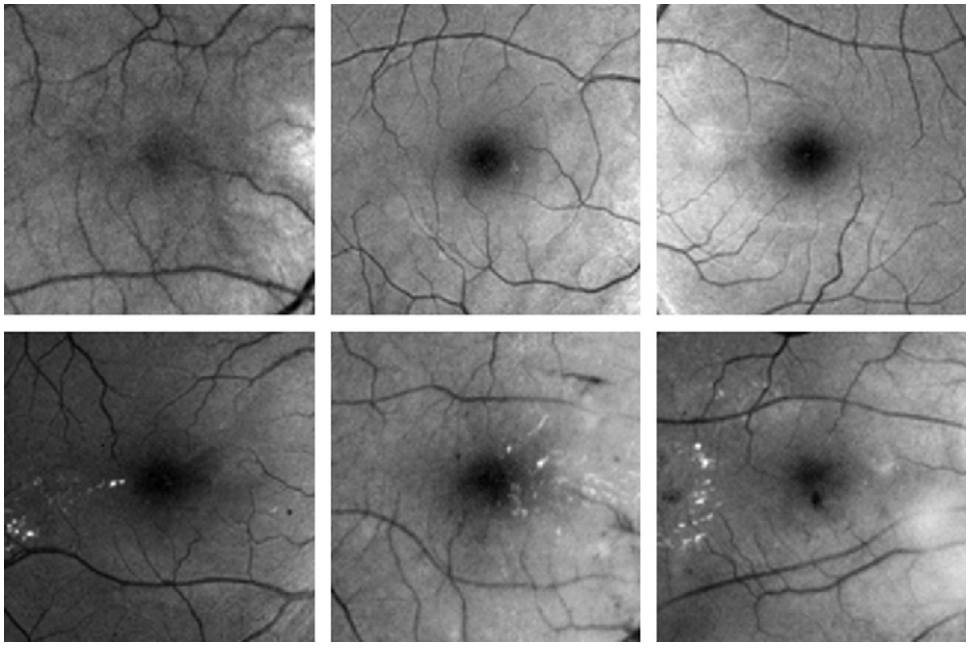

Fig. 10. Sample normal and abnormal near macular regions in color fundus images (green channel). Top row: normal ROI; bottom row: abnormal ROI. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

classification as it provides greater class separation. The poor performance of TOP descriptor is due to its lack of invariance to shift operations in images, and in the case of LBP, its inability to capture the subtle difference of texture (in the GMP images) between benign and malignant classes. The last plot attests to the importance of induced motion as its removal adversely affects the class separability.

6.2. Detection of macular edema in color fundus images

6.2.1. Dataset

MESSIDOR: The dataset comprises of color fundus images with associated annotation (at the image level) indicating the severity of macular edema. Disease severity is measured in terms of the extent of exudates near the macula region (see Fig. 4). About 522 images were used for detection of these disease symptoms. They exhibited wide variations in illumination, pigmentation and hence pose a good challenge. These images were divided into a set of 122 normal images picked randomly for training and 400 images for testing the presence of macular edema. Out of 400 test images, only 57 images show any signs of disease while a large proportion of the dataset (343 images) is normal which is typically the case in screening scenario.

Circular ROIs of 2 optic disk diameter were selected from the center of macula for all images. Only the green channel was used for processing as it provides maximum contrast. There is high variability in the structure as well as illumination of these ROIs within and across the two classes, namely, normal and abnormal (see Fig. 10). The task for the classifier was to detect whether a given ROI is normal or not.

6.2.2. Determining motion patterns

As in the case of lesion classification in mammograms, the optimal motion parameters were found using three normal and three abnormal ROIs for retina. The discriminability between the normal and abnormal images, measured using entropy difference in the gradient of the GMP, was found to be maximum for rotational GMP ($\theta_0 = 10^\circ$) with maximum as collating function.

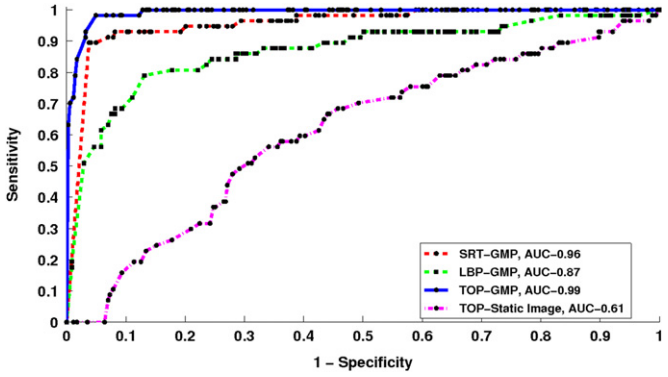

Fig. 11. ROC curves: classification performance for detection of macular edema.

