

Retinal Image Classification for the Screening of Age-related Macular Degeneration

Mohd Hanafi Ahmad Hijazi, Frans Coenen and Yalin Zheng

Abstract Age-related Macular Degeneration (AMD) is the most common cause of blindness in old-age. Early identification of AMD can allow for mitigation (but not cure). One of the first symptoms of AMD is the presence of fatty deposits, called *drusen*, on the retina. The presence of drusen may be identified through inspection of retina images. Given the aging global population, the prevalence of AMD is increasing. Many health authorities therefore run screening programmes. The automation, or at least partial automation, of retina image screening is therefore seen as beneficial. This paper describes a Case Based Reasoning (CBR) approach to retina image classification to provide support for AMD screening programmes. In the proposed approach images are represented in the form of spatial-histograms that store both colour and spatial image information. Each retina image is represented using a series of histograms each encapsulated as a time series curve. The Case Base (CB) is populated with a labelled set of such curves. New cases are classified by finding the most similar case (curve) in the CB. Similarity checking is achieved using the Dynamic Time warping (DTW).

Mohd Hanafi Ahmad Hijazi

Department of Computer Science, The University of Liverpool, UK e-mail: m.ahmad-hijazi@liverpool.ac.uk

Frans Coenen

Department of Computer Science, The University of Liverpool, UK e-mail: coenen@liverpool.ac.uk

Yalin Zheng

Ophthalmology Research Unit, School of Clinical Sciences, The University of Liverpool, UK e-mail: yzheng@liverpool.ac.uk

1 Introduction

Age-related Macular Degeneration (AMD) is the leading cause of blindness in people over 50 years of age. It is caused by damage to the macula, a small area on the human retina that is responsible for seeing fine detail and colour [20]. Although there is no cure for AMD, the condition can be mitigated against in the event of early detection. One of the first symptoms of AMD is the presence of fatty deposits, called *drusen*, on the retina. These can be detected by inspection of retina images routinely collected within screening programmes. This image inspection is conducted manually by trained clinicians. This paper describes an image classification mechanism to (at least partially) automate the identification of drusen in retina images.

The main challenge of the retina image AMD classification problem is that it is often difficult to distinguish drusen from background noise. The need for appropriate image representations, to facilitate the application of data mining, has been identified as a generic challenge within the context of medical image classification [9, 19]. In the context of AMD screening “standard” object segmentation techniques were deemed to be unsuitable as the shape and size of drusen varies significantly from image to image and tends to “blur” into the background. A spatial-histogram [18, 26] based approach was therefore adopted, a technique that features the ability to maintain spatial information between groups of pixels [3]. A *region* based approach is advocated in this paper where by the images are subdivided into “areas” and histograms are generated for each. The histograms were conceptualised as time series where the X-axis represents the histogram “bin” number, and the Y-axis the size of the bins (number of pixels contained in each).

To facilitate the desired classification a Case Based Reasoning (CBR) approach was adopted [21], where-by a collection of labelled cases were stored in a repository. A new case to be classified (labelled) is compared with the cases contained in this repository and the label associated with the *most similar* case selected. Given that the histograms can be conceptualised as time series, a Dynamic Time Warping (DTW) technique [1, 25] was adopted to determine the similarity between “curves”. The principal contributions of the work described are:

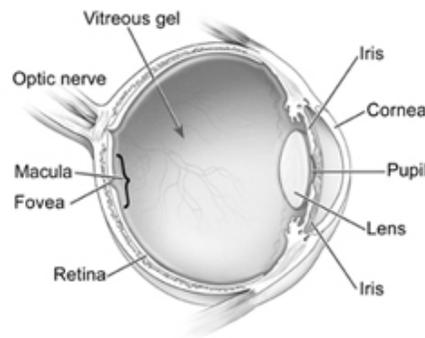
- A novel approach to AMD screening.
- A mechanism (that also has wider application) for classifying retina images for AMD without specifically identifying drusen.
- The use of regions in the representation to enhance the classification accuracy.
- An approach to CBR case similarity checking using a time series analysis technique.

The rest of this paper is organised as follows. Section 2 describes the application domain and Section 3 some relevant previous work. The screening process is described in Section 4. Section 5 and 6 provide further detail of how the retinal images are pre-processed and then transformed into the spatial-histogram (time series) representation. The specific classification technique used is described in Section 8, followed by an evaluation of the proposed approach in Section 9. Some conclusions are presented in Section 10.

