Automatic detection of the presence of bright lesions in color fundus photographs

M. Niemeijer*, M.D. Abràmoff** and B. van Ginneken*

* Image Sciences Institute, University Medical Center Utrecht, Utrecht, the Netherlands
** Department of Ophthalmology and Visual Sciences, The University of Iowa, Iowa City, USA

meindert@isi.uu.nl

Abstract: A novel scheme for the detection of bright lesions in digital color photographs is presented. A supervised system is used to assign a soft label to each pixel in a retinal image indicating the probability that it is a bright lesion pixel. Then, through a classifier combination scheme, the probability that bright lesions are present anywhere in the image is obtained. To counter spurious responses on certain retinal landmarks such as the optic disc, prior knowledge about the location of the major anatomical landmarks on the retina is integrated into the system. Different system setups are compared. The best performing system obtains an area under the ROC curve of 0.88.

Introduction

Screening for Diabetic Retinopathy (DR) can help prevent blindness and vision-loss in patients suffering from diabetes. Many screening programs use color fundus photographs, which allow imaging of the retina in a non invasive manner, for the detection of the symptoms associated with DR (see Figure 1). However, manual screening of large volumes of color fundus photographs by ophthalmologists is time consuming and expensive. Because approximately 90% of all imaged patients in a screening setting show no signs of the presence of DR [1], it would be interesting to develop a pre-screening system which will select the "suspect" images. Only these images would then be examined by an ophthalmologist.

The automatic DR pre-screening system should be able to recognize the characteristic signs of the presence of the disease. DR is a progressive disease characterized by different types of symptoms that manifest themselves on the retina. The first unequivocal sign of the presence of DR are the appearance of microaneurysms, visible as small red dots between the retinal vasculature. Next, larger hemorrhages appear and these are visible as larger red blots or smears between or connected to the vasculature. Finally, so called bright lesions appear (see Figure 2).

Previously, we published an algorithm which is able to detect the red lesions, i.e. microaneurysms and hemorrhages, in fundus images [2]. A successful system should also be able to detect the presence of bright lesions and this detection task will be the focus of this work. Screening data can contain three types of bright lesions. Two of these are related to DR, exudates and cottonwool spots, and one, drusen, is not. Exudates are intraretinal deposits of serum lipoproteins which leak from the vasculature. They usually have a high contrast with respect to the surrounding tissues. Cottonwool spots are ischemic areas of the retina that are colored brighter due to a lack of perfusion, they are normally larger than the exudates and have a low contrast with the background. Drusen are bright lesions which can have a high contrast and are usually clustered in larger groups. Examples of all three different types of lesions are shown in Figure 2. In this work we do not distinguish between the different types of bright lesions.

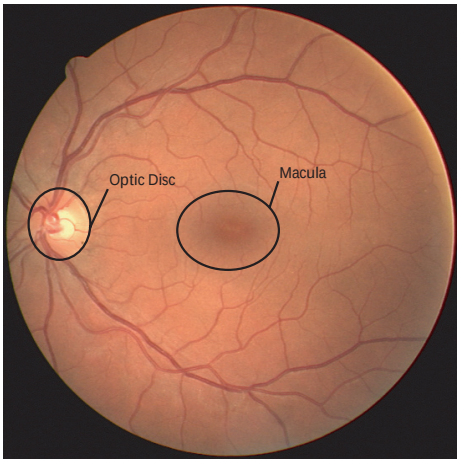

Figure 1: A typical color fundus photograph with the optic disc and macula indicated. Also visible is the vascular arch which is formed by the major vasculature flowing above and below the macula.

The detection of bright lesions in retinal images has proven a popular subject for research. Below we will give an overview of the previous work done towards the detection of bright lesions. Most of the presented work consider the detection of exudates only, while the detection of drusen or cottonwool spots has received less attention.

