Proceedings of the 29th Annual International
Conference of the IEEE EMBS
Cité Internationale, Lyon, France
August 23-26, 2007.

Contextual detection of ischemic regions in ultra-wide-field-of-view retinal fluorescein angiograms

E. Trucco, C. R. Buchanan, T. Aslam, B. Dhillon

***Abstract*— We report a novel prototype algorithm using contextual knowledge to locate ischemic regions in ultra-wide-field-of-view retinal fluorescein angiograms. We use high-resolution images acquired by an Optos ultra-wide-field-of-view (more than 200 degrees) scanning laser ophthalmoscope. We leverage the simultaneous occurrence of ischemia with a number of other signs, detected automatically, typical for the state of progress of the condition in a diabetic patient. The specific nature of ischemic and non-ischemic regions is determined with an AdaBoost learning algorithm. Preliminary results demonstrate above 80% pixel classification accuracy against manual annotations.**

I. Introduction and motivation

This paper reports a novel algorithm leveraging contextual knowledge to locate ischemic regions in ultra-wide-field-of-view retinal fluorescein angiograms (FAs).

Our contributions are a context-based detection scheme for retinal ischemic regions in FAs, a context framework based on a pathological (not just anatomical) knowledge, and the use of ultra-wide-field-of-view (UWFV) FAs.

Our main clinical motivation is proliferative diabetic retinopathy (PDR), a major cause of blindness in the developed world. The most prominent theories on its aetiology involve the development of ischemic areas in the retina [15]. These cause specific changes visible in images obtained by fluorescein staining of retinal vessels [4]. Other common conditions (e.g., retinal vein occlusions, arterial occlusions) involve development of ischemic areas. As PDR, they are identified by clinicians in fluosrescein angiography images and can be treated successfully by photocoagulation treatment by laser [5]. For all conditions it is essential to identify appropriate signs of ischemic disease, as laser treatment must be applied at an appropriate time to prevent neovascular complications [12].

Computer analysis and identification of ischemic areas would be useful in a variety of ways. Firstly, objective analysis of fluorescein angiograms can act as a useful marker for trials investigating treatments, e.g., for diabetes and vessel occlusions; current studies are hampered by the lack of objective outcome measures. Secondly, as exact aetiologies are still being investigated, quantitative characterization of retinal features together with accurate recording of patients may provide insight into the pathogenesis of potentially crippling disorders. Finally, in areas where expert ophthalmic opinion is not available, image analysis may help to identify patients at higher risk of developing neovascular complications requiring careful follow-up or laser treatment. Even experienced physicians may benefit from an objective, detailed and comprehensive analysis of fluorescein angiograms to complement their own suspicions and diagnosis.

To our best knowledge, the literature of retinal image analysis includes only a few papers on detecting ischemic and hypoxic retinal regions. A possible reason is that a characterization of ischemic regions in FAs in terms of own image properties, e.g., texture, shape or intensity, proves elusive. In contrast, the appearance of several retinal elements and pathology signs is rather well defined, e.g., the optic disk is generally elliptical; vessels are elongated structures with predictable intensity properties; leakages cause wide, bright areas in FAs. Jasiobedzky et al. [7], [8] segmented FAs frames into regions using morphology, then estimated the degree of perfusion as the local density of the microcapillary net. This was done by texture analysis as the microcapillaries could not be detected reliably. The authors report good agreements with clinician annotations on an unspecified number of FAs. No characterization of the data (e.g., image quality, presence of distracting signs or symptoms) was included. Conrath et al. [3] reported semi-automatic detection of avascular foveal zone (FAZ), which is enlarged in diabetic retinopathy, in FAs. They used 1280×1000 images without any signs possibly interfering with FAZ (exudates, laser scars, haemorrages, etc), and the commercially available ENVI geostatistics package. They concluded that FAZ detection was possible but improvements were necessary in image quality.

Such improvements include instruments enabling observation of the retinal periphery, where several diabetic signs appear first. Our work is based on FAs acquired with an OPTOS ultra-wide-field-of-view scanning laser ophthalmoscope (SLO), capable of field of view of more than 200 degrees at 3072×3900 with 256 grey levels, see Figure 1. [9], [10] are example of work using images from such sensors.

