Image Classification using Histograms and Time Series Analysis: A Study of Age-Related Macular Degeneration Screening in Retinal Image Data

Mohd Hanafi Ahmad Hijazi[1], Frans Coenen[1], and Yalin Zheng[2]

[1] Department of Computer Science, University of Liverpool, Liverpool, L69 3BX UK.
[2] Ophthalmology Research Unit, School of Clinical Sciences, University of Liverpool, Liverpool, L69 3GA UK.
{m.ahmad-hijazi,coenen,yalin.zheng}@liverpool.ac.uk
http://www.csc.liv.ac.uk

Abstract. An approach to image mining is described that combines a histogram based representation with a time series analysis technique. More specifically a Dynamic Time Warping (DTW) approach is applied to histogram represented image sets that have been enhanced using CLAHE and noise removal. The focus of the work is the screening (classification) of retinal image sets to identify age-related macular degeneration (AMD). Results are reported from experiments conducted to compare different image enhancement techniques, combination of two different histograms for image classification, and different histogram based approaches. The experiments demonstrated that: the image enhancement techniques produce improved results, the usage of two histograms improved the classifier performance, and that the proposed DTW procedure out-performs other histogram based techniques in terms of classification accuracy.

Keywords: Image mining, Medical image mining, Dynamic time warping, Image classification, Histogram based classification.

1 Introduction

There is much current interest within the data mining community in image mining [1, 2], especially medical image mining [3, 4]. This paper describes a histogram based approach to medical image mining, whereby the histograms are conceptualised as time series, and consequently time series analysis techniques may be applied to classify the images. The focus of the paper is the screening of colour fundus images to identify the possibility of a condition known as age-related macular degeneration (AMD) [5]. However, the described approach has much more general applicability.

Image mining can be undertaken in a number of ways, each requiring a different style of data input. The simplest approach is to express the image set as a set of tabular records, where each column (attribute) represents some image attribute, and then apply established tabular data mining techniques [2, 4]. The

issue here is the identification of the most appropriate image attributes [3, 6, 4]. Alternative representations encode the images as graphs where the nodes represent blocks of pixels of similar colour and the edges some relationship between the blocks [3]. In this case graph mining techniques can be applied. Another representation, and that advocated in this paper, is the histogram based approach whereby RGB (red, green and blue) and HSI (hue, saturation, intensity) values are represented as histograms.

In this paper histograms are conceptualised in terms of time series. By considering images in this manner time series analysis techniques may be applied. This paper advocates the use of Dynamic Time Warping (DTW) [7, 8], although other time series techniques may be used.

For histogram based medical image mining techniques (and many other mining techniques) to perform well it is desirable to first remove "noise" from the image set. In the case of our retina image set, the focus of this paper, we are interested in removing blood vessels from the image. The successful removal of noise is often helped by first applying some image enhancement process to the input data.

This paper is organised as follows: further details of the AMD background is presented in Section 2. Section 3 reports on some previous works on histogram based approaches, gives a brief review of popular image enhancement methods, and explanations of some general background concerning time series analysis, as applied to image mining, and DTW in particular. In Section 4 image enhancement is considered in the context of AMD; the results from a series of enhancement experiments are reported. The proposed AMD retina image screening process is outlined in Section 5, which includes details of our approach to: (i) noise reduction in the input image set and (ii) DTW. The experimental setup of the screening process, using real data, is reported in Section 6. The results are evaluated and discussed in Section 7. A summary, main findings and some conclusions are presented in Section 8.

