A HIERARCHICAL BAYESIAN CLASSIFICATION FOR NON-VASCULAR LESIONS DETECTION IN FUNDUS IMAGES

E. GRISAN* and A. RUGGERI*

* Department of Information Engineering, University of Padova, Italy

enrico.grisan@dei.unipd.it

Abstract: Due to its blood microcirculation, the retina is one of the first organ affected by hypertension and diabetes: retinal damages can lead to serious visual loss , that can be avoided by an early diagnosis. The most distinctive sign of diabetic retinopathy or severe hypertensive retinopathy are haemorrhages and microaneurysms (HM), hard exudates (HE) and cotton wool spots (CWS). Automatic detection of their presence in the retina is thus of paramount importance for assessing the presence of retinopathy., and therefore relieve he burden of images examination by retinal experts.

In this work we propose a simple and effective method to detect and identify these lesions in retinal images, by a two stage classifier. By considering a pixel-wise classification at the first stage and an object-wise classification at the second, it impose a hierarchy (or a multi-level) geometry to the classifier.

By using a Bayesian MAP classifier for the first stage and a simple linear discriminant for the second, it proved to achieve a sensitivity of 0.83, 0.71, 0.73 for the classes HE, HM, CWS respectively, with a specificity of 0.94, 0.99, 0.91. The area of the identified lesions shew a correlation with the clinical severity grading of 0.87, 0.83, 0.78.

Introduction

Retinopathies associated with systemic diseases such as diabetes and hypertension have an ever increasing importance as a cause of blindness and visual loss. Diabetic retinopathy is the first cause of blindness for people in working age in the United States, with all the consequent economic and social burdens.

The timely diagnosis and referral for management of diabetic retinopathy can prevent 98% of visual loss. Currently, a periodic dilated direct ophthalmoscopic examination seems the best approach for a screening with near universal coverage of the population at risk, despite the proved low sensitivity of direct ophthalmoscopy [1]. However, the number of ophthalmologist available is the limiting factor in initiating an ophthalmologist based screening [2].

With the increasing availability of digital fundus camera an automatic tool to recognize and grade non-vascular lesions in fundus images is highly required to relieve the burden of retina examinations from retina experts, and to envision the possibility of population screening for retinopathies.

The more distinctive lesions are hard exudates, cotton wool spots, haemorrhages and microaneurysms. Haemorrhages and microaneurysms can be grouped together, since they are often distinguishable only by their size. Automatic identification of lesions is made more difficult by the presence of luminosity variability within the retinal images, and by retinal pigmentation heterogeneity.

Several algorithms were proposed to identify microaneurysms in colour fundus images with various reported performance [3, 4, 5]. To the best of our knowledge, the highest overall performance was a sensitivity of 0.9 and a specificity of 0.8.Others focused on exudative lesions detection [6, 7] with the best performance obtaining a sensitivity of 0.92 and specificity 0.82.

Few attempts to identify different type of lesions within one algorithm have been reported. In [8], the lesions were divided into dark (HM) and bright (HE and CWS), yielding a sensitivity for bright lesions of 0.89 and for dark lesions of 0.78. In [9] both microaneurysms and exudates are searched in retinal colour images. Microaneurysms were identified with asensitivity and specificity reported of 0.77 and 0.88 respectively, whereas exudate detection obtained a sensitivity of 0.88 and a specificity of 0.99.

No reported work has tried to identify a set of lesions larger than microaneurysms and hard exudates. The only mention to classification of hemorrhages, microaneurysms, exudates and cotton wool spots is in [10], where different classification strategies are compared on manually segmented lesions.

We propose a method that makes use of the illumination corrected retinal images described in [11]. By imposing the first two moments to the image histograms, the colour content is an effective feature for the determination of the abnormal regions of the retina. Both normal fundus and lesions can only be described statistically, since their colour content is affected by noise and by an intrinsic variability of different retinas, and of different stage of evolution of the lesions.

