33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

Automatic Screening of Age-Related Macular Degeneration and Retinal Abnormalities

P. Burlina, D.E. Freund, B. Dupas, and N. Bressler

Abstract—We describe a novel approach for screening retinal imagery to detect evidence of abnormalities. In this paper, we focus our efforts on age-related macular degeneration (AMD), a pathology that may often go undetected in the early or intermediate stages, and can lead to a neovascular form often resulting in blindness, if untreated. Our strategy for retinal anomaly detection is to employ a single class classifier applied to fundus imagery. We use a multiresolution locally-adaptive scheme that identifies both normal and anomalous regions within the retina. We do this by using a hybrid parametric/non-parametric characterization of the support of the probability distribution of normal retinal tissue in color and intensity feature space. We apply this approach to screen for evidence of AMD on a dataset of 66 healthy and pathological cases and found a detection sensitivity and specificity of 95% and 96%.

I. RELEVANCE

We present an approach for automatically screening patients for evidence of retinal abnormalities. Among all pathologies, this paper is particularly focused on the case of AMD because of its prevalence and its clinical relevance: AMD is the leading cause of blindness in the US [1] and throughout much of the western world.

The advanced stage of AMD usually manifests itself by the development of choroidal neovascularization (CNV) which can lead to severely impaired central vision, impacting activities such as reading, driving, or face recognition [2]. In the United States alone, approximately 200,000 individuals over the age of 50 develop the advanced stage of AMD each year. Left untreated, approximately 70% of these cases develop substantial vision loss within 2 years in one eye [3], while over the next 5 years, approximately half of those who develop the advanced stage in one eye will develop the advanced stage in both eyes [1].

A major advance in the treatment of CNV in AMD uses anti-vascular endothelial growth factor (VEGF).

Manuscript received March 26, 2011.

P. Burlina is with The Johns Hopkins University Applied Physics Laboratory, Laurel, MD 20723, USA, and The Johns Hopkins University, Department of Computer Science, Baltimore, MD, USA. e-mail philippe.burlina@jhuapl.edu

D. E. Freund is with The Johns Hopkins University Applied Physics Laboratory, Laurel, MD 20723, USA. e-mail david.freund@jhuapl.edu.

B. Dupas is with Hôpital Lariboisière, Service d'Ophtalmologie. e-mail benedicte.dupas@gmail.com

N. Bressler is with The Johns Hopkins University School of Medicine Wilmer Eye Institute . e-mail nbressler@jhmi.edu

Recent clinical trials with anti-VEGF have demonstrated promise for mitigating the advance of the disease [3]. Thus, it is critical to identify in a timely manner those individuals most at risk for developing advanced AMD, specifically, individuals with the intermediate stage of AMD.

The main symptom of the intermediate stage AMD is the presence of small abnormal formations in the fovea called drusen (see Fig.1). The presence of drusen often causes no symptoms. That is, there is no loss of vision associated with the intermediate stage AMD and thus no outward signs to warn the individual of its presence [1]. For the hundreds of thousands of individuals affected with the neovascular form each year, early detection and treatment of the advanced neovascular stage could result in stabilization of vision at a very useful level for reading, driving, and other functions dependent on good central vision [4]. Currently, ophthalmoscopy of the retina or evaluation of fundus photographs remains the most effective method to identify drusen [1]. However, evaluating fundus images manually can be a tedious process requiring the expertise of an ophthalmologist or extensive training to understand the varying patterns recognized by an ophthalmologist. Therefore, there is a key need for automated visual diagnostic tools that can allow the detection of the intermediate stage AMD among a large pool of the at-risk population (millions of individuals) and facilitate longitudinal studies mapping the advance of the disease, thereby enabling timely intervention.