Using such images, we propose a *detection process based on co-occurring pathology signs* to characterize ischemic regions reliably. The basic idea is to leverage the simultaneous occurrence of ischemia with a number of other signs, typical for the state of progress of the condition in a diabetic patient. Our clinical authors have identified a set of signs supporting the hypothesis that a candidate region is ischemic;

E. Trucco is with the School of Computing, University of Dundee, UK, e.trucco@hw.ac.uk

C. R. Buchanan is with the School of Engineering and Physical Sciences, Heriot Watt University, EH14 4AS Riccarton, Scotland, UK and OPTOS plc, Carnegie Business Campus, Dunfermline, cbuchanan@optos.com

T. Aslam and B Dhillon are with the Princess Alexandra Eye Pavilion, Edinburgh, UK, Bal.Dhillon@luht.scot.nhs.uk

This work was supported by Knowledge Transfer Partnership grant KTP000684 of the UK Department of Trade and Industry.

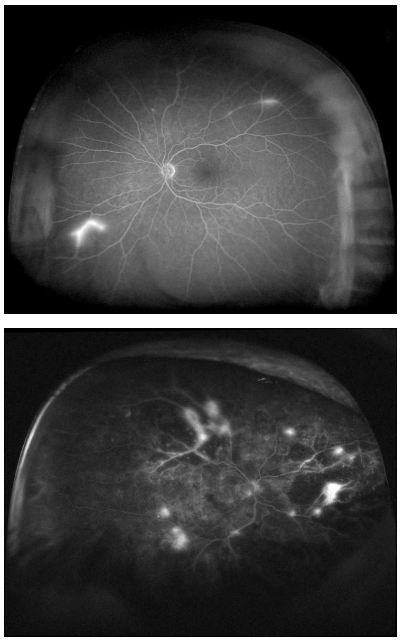

Fig. 1. Two UWFV frames, with (bottom) and without ischemic regions

the problem is then cast as one of combining evidence from multiple classifiers, each dedicated to a single sign. Evidence combination is performed by AdaBoost [13], [16], whereby a reliable committee classifier is built from a set of weak classifiers. We describe our current prototype using four signs (distance from vessel bifurcations, density of vessels, presence of leakages and microaneurysms), and initial results on UWFV FA sequences indicating the reliability of the context-based scheme.

We are not aware of contextual detection of ischemic regions in retinal image analysis. Previous investigations on the use of context for reliable location of retinal elements, e.g., in the presence of seriously degraded images, has concentrated on *anatomical knowledge*, for example the relative position of optic disk, macula and blood vessels [6], [9]. Our definition of context for ischemic regions extends this idea to the co-occurrence of pathological signs, based on *medical knowledge* about the progress of a condition (the set of co-occurring signs vary as the condition progresses). Context-based reasoning has also been reported for computer-assisted diagnosis STARE [11].

II. The context-based algorithm

A. Summary

The algorithm is visualised in Figure 2 and detailed in the following sections. The input is a sequence of FA frames. The output are the co-ordinates and approximate contours of regions classified as ischemic.

Detection is performed on a pixel-by-pixel basis, where we consider the likelihood that a particular pixel is within an ischemic region. This is achieved by a supervised learning algorithm (AdaBoost) using appearance, temporal and contextual features. The algorithm is trained with examples of ischemic and non-ischemic regions.

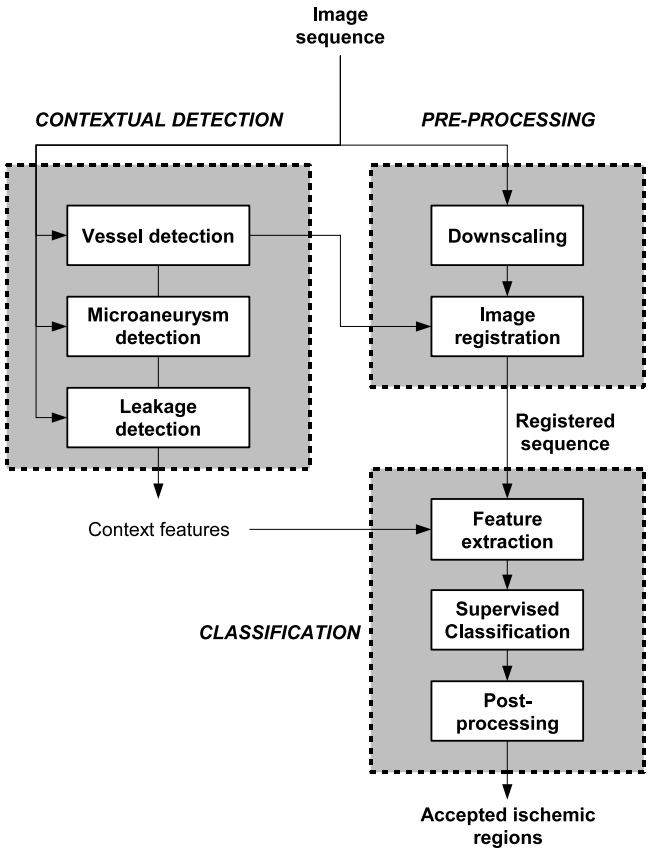

Fig. 2. System architecture.