2   Age-related Macular Degeneration

The macula is a small area located at the very centre of the retina, as shown in Figure 1(a) (indicated by a dashed circle). Dominated by cone photoreceptors, the macula allow people to see fine detail as well as colours. Sometimes the delicate cells of the macula become damaged and stop functioning. This leads to a number of eye disorders including age-related macular degeneration (AMD), where the macula degenerates with age [5]. AMD is often diagnosed, at its early stage, by the identification of drusen (yellowish-white subretinal deposits), through screening of patient retinal images. Drusen is usually the first clinical indicator of AMD. The severity of AMD is categorised as being either: *early*, *intermediate*, *advanced non-neovascular*, or *advanced neovascular* [5]. Each category is characterised by the existence of various sizes and shapes of drusen, pigment abnormality and/or other lesions. An example of a retina image that features drusen is given in Figure 1(b) (indicated by the white arrow). Drusen

itself are categorised as *hard* and *soft* drusen. Hard drusen have a well defined border, while soft drusen have boundaries that often blend into the background. The identification of features of AMD is thus not a straightforward process [9], and consequently hampers automated diagnosis of AMD at an early stage.
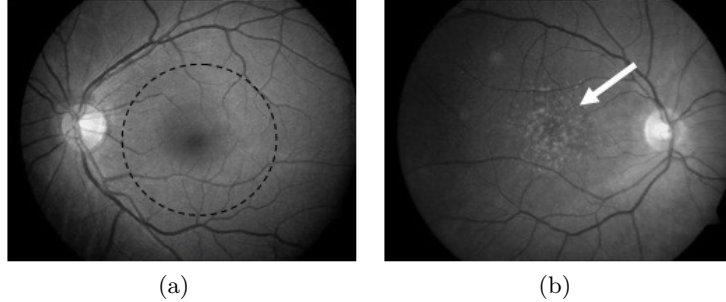

(a)   (b)

Fig. 1: (a) Normal and (b) AMD retinal images in greyscale

3   Previous Work

In this section an overview of relevant previous work is presented. The section is divided into three sub-sections. Sub-section 3.1 deals with histogram based approaches to image mining and especially medical image mining. Sub-section 3.2 deals with time series analysis, and especially DTW, in the context of medical image analysis. Sub-section 3.3 gives a brief overview of current image enhancement techniques.

3.1   Histogram Based Image Mining

Mining images according to content, in particular the image's colour distribution, is common in image mining [10–12]. The basic idea is to extract relevant information from the images (colour, texture, etc.) as feature vectors; which can then be represented in other forms, such as histograms, tables or graphs. Histograms tend to be used to signify colour and saturation information. Similarity measures may be used to measure "distance" between images [11, 12] for image retrieval, categorisation and classification.

Various similarity measures have been applied to histogram based image classification and retrieval techniques, and each has been empirically measured [13, 11]. Earth Mover's Distance (EMD) [14] computes the distance between two distributions, which are represented by user defined signatures. The minimum amount of work needed to transform one distribution into the other is used to measure how similar those distributions are. EMD has been reported to perform well for small sample sizes. Manhattan distance (L_1) has yielded the most effective common dissimilarity measures for histogram retrieval [13]. Euclidean Distance (L_2) is the most common metric used for calculating distance and is used in this paper as the base metric for performance evaluation.

3.2   Time Series Analysis for Image Mining

There is very little reported work on time series analysis for image mining. In [10] colour distribution was represented as a time series for image classification and clustering. Using a time series data representation, called Symbolic Aggregate approXimation (SAX) [15], and the K-nearest neighbour technique for classification. The results in [10] demonstrated a promising approach to time series analysis for image mining. The distinctions between this approach and that presented in this paper are in the time series data representation (SAX represents time series as a sequence of symbols) and the similarity measures adopted.