II. PREVIOUS WORK AND CHALLENGES

A fair amount of work has been devoted to the design of automated detectors for specific retinal pathologies such as diabetic retinopathy. However, despite its prevalence, much less has been done for AMD. Exploitation of digital fundus image processing for AMD was first reported by the study in [5]. Early AMD detection methods such as the study reported in [6] required user intervention. Recently, researchers have turned their attention to automated approaches: [7] used adaptive equalization and wavelets; [8] employed mathematical morphology on angiographic images with fluorescein injection; [9] used adaptive thresholding; [10] exploited a probabilistic boosting approach for the classification of non-homogeneous drusen textures; [11] used probabilistic modeling and fuzzy logic; [12] employed histogram normalization and adaptive

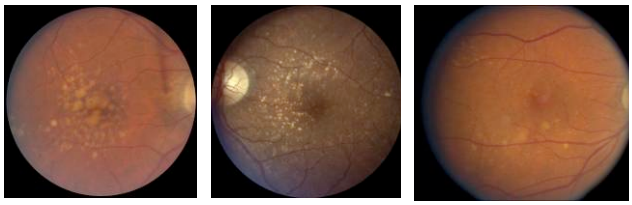

(a) (b) (c)

Fig. 1. Examples of pathological cases showing wide variation in drusen characteristics.

segmentation; finally, [13] exploited texture discrimination and the intensity topographical profile.

A common route employed by the most promising of the previously cited approaches consisted in using a two-class or multiclass classifier (drusen vs. vessels vs. retinal background tissue vs. other tissue). Because of variations in imaging conditions (fundus image quality, illumination, blur, background uniformity) and variations in patient specific appearance (variability in pigmentation and drusen appearance within and across subjects), it is difficult to identify stable image features characteristic of drusen that can be used to build a robust classifier that will perform reliably over a large dataset.

Because of this, we explore an alternate route, and investigate the use of a one-class classifier: we characterize the statistical distribution of 'normal' background retinal tissue, and search for areas exhibiting abnormalities. The salient features of the proposed algorithm are as follows: (a) Intensity, color, and gradient information is exploited. (b) A hybrid parametric Constant False Alarm Rate (CFAR) detector (for the fundus image intensity value) is used in conjunction with a non-parametric (adapted to color space features) CFAR detector based on Support Vector Machine (SVM). (c) The algorithm uses a multiscale and locally adaptive approach. In addition, our approach addresses other challenges in drusen detection including the presence of a background intensity gradient in retinal fundus imagery, the presence of various anatomical features (vessels, optical nerve, etc) and artifacts (flashes) resulting from specific Illumination conditions. To address these issues we develop additional processing stages that identify *appropriate* regions for training and testing and eliminate some of these spurious features.

III. TECHNICAL APPROACH

(A) Algorithm Description

This section describes the basic preprocessing steps. In the next section, we review the characterization of the support of the probability distribution of normal retinal tissue in color and intensity space.

Our ultimate goal in this preliminary processing phase is to find clean background training regions to train our single class classifier, and candidate testing regions on which to apply our classifier. Because the training is performed in-scene, this approach provides the basis of a patient-specific abnormality detector which (a) obviates the collection of an exhaustive training dataset, and (b) is patient specific, making this approach a good candidate for application over large datasets.

Figure 2 provides a block diagram of the basic algorithm. As shown in the figure, the algorithm has two distinct branches. As indicated by the horizontal dashed arrow, the right branch relies on results from the left branch, which is performed first. Although some of the blocks may be basic and self explanatory, for completeness, we will discuss each block separately and in turn.

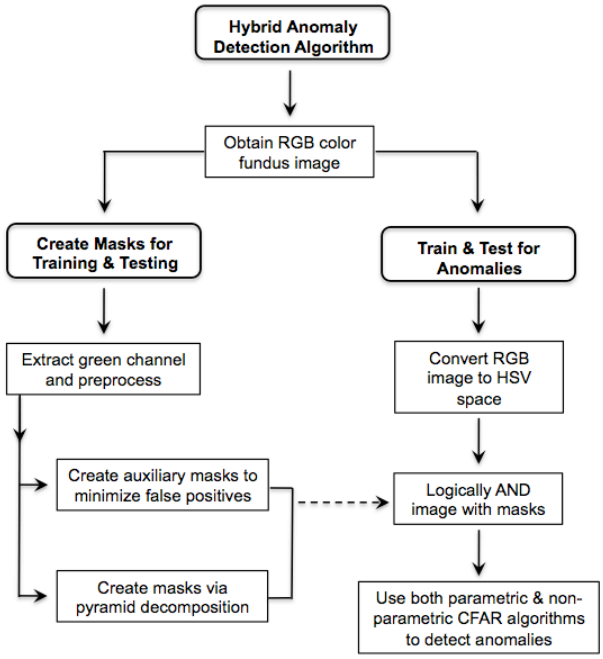

Fig. 2. Outline of anomaly detection algorithm.