B. Pre-processing

An FA image sequence $(I_1, \ldots I_n)$ has associated capture times $(T_1, \ldots T_n)$ from injection of the fluorescein dye. Pre-processing consists of vessel network location, downscaling and alignment. A standard, simple illumination normalization is also performed to compensate partially for effects due to patient motion or instrument adjustment.

1) Vessel network location: For each frame $(I_1, \ldots, I_n)$ vessels are located automatically and used for both image alignment and contextual detection. In brief, a vessel mask is constructed for each frame by applying matched filtering to enhance linear structures followed by grey-level thresholding. Bifurcations are located by skeletonisation and application of rules in an 8-connect pixel neighbourhood over the skeleton image [1].

2) Downscaling: To improve efficiency, frames are downscaled by a factor of four using bicubic interpolation (the native resolution of the sensor is 3072×3900, 8 bit/pixel). Fourfold downscaling does not destroy the intended targets (ischemic areas and context). In addition, the impact of sensor noise is reduced.

3) Frame alignement: The frame alignment algorithm considers each pair (I_i, I_{i+1}) of consecutive frames and matches control points (vessel bifurcations) between image pairs. Bifurcations are detected as sketched above. A bifurcation matching algorithm then searches for corresponding

pairs of control points. Correspondences are computed by contextual feature matching; attributes include branching angles, vessel widths, vessel cross-profiles and the displacement vectors to neighbouring bifurcations. Finally, a global affine transform identified by the correspondences is estimated by least squares and applied. Rigid registration provides an approximate solution sufficient for this initial study but which ought to be replaced by a more sophisticated scheme offering improved registration accuracy.

C. Feature definition

From the downscaled frame sequence we extract features of appearance, time and context associated with each pixel.

1) Appearance features: The appearance feature is simply the temporal average of the intensities at a specific pixel over the sequence. The intention is to capture areas that remain consistently dark over an FA session.

2) Temporal: The *time-intensity profile* is the function of pixel intensities over the whole sequence. We expect true ischemic regions to show little increase in intensity in comparison to areas of normal perfusion. To verify this, we consider averages of the time-intensity profile values in five adjacent temporal intervals (bins) spanning the sequence, and check the relative change between adjacent bins.

We observed that capture times vary for each session, i.e., the k-th frame (same position in sequence) of different sessions may be captured at different times from injection. Consequently, before further analysis we must obtain a measure of the signal at certain absolute times for comparison between sessions, and for this reason we perform binning of the frames by time stamp, not by frame number. The chosen time windows are (in seconds): before 30; 30-59; 60-149; 150-299; after 300. Each bin may contain from 3 to 6 frames approximately, depending on the sequence. For each bin b, an average value $x(b)$ is calculated. The result for each pixel is a 5-dimensional feature vector, $x(1),\ldots,x(5)$.

To capture the relative change in intensity between bins we simply calculate the differences between $x(b)$ and $x(b-1)$ for $b = 2,\ldots,b = 5$. As a result, the temporal feature is a 9-dimensional vector (5 bins, 4 differences).

3) Contextual features: The contextual features incorporate information of surrounding vasculature and pathologies; namely the vessel mask, bifurcation points, areas of hyperfluorescent leakage and microaneurysms. The associated detectors cannot be detailed for reasons of space; in essence, vessels and bifurcations are detected as sketched in Section II-B.1; leakages by spatio-temporal classification of intensity profiles [2], and microaneurysms by matched filtering, thresholding, region growing and analysis of size, shape and intensity, following [14].

For each candidate pixel the final 14-dimensional vector feature is a combination of appearance, temporal and context features.

4) Classification: We address this problem as a binary classification task, where a pixel (feature vector) is assigned to one of two classes, ischemic or non-ischemic. Ground truth data with ischemic regions were marked out for training. Accordingly, the training data consists of a set of feature vectors $\mathbf{x}_1,\ldots,\mathbf{x}_N$ with corresponding target labels $t_1,\ldots,t_N$ where $t_j \in \{-1,+1\}$.