DTW is a technique for measuring the similarity between two time series. It has been most commonly used in time series analysis [7, 8], but can also be applied to other domains [16]. DTW uses a dynamic programming approach to align two time series and then generates a *warping path* that maps (aligns) the two sequences onto each other. To map two time series T and S, of length n and m respectively, where $T = t_1, t_2, ..., t_n$ and $S = s_1, s_2, ...s_m$ a n-by-m matrix will be formed, where the (i^{th}, j^{th}) grid point corresponds to the alignment or distance between two points t_i and s_j. The warping path, W, is then the set of matrix elements that defines a mapping between T and S, defined as $W = w_1, w_2, ...w_K$, where $max(m, n) \leq K < m + n - 1$. The distance $d(t_i, s_j)$ between two points t_i and s_j is used to identify potential warping paths. There are many distance measure that may be used, the most common is the Euclidean distance, and this is the measure used in this paper. Thus:

$$d(t_i, s_j) = w_k = (t_i - s_j)^2 \tag{1}$$

The minimal warping path is selected by calculating the minimum cumulated distance between T and S as:

$$DTW(T, S) = min \left[\sqrt{\sum_{k=1}^{K} w_k} \right] \tag{2}$$

Figure 2(a) shows an example of the comparison of two time series, the generated warping path is given in Figure 2(b). A perfect match between the two time series will produce a direct diagonal line between the two corners of the grid.

3.3   Image Enhancement

Image enhancement is an important pre-processing step for most image mining applications. Many such techniques are reported in the literature [17, 2, 18, 19]. The most common enhancement in the context of histograms is Histogram Equalisation (HE), commonly used to enhance the contrast of an image. HE "spreads out" the most frequent intensity values to produced a better distributed histogram. Through this transformation, the contrast of an image is improved globally; unfortunately the enhancement may result in some bright parts of the
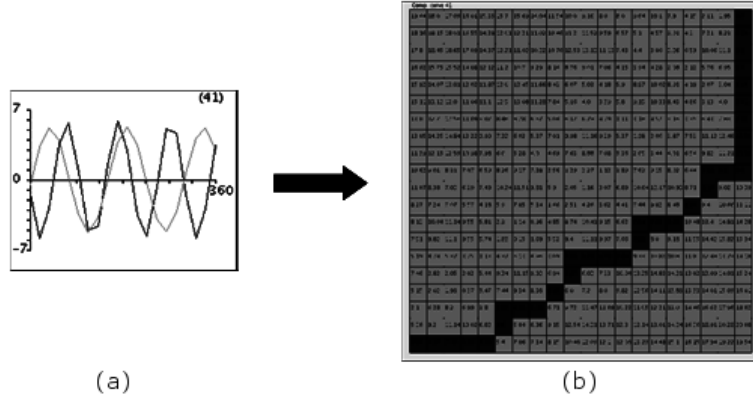
Fig. 2: (a) Two time series for comparison, (b) generated warping path

image being enhanced to the extent that they are "over exposed" and edges become less distinct. To overcome this problem Contrast Limited Adaptive HE (CLAHE) was introduced [19]. CLAHE computes several histograms that correspond to different sections of an image and equalise each histograms independently. Other enhancement technique exploit the correlation of low frequency coefficients with the illumination variation called Adaptive Histogram Equalisation with Rescaled Low Frequency (AHERLF) [17]. Using AHERLF the image contrast was first stretched using HE, and then transformed into a frequency based representation (coefficients) by means of Discreet Cosine Transform (DCT). The low coefficients, which were directly related to the illumination variations, were then rescaled by dividing by a constant. The approach [17] results in a more uniform illumination image with a better visualisation. A comparison of the above techniques is given in section 4.

An approach to combining the green and red channel histograms for retinal image enhancement has been reported in [18]. The Green channel has the highest contrast between retinal objects (blood vessels, fovea and etc.) and the background, while the red channel is much brighter and thus may improve the visualisation of dark area in the green channel image. In [18] the histogram matching was used to modify the Green channel histogram and consequently produced a histogram that displays the advantages of both channels.