2 Age-related Macular Degeneration

The work described in this paper is focused on the classification of retinal images, in particular the identification of age-related macular degeneration (AMD). Figure 1 illustrates a typical cross sectional view of the eye. The eye consists of the cornea, iris, pupil, lens, vitreous humour and the retina. As shown in Figure 1, centred at the fovea, the macula is a small area at the centre of the retina. It contains the densest photoreceptors and provides “central vision” and “colour vision”. Central vision is essential for humans to see fine detail as required by daily tasks such as reading and writing. Sometimes the delicate cells of the macula become damaged and stop functioning properly. There are various conditions for this to occur amongst which AMD is the leading cause of irreversible

Fig. 1 Cross sectional view of the eye National Institutes of Health (NIH), National Eye Institute (NEI), US (<http://www.nei.nih.gov/>).



Early diagnosis of AMD is achieved by the identification of *drusen* [20, 8], yellowish-white sub-retinal fatty deposits, by screening patient retinal images. The severity of AMD can be categorised into three classes: *early*, *intermediate*, and *advanced*. AMD can be either *non-neovascular* or *neovascular* [8]. Early AMD is characterised by the existence of several small ($63\mu\text{m}$ in diameter) or a few medium (63 to $124\mu\text{m}$) sized drusen or retinal pigmentary abnormalities. The presence of at least one large ($124\mu\text{m}$) and numerous medium sized drusen, or geographic atrophy, that does not extend to the centre of the macula, characterises intermediate AMD. Advanced non-neovascular (dry) AMD exists once the drusen has reached the center of the macula. *Choroidal neovascularisation* characterizes advanced neovascular (wet) AMD. The drusen itself is often categorised as *hard* and *soft* drusen. Hard drusen have a well defined border, while soft drusen have boundaries that often blend into the retinal background. Figure 2(a) shows an example of normal retinal image with the macula circled. A retina image that features drusen is given in Figure 2(b) (drusen indicated by a white arrow). The classification of AMD images by means of drusen identification is thus not a straightforward process. Most of the previous works have focused on automatic drusen segmentation [4, 13, 22, 23, 29] as opposed to AMD classification. The work proposed here however approaches the AMD screening problem without the need for identification of the physical existence of drusen and aims to classify images as either “AMD” or “non-AMD”.

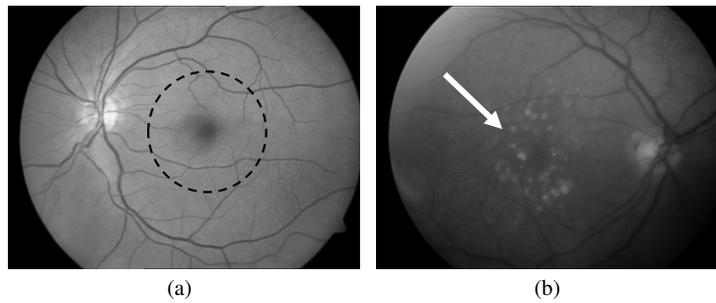


Fig. 2 Illustration of fundus images in grayscale: (a) Normal and (b) AMD.

3 Previous Work

The earliest work reported in the literature concerning drusen detection is that of Sbeh et al. [30] who used mathematical morphology to identify brightest points to detect drusen. More recent work [4] used a wavelet analysis technique to extract drusen patterns, and multi-level classification (based on various criteria) for drusen categorisation. Other works on the identification of drusen in retina images has focuses on segmentation coupled with image enhancement approaches [22, 23, 29]. Rapantzikos et al. [29] adopted a multilevel histogram equalisation to enhance the image contrast followed by drusen segmentation, in which two types of threshold, global and local, were applied to retinal images. Köse et al. [22, 23] proposed two approaches involving *inverse* drusen segmentation within the macular area. A region growing technique was used to identify “healthy” pixels by applying a threshold on the colour intensity levels [22]. Once this was done, the inverse of the segmented image was used to generate the segmentation of the drusen. A similar inverse segmentation approach, supported by statistical information, was adopted in [23]; where healthy *Characteristic Images* (CIs) were compared to new *Sample Images* (SIs) and a predetermined threshold is applied to classify SI. In [13] another approach, based on a non-parametric technique for anomaly detection, was described that uses a Support Vector Data Description (SVDD) to segment anomalous pixels.

