

# Classification and Localisation of Diabetic-Related Eye Disease

Alireza Osareh<sup>1</sup>, Majid Mirmehdi<sup>1</sup>, Barry Thomas<sup>1</sup>, and Richard Markham<sup>2</sup>

<sup>1</sup>Department of Computer Science, University of Bristol, Bristol, BS8 1UB, U.K.  
{osareh, majid, barry}@cs.bris.ac.uk

<sup>2</sup>Bristol Eye Hospital, Bristol, BS1 2LX, U.K.  
markham@gifford.co.uk

**Abstract.** Retinal exudates are a characteristic feature of many retinal diseases such as Diabetic Retinopathy. We address the development of a method to quantitatively diagnose these random yellow patches in colour retinal images automatically. After a colour normalisation and contrast enhancement pre-processing step, the colour retinal image is segmented using Fuzzy C-Means clustering. We then classify the segmented regions into two disjoint classes, exudates and non-exudates, comparing the performance of various classifiers. We also locate the optic disk both to remove it as a candidate region and to measure its boundaries accurately since it is a significant landmark feature for ophthalmologists. Three different approaches are reported for optic disk localisation based on template matching, least squares arc estimation and snakes. The system could achieve an overall diagnostic accuracy of 90.1% for identification of the exudate pathologies and 90.7% for optic disk localisation.

## 1 Introduction

Diabetic-related eye diseases affect the blood vessels in the retina and are the most common cause of blindness. In non-proliferative retinopathy damaged retinal vessels leak fatty and protein-based particles referred to as exudates (EXs). These appear as spatially random yellow pathologies with varying sizes and shapes as shown in Figure 1(c). When the EXs accumulate in the central part of the retina, the visual acuity is at risk [1]. Therefore, it is necessary to have a mass-screening retinal examination system for the early detection of diabetic-related retinal diseases and to assist the ophthalmologist to make the diagnosis faster and more efficiently. The aim of our work is to develop a decision support system that could be installed at GP clinics and high-street opticians alike that can act as an early warning system. In this paper, we report on our work towards this goal. We identify the EX regions of the colour retinal image using Fuzzy C-Means clustering and neural network classification following some key pre-processing steps. During an initial segmentation stage, the optic disk is also segmented as (one or more) candidate EX regions due to the similarity of its colour to the yellowish EXs. Far from being a disadvantage, the benefits of this are two-fold. Since optic disk localisation in retinal image analysis is of critical importance, the segmentation process provides us with an opportunity to identify the optic disk using a separate processing step. In turn, this allows us to remove the optic disk from the candidate regions for both better training and testing of a classifier for

EX region classification. Our experimental results indicate that we are able to achieve 90.1% accuracy, 93.4% sensitivity and 82.7% specificity in identifying the EX pathologies and 90.7% accuracy in localising the optic disk boundary.

Few investigations in the past have identified retinal EXs based on colour information. In many of the cases reviewed below only sensitivity and specificity figures were supplied. Philips et al. [2] separated all bright objects (including EXs) from the dark objects and background after noise removal, contrast enhancement, and simple thresholding with a sensitivity of 87%. Gardner et al. [3] used an artificial neural network for identification of EXs by classifying regions of size 20x20 pixels in greylevel images where the authors reported a sensitivity of 93.1%. However, this was the result of classifying the whole 20x20 region rather than a pixel-level classification. Sinthanayothin [4] applied a recursive region growing technique based on a manually-selected threshold value in greylevel images. The author reported 88.5% sensitivity and 99.7% specificity, but this was measured on a 10x10 window. Wang et al. addressed the same problem in [5] by using a Bayesian statistical classifier, based on colour features, to classify each pixel into lesion or non-lesion classes. However, other yellowish lesions (e.g. cotton-wool spots) were classified incorrectly at the same time. They reported 100% sensitivity and 70% specificity, but this was measured purely on whether EXs were present anywhere in the image. They did not report their quantification of the EX regions on a pixel-by-pixel level accuracy. One novelty of the proposed method here is that we locate EXs at pixel resolution in colour images and evaluate the performance using manually labelled groundtruth produced by an expert.

There have also been relatively few works for the detection of the optic disk. Liu et al. [6] applied an edge detection stage followed by a circular Hough transform to locate the disk. This approach failed when there were large areas of EXs interference in the image or when the boundary between the optic disk region and the background was not well defined. Li and Chutatape [7] produced a training set using the brightest pixels that were firstly clustered as candidate optic disk regions. Then principle component analysis was applied to project a new image to the 'disk space' specified by the eigenvectors of the training images and calculating the distance between the image and its projection. Sinthanayothin et al. [8] used an 80x80 subimage to evaluate the intensity variance of adjacent pixels, marking the point with the largest variance as the optic disk location. This technique could fail when there is a large area of EXs in the retinal image. The methods introduced in both [7] and [8] produce an optic disk centre but do not address optic disk boundary identification. The optic disk boundary was found in Mendels et al. [9], using mathematical morphology and an active contour approach in greylevel images. We have drawn inspiration from [9] and propose a similar snake-based method. However, our proposed method is implemented on colour images pre-processed using colour-based mathematical morphology. Additionally, our snake is automatically initialised.