4 Eliminating Noise

The quality of data mining results are typically detrimentally affected by the existence of noise in data [20]. In the case of image mining we typically wish to remove features in the input image set that are not considered relevant. With respect to the AMD screening application it was desirable to "remove" blood vessels from the retina image data. The segmentation of retinal blood vessels was conducted using the 2D matched filters proposed in [21].
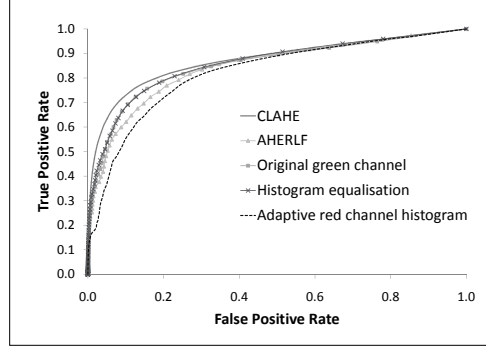
Fig. 3: ROC curves illustrating the performance of the vessel segmentation by using different image enhancement techniques

The identification and removal of noise can best be facilitated by first enhancing the images. The authors conducted a set of experiments to identify the most appropriate of the enhancement techniques identified above, compared with no enhancement, to support the identification of blood vessels in retina images. For evaluation purpose, the Receiver Operating Characteristic (ROC) curve [22] was used (Figure 3). The area under curve that corresponds to the overall performance on retinal vessel segmentation is shown in Table 1. From the table it can be seen that the best performance was achieved using CLAHE, this technique was therefore adopted to enhance the retina images.

Table 1: Area under curve (AUC) on vessel segmentation performance calculated for each enhancement technique on green channel

Image enhancement method	AUC
CLAHE	0.9180
AHERLF	0.8990
Green channel	0.9117
Histogram equalisation	0.9054
Adaptive red channel histogram	0.8941

5   AMD Screening

As noted above DTW provides a technique for comparing two curves. This can be fruitfully adapted for the purpose of data classification following the Case Based Reasoning (CBR) paradigm. Using this paradigm a new case is classified according to its similarity with a set of known pre-classified cases stored in a *Case Base*. With respect to the AMD screening application described here a set of pre-labelled retina images was used to form the "case base". New "unseen"

images could then be classified according to the "nearest match" within the case base, the class of the most similar case found in the case base being the class of the new case.

The proposed AMD image classification process is outlined in Figure 4. To represent each image, the green channel (RGB colour model) and saturation component (HSI colour model) histograms were extracted. The green channel was selected because it displays maximum contrast [9], and has the best discriminatory power between retinal anatomy and the retinal background [21]. In the context of AMD the green channel has been shown to produce good and consistent classification performance compared to other RGB channels and HSI components [16], particularly with respect to "normal" images. The saturation component was chose due to its good performance (in particular to identify AMD images) in AMD classification as shown in [16].

Prior to the generation of the histograms the images were pre-processed. First, CLAHE image enhancement was applied to each colour image to emphasise the contrast of the image and the "visibility" of retinal blood vessel edges. Next, the segmentation of retinal vessels was conducted using 2D matched filters [21] on the green channel. This was followed by the extraction of green channel and saturation component information, which was then represented in the form of histograms. The retinal blood vessel pixels were subsequently removed by subtracting the vessels intensity value from the generated histograms. The 'cleaned' green channel and saturation component histograms, referred to as curve hereafter, formed the "case base" (C) comprising: green ($G = g_0, g_1, ..., g_I$) and saturation ($S = s_0, s_1, ..., s_I$) curves (where I is the number of images and $G, S \in C$). New images to be classified formed a second "case base", $\bar{C}$, comprising: green ($\bar{G} = \bar{g}_0, \bar{g}_1, ..., \bar{g}_J$ and saturation, $\bar{S} = \bar{s}_0, \bar{s}_1, ..., \bar{s}_J$ curves (where J is the number of images and $\bar{G}, \bar{S} \in \bar{C}$). In the classification stage, each curve in $\bar{G}$ is compared with the content of G using DTW. A list of the n most similar g, $sim(\bar{g})$ is produced for each $\bar{g}$:

$$sim(\bar{g}_j) = \left\{ \left(g_0, \delta_0^j \right), ..., \left(g_n, \delta_n^j \right) \right\} \tag{3}$$