We believe the earliest work on automatic exudate detection was done by Akita and Kuga[3]. They extracted all edges in a retinal image using a standard edge detector. All vessel edges were removed using a previously made vessel centerline segmentation. Then the remaining edges were subjected to an algorithm which found the edges that formed a circular object. Measurements

ISSN: 1727-1983

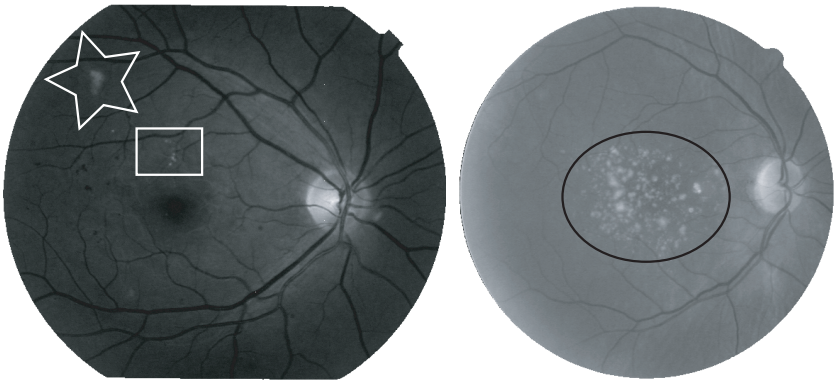

Figure 2: Examples of the three types of bright lesions. Left: Indicated with the star is a cottonwool spot, the rectangle indicates a group of exudates. Right: The ellipse contains a group of drusen.

on the inside and outside of the object then determined whether it was a hemorrhage or an exudate. No quantitative analysis was performed.

Gardner et al.[4] subdivided an image into 20×20 pixel grids. A neural network was trained using the 400 raw RGB values from each of the grids. Next, the trained neural network was used to classify grids in a test set. The authors report a per image sensitivity and specificity for the detection of exudates of 89% and 84.5% respectively.

A supervised method for exudate detection is presented by Wang et al.[5]. The approach starts by adjusting the brightness of the retinal image using a brightness transform function. This equalizes the brightness across the image and helps to detect fainter exudates. Next, a minimum distance discriminant function is trained using pixels from a training set. Using the discriminant function exudate pixels in new images are detected. Finally, the output of the classifier is cleaned up using a moving window approach. The authors report a per image sensitivity of 100% and a specificity of 70% on a set of 154 images.

Hsu et al.[6] propose an approach in which pixels are clustered according to their gray value. The clusters are then classified into normal or exudate according to a set of rules based on the distance to the vasculature as well as the gray values inside the cluster. A per image sensitivity of 100% and a specificity of 74% on 543 images is reported.

Sinthanayothin et al.[7] present an approach in which the exudates are segmented using a recursive region growing procedure. This procedure is applied after preprocessing of the image to enhance the contrast between exudates and the image background. The result is a set of grown regions representing the different objects in the image. A threshold on the gray levels in the regions is applied, next the regions belonging to the optic disk are removed. The remaining regions are labelled as exudate. A sensitivity of 88.5% and 99.7% specificity for 10×10 pixel image patches is reported.

A morphological approach to exudate detection is presented by Walter et al.[8]. Through the use of several morphological operations a segmentation of the exudates is accomplished. The reported performance is on a per pixel basis, 92.8% sensitivity and a mean predictive value of 92.4%.

Osareh et al.[9] propose another supervised approach. First, the contrast of the image is equalized after which the pixels in the image are clustered using Fuzzy C-Means clustering based on their color. A neural network is trained using example clusters which have been labelled by an ophthalmologist. The reported results on a per image basis are a sensitivity of 95% and a specificity of 88.9%.

Li et al.[10] propose a method which uses a model of the position of the anatomy on the retina to find the fovea and eliminate the optic disk. Next the image is blurred using a mean squared Wiener filter followed by a combined region growing and edge detection method to segment candidate regions. A per image sensitivity of 100% with a specificity of 71% is reported on a small test set.

In this work we propose a supervised pixel classification approach in which each pixel in a retinal image is assigned a probability of it being a bright lesion pixel. Then, by way of a classifier combination scheme, each of the individual pixel probabilities in an image are combined into one probability for the entire image. The final probability indicates whether it is likely or not the image contains bright lesions.

Methods

Generally, in the development of a Computer Aided Diagnosis (CAD) algorithm there are a number of steps which seem to be present in most approaches as described in the literature. First, there is often a pre-processing step to correct for some characteristics of the image which are not beneficial to the automatic analysis and interpretation of the image. In retinal imaging the illumination of the retina is often non-uniform resulting in varying contrast and colors over an image. Next, candidate objects are extracted using a basic technique which should be highly sensitive so as to detect every possible object in the image. Obviously, this often results in a low specificity. The used candidate detection technique should depend on the

ISSN: 1727-1983

type of objects one is looking for. Finally, the candidate objects are classified using a statistical classifier or a rule based system into two or more classes of objects. If a certain number of detections are made in an image, it can be labelled as "suspect".