In this study, we used 60 colour retinal images obtained from a non-mydratric retinal camera with a 45° field of view. The image resolution was 760x570 at 24bit RGB. This paper is organised as follows. Section 2 describes our proposed automatic method for identification of the EX pathologies in colour retinal images. In Section 3, three different approaches for identification of the optic disk are presented. Results

are outlined as the algorithms are described, but in Section 4 we present further results to compare the proposed methods with [9]. The paper is concluded in Section 4.


## 2 Detecting Exudates

We put our data through two pre-processing steps before commencing the detection of EXs. The retina's colour in different patients is variable being strongly correlated to skin pigmentation and iris colour. The first step is therefore to normalise the retinal images. We selected a retinal image as a reference and then used histogram specification [10] to modify the values of each image in the database such that its frequency histogram matched the reference image distribution. Figure 1(c) shows the result of this normalisation. In the second pre-processing step, the contrast between the EXs and retina background is enhanced to facilitate later segmentation. The contrast is not sufficient due to the internal attributes of lesions and decreasing colour saturation, especially in the areas around the retina. We applied local contrast enhancement [8] to distribute the values of pixels around the local mean. In this way a pixel  $p$  in the centre of a small running window  $w$ , was changed to a new value  $p_n$ :

$$p_n = 255 * \left( \frac{[\phi_w(p) - \phi_w(\text{Min})]}{[\phi_w(\text{Max}) - \phi_w(\text{Min})]} \right) \quad (1)$$

$$\text{where } \phi_w(p) = \left[ 1 + \exp\left(\frac{\mu_w - p}{\sigma_w}\right) \right]^{-1} \quad (2)$$

$\text{Max}$  and  $\text{Min}$  are the maximum and minimum intensity values in the whole image, while  $\mu_w$  and  $\sigma_w$  indicate the local window mean and standard deviation. The exponential produces significant enhancement when the contrast is low ( $\sigma_w$  is small), while it provides little enhancement if the contrast is high ( $\sigma_w$  is large). Figure 1(d) shows the result of local contrast enhancement on a typical retinal image.



**Fig. 1.** Colour normalisation and local contrast enhancement: (a) reference image, (b) typical retinal image (including EXs), (c) colour normalised version, (d) after contrast enhancement.

### 2.1 Colour Image Segmentation

The ideal segmentation of an image is usually application-dependent. Unlike hard segmentation methods, which force pixels to belong exclusively to one class, colour Fuzzy C-Means (FCM) clustering allows pixels to belong to multiple classes with varying degrees of membership [11-12]. The segmentation approach is based on a coarse and a fine stage. The coarse stage is responsible for evaluating Gaussian-smoothed histograms of each colour band of the image in order to produce an initial classification into a number of classes (call it  $K$ ). Then, in the fine stage, FCM clustering is applied to find an optimum solution for the following objective function, with respect to  $U$  (a fuzzy  $K$ -partition of the data set) and  $V$  (a set of  $K$  prototypes):

$$J(U, V) = \sum_{j=1}^N \sum_{i=1}^K (\mu_{ij})^q \|X_j - V_i\| \quad ; \quad K \leq N \tag{3}$$

where  $q$  is the weighting exponent which controls the fuzziness of the resulting clusters,  $X_j$  is the  $j$ th feature vector,  $V_i$  is the centroid of the  $i$ th cluster,  $N$  is the number of data points, and  $\mu_{ij}$  is the degree of membership of  $X_j$  to cluster  $i$ . These memberships must lie between 0 and 1, and  $\mu_{ij}$  must sum to 1 for all  $i$ . Then the fuzzy partitioning is carried out through an iterative optimisation of (3) using the following algorithm:

- 1) Choose primary centroids (prototypes)
- 2) Compute the degree of membership of all feature vectors in all the clusters:


$$\mu_{ij} = \left( \sum_{k=1}^K \left( \frac{\|X_j - V_i\|^2}{\|X_j - V_k\|^2} \right)^{\frac{2}{q-1}} \right)^{-1} \tag{4}$$

- 3) Compute new centroids:  $V_i' = \left( \sum_{j=1}^N (\mu_{ij})^q X_j \right) \left( \sum_{j=1}^N (\mu_{ij})^q \right)^{-1}$  (5)

and update the membership values  $\mu_{ij}$  to  $\mu'_{ij}$  according to (4).