Table 3
Performance comparison for Edema detection on MESSIDOR database.

Method	No. of images	Sensitivity	Specificity	AUC
Barriga et al. [26]	400	1	–	**0.98**
Proposed method	522	1	0.87	**0.99**

6.2.3. Results

Receiver operating characteristics: The GMP for a given retinal image was derived using induced rotation with $\theta_0 = 10^\circ$, $\theta_k = 360^\circ$ and *maximum* as the collation function. The ROC plot for macular edema detection is shown in Fig. 11. The value of K (KNN DD classifier) was taken to be 5 while computing the sensitivity and specificity. From these plots it can be observed that (a) the detection performance is best with features computed with induced motion as the curves for GMP-X where X denotes different descriptors, lie above the curve for the static case; (b) the ROC for GMP-TOP outperforms all other descriptors. The AUC is nearly 1 for the best performance and close to 0.5 for the static-TOP.

Table 3 shows the performance comparison for macular edema detection on the Messidor dataset. Ref. [26] report a high

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020

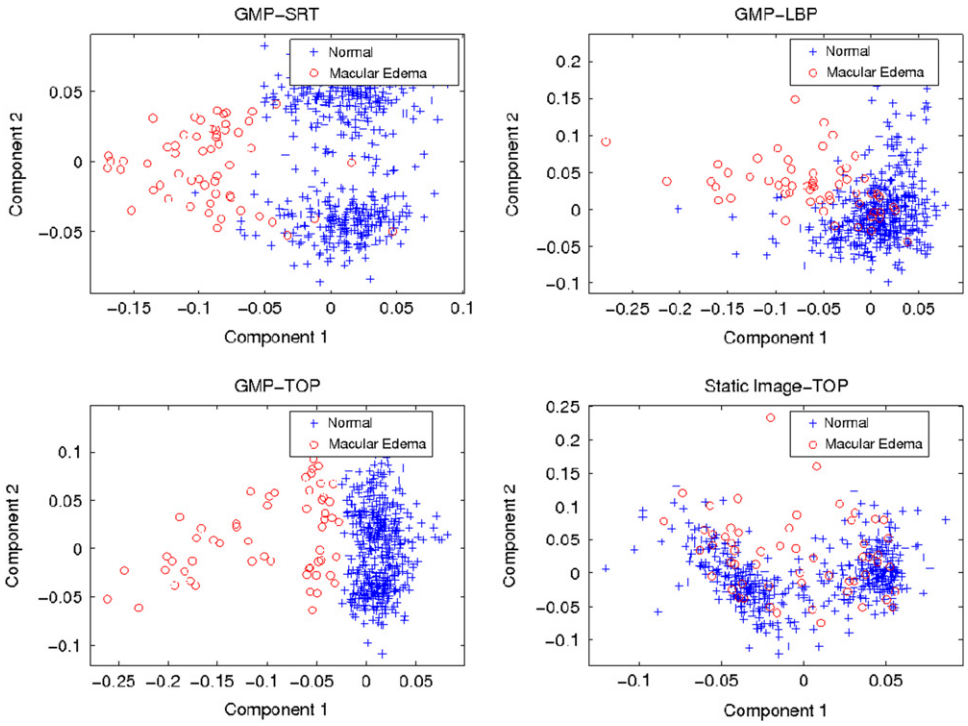

Fig. 12. Feature scatter plot for macular edema.

performance on abnormality detection for only the severe (and not the moderate) cases of clinically significant macular edema. Our figures are reported for both moderate and severe cases. Based on the reported figures for our method, we conclude that the performance has actually improved since the larger dataset we considered had a higher number of normals despite which the AUC marginally improved. It was also noted that when least number of ROIs are misclassified, a specificity of 0.95 was achieved for the sensitivity of 0.98.

Scatter plots: As in the previous case study, the feature subspace for different descriptors were analyzed in lower 2-dimensional feature space using scatter plots.

The scatter plots in Fig. 12 for the Messidor dataset shows that GMP-TOP (with induced rotation) is able to generate a well-separated, compact normal subspace in the case of macular edema. LBP fails to capture the normal subspace as the granularity of the resultant texture is non-uniform. SRT performs relatively better than LBP due to the smearing effect of abnormalities in the radial direction in rotational motion. The last scatter plot in this figure shows the effect of removing the induced motion for the TOP descriptor: the two classes are no longer well separated. Hence, it can be concluded that GMP plays a vital role in capturing the difference between normal and abnormal.

From the assessment results of the two case studies, we conclude that GMP is indeed effective for generating a compact, well-separated normal subspace.