$$\delta_n^j = DTW \left(\bar{g}_j, g_i \right) \tag{4}$$

$DTW(\bar{g}, g)$ is the minimal warping path or distance of the green channel "new case" curve, $\bar{g} \in \bar{G}$ and its most similar green channel "case base" curve, $g \in G$ (equation 2), $0 \leq i < I$ and $0 \leq j < J$. Each $\bar{g} \in \bar{G}$ is then classified according to the class associated with its most similar $g \in G$. The classifier will sort the curves in $sim(\bar{g}_j)$ in ascending order of similarity. Where several $g \in G$ have similar δ values a similar distance measuring process is applied to $\bar{S}$ and S. This is undertaken where the distance between two δ values is less than a user specified threshold, $diff$:

$$diff_j = \delta_0^j \times (1 + \alpha) \tag{5}$$

where α is a predefined constant (set to 0.3 in this paper). This step will produce a list of saturation curves distance, $sim(\bar{s})$, for each $\bar{s} \in \bar{S}$ of length m (m is determined by the number of δ that is less than $diff$):

$$sim\left(\bar{s}_j\right) = \left\{\left(s_0, \bar{\delta}_0^j\right), ..., \left(s_m, \bar{\delta}_m^j\right)\right\} \qquad (6)$$

$$\bar{\delta}_m^j = DTW\left(\bar{s}_j, s_i\right) \qquad (7)$$

Finally, we can work out the list of m most similar cases, $c \in C$ for each $\bar{c} \in \bar{C}$ as follows

$$sim\left(\bar{c}_j\right) = \left\{\left(c_0, \mu_0\right), ..., \left(c_m, \mu_m\right)\right\} \qquad (8)$$

$$\mu_m = \frac{1}{2}\left(\delta_m^j + \bar{\delta}_m^j\right) \qquad (9)$$

Each "new case", $\bar{c} \in \bar{C}$ will then be classified as belong to the same class of its most similar curve in the "case base", $c \in C$.
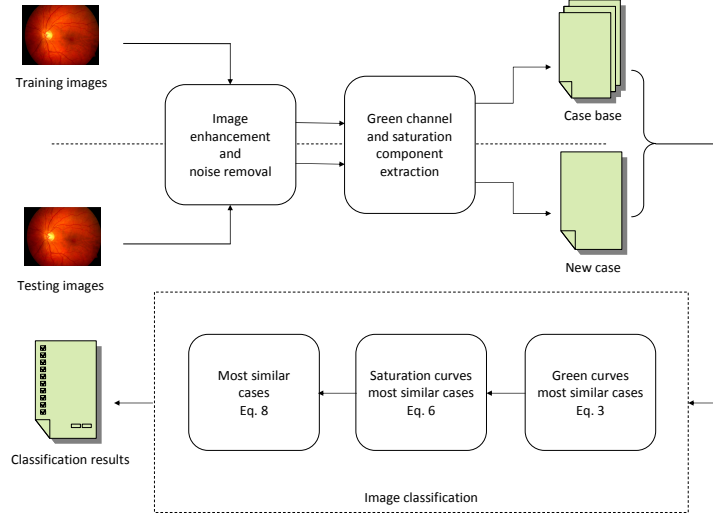

Fig. 4: Proposed AMD screening process

6   Experimental Setup

In this study 144 hand labelled images were used, of which 86 featured AMD that were collected as part of the ARIA project[3]. Ten-fold Cross Validation (TCV) was applied in all experiments (where one tenth of the images were taken as a testing set and the rest for training in each fold). The aims of the experiments were: (i) to investigate the classification performance, using images enhanced with CLAHE and noise removed, against the raw images; (ii) to evaluate the

---

[3] http://www.eyecharity.com/aria_online/

performance of combining two different histograms, green and saturation histograms, on image classification; and (iii) to investigate how well the DTW approach operated with respect to other histogram based approaches (L_1, L_2, and EMD).