There has been very little reported work on the application of image mining techniques for AMD screening. The existing work (see above) has been mostly focuses on the segmentation/identification of drusen. Of the reported work that the authors’ are aware of, only two reports [4, 13] extend drusen detection and segmentation to distinguish retinal images with and without AMD features. However, all the previous work is focused on the detection of drusen using segmentation, a challenging task given the inconsistent visual appearance of drusen and other lesions. The clarity, colour, luminosity and texture of images are affected by several factors during the image acquisition process, such as involuntary eye movement and the media opacity of the subject.

The distinction between the work described here and previous approaches is that we make no attempt to locate and isolate (segment) drusen within retinal images. Instead, we extend the uses of individual colour channel histograms [16] to a spatial-histogram based approach that obviates the need for accurate segmentation

of drusen. Spatial-histograms extend the concept of simple colour histograms by including spatial pixel information [3, 33, 35] and have been shown to perform well in region-based tracking [3], object detection [35] and image retrieval [33].

Space limitations preclude an overview of CBR. However CBR is a well established AI technique with an associated, well established, body of literature. Recommended reference works include [24] and [21]. For a review of the application of CBR in medical domains interested readers are referred to [17] or [2].

4 The AMD Screening Process

An overview of the proposed retinal image classification, to identify AMD, is presented in this section. The approach can be viewed as consisting of two stages, (i) Case Base (CB) generation and (ii) image classification. A block diagram outlining the process is given in Figure 3 (the directed arcs indicate process flow). In the figure the two stages are delimited by dashed boxes. The case base generation process commences at the top left of the figure, while the classification process at the bottom left.

CB generation comprises three sub-stages: (i) image preprocessing, (ii) histogram generation and (iii) feature selection. CB generation commences with a training set of pre-labelled images which are preprocessed as follows:

1. **Image Enhancement:** Normalisation and enhancement of the image contrast. Colour normalisation is applied first, followed by illumination normalisation and then contrast enhancement to increase the “visibility” of the main retinal anatomy (blood vessels, etc.).
2. **Object Segmentation:** Identification of the main retinal structures.
3. **Noise Reduction:** Removal of blood vessel pixels from the retina images.

The image pre-processing is described in further detail in Section 5.

The next step is to generate the spatial-histograms. In order to make the representation more tractable, colour quantisation was applied to the preprocessed images to reduce the overall dimensionality (number of colours). To generate the histograms the quantised colour retinal images were first partitioned into nine *regions* and then spatial-histograms were extracted for each region. The idea here is that the presence of drusen is often regionalised and consequently we may be more interested in some regions than others. Section 6 gives more detail of the technique used to generate the spatial-histograms.

During feature selection the spatial-histograms (regions) that feature the best discriminatory power (in the context of AMD classification) are identified. The regions are ranked according to their discriminatory power and the top T selected. This process also ensured that the size (number of pixels) of each region/histogram does not bias the resulting classification. The feature selection was conducted using a *class separability* measure which was applied to the collection of histograms representing each retina image and the most appropriate histograms selected. The selected

spatial-histograms were then combined and stored in the form of time series curves (one per image). The feature selection process is discussed in further detail in Section 7.

The image classification task is detailed in Section 8.