7. Discussion and conclusion

The key idea that has been proposed in this paper uses the spatio-temporal collation of information that is found (at the receptor level) in low-level vision, as an inspiration to generate a representation for a normal subspace. The receptor response is simulated with the GMP pattern generated by inducing motion to a given image. An abnormality is always with reference to a context which is deemed to be normal. Hence, collation of intensities from induced motion ensures that the required context (normal) is captured and the disease related deviations are enhanced. The choice of the type of induced motion and the collation function depends solely on the properties of the selected ROI and the characteristics of disease deviations observed.

The case of collation of information from rotational motion over the ROI can also be explained in terms of integration of intensities over a *rotation group* of the ROI. Deriving image features by normalization or group integration has been attempted [27] for representing rotationally invariant features. In this work, however, the rotational subgroup generated by inducing rotations is used to enhance the deviation between the normal and abnormal ROI.

We have established that features computed over GMP are suited for learning the normal subspace by experimentally validating the performance of the method for two popular screening problems, classification of lesions in breast cancer and detection of macular edema (a sub-problem of diabetic retinopathy) in retina. For breast cancer, an ROI is considered, whose malignancy is of interest whereas in macular edema, the disease manifests as multiple, subtle/high contrast lesions within the ROI. Relevant GMP after evaluating the best motion parameter values are derived for both the problems separately. The experiments on two publicly available representative datasets and its results are compared with the earlier methods in the literature. High performance of even a naive single class classifier (KNN DD) with the proposed moment patterns demonstrates its effectiveness. A key strength of the proposed method is that it does not require any (a) additional preprocessing for enhancing image quality or removal of background structures, like blood vessels near macula or (b) segmentation of the individual lesions for disease classification. The current method is promising in terms of its applicability to CAD systems as it is able to achieve high sensitivity and specificity.

There is scope for future work in different avenues. The theoretical analysis of different collation functions and their effect on the contextual information captured by GMP is one avenue. In addition to functions like mean and maxima considered in this

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020

work, the effect of other more complex functions, for example, higher order moments and other non-linear functions can be explored. Further, we can seek to find a systematic way in which the best motion parameters can be determined for a given problem, via a simulation of normal and abnormal tissues.

References

[1] H.D. Cheng, X. Cai, X. Chen, L. Hu, X. Lou, Computer-aided detection and classification of microcalcifications in mammograms: a survey, Pattern Recognition 36 (12) (2003) 2967–2991.

[2] B. Zhang, X. Wu, J. You, Q. Li, F. Karray, Detection of microaneurysms using multi-scale correlation coefficients, Pattern Recognition 43 (6) (2010) 2237–2248.

[3] M. Siyal, L. Yu, An intelligent modified fuzzy c-means based algorithm for bias estimation and segmentation of brain MRI, Pattern Recognition Letters 26 (13) (2005) 2052–2062.

[4] A.R. Domínguez, A.K. Nandi, Toward breast cancer diagnosis based on automated segmentation of masses in mammograms, Pattern Recognition 42 (6) (2009) 1138–1148.

[5] D.T. Gering, W.E.L. Grimson, R. Kikinis, Recognizing deviations from normalcy for brain tumor segmentation, in: Proceedings of the Fifth International Conference on Medical Image Computing and Computer-Assisted Intervention-Part I, MICCAI '02, Springer-Verlag, 2002, pp. 388–395.

[6] H. Nosato, H. Sakanashi, M. Murakawa, T. Higuchi, N. Otsu, K. Terai, N. Hiruta, N. Kameda, Histopathological diagnostic support technology using higher-order local autocorrelation features, in: ECSIS Symposium on Bio-inspired, Learning, and Intelligent Systems for Security, 2009, pp. 61–65.

[7] G. Erus, E. Zacharaki, N. Bryan, C. Davatzikos, Learning high-dimensional image statistics for abnormality detection on medical images, Computer Vision and Pattern Recognition Workshops (2010) 139–145.

[8] S. Philip, A. Fleming, K. Goatman, S. Fonseca, P. McNamee, G. Scotland, G. Prescott, PF. Sharp, J. Olson, The efficacy of automated disease/no disease grading for diabetic retinopathy in a systematic screening programme, British Journal of Ophthalmology 91 (2007) 1512–1517.

[9] Z. Zivkovic, F. van der Heijden, Efficient adaptive density estimation per image pixel for the task of background subtraction, Pattern Recognition Letters 27 (7) (2006) 773–780.

[10] X. Gao, T. Boult, F. Coetzee, V. Ramesh, Error analysis of background adaption, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, vol. 1, 2000, pp. 503–510.

[11] A. Rajagopalan, R. Chellappa, N. Koterba, Background learning for robust face recognition with pca in the presence of clutter, IEEE Transactions on Image Processing 14 (6) (2005) 832–843.