i) Obtain RGB color fundus image: Color is an important visual cue in finding retinal anomalies as well as characterizing drusen. Our algorithm is designed to take advantage of this fact. Thus, the first step is to acquire a red, green, blue (RGB) color image of the retina using a standard fundus camera. Although digitized fundus images are always square, they sometimes are of different sizes. Thus the RGB images are resized to a uniform 1000 x 1000 size.

ii) Extract the green channel and preprocess: The green channel is extracted and used for analysis.

iii) Create auxiliary masks: In this step, two logical masks are created to minimize the detection of false positives: A "border" mask to eliminate the large circular edge (see Fig.1) and all pixels exterior to it and an "optic nerve" mask to eliminate the region occupied by the optic nerve (see Fig.1b). This step depends on *a priori* knowledge of the specifics of the image acquisition process in step *i)* above. For example, if the magnification is such that the entire image is comprised of retina (i.e. no edge and no exterior black pixels) then clearly no border mask is necessary. We note that the three images shown in Fig.1 are in a standard format and are completely representative of all the images analyzed in this study. There have been numerous studies devoted to accurately locating and sizing the optic nerve in

fundus images. However, we found that using basic image processing steps were sufficient for the purpose of eliminating the area around the optic nerve in order to avoid spurious detections. In particular, we perform radiometric adjustment by doing histogram equalization after excluding the top and bottom 1% of intensity histogram. Next the mean pixel value in the image is set to 127 and then this image is filtered with a 5 x 5 median filter. After filtering, the image is thresholded, the largest connected component is found, and then this area is dilated.

iv) Create masks via pyramid decomposition: In this step, two logical masks are created for the ultimate purpose of (a) finding 'clean' training background areas devoid of blood vessels, drusen, and other features or artifacts (thereby minimizing outliers) and (b) finding testing areas likely to contain abnormalities (thereby minimizing false positive detections).

We use a multiresolution pyramid decomposition process on the green channel image to create the training and testing masks. Specifically, we look for regions of high gradient magnitude indicative of high-frequency components such as imaging artifacts or vessels. The gradient magnitude is calculated at each level (scale) of the pyramid and binary masks are then obtained by thresholding such that regions of high gradient magnitude (indicating the presence of an edge) are assigned a value of one. The union of these binary images is accumulated at each level of the image pyramid, yielding an intermediate binary mask.

The testing mask is obtained by morphologically closing the intermediate binary mask. The training mask is obtained by dilating the intermediate binary mask and then taking the logical complement. Note that pixels with a value of one in the training mask represent regions with low frequency components (i.e. regions devoid of edges and artifacts) and presumably correspond to normal retinal tissue.

v) Convert RGB image to HSV: The original color fundus image is converted from RGB to Hue, Saturation, and Value (HSV) space. The HSV vectors are used for training and anomaly detection testing.

vi) Logically AND HSV image with training & testing masks: The next step is to logically AND the training mask with the HSV image and split the subsequent image into local sub-images (100x100 windows) for training. This is to provide local adaptation and address the presence of smoothly varying gradient intensity often present in fundus imagery.

vii) Use both parametric and non-parametric CFAR algorithms to detect anomalies: Since the H and S space is not a Euclidian space, the H and S channels are used for training a non-linear non-parametric (i.e. SVM) CFAR detector to find anomalies based on color. The value V is instead used for training a traditional single-sided CFAR detector, to find anomalies based on intensity. The traditional CFAR detector is trained on the background V value and splits regions into low V (generally low intensity corresponding to vessels, or hyper-pigmentations) and high V values (corresponding to drusen). Once training is complete, the anomaly detector is applied to candidate testing regions obtained using the testing mask. The process is applied to local 100x100 tiles. Pixels found to be anomalous under both CFAR detectors are then combined and the image is reconstructed from its sub- windows.