In order to cope with this variability, a two stage approach has been devised. The first stage has the aim of segmenting the candidate lesions from the normal fundus, assign-

ISSN: 1727-1983

ing to each pixel a probability of being lesion or retinal background. This stage should be very sensitive, relaxing the requirements of specificity. The second stage, corresponding to a higher level of abstraction, classifies the candidate regions, identified in the first stage, into the lesions of interest. This is achieved through a supervised classifier, which looks at region-wide features, instead of pixel-wise.

Materials

From a pool of 50^o fundus images acquired on slide, and subsequently digitized with a 1378 dpi resolution, true colour, 60 images were chosen by three ophthalmologists, to be used as a training set (DB60). Among these images, 10 presented some non-vascular lesions: 6 HM, 7 CWS, 6 HE.
A second set (DB200) of 200 images were chosen by an ophthalmologist, to be used for validating the proposed method. In this set, 47 showed non-vascular lesions: 33 HM, 22 CWS, 31 HE.
All images were graded, with a separate severity score for the lesions considered (HM, CWS and HE).
We suppose that retinal vessels and optic disc have been previously extracted by one of the available algorithms (e.g. [12, 13] for vessels and [14, 15] for optic disc), so that the region of the retina covered by vessels and optic disc can be excluded from the non-vascular identification algorithm without processing.

Methods

Bayesian Classification of Suspect Areas

This is the first stage of the lesion identification procedure. In order to provide a high sensitivity to the lesions of interest (classes), the fundus image can be divided into three types of areas. Those that appear darker than the normal fundus ω_1, which correspond to possible hemorrhages and microaneurysms, those that appear brighter than the retinal background ω_3, which correspond to possible hard exudates and cotton wool spots, and those belonging to the normal fundus ω_2. The colour content of each pixel of the image is assumed to be sufficient to identify the three classes. Given the image I, the colour content of the pixel at position (x,y) can be described by the vector $\mathbf{f}$ containing the intensity values of the three colour channel:

$$\mathbf{f}(x,y) = [I_{red}(x,y), I_{green}(x,y), I_{blue}(x,y)] \qquad (1)$$

Given the a priori probability $P(\omega_i)$, $i = 1,2,3$, of the three types of regions, and the conditional probability densities $p(\mathbf{f}|\omega_i)$, describing the statistical colour distribution for the three classes, the probability that a pixel $I(x,y)$ belongs to the class ω_i can be computed by means of the Bayes theorem:

$$P(\omega_i|\mathbf{f}) = \frac{p(\mathbf{f}|\omega_i)\cdot P(\omega_i)}{p(\mathbf{f})} \qquad i = 1,2,3 \qquad (2)$$

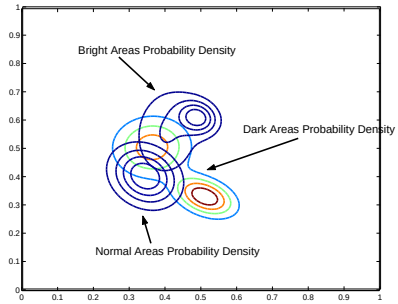

(a)

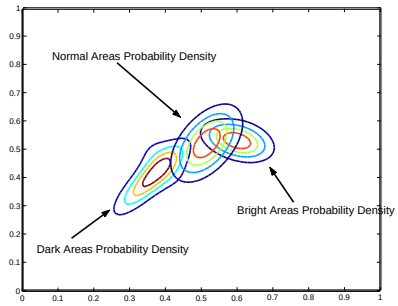

(b)

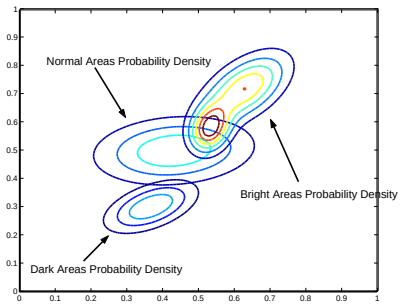

(c)