- 4) if  $\| \mu_{ij} - \mu'_{ij} \| < \xi$ , where  $\xi$  is a termination criterion between 0 and 1, stop, otherwise continue from step 3. We fixed  $q=2$  and  $\xi=0.5$  for all experiments and the algorithm was iterated until FCM could distinguish three different clusters. Figure 2(b) illustrates an example result after this colour segmentation stage. Figure 2(c) shows the candidate EX regions overlaid on the original image.

We performed K-means clustering on the image as an alternative segmentation technique but with limited success. We also experimented with the method described by Comaniciu and Meer [13], but for our images the results were inconsistent and unsatisfactory. The FCM approach was straightforward to implement, fast and, had fixed parameters, but most importantly it allowed us to segment colour images. To assess the accuracy of the proposed segmentation technique, an expert clinician marked the EX regions in all 60 of our retinal images. Accurate, manual, pixel by pixel registration of small pathologies like EXs is very difficult due to the wide




**Fig. 2.** Colour image segmentation: (a) a typical normalised image, (b) FCM segmented image, and (c) candidate EX regions overlaid on the original image.

variability in their colour. FCM could achieve 98% accuracy compared to the groundtruth for identification of EXs. Only extremely faint EXs were not identified. Also a simple and effective overlap measure of the match between two regions was used to investigate the accuracy of the FCM in producing an EX contour:

$$M = \frac{N(R \cap T)}{N(R \cup T)} \quad (6)$$

where  $R$  and  $T$  are two regions being compared and  $N(\cdot)$  is the number of pixels in the set. Using this measure FCM could achieve 96.4% accuracy compared to the groundtruth. Figure 3 shows typical regions used for applying the metric in (6), i.e. the groundtruth and the segmentation result.



**Fig. 3.** Segmentation performance evaluation: (a) overlaid candidate EXs, (b) a close-up of some EXs (c) groundtruth of EXs as determined by a clinician, (d) FCM result.

After the candidate EX regions are located, they must be classified. To perform this accurately, it is crucial to extract efficient features that produce the best class separability. It is worth noting that after FCM, false positive EX candidates arise due to both general reflections in the retina and the similar characteristics of the optic disk. The optic disk regions were removed prior to classification using the automatic optic disk localisation method described later in Section 3.

### 2.2 Feature Extraction

Clinically, ophthalmologists use colour to differentiate various pathological conditions. Similarly coloured objects like cotton-wool spots and EXs are differentiated with further features such as size, edge strength and texture. Indeed, we found features such as average intensity and standard deviation of intensity were enough to encode appearance in terms of brightness and texture. After a comparative study of our training dataset (more in section 2.4), to investigate the relative discriminative power of a small set of features for use in the classification process, ten were selected:  $R_1, G_1, B_1$  (representing mean RGB values inside a region),  $R_2, G_2, B_2$  (representing mean RGB values outside a region), *region size*, *mean* and *standard deviation* of intensity, and *edge strength*. The within-class matrix ( $S_w$ ) and between-class scatter matrix ( $S_b$ ) were computed. The value  $J = trace(S_b/S_w)$  was used as a measure of feature-set efficiency. This metric increases when between-class scatter is larger or within-class scatter is smaller. Therefore the feature subset, which yields the largest  $J$ , is considered to be the best subset to use for classification.

**Table 1.** Discriminative importance of the feature set employed.

Number of Features	Features in subset										Metric $J$
	$G_1$	$G_2$	$B_1$	$B_2$	Average Intensity	$R_1$	$R_2$	Standard Deviation	Edge Sharpness	size	
1	*	-	-	-	-	-	-	-	-	-	0.5675
2	*	*	-	-	-	-	-	-	-	-	1.1346
3	*	*	*	-	-	-	-	-	-	-	1.6752
4	*	*	*	*	-	-	-	-	-	-	2.2155
5	*	*	*	*	*	-	-	-	-	-	2.7404
6	*	*	*	*	*	*	-	-	-	-	3.2340
7	*	*	*	*	*	*	*	-	-	-	3.7187
8	*	*	*	*	*	*	*	*	-	-	4.1455
9	*	*	*	*	*	*	*	*	*	-	4.5161
10	*	*	*	*	*	*	*	*	*	*	4.7056

Table 1 represents the relative importance of the selected features, as the number of features allowed is increased. A ‘\*’ indicates that a feature was used and a ‘-’ indicates that it was not. As expected, features that provide colour information seem to contribute significantly more than the other features. The class separability can naturally improve by including additional features, but at the expense of extra features and classifier complexity.