(B) Characterization of the probability density support region of normal retinal tissue

This section details the anomaly detection. The idea of anomaly detection is to characterize the probability distribution that explains the training samples exemplifying normal retinal tissue. One starts with a training vector set T composed of L vectors drawn from an underlying probability distribution P of normal background retinal tissue pixels:

$$T = \left\{ x_i \in \mathbf{R}^N, i = 1,2,...L \right\}. \quad (1)$$

For our problem, N equals 2 since $\mathbf{x}_i$ consist of the Hue (H) and Saturation (S) values. Since we adopt a CFAR approach, the support region in feature space of the probability distribution P is found by using a constrained optimization approach developed in [14],[15]: this approach seeks to construct a support region R in feature space in such a way that that the probability that a point $\mathbf{x}$ drawn from P will have a specified and bounded probability (a false alarm rate) to lie outside of R. As described shortly, this false alarm rate will be denoted by a parameter υ.

A Gaussian Radial Basis Function (RBF) kernel is used,

$$K(\mathbf{x},\mathbf{y}) = \exp(-\|\mathbf{x}-\mathbf{y}\|^2/\sigma^2) = \langle \Phi(\mathbf{x}), \Phi(\mathbf{y}) \rangle \quad (2)$$

in place of the traditional linear inner product $\langle \mathbf{x}, \mathbf{y} \rangle$. Doing so implicitly imparts a specific geometric interpretation to the subsequent steps. In particular, it maps the exemplars in T from an N-dimensional input space (here N=2) into an infinite-dimensional feature space. In the feature space, the exemplars reside on an infinite-dimensional hypersphere of unit radius and we then seek a hyperplane defined as

$$\langle \mathbf{w}, \Phi(\mathbf{x}) \rangle = \rho \quad (3)$$

which maximally separates the origin from the training exemplars in T. That is, maximizes the distance $\rho/\|\mathbf{w}\|$ from the origin to the separating hyperplane. Based on this, the decision region for testing becomes

$$R = \left\{ \mathbf{x} : \mathrm{sgn}(\langle \mathbf{w}, \Phi(\mathbf{x}) - \rho \rangle) \geq 0 \right\} \quad (4)$$

where sgn(.) denotes the sign function. To complete the specification of R, a certain fraction, n, of the vectors in T are allowed to violate the original constraint that they must lie within the half space R. That is, slack variables, ξ_i, are introduced such that the vectors in T are allowed to satisfy: $\langle \mathbf{w}, \Phi(\mathbf{x}_i) \rangle \geq \rho - \xi_i$, for ξ_i non-negative. Doing so provides some robustness to the incorrect inclusion of outliers such as

imaging artifacts during the training mask definition phase described earlier.

These specifications can be summarized in a constrained minimization of a quadratic objective function [14]:

$$\min_{w\in F;\,\xi\in R^L;\,\rho\in R} \frac{1}{2}\|\mathbf{w}\|^2 + \frac{1}{\upsilon L}\sum_i \xi_i - \rho \qquad (5)$$

subject to constraints

$$\langle \mathbf{w}, \Phi(\mathbf{x}_i)\rangle \geq \rho - \xi_i, \text{ with } \xi_i \geq 0\,. \qquad (6)$$

Note, when υ is small, the penalty imposed on the points violating the separating plane increases.

To solve this constrained problem, a Lagrangian is defined as

$$L(\mathbf{w},\xi,\rho,\alpha,\beta) = \frac{1}{2}\|\mathbf{w}\|^2 + \frac{1}{\upsilon L}\sum_i \xi_i - \rho + A(\mathbf{w},\xi,\rho,\alpha,\beta,\mathbf{x}) \qquad (7)$$

where the function A is defined by

$$A(\mathbf{w},\xi,\rho,\alpha,\beta,\mathbf{x}) = -\sum_i \alpha_i\left(\langle \mathbf{w},\Phi(\mathbf{x})\rangle - \rho + \xi_i\right) - \sum_i \xi_i\beta_i \qquad (8)$$

Invoking the Kuhn-Tucker conditions [15], it follows that a certain number of the Lagrange multipliers, α_i, are non-zero. These non-zero multipliers correspond to exemplars that lie exactly on the hyperplane boundary of R (the support vectors). Taking the first order optimality condition and setting the partial derivative of the Lagrangian with respect to *(***w***,x,r,a,b)* to zero, yields the following properties:

$$\mathbf{w} = \sum_i \alpha_i \Phi(x_i) \;\; and \;\; \sum_i \alpha_i = 1. \qquad (9)$$