Figure 1: Slices of the probability density functions of the three classes of interest

$$p(\mathbf{f}) = \sum_{i=1}^{3} (p(\mathbf{f}|\omega_i)\cdot P(\omega_i)) \qquad (3)$$

This formulation allows the straightforward description of a space variant a priori probability. For each pixel position (x,y), the probability will be $P(\omega_i,x,y)$. Previously identified structures, as vessels and optic disc, can then be included in (2) through the *a priori* probability, forcing the pixels belonging to them to be considered normal fundus, excluding them from possible abnormal areas.
After the three conditional *a posteriori* probabilities $P(\mathbf{f}|\omega_i)$ are evaluated, a pixel is assigned to one of the three classes with a maximum a posteriori (MAP) procedure. A set of ten images, not belonging to the DB60 nor to the DB200, have been used to estimates the probabilities and the probability densities needed to evaluate the *a posteriori* probabilities of (2). The lesions present in these images were manually segmented, and the intensity values of the three colour channels of the identified pixels have been pooled to estimate $p(\mathbf{f}|\omega_i)$. The colour probability density for each class has been modelled with a multivariate mixture of Gaussian. Normal fundus have

ISSN: 1727-1983

shown a monomodal density, thus it is described by a single Gaussian, whereas both dark and bright classes shew a bimodal density, thus modelled with the combination of two Gaussian curves. Since significant correlation was found among the colour channels for the three classes, a full covariance matrix is used in the gaussian mixture model. The a priori probabilities $P(\omega_i)$ have been empirically set. Since lesions are uncommon but in the more severe stages of retinopathy, the occurrence (relative frequency) of lesion pixels in the images would be excessively low, providing an unsatisfactory sensitivity:

$$P(\omega_1) = 0.01 \quad (4)$$
$$P(\omega_2) = 0.98 \quad (5)$$
$$P(\omega_3) = 0.01 \quad (6)$$

In Fig. 1 three slices along the blue component plane of the probability density functions are shown, for increasing values of the blue component.

Blobs Analysis and Features Extraction

The result of the procedure described in the previous section is a sparse set of pixels for the two classes of interest, that not convey any information on spatial relationship or global object features. Linking single pixels to the lesion they belong to would not be an issue if the proposed algorithm were 100% specific. However, since pixel-wise classification is not very specific, a refinement is needed: this is done by introducing a further level of abstraction. Every set of connected pixels with the same classification represent a separate object, that is a candidate lesion.
Small set of pixels classified as candidate lesions were excluded from the subsequent analysis. The ratio for doing this is that isolated pixels are more likely to be noise overlapped to normal fundus, and then misclassified, rather than tiny lesions. A pair of cascaded morphological operators well suit the task, providing both the removal of isolated areas of dimension comparable to the morphological kernel, and also the filling of small holes in connected regions. After the pixel classification has been cleaned by this filtering, all connected components are extracted separately, and each represent a candidate lesion.
In order to distinguish the different lesions, and the lesions from false positives, is necessary to provide a description of every identified region, to provide the most discriminating factors. This is not a trivial task, given the extreme variability in appearance of the same type of lesions (intra-class variability), and the close similarity of lesions to variations in pigmentation or to other lesions (inter-class variability).
The regions can be described roughly by means of four types of features. The first are chromatic features, which statistically describe the colour aspect of the region.To enhance the chromatic information, the HSV colour space in addition to the RGB available intensities has been used. This is supposed to be the most discriminant space for perceptually different colours [16]. The second set of features provides a geometrical characterization of the region. It evaluates both perimeter and area features, in order to evaluate the compactness, but also aspect ratios, principal axes, eccentricity. Also the standard deviation of the radius is evaluated: given the region centroid, the variation of the distance between the region border and the centroid for every perimeter point, should give a measure of the smoothness of the region borders and of its roundness.
Then, on one hand the region should have a definite differentiation from the neighborhood, on the other the characteristics of features on the border may provide a hint on the particular lesion to be classified. Three areas are defined in each candidate lesion: the internal part B_i, the internal border ∂B_i and the external border ∂B_o, and they are shown in Fig. 2. They are easily defined using the morphological operations of dilation and erosion, with a kernel of suitable dimension. B_i is simply the eroded version of the original region B, the internal border is what remains by subtracting form B the internal part B_i, and the external border is what remains after the subtraction of B from its dilation.
The fourth set of features is composed by colour gradi-