In this general CAD setup there is a progression in scale. In the beginning one starts with a collection of pixels and locations which together form the image. Then, after candidate detection, clusters of pixels have been found to form the candidate objects. Finally, the classification allows progression to the image level, where a judgement is made concerning the entire image. In contrast, our approach takes a set of local judgements, i.e. on the pixel level, and combines these to directly form a judgement concerning the complete image. Thus we skip the candidate object detection and classification steps.

For our approach we do not pre-process the image, only the unprocessed green plane of the RGB color image is used. Preliminary experiments showed that simple preprocessing techniques such as high pass filtering by subtraction of a blurred version of the image, which is a popular method in the retinal image analysis literature [11, 6, 12], didn't improve results.

In the first step of the method a classifier first assigns a soft label p_i to each pixel i in the field of view (FOV) indicating the probability that it is a bright lesion pixel. To train the classifier we use a set of pixels extracted from a training set in which each pixel has been labelled either "bright lesion" or "background". Each of these pixels i and its direct neighborhood is described using a rich set of local image structure features. In this case these are the responses of the Gaussian filter and its derivatives up to and including second order (i.e. $L, L_x, L_y, L_{xx}, L_{yy}, L_{xy}$) at scales $\sigma = 1, 2, 4, 8, 16$ in addition to the gray value of the pixel under consideration. Together this makes a total of 31 features, 6 derivatives times 5 scales plus one gray value.

For the classification a k-Nearest Neighbor (kNN) classifier was used. This classifier can assign a posterior probability, or soft label, p_i to an unseen pixel i by

$$p_i = \frac{n}{k} \tag{1}$$

where n is the number of bright lesion samples amongst the k nearest neighbors in feature space of i.

After training the classifier and applying it to a previously unseen image there often is some spurious response in the optic disc (OD) region of the image. As the OD (see Figure 1) is a bright circular object, the pixels inside it locally appear similar to those in bright lesions. As we use a set of multi-scale features and OD pixels have been used for training the system, it is able to tell the difference between a smaller bright lesion and the larger OD for most pixels in the OD. However, as vessels run over the OD smaller areas of the OD are "separated" from the main disc area (see Figure 3). Two other sources of spurious response are the nerve fibers running through the retina and flash light reflections. In a large part of the retina the nerve fibers are not visible but near the major

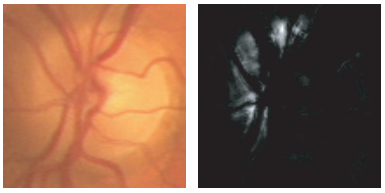

Figure 3: Spurious responses inside the optic disc.

vasculature of the vascular arch they often are (see Figure 4). This effect is enhanced by the flash light which reflects off of the major vessels which lie in a tube-like fashion on top the retina.

In order to suppress the spurious responses we used an automatic retinal segmentation system that we previously developed [13]. This system is able to fit a shape model to a retinal image thereby finding the major anatomical landmarks, i.e. the OD, the macula and the vascular arch. Now, this model is fitted to all training and test images and the distance to that part of the shape model which defines the OD and the major vascular arch (see Figure 5) is added as another feature. This brings the total number of features to 32.

Now, after applying the described system to all the pixels in an image we obtain a p_i for every pixel. One could view each of these p_i as a local decision on whether the image contains bright lesions. To combine these local decisions into a global probability P we use a static classifier combination scheme. Examples of well known combination schemes are for example the minimum, maximum, average, product and median rule. The usage of these rules is obvious, e.g. for the minimum rule the minimum p_i in the image is selected or for the product rule the product of all p_i should be used.

Since there is always a majority of pixels with $p_i = 0$ and even in images without bright lesions there are always a few pixels with $p_i > 0$ these standard combination rules do not work. Therefore we use a quantile based combination scheme which is a generalization of the median combination rule[14]. This scheme, given n soft labels is given by

$$P = \frac{\mathbf{p}_{(\lfloor rn \rfloor)} + \mathbf{p}_{(\lceil rn \rceil)}}{2} \tag{2}$$

where $r \in (0,1)$, $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ indicate the largest integer which is smaller than $\cdot$ and the smallest integer which is

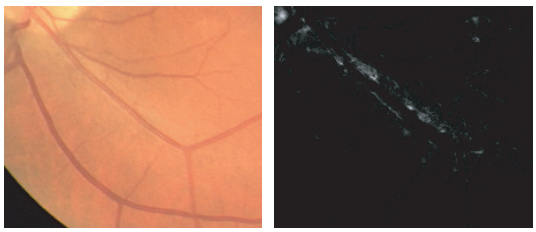

Figure 4: Spurious responses alongside the major vasculature.