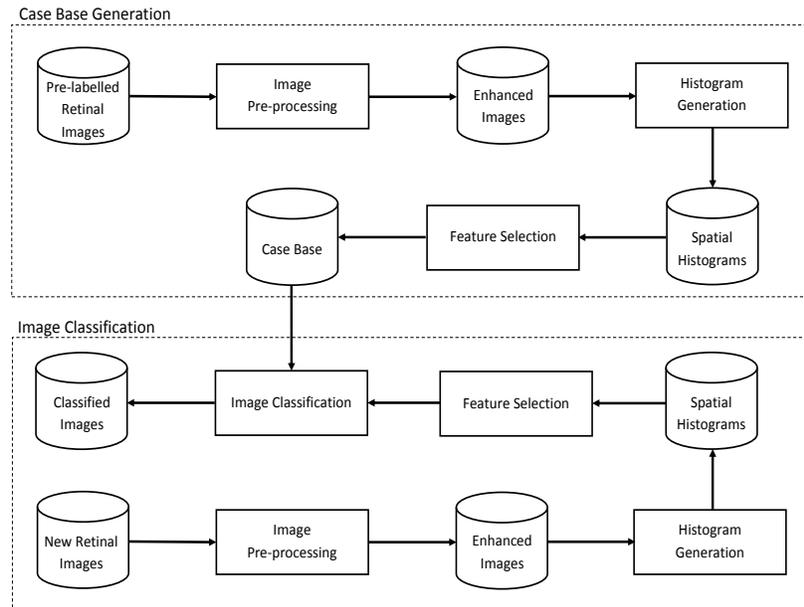


Fig. 3 Block diagram of the proposed retinal images screening system

5 Image Pre-processing

This section describes the image pre-processing steps required to represent images into meaningful forms for image mining. The image pre-processing consists of two steps: (i) image enhancement and (ii) segmentation of anatomic structures to identify retinal blood vessels.

5.1 Image Enhancement

The quality of the retinal images is often severely affected by factors such as: colour variance and non-uniform illumination [11, 27], which are difficult to control. In the context of AMD screening this will lead to difficulties in the detection of drusen, and hamper the associated identification and localisation of retinal common struc-

tures such as retinal blood vessels. Thus, colour and illumination normalisation, and contrast enhancement are important.

Due to the colour variation between different retinal images, colour normalisation must be performed prior to image enhancement. To normalise the colours featured in retinal images a *histogram specification* approach was applied [14]. First, a reference image that represents the best colour distribution and contrast is selected by a trained clinician. Then, the Red-Green-Blue (RGB) colour histograms of the reference image are generated. Finally, the RGB histograms of other images are extracted and each of these histograms is tuned to match the reference image histograms.

Once the colour is normalised, illumination normalisation is applied so as to reduce the luminosity variations on the image. An approach, to estimate the luminosity and contrast variability of the retinal image based on the image background colour, proposed by Foracchia et al. [11] was adopted. This approach estimates the original image, \bar{I} , as follows:

$$\bar{I}(x,y) = \frac{I(x,y) - \bar{L}(x,y)}{\bar{C}(x,y)}, \quad (1)$$

where I is the observed image, and \bar{L} and \bar{C} are the estimations of luminosity and contrast, calculated in the neighbourhood N of each pixel. One drawback of this approach is that drusen that are larger than the window size N , used for the estimation, are smoothed in the normalisation process. However, the authors found that this disadvantage could be limited by setting the \bar{C} value to 1 there by excluding the contrast estimation. Contrast normalisation was then conducted using Contrast Limited Adaptive Histogram Equalisation (CLAHE) as described in [36].

5.2 Objects Segmentation

The presence of retinal anatomies, such as blood vessels and the optic disc, sometimes hampers the detection of drusen. The authors' own experiments have indicated that the removal of blood vessel pixels from retina images can improve classification accuracy [16]. This has also been observed more generally by other researchers in the field ([23, 28, 29]).

To segment the retinal blood vessels 2-D Gabor wavelet filters [31] were applied. A pixel is classified as *vessel* or *non-vessel* by means of a Bayesian classifier with a class-conditional probability density function, generated using the Gaussian mixture model. As a result a "retinal vessels" binary representation is generated for each image which is then applied as a "mask" to the enhanced retinal images and consequently the blood vessels pixel values replaced with a "null" value.

The optic disc was however left untouched as experiments conducted by the authors, reported in [16], indicated that removal of the optic disc only results in increased accuracy with respect to a minority of retina images and decreases accuracy with respect to the majority.

6 Spatial Histogram Generation

Colour histograms have been widely used as a simple way of representing images for object identification and retrieval [5, 32]. The main advantage is their robustness against object changes in terms of shape and position within images. The main disadvantage is the loss of spatial information between pixels and colours, thus images with similar histograms may have very different appearances [33, 35]. In some images, the colour distribution of pixels at different sections of an image may be an essential feature that should be included in the image representation. In the context of AMD classification there are a significant number of cases where the AMD images have almost similar colour histograms to the normal ones. The fact that drusen pixel colours are very similar to the colours of pixels adjacent to the retinal blood vessels boundaries (as well the optic disc), may thus lead to classification errors. A spatial-histogram [18, 26] representation was therefore adopted.