AdaBoost [13], [16], an adaptive boosting scheme, is used to combine a series of weak classifiers into a powerful committee classifier. For this study each weak classifier is a decision stump, equivalent to a simple threshold on one of the input variables.

5) Training: During training AdaBoost creates a series of weak classifiers over a number of rounds $m = 1,\ldots,M$. For each round a distribution of weights is updated indicating the importance of examples in the data set where the weights of incorrectly classified examples are increased. In this way the next classifier focuses more on those examples. The result is a series of weak classifiers $h_1,\ldots,h_M$ and corresponding weighting coefficients $a_1,\ldots,a_M$, with greater weights assigned to more accurate classifiers. Effectively we employ AdaBoost to select the most informative criteria with which to distinguish between the two classes given the features supplied.

6) Classifying pixels in novel sequences: To classify pixels in an unseen test sequence, a set of feature vectors $\mathbf{y}_1,\ldots,\mathbf{y}_N$ is extracted, applying the aforementioned methods to detect vessels, microaneurysms and areas of fluorescein leakage. This data is supplied to the trained classifier, which queries each of the weak learners and combines their outputs. The result is a prediction $H(\mathbf{y_i})$ for each test example $\mathbf{y}_i$, where the sign denotes the class as positive or negative:

$$H(\mathbf{y}) = \text{sign}\left(\sum_{t=1}^{T} \alpha_t h_t(\mathbf{y}) - \theta\right) \tag{1}$$

Changing the parameter θ (normally set to zero) makes the prediction more or conservative or more speculative.

7) Post-processing segmentation: Classification results are recorded in a final image (map), where regions of positive predictions are considered as regions of ischemia. Morphological closing is employed to minimize small areas considered to be unreliable.

Accepted regions can be characterised by area and level of severity defined by the magnitude of their score from Eq. (1). Boundary tracing is then applied to the accepted regions and a list of region coordinates scaled to match the size of the input image is the final output.

Notice that we concentrate on reliable detection at this stage, meaning the existence and approximate position of ischemic regions, not on their exact extent and shape. Exact characterization is however already computed by the scheme components, but requires refinements to yield sufficient reliability.

III. Experimental results

All code was written in Matlab on a 1.6 GHz PC running under Windows XP.

For the purpose of this study we acquired five anonymous UWFV FA sequences (three with ischemia and other pathologies and two absent of ischemia). The sequences contain 20

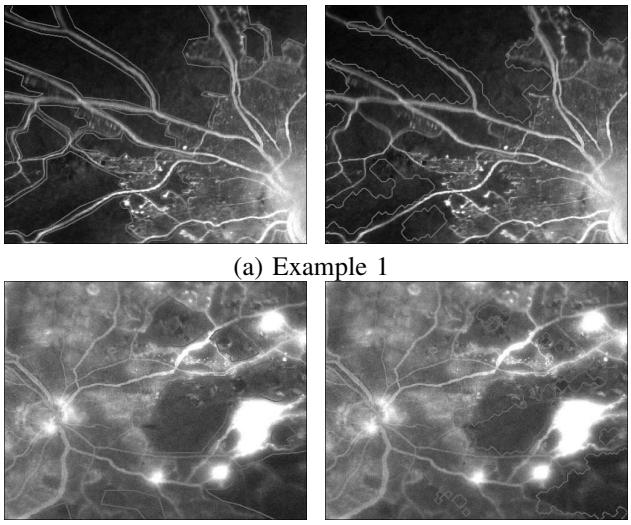

(a) Example 1

(b) Example 2

Fig. 3. Two ischemic examples showing (left) manual annotations, (right) overlay of final predictions

to 40 frames captured over approximately six minutes from injection of the fluorescein dye. Ground truth data for all sessions was obtained by manually tracing ischemic regions on registered frames. Care was taken to avoid including vessels in the ischemic regions.

This dataset was used to train and test the classifier using a *leave-n-out* procedure so that training was performed on four sequences and testing performed on the remaining held-out sequence. Training of the classifier proceeds as described in Section II-C.5. We observed that good generalisation can be achieved by approximately 200 training rounds.

Inspection of the results for the two healthy retinal sequences reveals no false alarms for either sequence even in the presence of leakages and microaneurysms. Two examples of final detections are presented (Figure 3) and compared against ground-truth data. Approximately 80% of the pixels were classified correctly in our experiments; predictions tend to be conservative. Higher sensitivity can be achieved by decreasing the parameter θ in Eq. (1) but this may lead to an unacceptable increase in false positives. In our experiments we used a narrow range values between -1 and 0. In a clinical screening context it may be preferable to decrease the rate of false negatives, as it is better to raise a limited number of false alarms (to be checked by clinicians) than missing some pathology altogether.