ISSN: 1727-1983

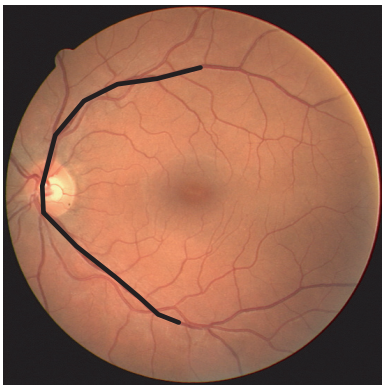

Figure 5: Part of the fitted shape model shown in black. The distance to the black line is used as a feature to eliminate spurious responses in the optic disk and along the vascular arch.

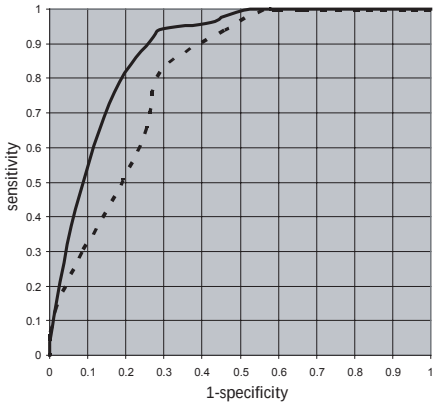

Figure 6: The ROC curves of both systems. System 1 is given by the dotted line while system 2 is given by the solid line.

bigger than $\cdot$ respectively and $\mathbf{p}_{(i)}$ indicates the ith order statistic. The ith order statistic can be calculated by sorting all n posterior probabilities p_i in ascending order and selecting the ith entry in this list. Note that if $r = 0.5$ one obtains the median rule. After preliminary experiments it was decided to set $r = 0.96$.

Materials

In the development and testing of the system a set of 153 different macula-centered digital color fundus photographs were used. These were obtained from a DR screening program in the Netherlands. All images were JPEG compressed and all identifyable patient data was removed. JPEG compression is not ideal from an image processing point of view but is standard practice in many screening settings. Image size was 768×584 with a circular field of view with a diameter of 540 pixels. Because the proposed algorithm is supervised we split the dataset into a training and a test set. A total of 44 images were used for training the system while the remaining 109 images were used for testing.

An experienced ophthalmologist marked all bright lesion pixels in the training set. All images in the training set contained at least one bright lesion. In the test set every image containing bright lesions, 59 of the 109, was labelled as such by the ophthalmologist. This allowed us to train the system using the marked pixels from the training set and to do a per image evaluation on the test set.

Experiments and Results

The described system was run twice on the complete test set. Once without the distance to the vascular arch and OD added as a feature (System 1) and once with this extra distance feature (System 2). After obtaining P for every image in both experiments we can threshold it to vary the sensitivity and specificity of both methods. The final ROC curves [15] are given in Figure 6. The area under the curve for System 1 and 2 were 0.80 and 0.88 respectively. Some result images are shown in Figure 7.

Discussion

The relatively simple scheme as proposed in this work produces promising results and shows it is possible for this application to go from pixel level decisions directly to decisions on an image level. Also we have shown how prior information derived from a segmentation of the major anatomical landmarks on the retina can be used within this scheme to boost its performance.

A problem with the current system setup is that we do not know anything about the type of lesions that are present in the image. As was said in the introduction these bright lesions can be of three types, one of which is not connected to DR at all. This suggests an expansion of the system so that it is able to distinguish the different types of lesions from each other. We have already done some work on the classification of lesions into one of the three classes [16]. Combination of that method with the automatic detection technique detailed in this paper will be a logical next step.

Another element which could be improved is the fact that the system does not use any color information at the moment. Previous publications have shown the importance of color for bright lesion detection and classification [17] so it is very well possible the performance of the system could be improved given the use of the color information in the other two color planes.

Conclusions

An automatic system that determines whether a fundus photograph contains bright lesions has been presented. The system shows reasonable performance on a large test set of screening images.