The spatial information of an image can be generated by preserving the objects texture and shape using templates [35], as well as by partitioning the image into regions based on the chosen colour values and keep the regions location for each of the chosen colours [18]. The utilisation of texture and shape to extract spatial information is hampered by the nature of the AMD featured images where no common textures and shapes exist, other than the main retinal structures. Therefore, a method to generate colour distribution for each region [33] has been applied in the work described here as it is conjectured that the similar regions of two different classes of retinal images will have different colour distribution. The generation of spatial-histogram consisted of several steps. First, the number of colours was reduced to make the computational cost more feasible. The minimum variance quantisation technique [34], with dithering [10] (implementation using Matlab¹ function *rgb2ind*), was used to reduce the image colours to C colours. A careful selection of C value is essential as it will affect the quality of the generated histograms, as shown in Section 9. The colour quantisation was applied on the global colour space, instead of local, in order to standardise the colour mapping. Thus, all images referenced a similar colour map.

Once the colour quantisation was complete each image was partitioned into N similar sized regions, $\mathcal{R} = \{r_1, r_2, \dots, r_N\}$, and a spatial-histogram generated for each. The set of spatial histograms for a given image m is defined as:

$$h_m = \{sh_1^m, sh_2^m, \dots, sh_N^m\} \quad (2)$$

where sh_n^m is the spatial-histogram generated for region n , ($1 \leq n \leq N$) in image m with size of C bins. The histogram value for colour c in histogram sh_n^m is then given by:

$$sh_n^m(c) = \alpha \quad (3)$$

¹ <http://www.mathworks.com>

where α is the c -th bin count in region n of image m , and ($0 \leq c < C$). The size of each image spatial-histograms, h_m , for an image m is equivalent to $C \times N$; the number of colours, C , multiplied by the number of regions, N . The complete set of histograms representing an image set is then defined as $H = \{h_1, h_2, \dots, h_M\}$, where M is the number of images.

7 Feature Selection

Feature selection is a process to reduce the number of features contained in a feature space by removing irrelevant or redundant features[6, 7, 12]. By selecting only those features that have a strong discriminatory power between classes, the computational cost of classification can be considerably reduced while at the same time maximising classification accuracy [6]. Common feature selection techniques[7, 12] include the χ^2 measure, mutual information, Odds Ratio and Principal Component Analysis.

With respect to the AMD screening process described here a class separability method [6] that estimates the effectiveness of a features ability to distinguish between classes using the Kullback-Leibler (KL) distance measure was adopted. This was a two stage process. First an average *signature*, γ_n , histogram was generated for each region with respect to each class as follows:

$$\gamma_n^a = \frac{1}{p} \sum_{j=1}^p sh_n^j \quad (4)$$

where n is the region identifier, a is a class label and p is the number of training set images labelled as class a . The *class separability*, $dist_n$, is then calculated by:

$$dist_n = \sum_{a=1}^d \sum_{b=1}^d \delta_n(a, b) \quad (5)$$

where d is the number of classes and $\delta_n(a, b)$ is the KL distance between histograms of γ_n corresponding to classes a and b described as:

$$\delta_n(a, b) = \sum_{i=1}^c p_n(\gamma_n^a(i)) \log \left(\frac{p_n(\gamma_n^a(i))}{p_n(\gamma_n^b(i))} \right) \quad (6)$$

where c is the number of bins or colours in the histograms, and $p_n(\gamma_n^a(i))$ is the probability that the n -th feature takes a value in the i -th bin of the signature spatial-histogram γ_n given a class a . The probability, p_n was calculated by dividing each bin count of γ_n by the total number of elements in γ_n .

The features are then sorted in descending order of $dist_n$; the top T features with the highest $dist_n$ provide the best separation between classes and are therefore selected. However, the selection of value of T is domain dependent and for the work described here, $T = 5$ consistently produced the best result as shown in Section 9. The other regions were omitted from further processing. Thus, the size of h_m has

been reduced to only $C \times T$. These histograms then make up the CB for the CBR process.