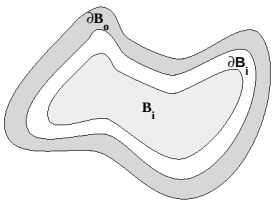

Figure 2: A region divided into its internal part B_i, internal border ∂B_i and external border ∂B_o

ent of the regions. The gradients are computed by looking at intensity variation between growing region perimeters, starting from the centroid. Defining a sequence of boundaries by iteratively removing all perimeter points from the region, a sequence of vectors $\{\mathbf{n}_1 \dots \mathbf{n}_j\}$ normal to each boundary is computed. The defined gradient is therefore the intensity variation along the normal direction from the region centroid to its boundary. Sharp margins and soft margins are likely to be identified via a combination of internal border and external border features.

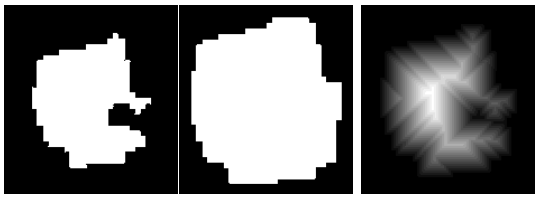

(a) Original region (b) Dilated Region (c) Border Sequence

Figure 3:

ISSN: 1727-1983

Feature Selection

There are several reasons that suggest to keep as low as possible the number of features with which classify the candidate regions. The first is computational complexity. Another is that the increase in complexity in computing a larger number of features is not always matched by an increase in discriminatory power, because of the possible redundancy and correlation among features. The most important reasons lie however in the generalization power of a classifier. Since the number of features can be considered as the number of free parameters of the classifier, the smaller the ratio between this number and the cardinality of the training set, the best the performance on the training set but the less robust the classifier.
Given the large number of features that can be computed in the framework of the previous section, there is the need to select the most significant ones, so as to maintain the number of features used by the classifier small , and at the same time keep the discriminating power elevated.
In order to perform this feature selection, it is necessary to have a measure of the discrimination ability of a set of features. The most common measures involve the evaluation of the intra-class scatter matrix S_w and the inter-class scatter matrix S_b. Given M classes, a set of features vectors **x**, each one belonging to one and only one class, the intra-class scatter matrix is:

$$S_w = \sum_{i=1}^{M} P(\omega_i)\Sigma_i \tag{7}$$

$P(\omega_i)$ is the a priori probability of the class ω_i, and Σ_i is its covariance matrix:

$$\Sigma_i = E\left[(\mathbf{x}-\mu_i)(\mathbf{x}-\mu_i)^T\right] \tag{8}$$

The inter-class scatter matrix is:

$$S_b = \sum_{i=1}^{M} P(\omega_i)(\mu_i-\mu_0)(\mu_i-\mu_0)^T \tag{9}$$

with μ_0 being the global mean:

$$\mu_0 = \sum_{i=1}^{M} P(\omega_i)\mu_i \tag{10}$$

Various combination of these two matrices have been proposed to evaluate at the same time the closeness of samples coming from the same class, and the separation of the classes. The main disadvantage of these criteria is that they do not have any relationship with the Bayesian classification error. On the contrary, the probability of error in classifying a point between two classes ω_1 and ω_2 is bounded by the Chernoff bound:

$$\varepsilon \leq P(\omega_1)^{1-s}P(\omega_2)^s \int p(x|\omega_1)^{1-s}p(x|\omega_2)^s dx \tag{11}$$