ISSN: 1727-1983

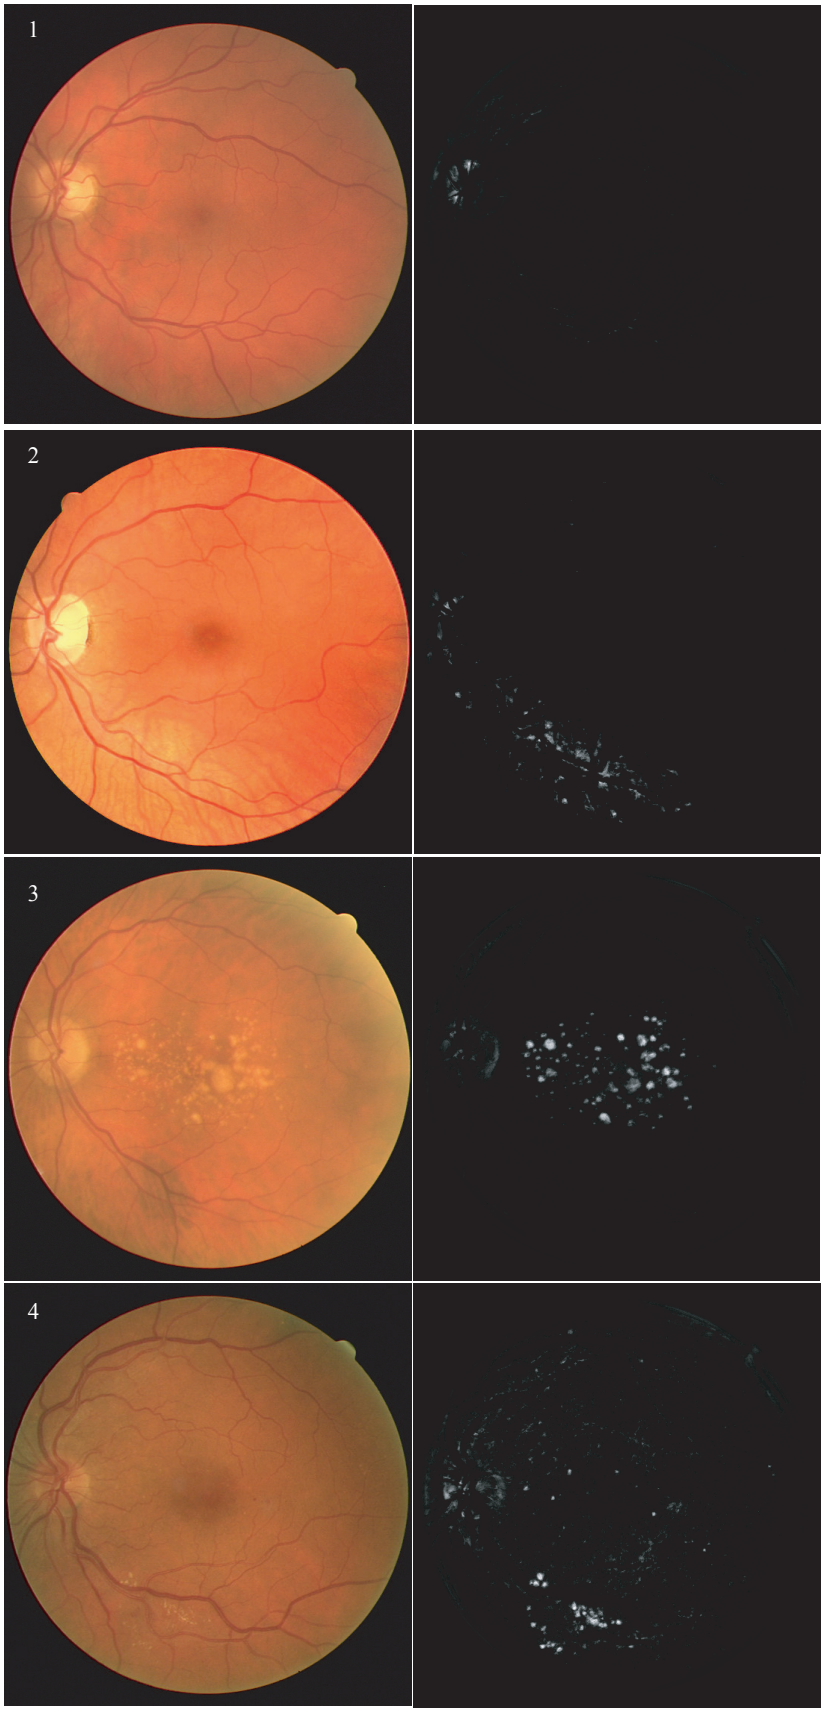

Figure 7: Examples of the soft labels assigned to images pixels by System 2 (i.e. the best performing system). Images 1 and 2 show normal, healthy retinas. Image 1 shows some response on the OD, but the probabilities never are higher than 0.6. In image 2 erroneous responses are generated on an area of the retina which is very bright. Images 3 and 4 are examples of the response of the system to retinas with visible lesions.