8 Retinal Image Classification using CBR and DTW

Given a new set of images produced during an AMD screening process these may be classified using the CB developed as described in the foregoing subsections. As noted above the histograms in the CB may be viewed as time series. Similarity checking may therefore be conducted using time series analysis techniques. For the AMD screening a Dynamic Time Warping (DTW) technique [1, 25] was adopted. DTW is a time series analysis technique that measures the distance between two time series through the generation of a *warping path* between these sequences. Given two time series, $T = \{t_1, t_2, \dots, t_m\}$ and $\bar{T} = \{\bar{t}_1, \bar{t}_2, \dots, \bar{t}_n\}$, a matrix of size $m \times n$ will be formed. The distance between t_i and \bar{t}_j , $d(t_i, \bar{t}_j)$, where $0 \leq i < m$ and $0 \leq j < n$ for all i and j is computed using the Euclidean distance similarity measure (other similarity measure methods can also be applied). The minimal warping path is computed by summing up the minimal d for each matrix grid points thus giving a distance between T and \bar{T} . More details of the DTW approach with respect to retinal image classification can be found in [15].

9 Evaluation

To evaluate the AMD screening process a collection of 144 retinal images, acquired as part of the ARIA² project, were used. The collection was manually pre-labelled, included 86 AMD images and 58 non-AMD images. The experiments described in this section evaluate the performance of the proposed approach. Three metrics are used for evaluation purposes: *Specificity*, *Sensitivity* and *Accuracy*. All experiments were conducted using Tenfold Cross Validation (TCV) whereby the dataset was randomly divided into equal sized “tenths”; and on each TCV iteration, one tenth was used as the test set while the remainder was used as the training set. The objectives of the experiments may be summarised as follows and is described in the following subsections:

1. **Number of Bins Parameter:** To determine the minimum number of bins for the histograms, with respect to colour quantisation, such that classification accuracy would not be adversely affected.
2. **T Parameter Identification:** To determine the most appropriate setting for the T parameter, the threshold that determines the number of regions to be included in the final representation during feature selection.

² http://www.eyecharity.com/aria_online

9.1 Number of Bins Parameter

The first set of experiments was designed to determine the number of output bins for colour image quantisation. The aim was to determine the least number of bins while maintaining classification accuracy. Experiments using 32, 64, 128, and 256 bins were conducted (but without the region concept). Table 1 shows the classification results obtained. The results clearly indicate that the overall classification accuracy is relative to the number of bins up to 128. This was expected as low numbers of colour bins will tend to group different coloured pixels in to the same bin, and consequently reduce the discriminative power of the colour representation.

Table 1 Classification results for a range of colour quantisation output bins

Bins	Specificity (%)	Sensitivity (%)	Accuracy (%)
32	53	74	66
64	69	67	68
128	55	81	71
256	52	84	71

9.2 T Parameter Identification

The results presented in the foregoing were generated by setting the number of regions parameter to one. The experiment described in this sub-section consider the effect of using regions, as opposed to the entire image, and how many regions should be considered. For this purpose, the retinal images were partitioned into $3 \times 3 = 9$ regions. The number of regions however could be tailored depending on the problem domain. Spatial-histograms were then generated as described in Section 6. Bin parameter values of 32, 64 and 128 were used; the 256 output bin was omitted from further analysis as it did not give any significant improved performance over the 128 bin threshold and also because it would introduce a significant computational overhead. The retinal image classification was performed using the top- T regions that had the highest discriminatory capability.

Table 2 Classification results for AMD classification using 32 colour output bins with various T values

T	SH-dimension	Specificity (%)	Sensitivity (%)	Accuracy (%)
1	32	71	72	72
2	64	71	65	67
3	96	68	75	72
4	128	68	75	72
5	160	74	79	77
6	192	71	72	72
7	224	69	74	72
8	256	66	74	71
9	288	71	76	74

Table 3 Classification results for AMD classification using 64 colour output bins with various T values

T	SH-dimension	Specificity (%)	Sensitivity (%)	Accuracy (%)
1	64	59	70	65
2	128	66	68	67
3	192	64	71	68
4	256	69	68	69
5	320	70	74	73
6	384	69	73	71
7	448	68	70	69
8	512	69	67	68
9	576	69	76	74

Table 4 Classification results for AMD classification using 128 colour output bins with various T values

T	SH-dimension	Specificity (%)	Sensitivity (%)	Accuracy (%)
1	128	61	69	65
2	256	61	75	69
3	384	59	81	72
4	512	65	80	74
5	640	67	80	75
6	768	64	78	72