$$s \in [0,1] \tag{12}$$

The optimum value of s, which provides the minimum error, is not easy to find analytically, so that the value $s=0.5$ is often chosen to provide an upper bound, even if not the lowest [17]. In this case, the value $\mu(\frac{1}{2})$ is called the Bhattacharayya distance, and it can be used as a criterion for class separability. For two normal distribution with mean μ_i and covariance Σ_i, $i=1,2$, the distance can be evaluated in closed form:

$$\begin{aligned}\mu(\tfrac{1}{2}) &= \tfrac{1}{8}(\mu_1-\mu_2)'\left(\frac{\Sigma_1+\Sigma_2}{2}\right)^{-1}(\mu_1-\mu_2) \\ &+ \quad \tfrac{1}{2}\ln\left(\frac{|0.5(\Sigma_1+\Sigma_2)|}{|\Sigma_1|^{1/2}+|\Sigma_2|^{1/2}}\right)\end{aligned} \tag{13}$$

This distance is used to iteratively pick the most discriminating feature among in the set of available features, given the features already chosen. In this way 29 features have been chosen to build the classifier for the bright classes and 30 for the classification of dark lesions.

Linear Classifier

The classifier that is build from the training set and the chosen features is based on the Linear Discriminant Analysis (LDA). Given a certain numbers of classes with supposedly different characteristics, LDA is a method for linearly mapping the high dimensional characteristics vector in a lower dimensional space, which maximize the separation between classes, supposing their distribution normal. LDA is based on the maximization of a function J that is an indicator of the class separation. Given N samples, M classes and a 1xm vector f of features, the function J considered is:

$$J_1 = trace(S_w^{-1}S_b) \tag{14}$$

with S_w and S_b defined by (9) and (7) respectively. The mxm matrix $S_w^{-1}S_b$ has rank $M-1$. A transformation matrix preseserving the value of J is the matrix **C** having on its columns the (M-1) non trivial eigenvectors of $S_w^{-1}S_b$.

Blob Classification

In the case of lesion classification, both dark and bright lesions may belong to one of three classes to be distinguished. Bright lesions have to be separated in hard exudates, cotton wool spots and false positives. Dark lesions have to be distinguished in hemorrhages, vessels and false positives. Vessel have been incorporated in the classification of dark region separately from false positives, because of their features, separate from that of retinal background: consider them in a unique class together with the false positives would have resulted in poor classification performances.
Given the vector $\mathbf{f}_{ij}$ of the features belonging to blob j classified as class i, the transformed features vector is obtained through the linear transformation:

$$\mathbf{y}_{ij} = \mathbf{C}\mathbf{f}_{ij} \tag{15}$$

For each class ω_i of numerosity N_i, the sample mean and covariance matrix, μ_i and Σ_i, are evaluated from the transformed features of the training set:

$$\mu_i = \tfrac{1}{N_i}\textstyle\sum_{j=1}^{N_i}\mathbf{y}_{ij} \tag{16}$$

ISSN: 1727-1983

Table 1: Results for identification of images with non vascular lesions in DB60

	HE	CWS	HAE
True Positives	5	4	3
False Positives	1	1	1
False Negatives	1	2	2
Sensitivity	0.83	0.66	0.6
Specificity	0.98	0.98	0.98

$$\Sigma_i = \frac{1}{N_i - 1} \sum_{j=1}^{N_i} (\mathbf{y}_{ij} - \mu_i)'(\mathbf{y}_{ij} - \mu_i) \quad (17)$$

These means and covariances are used to evaluate the Mahalanobis distance from each class ω_i of the transformed features vector $\mathbf{y}$ of an unknown region:

$$d_i = (\mathbf{y} - \mu_i)' \Sigma_i^{-1} (\mathbf{y} - \mu_i) \quad (18)$$

The region has the highest probability of belonging to the class with the minimum distance, and the probability is proportional to the distance d_i

(a) Original image with various non vascular lesions

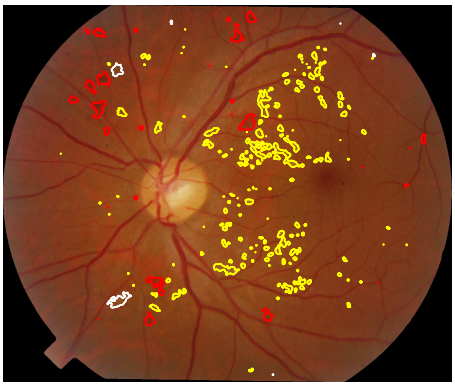

(b) Identified lesions: hemorrhages (red), cotton wool spots (white), and hard exudates (yellow)

Figure 4: Lesion detection and classification in a sample image.