We believe some error and the tendency for conservative predictions is largely due to the limited training data in this initial study. We must also consider that ground truth boundaries are rather subjective. Other errors are due to image acquisition (e.g., incorrect patient position, eyelids).

IV. Discussion and future work

We have presented a context-based algorithm detecting ischemic regions in UWFV fluorescein angiograms. Our main contributions are a context-based detection scheme for retinal ischemic regions in FAs, a definition of context based on a pathological (not just anatomical) knowledge, and the use of UWFV FAs, not yet common in the literature but very promising as showing large parts of the retinal periphery. Experimental results suggest that our scheme is a very promising solution for the detection of ischemic regions, the image appearance of which, *per se*, is not easily characterized. Further, more extensive tests are planned to assess the power of our scheme, for which purpose a higher number of UWFV FA sequences are being collected.

This work opens up a wealth of future investigations. These include improvements to algorithm components, e.g., registration and more contextual pathology signs, as well as more fundamental issues like a comparison of medical and automatic judgements about the relative importance of contextual pathology signs.

References

[1] H. Azegrouz and E. Trucco. Max-min central vein detection in retinal fundus images. *Proc. IEEE Int Conf on Image Processing (ICIP'06)*, pages 1925–1928, 2006.

[2] C. Buchanan and E. Trucco. Automated detection of hyper-fluorescent leakage in retinal angiographic sequences by temporal analysis. In *Proc. IEEE Int. Conf. on Image Processing (ICIP 2007), to appear*, 2007.

[3] J Conrath, O Valat, R Giorgi, M Adel, and D Raccah. Semi-automated detection of the foveal avascular zone in fluorescein angiograms in diabetes mellitus. *Clinical and Experimental Ophthalmology*, 34:119–123, 2006.

[4] Early Treatment Diabetic Retinopathy Study Research Group. Fluorescein angiographic risk factors for progression of diabetic retinopathy. *Ophthalmology*, 98:834–840, 1991.

[5] The Diabetic Retinopathy Study Research Group. Preliminary report on effects of photocoagulation therapy. *Am. Journ. Ophthalm.*, 81:383–396, 1976.

[6] A. Hoover and M. Goldbaum. Locating the optic disk in a retinal image using the fuzzy convergence of blood vessels. *IEEE Trans. on Medical Imaging*, 22(8):951–8, 2003.

[7] P. Jasiobedzki, D. McLeod, and C.J. Taylor. Detection of non-perfused zones in retinal images. In *Proc. Fourth Annual IEEE Symposium Computer-Based Medical Systems*, pages 162–169, 1991.

[8] P. Jasiobedzki, C.J. Taylor, and J.N.H. Brunt. Automated analysis of retinal images. *Image and Vision Computing*, 11(3):139–144, 1993.

[9] P. Kamat and E. Trucco. Locating the optic disk in retinal images via plausible detection and constraint satisfaction. In *Proc. IEEE Int Conf on Image Processing (ICIP'04)*, volume CD 1. IEEE, 2004.

[10] A Manivannan, J Plskova, A Farrow, S McKay, P F Sharp, and J V Forrester. Ultra-wide-field-of-view fluorescein angiography of the ocular fundus. *American Journ of Ophthalmology*, 140:525–527, 2006.

[11] R. Pai, A. Hoover, and M. Goldbaum. Automated diagnosis of retinal images using evidential reasoning. In *15th Int Conf on Systems Engineering*, 2002.

[12] The Central Vein Occlusion Study Group N report. A randomized clinical trial of early panretinal photocoagulation for ischemic central vein occlusion. *Ophthalmology*, 102:1434–1444, 1995.

[13] R.E. Schapire. The boosting approach to machine learning: An overview. *MSRI Workshop on Nonlinear Estimation and Classification*, 2002.

[14] T. Spencer, J.A. Olson, K.C. McHardy, P.F. Sharp, and J.V. Forrester. An Image-Processing Strategy for the Segmentation and Quantification of Microaneurysms in Fluorescein Angiograms of the Ocular Fundus. *Computers and Biomedical Research*, 29(4):284–302, 1996.

[15] D. Verma. Pathogenesis of diabetic retinopathy: the missing link? *Medical Hypotheses*, 41:205–210, 1993.

[16] P. Viola and M. Jones. Robust real-time face detection. *Int. Journ. of Computer Vision*, 63(2):137–154, 2004.