Comparisons with various T values are reported in Tables 2, 3 and 4 shows that the average classification results obtained using 32, 64 and 128 sized bins respectively compared to a range of T parameter values. The *SH dimension* column indicates the total number of bins (dimensions) in the spatial-histogram representation (calculated by multiplying the Bin parameter by the T parameter). Inspection of the results indicates that there is a tendency for best results to be produced when $T = 5$, although the evidence is not conclusive. In Table 2 (32 bins) the best results were obtained when $T = 5$, with an overall accuracy of 77%. Similar results are shown in Table 4 (128 bins) with the best overall accuracy of 75% when $T = 5$. The results in Table 3 however performed best with $T = 9$ with overall accuracy of 74%, although a setting of $T = 5$ also produced good results. The best specificity of 74% was recorded with $T = 1$ and 32 colour bins, and the best sensitivity of 81% with $T = 3$ and 128 colour bins. One interesting observation is that specificity tends to increase as the number of colour bins decreases. This may be because a low number of colour bins gives lower colour variation.

Overall the results demonstrate that by using only some portion of the images a comparative or better classification result is generated than when using the entire image. The results in Table 4 contains only six T values (1 to 6) as the machine memory required for the classification process increases quadratically with the size of the colour bins. Thus, the authors have decided to stop the process at $T = 6$ because of: (i) the computational complexity when comparing two spatial-histograms of $T = 6$ with 128 colour output bins is $\mathcal{O}(n^2)$, the time complexity is more than two orders of magnitude compared to the best results recorded in the experiment ($T = 5$

and 32 colour output bins), and (ii) as indicated by Table 2 and 3 performance will most probably decrease as the size of the spatial-histograms increases.

10 Conclusion

An approach of retinal image classification for AMD screening has been described. The images were represented in the form of spatial-histograms that stored the colour information of the images, while maintaining the spatial information of each colour value. A feature selection strategy, to identify regions in an image that have strong discriminative power to separate classes, was applied to remove irrelevant features, as well as reducing the overall computational cost. The experiments described show both promising and interesting results. Best performance was achieved with a low number of colour bins (32) and a T parameter (number of regions) of 5.

References

1. D. J. Berndt and J. Clifford. Using dynamic time warping to find patterns in time series. In *AAAI Workshop on Knowledge Discovery in Databases*, pages 229–248, 1994.
2. I. I. Bichindaritz and C. C. Marling. Case-based reasoning in the health science: What's next? *Artificial Intelligence in Medicine*, 36(2):127–135, 2006.
3. S. T. Birchfield and S. Rangarajan. Spatial histograms for region-based tracking. *ETRI Journal*, 29(5):697–699, 2007.
4. L. Brandon and A. Hoover. Drusen detection in a retinal image using multi-level analysis. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 618–625. Springer-Verlag, 2003.
5. R. Brunelli and O. Mich. Histograms analysis for image retrieval. *Pattern Recognition Letters*, 34:1625–1637, 2001.
6. E. Cantu-Paz. Feature subset selection, class separability, and genetic algorithms. In *Proceedings of Genetic and Evolutionary Computation Conference*, pages 959–970, 2004.
7. E. Cantu-Paz, S. Newsam, and C. Kamath. Feature selection in scientific applications. In *Proceedings of 2004 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 788–793, 2004.
8. P. T. V. M. de Jong. Age-related macular degeneration. *The New England Journal of Medicine*, 355(14):1474–1485, 2006.
9. U. M. Fayyad, P. Smyth, N. Weir, and S. Djorgovski. Automated analysis and exploration of image databases: Results, progress, and challenges. *Journal of Intelligent Information Systems*, 4:7–25, 1995.
10. R. W. Floyd and L. Steinberg. An adaptive algorithm for spatial greyscale. *Society for Information Display*, 17(2):75–77, 1976.
11. M. Foracchia, E. Grisan, and A. Ruggeri. Luminosity and contrast normalization in retinal images. *Medical Image Analysis*, 9:179–190, 2005.
12. G. Forman. An extensive empirical study of feature selection metrics for text classification. *Journal of Medical Learning Research*, 3:1289–1305, 2003.
13. D. E. Freund, N. Bressler, and P. Burlina. Automated detection of drusen in the macula. In *Proceedings of the Sixth IEEE International Conference on Symposium on Biomedical Imaging: From Nano to Macro*, pages 61–64, 2009.
14. R. C. Gonzalez and R. E. Woods. *Digital image processing*. Pearson Prentice Hall, 2008.