### 2.3 Classification

In this study we investigated various classifiers including Linear Delta Rule (LDR), K-Nearest Neighbours (KNN), Quadratic Gaussian classifier (QG), and Neural Networks (NN) [14]. We opted for a NN due to its superior performance, however a comparison with the other methods will be presented later. FCM processing of our database of our 60 colour images resulted in 4037 segmented regions. These were

labelled by a consultant ophthalmologist as EX or non-EX to create a fully labelled groundtruth dataset. Then, they were divided into a training set, a validation set and a test set in a 64:10:26 ratio. Therefore, the training and validation set contained around 3000 candidates, of which 1205 were labelled as EX. The remaining 1037 (of which 417 were labelled as EXs) were used as an independent testing set. A 10-fold cross-validation was used for the network generalisation. The three-layer perceptron NN employed had a 10-node input layer corresponding to our feature vector. We experimented with a hidden layer with a range of 2 to 35 hidden units to find the optimum architecture. A single output node gave the final classification probability. The network was trained using two different learning methods, standard Back-Propagation (BP) and Scaled Conjugate Gradient (SCG) descent [15]. The classifier error was calculated using regions from validation set after each iteration of training. Training was terminated when this error began to rise as it signified over-training. The NN performance was measured using the previously unseen regions from the test set in the usual terms of detecting the presence or absence of a pathology in an image, i.e. *sensitivity* and *specificity*. Sensitivity is the ratio of the true positive (TP) decisions to all positive decisions, while specificity is the ratio of true negative (TN) decisions to all negative decisions. Another reported measure, the accuracy, is the ratio between the total numbers of correctly classified instances to all the instances that exist in the test set. Table 2 summarises the results obtained on the test set where the best classifiers, including 30 hidden units for SCG and 10 for BP learning method, have been chosen on the basis of validation error. These results are just a selection out of a number of configurations used for training the classifiers. For each network the optimum threshold value giving the best balance between sensitivity and specificity is shown.

**Table 2.** Performance comparison of different neural network based classifiers.

Classifier	Output Threshold	Overall Accuracy	Sensitivity	Specificity
<b>SCG (10)</b>	0.4	81.2%	70.3%	91.4%
<b>SCG (20)</b>	0.5	84.0%	84.0%	82.8%
<b>SCG (30)</b>	0.3	90.1%	93.4%	82.7%
<b>BP (10)</b>	0.3	89.6%	89.0%	89.8%
<b>BP (20)</b>	0.5	82.0%	77.5%	84.6%
<b>BP (30)</b>	0.5	87.0%	80.3%	91.7%

The classical tool to achieve tradeoffs between sensitivity and specificity criteria is the Receiver Operating Characteristic (ROC) curve [16]. This curve is typically plotted with the TP fraction against the FP fraction. The bigger the area under the ROC curve ( $A_c$ ) the higher the probability of making a correct decision. Figure 4 compares the behaviour of the NN classifiers for the full range of output threshold values. The best performance for the BP network was achieved using 10 hidden units with  $A_c = 0.897$ . Similarly, for the SCG network, 30 hidden units achieve the best performance with  $A_c = 0.916$ . These show that these performances are very close and comparable. These conclusions also match the results shown in Table 2.

Table 3 illustrates comparative classification performances for three other classifiers along with the best results from the NN. These were the LDR, KNN and QG classifiers. The combination of the selected features provides a good classification performance for all the classifiers with the KNN classifier performing the best out of the three new methods. Overall, the classification analysis indicates that the best suitable classifier for distinguishing the EX pathologies from the other non-EXs is a NN classifier based on the SCG learning method.

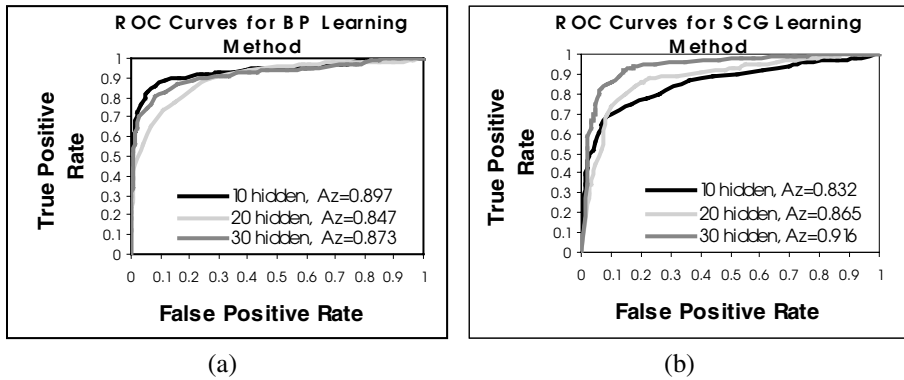


Fig. 4. NN classifier ROC curves analysis, (a) BP learning method, (b) SCG learning method.