ISSN: 1727-1983

References

[1] M.D. Abràmoff and M.S. Suttorp. Web-based screening for diabetic retinopathy in a primary care population: the Eye Check project. *Telemedicine and e-Health*, 2005. In press.

[2] M. Niemeijer, B. van Ginneken, J.J. Staal, M.S.A. Suttorp-Schulten, and M.D. Abràmoff. Automatic Detection of Red Lesions in Digital Color Fundus Photographs. *IEEE Transactions on Medical Imaging*, 24(5):584–592, May 2005.

[3] K. Akita and H. Kuga. A computer method of understanding ocular fundus images. *Pattern Recognition*, 15(6):431–443, 1982.

[4] G.G. Gardner, D. Keating, T.H. Williamson, and A.T. Elliot. Detection of diabetic retinopathy using neural network analysis of fundus images. *British Journal of Ophthalmology*, 80(11):937–948, November 1996.

[5] 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 (CVPR)*, 2000.

[6] W. Hsu, P. M. D. S. Pallawala, M. L. Lee, and K. G. Au-Eong. The role of domain knowledge in the detection of retinal hard exudates. In *Proceedings of IEEE Computer Vision and Pattern Recognition*. IEEE Computer Vision and Pattern Recognition, 2001.

[7] C. Sinthanayothin, J.F. Boyce, T.H. Williamson, H.L. Cook, E. Mensah, S. Lal, and D. Usher. Automated detection of diabetic retinopathy on digital fundus images. *Diabetic Medicine*, 19:105–112, 2002.

[8] T. Walter and J.C. Klein. A computational approach to diagnosis of diabetic retinopathy. In *6th Conference on Systemics, Cybernetics and Informatics (SCI)*, July 2002.

[9] A. Osareh, M. Mirmehdi, B. Thomas, and R. Markham. Automated identification of diabetic retinal exudates in digital colour images. *The British Journal of Ophthalmology*, 87(10):1220–1223, October 2003.

[10] H. Li and O. Chutatape. Automated feature extraction in color retinal images by a model based approach. *IEEE Transactions on Biomedical Engineering*, 51(2):246–254, 2004.

[11] T. Spencer, J.A. Olson, K.C. McHardy, P.F. Sharp, and J.V. Forrester. An image-processing strategy for the segmentation and quantification in fluorescein angiograms of the ocular fundus. *Computers and Biomedical Research*, 29:284–302, 1996.

[12] M. Foracchia, E. Grisan, and A. Ruggeri. Luminosity and contrast normalization in retinal images. *Medical Image Analysis*, 9(3):179–190, 2005.

[13] M. Niemeijer, B. van Ginneken, F. ter Haar, and M.D. Abràmoff. Automatic detection of the optic disc, fovea and vascular arch in digital color photographs of the retina. In W.F. Clocksin, A.W. Fitzgibbon, and P.H.S. Torr, editors, *Proceedings of the British Machine Vision Conference*, pages 109–118. BMVA, 2005.

[14] M. Loog and B. van Ginneken. Static Posterior Probability Fusion for Signal Detection. In J. Kittler, M. Petrou, and M.S. Nixon, editors, *Proceedings of the 17th International Conference on Pattern Recognition*, pages 644–647. IEEE, 2004.

[15] C.E. Metz. ROC Methodology in Radiologic Imaging. *Investigative Radiology*, 21(9), 1986.

[16] M. Niemeijer, B. van Ginneken, M. Sonka, and M.D. Abramoff. Automated Classification of Exudates, Cottonwool Spots and Drusen From Retinal Color Images for Diabetic Retinopathy Screening. In *Association for Research in Vision and Ophthalmology (ARVO)*, 3468, page 150. The Association for Research in Vision and Ophthalmology, 2005.

[17] M.H. Goldbaum, N.P. Katz, M.R. Nelson, and L.R. Haff. The discrimination of similarly colored objects in computer images of the ocular fundus. *Investigative Ophthalmology & Visual Science*, 31(4):617–623, April 1990.

ISSN: 1727-1983