15. M. H. A. Hijazi, F. Coenen, and Y. Zheng. A histogram based approach for the screening of age-related macular degeneration. In *Medical Image Understanding and Analysis 2009*, pages 154–158. BMVA, 2009.
16. M. H. A. Hijazi, F. Coenen, and Y. Zheng. Retinal image classification using a histogram based approach. In *Proc. International Joint Conference on Neural Networks*, pages 3501–3507. IEEE, 2010.
17. A. Holt, I. Bichindaritz, R. Schmidt, and P. Perner. Medical applications in case-based reasoning. *The Knowledge Engineering Review*, 20:289–292, 2005.
18. W. Hsu, S. T. Chua, and H. H. Pung. An integrated color-spatial approach to content-based image retrieval. In *Proceedings of the Third International Conference on Multimedia*, pages 305–313, 1995.
19. W. Hsu, M. L. Lee, and J. Zhang. Image mining: Trends and developments. *Intelligent Information Systems*, 19(1):7–23, 2002.
20. R. D. Jager, W. F. Mieler, and J. W. Mieler. Age-related macular degeneration. *The New England Journal of Medicine*, 358(24):2606–2617, 2008.
21. J. Kolodner. *Case-based reasoning*. Morgan Kaufmann, 1993.
22. C. Köse, U. Şevik, and O. Gençaliöğlü. Automatic segmentation of age-related macular degeneration in retinal fundus images. *Computers in Biology and Medicine*, 38:611–619, 2008.
23. C. Köse, U. Şevik, and O. Gençaliöğlü. A statistical segmentation method for measuring age-related macular degeneration in retinal fundus images. *Journal of Medical Systems*, 34(1):1–13, 2008.
24. D. B. Leake. *Case-based reasoning: Experiences, lessons and future directions*. AAAI Press/MIT Press, 1996.
25. C. S. Myers and L. R. Rabiner. A comparative study of several dynamic time-warping algorithms for connected word recognition. *The Bell System Technical Journal*, 60(7):1389–1409, 1981.
26. B. C. Ooi, K-L. Tan, T. S. Chua, and W. Hsu. Fast image retrieval using color-spatial information. *The International Journal of Very Large Data Bases*, 7(7):115–128, 1998.
27. A. Osareh. *Automated identification of diabetic retinal exudates and the optic disc*. PhD thesis, University of Bristol, UK, 2004.
28. N. Patton, T. M. Aslam, and T. MacGillivray. Retinal image analysis: Concepts, applications and potential. *Progress in Retinal and Eye Research*, 25:99–127, 2006.
29. K. Rapantzikos, M. Zervakis, and K. Balas. Detection and segmentation of drusen deposits on human retina: Potential in the diagnosis of age-related macular degeneration. *Medical Image Analysis*, 7:95–108, 2003.
30. Zakaria Ben Sbeh, Laurent D. Cohen, Gerard Mimoun, and Gabriel Coscas. A new approach of geodesic reconstruction for drusen segmentation in eye fundus images. *IEEE Transactions on Medical Imaging*, 20(12):1321–1333, 2001.
31. J. V. B. Soares, J. J. G. Leandro, R. M. Cesar Jr., H. F. Jelinek, and M. J. Cree. Retinal vessel segmentation using the 2-d gabor wavelet and supervised classification. *IEEE Transactions on Medical Imaging*, 25(9):1214–1222, 2006.
32. M. J. Swain and D. H. Ballard. Color indexing. *International Journal of Computer Vision*, 7(1):11–31, 1991.
33. H-C. Wu and C-C. Chang. An image retrieval method based on color-complexity and spatial-histogram features. *Fundamenta Informaticae*, 76:481–493, 2007.
34. X. Wu. *Graphic Gems II*, chapter Efficient statistical computations for optimal color quantization, pages 126–133. Elsevier Science and Technology, 1991.
35. H. Zhang, W. Gao, X. Chen, and D. Zhao. Object detection using spatial histograms features. *Image and Vision Computing*, 24:327–341, 2006.
36. K. Zuiderveld. *Contrast limited adaptive histogram equalization*, pages 474–485. Academic Press Graphics Gems Series. Academic Press Professional, Inc., 1994.