Results

The results of the algorithm on the images belonging to DB60 database are summarized in Tab. 1 and in Tab. 2 for those belonging to DB200. Since the evaluation of sensitivity and specificity on single lesion requires a carefully identified ground truth, it has not been possible to perform such an evaluation. The performance evaluation is therefore on the ability of the algorithm to correctly identify whole images with particular lesions.
Nevertheless, the ability of the algorithm to estimate the severity of the non-vascular lesions can be evaluated by looking at the correlation between the estimated area of the retina covered by specific lesion, and the related ophthalmologist grading. This can be quantitatively assessed with the measure of correlation, reported in Tab. 3, and visually evaluated in Fig. 5, Fig. 7 and Fig. 6

Table 2: Results for identification of images with non vascular lesions in DB200

	HE	CWS	HAE
True Positives	22	15	22
False Positives	1	17	1
False Negatives	7	7	9
Sensitivity	0.75	0.68	0.71
Specificity	0.99	0.90	0.99

Table 3: Correlation between the estimated area covered by lesions and the ophthalmologist grading, on the images of the DB60

	HE	CWS	HAE
Correlation	0.87	0.78	0.83

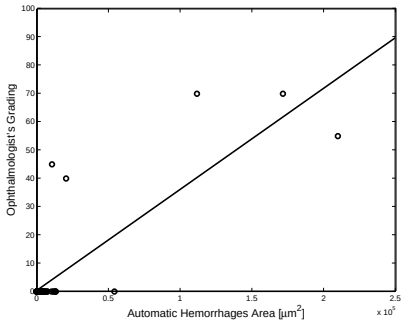

Figure 5: Total area of the identified hemorrhages versus ophthalmologist grading

Conclusions

A new method for identifying and distinguish between retinopathy lesions has been presented. Its simplicity makes it appealing when very fast computation is required.
On the side of excellent specificity and good sensitivity, the method proved its ability to correctly grade the severity of the lesions in the retina,showing a good correlation with the ophthalmologist perception.
The lower sensitivity among the lesions is obtained for hemorrhages. This is due by the difficulty in distinguish-

ISSN: 1727-1983

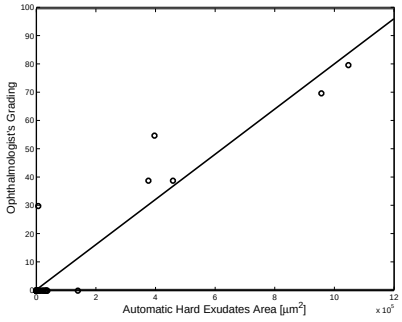

Figure 6: Total area of the identified hard exudates versus ophthalmologist grading

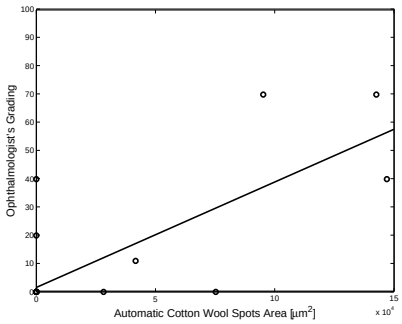

Figure 7: Total area of the identified cotton wool spots versus ophthalmologist grading

ing faint haemorrhages from pigmentation variation or shades by the first stage Bayesian classifier.

References

[1] H. M. HERBERT, K. JORDAN, and D. W. FLANAGAN. Is screening with digital imaging using one retinal view adquate? *Eye*, 17(3):497–500, 2003.