[12] C. Stauffer, W. Grimson, Learning patterns of activity using real-time tracking, IEEE Transactions on Pattern Analysis and Machine Intelligence 22 (8) (2000) 747–757.

[13] C.R. Wren, A. Azarbayejani, T. Darrell, A.P. Pentland, Pfinder: real-time tracking of the human body, IEEE Transactions on Pattern Analysis and Machine Intelligence 19 (1997) 780–785.

[14] D.E. Freund, N. Bressler, P. Burlina, Automated detection of drusen in the macula, in: Proceedings of the Sixth IEEE International Conference on Symposium on Biomedical Imaging, 2009, pp. 61–64.

[15] M. Heikkilä, M. Pietikäinen, C. Schmid, Description of interest regions with local binary patterns, Pattern Recognition 42 (3) (2009) 425–436.

[16] L. Zheng, D. Shi, Advanced radon transform using generalized interpolated fourier method for straight line detection, Computer Vision and Image Understanding 115 (2) (2011) 152–160.

[17] N.V.K. Medathati, J. Sivaswamy, Local descriptor based on texture of projections, in: Proceedings of the Seventh Indian Conference on Computer Vision, Graphics and Image Processing, ICVGIP '10, 2010, pp. 398–404.

[18] H. Cheng, X. Shi, R. Min, L. Hu, X. Cai, H. Du, Approaches for automated detection and classification of masses in mammograms, Pattern Recognition 39 (4) (2006) 646–668.

[19] S. Ravishankar, A. Jain, A. Mittal, Automated feature extraction for early detection of diabetic retinopathy in fundus images, in: IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2009, pp. 210–217.

[20] J. Suckling, J. Parker, D.R. Dance, et al., The mammographic image analysis society digital mammogram database, in: Proceedings of the 2nd International Workshop on Digital Mammography, vol. 1069, 1994, pp. 375–378.

[21] MESSIDOR, URL 〈http://messidor.crihan.fr/index-en.php〉, 2011.

[22] N. Mudigonda, R. Rangayyan, J. Leo Desautels, Detection of breast masses in mammograms by density slicing and texture flow-field analysis, IEEE Transactions on Medical Imaging 20 (12) (2001) 1215–1227.

[23] I. Buciu, A. Gacsadi, Directional features for automatic tumor classification of mammogram images, Biomedical Signal Processing and Control 6 (4) (2011) 370–378.

[24] L.d.O. Martins, A.M.d. Santos, A.C. Silva, A.C. Paiva, Classification of normal, benign and malignant tissues using co-occurrence matrix and Bayesian neural network in mammographic images, in: Ninth Brazilian Symposium on Neural Networks, 2006, pp. 24–29.

[25] R. Mousa, Q. Munib, A. Moussa, Breast cancer diagnosis system based on wavelet analysis and fuzzy-neural, Expert Systems with Applications 28 (4) (2005) 713–723.

[26] E. Barriga, V. Murray, C. Agurto, M. Pattichis, W. Bauman, G. Zamora, P. Soliz, Automatic system for diabetic retinopathy screening based on am-fm, partial least squares, and support vector machines, in: IEEE International Symposium on Biomedical Imaging, 2010, pp. 1349–1352.

[27] M. Reisert, H. Burkhardt, Invariant features for 3d-data based on group integration using directional information and spherical harmonic expansion, in: 18th International Conference on Pattern Recognition, vol. 4, 2006, pp. 206–209.

K. Sai Deepak is a graduate student at the center of visual information technology at IIIT Hyderabad. He is currently working on disease detection, namely diabetic retinopathy from retinal images. Prior to joining IIIT, Hyderabad, he has worked with Software Engineering and Technology Labs (Infosys) on medical image analysis for 3 years. His broad interests are in the areas of medical image analysis and pattern recognition.

N.V. Kartheek Medathati is a graduate student at the center of visual information technology at IIIT Hyderabad. He also finished his BTech in Electronics and Communication from the same university. He is currently working on developing biologically inspired algorithms for saliency detection in images. His interests are in developing image analysis systems based on the principles of computational visual attention.

Jayanthi Sivaswamy did her MS and PhD in Electrical Engineering from Syracuse University. She is currently a professor at IIIT Hyderabad. Prior to joining IIIT-H in 2001, she was with the University of Auckland, New Zealand, for 8 years. Her broad research interests are biological and computer vision. Bio-inspired algorithm development and CAD tool development based on automated medical image analysis form the thrust of her current work.

Please cite this article as: K.S. Deepak, et al., Detection and discrimination of disease-related abnormalities based on learning normal cases, Pattern Recognition (2012), http://dx.doi.org/10.1016/j.patcog.2012.03.020