Therefore the optimal vector $\mathbf{w}$ is the weighted average (center of mass) of the support vectors. Finally, plugging back the last equation, in the definition of R, we get:

$$R = \left\{\mathbf{x} : \sum_i \alpha_i K(\mathbf{x}_i,\mathbf{x}) \;\geq\; \rho\right\}. \qquad (10)$$

The decision region R has a probabilistic interpretation [16] if one considers the sum of weighted $K(\mathbf{x}_i,\mathbf{x})$ as being comprised of a mixture of Gaussians and therefore R as region resulting from a CFAR detection rule.

The next step in the anomaly detection is to find anomalies based on V. This step is accomplished by adopting an additional CFAR test on the intensity. The intensity (i.e. the value, V, in HSV) is assumed Gaussian. The mean and variance are computed and a one-sided detection region R_V is defined as

$$R_V = \left\{V : |V - \bar{V}| \geq T\sigma_V\right\}. \qquad (11)$$

Because drusen almost invariably have a higher intensity (i.e. V) compared to other anomaly classes, R_V ,will contain pixels associated with the drusen. The final detection region is then taken as the intersection of R and R_V.

IV. RESULTS

We applied our retinal anomaly detection to the problem of screening for intermediate stage AMD (see Figure 3). For testing, after approval by The Johns Hopkins University School of Medicine's Institutional Review Board, we collected 66 color fundus images from a clinical dataset. 27 of the images were of a normal (i.e. no drusen) fundus and the remaining 39 contained drusen. The decision to classify the patient as normal or pathological was based on the number or percentage of pixels found to be abnormal and

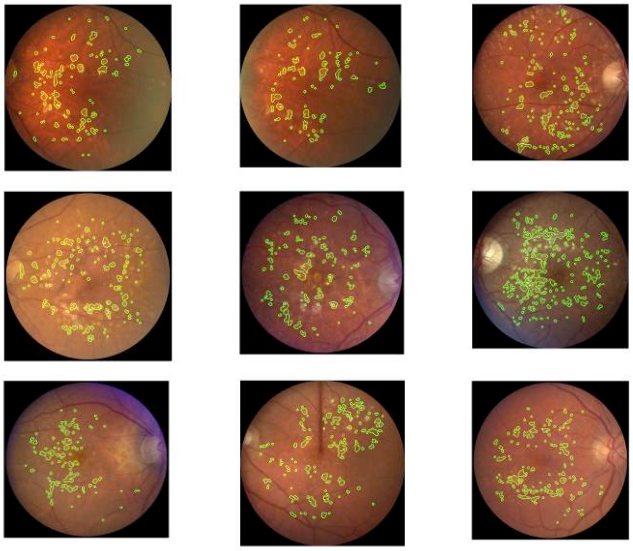

Fig. 3 Example of detection of drusen.

presumably drusen. This percentage threshold was kept at 3% for all our experiments.

The images were collected from 38 subjects overall. Of these 38 subjects, most were Caucasian and some were Asian or African American subjects whose fundus images tend to be darker due to the higher melanin content in the retinal pigment epithelium. The images were acquired with a Zeiss Fundus camera FF4 after pupil dilation by the Wilmer Eye Institute Photography Department. A 40° field of view (FOV) image centered on the macula was used in all cases.

Of the 38 subjects used, 15 were healthy and 23 were diagnosed for various stages of AMD. The majority of AMD diagnosed fundus images were graded as the intermediate stage AMD, with some early stage AMD and a few advanced cases. In all, the dataset represented 16 unique eyes for normal subjects and 24 unique eyes for pathological

subjects. In some instances stereo fundus images were obtained, yielding two distinct views of the same eye.

The images were classified by experienced graders of AMD from the JHU Wilmer Eye Institute who have participated in grading AMD for NIH-sponsored trials.
Based on the ground truth, as determined by the graders, we found about 95% sensitivity, 96% specificity, and positive predictive value (PPV) equal to 97%, and negative predictive value (NPV) of 92%. A comparison with recently published methods cited above that were tested on datasets of size commensurate with our dataset shows that our method exceeds the performance reported by others.