The constant α (equations 5) was set to 0.3 (determined through a series of tests), while the number of most similar green curves n (equations 3) was set to 5. Three evaluation metrics were utilised to measure the classification performance: sensitivity, specificity and accuracy. Overall accuracy was used to measure the overall performance of the classifier in terms of classifying retinal images correctly according to their class.

7   Results and Discussions

In this section the results of the three sets of experiments, introduced above, are presented and discussed.

7.1   Performances of Enhanced Images for AMD Classification

Table 2 shows the effect of image contrast enhancement and noise removal, on the classification result. The proposed AMD screening approach, for each green (equations 3) and saturation (equations 6) curve, was applied separately. From the table it can be seen that the application of noise reduction produced superior results, an average improvement of 4% per evaluation metric compared to the raw images for green channel curves, while saturation curves recorded slight improvements only on both the sensitivity and accuracy.

Table 2: Ten-fold Cross Validation (TCV) results of applying image pre-processing on retinal images for AMD classification

	Original images		Pre-processed images	
	Green	Saturation	Green	Saturation
Specificity (%)	57	55	62	55
Sensitivity (%)	69	81	72	83
Accuracy (%)	64	71	68	72

The pre-processing (image enhancement and noise removal) of the retinal images drastically changed most of the histogram curves. Overall, the curves became smoother and more consistent according to their classes, this was particularly so in the case of the normal control images. A more distinctive pattern could also be observed between classes. Enhancement significantly altered the time series analysis output in all cases. It is suggested that the better performance, resulting from the "smoothed" curves, has contributed to a better DTW.

Figure 5 shows examples of the green histogram curves, and the best matching curve, for three selected images (each column represents an image). Two of the images featured AMD, one did not. Figure 5(a-c) presents the "raw" curves

without enhancement. Figure 5(g-i) shows the same curves after the application of the proposed enhancement and noise reduction techniques. From the figure it can be seen that there are significant differences between the two sets of curves.

Applying the proposed DTW technique to the raw images given in Figure 5(a-c) resulted in misclassifications in all cases. The closest matches in each case are presented in Figure 5(d-f). After pre-processing, curves (g) and (i) were classified correctly, although curve (h) remained misclassified. Careful analysis of the results suggests that the application of enhancement and noise reduction techniques to the images has a significant effect on the nature of the histogram curves and consequently on the classification accuracy.
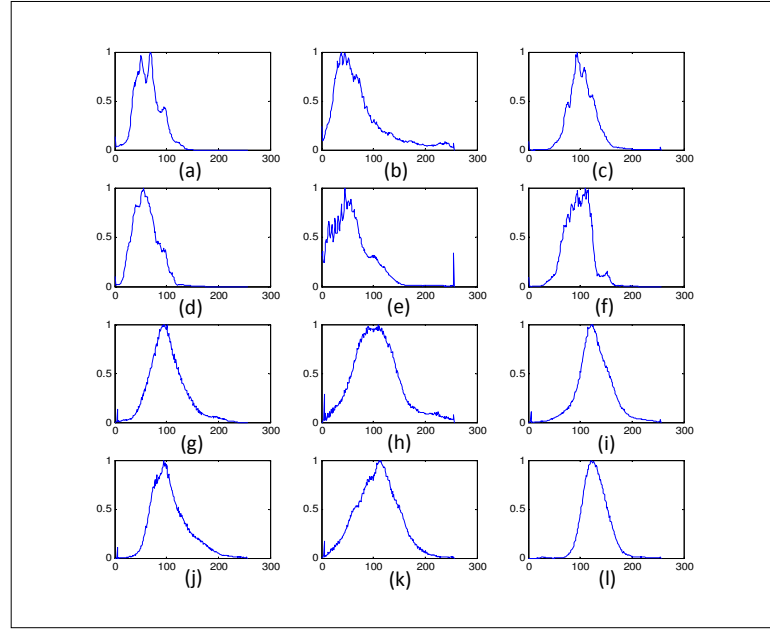