Table 3. Comparison of the performance of different classifiers.

Classifier	Accuracy (%)	Sensitivity (%)	Specificity (%)
SCG (30)	90.1	93.4	82.7
BP (10)	89.6	89.0	89.8
LDR	81.8	65.3	93.7
KNN (K=4)	87.1	80.2	93.1
QG	81.3	64.4	93.0

### 3 Locating the Optic Disk


The optic disk is the entrance region of blood vessels and optic nerves to the retina. Localising it is crucial, for example, to identify anatomical and pathological parts in retinal images (e.g. the fovea), for blood vessel tracking, and as a reference length for measuring distances in retinal images. The optic disk is a brighter region than the rest of the choroids due to the absence of retina layer and its shape is approximately round. Another motivation for identification of the optic disk is to remove it as a (false positive) candidate EX region. We investigated three different approaches: template matching, least squares regression arc estimation, and snakes. The first strategy provided an approximate location of the optic disk centre, while the second method estimated the optic disk as a circle. Either of these were enough to remove the



optic disk as a candidate EX region. However, as a precise localisation is of more use in clinical analysis, the snake technique was developed to produce a more accurate boundary.

### 3.1 Template Matching

We generated a 110x110 template by averaging the optic disk region in 16 normalised retinal images selected randomly. The normalised correlation coefficient [10] was evaluated to present an indication of the match between the template image and each individual pixel in the image under consideration. We applied the method to our 60 retinal images. In all cases the template image could provide a good approximate match of the optic disk centre even when large, albeit scattered, areas of similarly coloured EX pathologies existed in the image. Figure 5 shows the template and three typical results. The main weakness of this approach is that the largest correlation coefficient value does not necessarily correspond to the true optic disk centre as can be seen easily in Figure 5.




**Fig. 5.** (a) Close up view of the template image (b-d) examples of optic disk centre estimation using template matching.

### 3.2 Least Squares Regression Arc Estimation

In this section, we report on a faster and more accurate estimation of the optic disc region. After the FCM segmentation stage, we performed connected component labelling [17] to identify the size and extent of each separate region (Figure 6(b)). The optic disc region is often fragmented into smaller regions due to blood vessels running across it, including a main large region on the temporal side and some smaller ones on the nasal side. We used a very conservative connected region size threshold of 1000 to select all possible candidate regions. This step removed small minor regions and preserves computation time, but was not strictly necessary since our application is not time-critical. Next, we selected the best optic disk region candidate by applying a compactness criterion [17] that rejected all EX regions, since they have an irregular shape and are non-circular. In contrast a real candidate optic disk region exhibits a more regular shape like the fragmented half-circle in Figure 6(b). We then used the boundary of the candidate object to apply an iterative algorithm [18] for estimating the centre of a circular arc and its radius. This algorithm (see Appendix A) is based on minimization of the error between a set of pixels (arc points) and the estimated arc using least squares regression, with an area-based measure of best fit. Figure 6 illustrates the results of locating the optic disk based on arc estimation.

The method was applied to 16 retinal images and could achieve acceptable estimation of the optic disk region which will be reported later in Section 4. This simple approach is more than adequate to enable us to locate the optic disc approximately and exclude it from the candidate EX regions found by FCM. In turn, that allows us to train the network for EX classification more accurately. However, it still only approximates the boundary of the optic disk. Indeed, in some eye diseases, it is necessary to identify the exact boundary of the optic disk.



**Fig. 6.** (a) Typical retinal image including EXs, (b) connected component labelling of EX candidates after FCM, (c) optic disk candidate partial boundary, (d) estimated optic disk area.

### 3.3 Accurate Localisation of the Optic Disk Using Snakes

In [9], Mendels et al. introduced a method based on greylevel morphological smoothing of the optic disk region followed by a manually-placed snake implementation to localise the optic disk. Here, we use a very similar approach, except we introduce two novel improvements. Firstly, we place the initial snake automatically using our earlier levels of processing. Secondly, we improve on the morphological step by using colour morphology which results in a more homogeneous inner disk area and aids the accuracy of the optic disk localisation. The initial contour for a snake must be close to the desired boundary otherwise it can converge to the wrong resting place. We can use either of the methods described in sections 3.1 and 3.2 to automatically position an initial snake. In all our experiments this was performed successfully. In general, a snake is a set of points initially placed near the contour of interest, which are gradually brought closer to the exact shape of the desired region in the image [19]. This is carried out through iterative minimization of an energy function comprising an internal and an external term:


$$E_{snake} = \int_0^1 \frac{1}{2} (\alpha |X'(s)|^2 + \beta |X''(s)|^2) + E_{ext}(X(s)) ds \tag{7}$$