V. CONCLUSION

A novel hybrid parametric/non-parametric CFAR anomaly detector was presented and applied to the detection of abnormalities in retinas using fundus imagery. In this paper, we applied the technique to the detection of drusen which are indicative of AMD. Early detection of drusen may have substantial impact on the prevention of blindness due to AMD [4].

The results presented here are still preliminary and the data set used to test our algorithm was relatively small and needs to be increased before definitive conclusions can be made regarding the robustness of the method. Our goals are to continue characterizing and refining our algorithm on expanding datasets to eventually allow for large-scale deployment. Furthermore, we plan to construct ROC curves in order to quantify algorithm performance. In addition, we will expand the algorithm to be applicable to image modalities other than standard fundus images (e.g. Optical Coherence Tomography (OCT) images).

Finally, we emphasize here that the algorithm we have presented should be viewed as an anomaly detector. Although we have explicitly applied it to the specific problem of drusen detection in early or intermediate stage AMD, it is not limited to just this application. It is applicable to screening for other retinal pathologies which entail the formation of abnormal tissue or bleeding such as diabetic retinopathy or geographic atrophy.

REFERENCES

[1] Bressler, N.M., “Age-related macular degeneration is the leading cause of blindness”, JAMA, 291:1900-1, 2004.

[2] Patients’ Perceptions of the Value of Current Vision: SST report No. 4. Am J Ophthalmol.: p. 91:108, 2004.

[3] Bressler, N.M. and Bressler, S.B.. “Photodynamic therapy with verteporfin (Visudyne): impact on ophthalmology and visual sciences”, Invest Ophthalmol Vis Sci. 41(3): p. 624-8, 2000.

[4] AREDS 1 report number 8 Arch ophth, 119 (10) : 1417-36, 2001.

[5] Mimoun, G. “Intéret de l’image numérisée dans le diagnostic et la classification des drusen maculaires,” Soc. Fr. Ophtalm, vol. 4, 1989.

[6] Smith, R.T. , Chan, J.K. , Nagasaki, T. J.R. Sparrow, I. Barbazetto, C.C.W. Klaver, “A digital method of drusen quantification based on the geometry of fundus reflectance”, in Proc. IEEE Bioengineering conf., 2003.

[7] Brandon, L. and A. Hoover, “Drusen detection in a retinal image using multi-level analysis”, MICCAI 2003, LNCS 2878, ed. R.E. Ellis and T.M. Peters. Springer-Verlag. 618-625, 2003.

[8] Ben Sbeh, Z.B. Cohen, L.D. Mimoun, G. Coscas, G.and Soubrane, G. "An adaptive contrast method for segmentation of drusen", in Proc. ICIP, 1997.

[9] Rapantzikos, K., Zervakis, M. and Balas, K., “Detection and segmentation of drusen deposits on human retina: Potential in the diagnosis of age-related macular degeneration”, in Medical Image Analysis, Vol. 7, Issue 1, 95-108, 2003.

[10] Lee, N.; Laine, A.F.; Smith, T.R., "Learning non-homogenous textures and the unlearning problem with application to drusen detection in retinal images", in Proc. ISBI, 2008.

[11] Thdibaoui, A. , Rajn, A. Bunel, P. , “A fuzzy logic approach to drusen detection in retinal angiographic images”, in Proc. ICPR, 2000.

[12] Checco, P., Corinto, F., "CNN-based algorithm for drusen identification", Proc IEEE Int. Symp. Circ. and Syst., 2006.

[13] Parvathi, S.S.; Devi, N., "Automatic Drusen Detection from Colour Retinal Images", Proc. Int. Conf. on Comp. Intelligence and Multimedia Applications, 2007.

[14] Scholkopf, B., Platt, J., Shawe-Taylor, J., Smola, A. J. and Williamson, R.C., “Estimating the support of a high-dimensional distribution”, Neural Computation 13, 1443–1471, 2001.

[15] Vapnik, V.N., *Statistical Learning Theory*, John Wiley & Sons, Inc. (New York), 1998.

[16] Freund D. E., Burlina P., Banerjee A., and Justen E., “Comparison of Kernel Based PDF Estimation Methods”, Proc. SPIE Automatic Target Recognition XIX, Vol. 7335, 733508, 2009.