Fig. 5: Sample green channel curves: (a-c) raw image histograms, (d-f) most similar case for curves a to c, (g-i) enhanced and noise removed histograms, (j-l) most similar case for curves g to i

7.2   Performances Using Both Green and Saturation Histograms

Both the green and saturation curves produced good performances. However closer inspection of the results revealed that the difference between the best fit and the second best fit was sometimes marginal, thus calling into question the robustness of the approach. It was also noted that in the case of the green channel histograms the distinction between the best fit and the second best fit tended to be larger than in the case of the saturation histograms (but not in all cases). It was therefore deemed appropriate to combine the evidence of both types of

histogram (the process described in Section 5) so that a more robust classification would result. Recall from Section 5 that the saturation histogram is only considered where the difference between the top two classes using green channel histograms disagree and the difference in similarity values (δ) is less than $diff$. From Table 3 it can be seen that the use of both sets of histograms (where appropriate) produced better results than when using the green histograms alone; and comparable with the saturation histograms results, but with a much greater degree of "confidence" in the end results.

Table 3: Ten-fold Cross Validation (TCV) results of using green and saturation curves for AMD classification

	Green	Saturation	Green and saturation
Specificity (%)	62	55	62
Sensitivity (%)	72	83	79
Accuracy (%)	68	72	72

7.3 Comparison of Performances of Various Histograms Based Image Classification Techniques for AMD Classification

Table 4 presents a comparison of classification accuracy using the proposed DTW technique and L_1, L_2, and EMD to classify retinal images using the pre-processed and original images. TCV was again used throughout. Using DTW' and L_1' with raw images produced a superior specificity compared to others. DTW however performed exceptionally well with respect to sensitivity and overall accuracy (79% and 72% each). The other techniques (L_1, L_2, and EMD), applied to raw and enhanced images, produce mixed results as shown in the Table.

Table 4: Ten-fold Cross Validation (TCV) Results using DTW, L1, L2 and EMD for retinal images classification

	DTW	L_1	L_2	EMD	DTW'	L_1'	L_2'	EMD'
Specificity (%)	62	62	61	55	*64*	*64*	56	52
Sensitivity (%)	*79*	76	74	67	75	76	73	69
Accuracy (%)	*72*	70	69	62	71	71	66	62

The results presented in Table 4 indicate that the ability of DTW to find dissimilarities between two time series by calculating the shortest path between points in the time series data may have contributed to the results produced. Other techniques that calculate point to point distances between two curves are not good enough to classify complex and non-uniform curves, like does used to represent retinal images in this paper. DTW measures the distance between a point in a time series curve to all points in the other time series curve and selects the shortest distance. From Table 4 the best performance recorded was 79% sensitivity.

8   Conclusions

In this paper, an image classification technique, founded on a histogram based representation and DTW was described. The focus of the work was the classification of retinal image data to provide an AMD screening service. In the case of the AMD application the images were represented in terms of their green channel and saturation component histograms. Experiments were conducted to compare the legitimacy of the technique with respect to: (i) image enhancement and noise reduction, (ii) combination of two different histograms, and (iii) other histogram based techniques. From the experimentation it was found that noise reduction (removal of blood vessels in the case of the AMD data) and image enhancement produced better classification results. The results were improved further by using both the green and saturation histograms. The experiments also indicated that the proposed DTW technique, combined with the proposed enhancement strategy, produced the best classification results compared to other histogram based techniques. The best result achieved by the proposed approach was a sensitivity of 79%. It is worth noting that in other work [23], the mean sensitivity achieved through manual graders observation on different sets of retinal images was 86%. For future works, we intend to investigate further noise removal, as well as focusing the screening process on only the central area of the retinal image (the Macula). We are also aware of the issues of inadequate dataset size, therefore, effort is being made to collect more examples of both AMD and normal images.