where  $X(s)$  is a parametric curve and the coefficients  $\alpha$  and  $\beta$  control the snake's tension and rigidity respectively. The last term  $E_{ext}(X(s))$  is the external energy and can be defined by an edge map, for example  $f(x,y) = |\nabla[G_\sigma(x,y) * I(x,y)]|^2$ , where  $I$  is the image, and  $G$  is a Gaussian smoothing filter with standard deviation  $\sigma$  and  $\nabla$  and  $*$  are the gradient and convolution operators respectively. The disadvantage of this external energy field is that it may lead to many local energy minima that do not represent the desired boundary. Snakes also have difficulties progressing into

boundary concavities. To overcome these problems, we used the Gradient Vector Flow (GVF) field [20] which replaces the standard formulation of the external energy with a GVF field,  $V(x, y) = (u(x, y), v(x, y))$ . This vector field is evaluated by applying generalised diffusion equations to both components of the gradient of an image edge map. Then, the external energy is defined to minimise the following energy function:

$$\gamma = \iint \mu(u_x^2 + u_y^2 + v_x^2 + v_y^2) + |\nabla f|^2 |V - \nabla f|^2 dx dy \tag{8}$$

where  $\mu$  is a parameter that determines the trade-off between smoothness of the vector field and how much the snake is attracted to the edges. It should be set according to the amount of noise in the image (more noise, bigger  $\mu$ ). In optic disk images, a snake would be significantly affected by the blood vessels present in the optic disk region. Hence, a preparation step is necessary. Morris and Wood [21] used a median filter with size 52x52. This technique is slow and would heavily blur the boundaries of the disk. Mendels et al. [9] applied a closing grey morphology operation to smooth the vascular structures while keeping the actual edges of the optic disk. A symmetrical disk-structuring element of size 13 was used since the blood vessels are not wider than 11 pixels. An example of this is shown in Figure 7(b).



**Fig. 7.** Morphology closing results: (a) close-up of typical optic disk region, (b) greylevel morphological closing, and (c) proposed colour morphological closing.

As Figure 7(a) indicates, all the vessel structures should be eliminated properly and without using morphological filtering it is unlikely that a snake can converge to the optic disk boundary appropriately. Using greylevel morphology, the operation can be applied to the intensity or lightness channel. In colour, the optic disk has quite a bright, sometimes near-saturated characteristic. More importantly, a pixel is multi-valued and must be treated as a vector. There has been relatively little research on colour morphology with the most significant and practical work being Hanbury and Serra [22]. We devised a more basic definition of colour morphology, in which the definitions of *maximum* and *minimum* operations on colour pixels reduce to the maximum and minimum of the set of vectors in the structuring mask. As we are looking to remove dark vascular regions against the bright, near-saturated optic disk region, we use as our point of reference the origin of the colour space to compare colour vectors against. Hence, the definitions for dilation ( $I_d$ ) and erosion ( $I_e$ ) in our simple colour morphology domain become:

$$I_d(x) = \{I(y) : I(y) = \max [I(z)] , z \in s_x \} \tag{9}$$

$$I_e(x) = \left\{ I(y) : I(y) = \min [I(z)] , z \in s_x \right\} \quad (10)$$


where  $s$  is the structuring mask. To perform such operations, a more metric colour space is appropriate and hence we used the  $L^*a^*b$  colour space. Figure 7(c) shows a closing operation using the same structuring mask size. This approach preserves the edges better and, as will be shown later, results in a more accurate localisation of the optic disk.

## 4 Final Processing and Comparison

After the colour morphological pre-processing step, the snake was applied. Several different values were tested for the regularisation parameter of the GVF and we found  $\mu = 0.27$  as the best for our images. The number of iterations during the evaluation of the GVF field was set to 40. We used  $\alpha = 0.7$  and  $\beta = 0.1$  which control the snake's tension and rigidity respectively. Finally, the number of iterations for the snake convergence was determined experimentally and set to 175. All these parameters were kept constant throughout the experiments. We initialised the snake as a circle with a centre at the point found through template matching and with a circular radius set to half of the width of the template. In a number of cases this initial snake intersected the real optic disk boundary, but the GVF snake has the ability to shrink or grow towards the final boundary. Figures 8(a)-8(f) illustrate the optic disk boundary localisation showing the initial and final snake. Figure 8(b) displays the morphological result, but the rest of the images show the snake overlaid on the original image for better visual analysis. Both full and close-up views are shown of the initial and final snake.