[2] L. VARMA, G. PRAKASH, and H. K. TEWARI. Diabetic retinopathy: time for action. no complacency, please! *Bullettin of the World Health Organization*, 80:419, 2002.

[3] G. E. ØIEN and P. OSNES. Diabetic retinopathy: Automatic detection of early symptoms from retinal images. In *Proceedings NORSIG-95 Norwegian Signal Processing Symposium*, September 1995.

[4] J. H. HIPWELL, F. STRACHAN, J. A. OLSON, K. C. MCHARDY, P. F. SHARP, and J. V. FORRESTER. Automated detection of microaneurysms in digital red-free photographs: a diabetic screening tool. *Diabetic Medicine*, 17(8):588–594, August 2000.

[5] G. YANG, L. GAGNON, S. WANG, and M. C. BOUCHER. Algorithm for detecting Micro-Aneurysms in low resolution color retinal images. In *Proceedings of Vision Interface 2001, Ottawa*, pages 265–271, June 2001.

[6] A. OSAREH, M. MIRMEHDI, B. THOMAS, and R. MARKHAM. Automatic recognition of exudative maculopathy usinf fuzzy C-Means clustering and neural networks. In *On-line Proceedings MIUA 2001*, 2001.

[7] H. WANG, W. HSU, K. G. GOH, and M. L. LEE. An effective approach to detect lesions in color retinal images. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2000, June.

[8] M. GOLDBAUM, S. MOEZZI, A. TAYLOR, S. CHATTERJEE, J. BOYD, E. HUNTER, and R. JAIN. Automated diagnosis and image understanding with object extraction, object classification, and inferencing in retinal images. In *1996 IEEE International Conference on Image Processing*, pages 695–698, 1996.

[9] C. SINTHANAYOTHIN, J. F. BOYCE, T. H. WILLIAMSON, H. L. COOK, E. MENSAH, and D. USHER. Automated detection of diabetic retinopathy of digital fundus images. *Diabetic Medicine*, 19(2):105–112, February 2002.

[10] B. M. EGE, O. K. HEJLESEN, O. V. LARSEN, K. MØLLER, B. JENNINGS, D. KERR, and D. A. CAVAN. Screening for diabetic retinopathy using computer based image analysis and statistical classification. *Computer Methods and Programs in Biomedicine*, 62:165–175, 2000.

[11] M. FORACCHIA, E. GRISAN, and A. RUGGERI. Luminosity and contrast normalization in retinal images. *Medical Image Analysis*, 9(3):179–190, June 2005.

[12] A. HOOVER, V. KOUZNESTOVA, and M. GOLDBAUM. Locating blood vessels in retinal images by piece-wise threshold probing of a matched filter response. *IEEE Transactions on Medical Imaging*, 19(3):203–210, March 2000.

[13] E. GRISAN, A. PESCE, A. GIANI, M. FORACHHIA, and A. RUGGERI. A new tracking system for the robust extraction of retinal vessel structure. In *The 26th International Conference of the IEEE Engineering in Medicine and Biology Society*, September 2004.

[14] HUIQI LI and O. CHUTATAPE. A model-based approach for automated feature extraction in fundus images. In *Proceedings of the Ninth IEEE International Conference on Computer Vision (ICCV 2003)*, pages 394–399. IEEE Computer Society, October 2003.

[15] T. WALTER, J.-C. KLEIN, P. MASIN, and A. ERGINAY. A contribution of image processing to the diagnosis of diabetic retinopathy - detection of exudates in color fundus images of the human retina. *IEEE Transactions on Medical Imaging*, 21(10):1236–1243, October 2002.

[16] G. PASCHOS. Perceptually uniform color spaces for color texture analysis: an empirical evaluation. *IEEE Transactions on image processing*, 10(6):932–937, June 2001.

[17] K. FUKUNAGA. *Introduction to statistical pattern recognition*. Academic Press, New York and London, 1972.

ISSN: 1727-1983