References

1. Hsu, W., Lee, M.L., Zhang, J.: Image mining: Trends and developments. Intelligent Information Systems **19** (2002) 7–23
2. Ordonez, C., Omiecinski, E.R.: Discovering association rules based on image content. In: IEEE Forum on Research and Technology Advances in Digital Libraries. (1999) 38–49
3. Elsayed, A., Coenen, F., Jiang, C., Garcia-Finana, M., Sluming, V.: Corpus callosum MR image classification. In: Proceedings of AI-2009, Springer London (2009) 333–346
4. Yu, X., Hsu, W., Lee, W.S., Lozano-Peres, T.: Abnormality detection in retinal images (2004)
5. Jager, R.D., Mieler, W.F., Mieler, J.W.: Age-related macular degeneration. The New England Journal of Medicine **358** (2008) 2606–2617
6. Hsu, W., Lee, M.L., Goh, K.G.: Image mining in IRIS: Integrated retinal information system. ACM SIGMOD Record **29** (2000) 593
7. Berndt, D.J., Clifford, J.: Using dynamic time warping to find patterns in time series. In: AAAI Workshop on Knowledge Discovery in Databases. (1994) 229–248
8. Keogh, E.J., Pazzani, M.J.: Derivative dynamic time warping. In: First SIAM International Conference on Data Mining. (2001)
9. Rapantzikos, K., Zervakis, M., Balas, K.: Detection and segmentation of drusen deposits on human retina: Potential in the diagnosis of age-related macular degeneration. Medical Image Analysis **7** (2003) 95–108

10. Al-Aghbari, Z.: Effective image mining by representing color histograms as time series. Journal of Advanced Computational Intelligence and Intelligent Informatics **13** (2009) 109–114
11. Conci, A., Castro, E.M.M.: Image mining by content. Expert System with Applications **23** (2002) 377–383
12. Foschi, P.G., Kolippakkam, D., Liu, H., Mandvikar, A.: Feature extraction for image mining. In: International Workshop on Multimedia Information Systems. (2002) 103–109
13. Brunelli, R., Mich, O.: Histograms analysis for image retrieval. Pattern Recognition Letters **34** (2001) 1625–1637
14. Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover's distance as a metric for image retrieval. International Journal of Computer Vision **40** (2000) 99–121
15. Lin, J., Keogh, E., Lonardi, S., Chiu, B.: A symbolic representation of time series, with implications for streaming algorithms. In: Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery, ACM (2003) 2–11
16. Hijazi, M.H.A., Coenen, F., Zheng, Y.: A histogram approach for the screening of age-related macular degeneration. In: Medical Image Understanding and Analysis 2009, BMVA (2009) 154–158
17. Hossain, M.F., Alsharif, M.R.: Image enhancement based on logarithmic transform coefficient and adaptive histogram equalization. In: International Conference on Convergence Information Technology, IEEE (2007) 1439–1444
18. Salem, N.M., Nandi, A.K.: Novel and adaptive contribution of the red channel in pre-processing of colour fundus images. Journal of the Franklin Institute **344** (2007) 243–256
19. Zuiderveld, K. Academic Press Graphics Gems Series. In: Contrast limited adaptive histogram equalization. Academic Press Professional, Inc. (1994) 474–485
20. Cios, K., Swiniarski, R., Pedrycz, W., Kurgan, L.: 2. In: Data. Springer US (2007) 27–47
21. Chaudhuri, S., Chatterjee, S., Katz, N., Nelson, M., Goldbaum, M.: Detection of blood vessels in retinal images using two-dimensional matched filters. IEEE Transactions on Medical Imaging **8** (1989) 263–269
22. Fawcett, T.: An introduction to ROC analysis. Pattern Recognition Letters **27** (2006) 861–874
23. Jain, S., Hamada, S., Membrey, W.L., Chong, V.: Screening for age-related macular degeneration using nonstereo digital fundus photographs. Eye **20** (2006) 471–475