Next, we evaluate the accuracy of our localisation compared to manually labelled groundtruth produced by a clinician (e.g. see Figure 8(g)). We also demonstrate the improvement obtained using the proposed simple colour morphological closing over greylevel morphology for this application. We applied closing in greylevel and  $L^*a^*b$  colour space to a set of 16 colour retinal images selected randomly. Then a GVF snake was automatically positioned to localise the boundary of the optic disk in each case. The same measure of accuracy as in (6) was used to compare a groundtruth region with that inside a snake. In this case  $R$  and  $T$  correspond to the groundtruth and the final snake-localised optic disk region respectively (e.g. Figure 8(h)).

Table 4 outlines the evaluated performance for all the 16 images as well as their overall average accuracy. The first column shows the accuracy when the boundary estimation method from section 3.2 was applied. The second and third columns show the accuracy of the snake after using greylevel morphology and the proposed colour morphology respectively. The least squares arc estimation method clearly performed worse than the snake (whatever the morphological pre-processing step) except in case 10. The general success of this method is still significant and is due to the precision of the regions segmented by the FCM algorithm. The best optic disk localisation performance was achieved in 14 of the 16 images by employing a snake after morphological operations in the colour space. Again, it should be taken into account that our clinician agreed that the final snake results represent an accurate boundary for almost all the images compared to his manually determined groundtruth.



**Fig. 8.** Optic disk localisation results, (a) a typical retinal image, (b) vessel elimination in  $L^*a^*b$ , (c) initial snake, (d) final snake, (e) close up of optic disk region and initial snake, (f) close up of final snake, (g) close up of hand-drawn groundtruth, and (h) overlay of (f) and (g).

**Table 4.** Performance evaluation for different optic disk approaches.

RETINAL IMAGE	LEAST SQUARES ARC ESTIMATION	GREY MORPHOLOGY BASED SNAKE	L*A*B MORPHOLOGY BASED SNAKE
1	68.00	89.40	<b>91.07</b>
2	78.66	92.14	<b>92.78</b>
3	82.73	91.89	<b>92.92</b>
4	84.24	92.66	<b>92.72</b>
5	78.80	86.40	<b>90.65</b>
6	83.37	88.70	<b>89.00</b>
7	89.21	<b>92.96</b>	92.16
8	85.70	86.14	<b>89.20</b>
9	80.37	87.74	<b>87.83</b>
10	<b>92.27</b>	89.30	89.72
11	83.08	85.26	<b>88.84</b>
12	84.67	90.41	<b>90.51</b>
13	73.77	87.00	<b>90.63</b>
14	86.48	88.03	<b>89.93</b>
15	81.42	77.93	<b>90.32</b>
16	63.54	86.57	<b>92.62</b>
<b>Overall Accuracy</b>	<b>81.02%</b>	<b>88.28%</b>	<b>90.68%</b>

## 5 Conclusion

In this study we investigated various classifiers including K-Nearest Neighbours, quadratic Gaussian classifier, Linear Delta Rule, and neural networks. We found a NN based on the SCG learning method could represent the best overall diagnostics with 90.1% accuracy, 93.4% sensitivity, and 82.7% specificity, where the trade-off between sensitivity and specificity was appropriately balanced for this particular problem. Due to the importance of optic disk localisation in retinal image analysis and in order to be able to separate the optic disk regions from the EX candidate region list, three different approaches were investigated. Template matching could only provided an approximate location of the optic disk centre, but the least squares regression arc estimation technique could localise the optic disk region to 81% accuracy. However, the third approach localised the optic disk boundary based on the active contours concept to 90.7% accuracy after applying a colour morphology based step.

Another useful by-product of the proposed FCM approach is that it also segments the blood vessels. We hope to turn our attention to these in our future work. The study presented here indicates that automated diagnosis of exudative retinopathy based on colour retinal image analysis is very successful in detecting EXs. Even though not all exudates may sometimes be found in a retinal image, it is important that some are found. This helps in bringing the problem to the attention of the ophthalmologist. Hence, the system could be installed as a mass-screening tool in optician shops and used by non-experts. We have shown that we can detect the majority of exudates, and in all cases when the disease was present, it was detected.

## APPENDIX A

This simple algorithm is taken from [18]. Suppose the coordinates of the data in the image are  $Q = \{q_i = (x_i, y_i), i = 1, \dots, N\}$ . In our case, these are the candidate optic disk points. The estimated circle is  $(x - x_c)^2 + (y - y_c)^2 = R^2$ , where  $(x_c, y_c)$  and  $R$  represent the centre and radius of the circle respectively. These parameters are estimated based on least squares regression as follows:

$$R^2 = \frac{1}{N} \left( \sum x^2 - 2 \sum x_c + N x_c^2 + \sum y^2 - 2 \sum y_c + N y_c^2 \right), \quad x_c = \frac{c_1 b_2 - c_2 b_1}{a_1 b_2 - a_2 b_1}, \quad y_c = \frac{a_1 c_2 - a_2 c_1}{a_1 b_2 - a_2 b_1}$$

where  $a_1 = 2 \left( \left( \sum x \right)^2 - N \sum x^2 \right)$ ,  $a_2 = 2 \left( \sum x \sum y - N \sum xy \right)$ ,  $b_1 = 2 \left( \sum x \sum y - N \sum xy \right)$

$$b_2 = 2 \left( \left( \sum y \right)^2 - N \sum y^2 \right), \quad c_1 = \left( \sum x^2 \sum x - N \sum x^3 + \sum x \sum y^2 - N \sum xy^2 \right)$$

$$c_2 = \left( \sum x^2 \sum x - N \sum y^3 + \sum y \sum y^2 - N \sum x^2 y \right)$$

## References

1. Ghafour, I, Allan, D., Foulds, W.: Common Causes of Blindness and Visual Handicap in West of Scotland. *British Journal of Ophthalmology* **67** (1983) 209-213
2. Philips, R., Forrester, J., Sharp, P.: Automated Detection and Quantification of Retinal Exudates. *Graefe's Archive for Clinical and Experimental Ophthalmology* **231** (1993) 90-94
3. Gardner, G.G., Keating, D., Williamson, T.H., Elliott, A.T.: Automatic Detection of Diabetic Retinopathy Using an Artificial Neural Network: A Screening Tool. *British Journal of Ophthalmology* **80** (1996) 940-944
4. Sinthanayothin, C.: Image Analysis for Automatic Diagnosis of Diabetic Retinopathy. PhD Thesis, King's College London (1999)
5. Wang, H., Hsu, W., Goh, K.G., Lee, M.L.: An Effective Approach to Detect Lesions in Colour Retinal Images. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2000) 181-187
6. Liu, Z., Chutatape, O., Krishnan, S.M.: Automatic Image Analysis of Fundus Photographs. In *19<sup>th</sup> Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (1997)
7. Li, H., Chutatape, O.: Automatic Location of Optic Disk in Retinal Images. *IEEE International Conference on Image Processing* (2001) 837-840
8. Sinthanayothin, C., Boyce, J., Williamson, C.T.: Automated Localisation of the Optic Disk, Fovea, and Retinal Blood Vessels from Digital Colour Fundus Images. *British Journal of Ophthalmology* **38** (1999) 902-910
9. Mendels, F., Heneghan, C., Thiran, J.P.: Identification of the Optic Disk Boundary in Retinal Images Using Active Contours. *Proceedings of the Irish Machine Vision and Image Processing Conference* (1999) 103-115
10. Jain, A.K.: *Fundamentals of Digital Image Processing*. Prentice-Hall Inc (1989)
11. Lim, Y.W., Lee, S.U.: On the Colour Image Segmentation Algorithm Based on the Thresholding and the Fuzzy C-Means Techniques. *Pattern Recognition* **23**, 9 (1990) 935-952
12. Bezdek, J., Keller, J., Krisnapuram, R., Pal, N.R.: *Fuzzy Model and Algorithms for Pattern Recognition and Image Processing*. Kluwer Academic Publishers (1999)
13. Comaniciu, D., Meer, P.: Robust Analysis of Feature Space: Color Image Segmentation. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (1997) 750-755
14. Fukunaga, K.: *Statistical Pattern Recognition*. Academic Press, Inc (1990)
15. Bishop, C.M.: *Neural Networks for Pattern Recognition*. Oxford University Press (1992)
16. Henderson, A.R.: Assessing Test Accuracy and its Clinical Consequences: A Primer for Receiver Operating Characteristics Curve Analysis. *Annals of Clinical Biochemistry* **30** (1993) 521-539
17. Sonka, M., Hlavac, V., Boyle, R.: *Image Processing, Analysis and Machine Vision*. Thomson Computer Press (1993)
18. Thomas, S.M.: A Simple Approach for the Estimation of Circular Arc Centre and its Radius. *Computer Vision, Graphics and Image Processing* **45** (1989) 362-370
19. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active Contour Models. *International Journal of Computer Vision* **1**, 4 (1987) 321-331
20. Xu, C., Prince, J.L.: Snakes, Shapes and Gradient Vector Flow. *IEEE Transactions on Image Processing* **7**, 3 (1998) 359-369
21. Morris, T., Wood, I.: The Automatic Extraction of the Optic Nerve Head. In *American Academy of Optometrists, Biennial European Meeting* (1994) 11-22
22. Hanbury, A., Serra, J.: Mathematical Morphology in the HLS Colour Space. *Proceedings of the 12<sup>th</sup> British Machine Vision Conference* (2001) 451-460