Proceedings of the 28th IEEE
EMBS Annual International Conference
New York City, USA, Aug 30-Sept 3, 2006

Locating the Optic Nerve in Retinal Images: Comparing Model-Based and Bayesian Decision Methods

Thomas P. Karnowski, *Member, IEEE*, V. Priya Govindasamy, *Member, IEEE,*
Kenneth W. Tobin, Jr., *Senior Member, IEEE,* Edward Chaum

***Abstract*—In this work we compare two methods for automatic optic nerve (ON) localization in retinal imagery. The first method uses a Bayesian decision theory discriminator based on four spatial features of the retina imagery. The second method uses a principal component-based reconstruction to model the ON. We report on an improvement to the model-based technique by incorporating linear discriminant analysis and Bayesian decision theory methods. We explore a method to combine both techniques to produce a composite technique with high accuracy and rapid throughput. Results are shown for a data set of 395 images with 2-fold validation testing.**

I. Introduction

The World Health Organization estimates that 135 million people have diabetes mellitus worldwide and that this number will increase to 300 million by the year 2025 [1]. More than 18 million Americans currently have diabetes and the number of adults with the disease is projected to more than double by the year 2050 [2]. Visual disability and blindness have a profound socioeconomic impact upon the diabetic population and diabetic retinopathy (DR) is the leading cause of new blindness in working-age adults in the industrialized world [3]. Thus, there is a significant need to develop inexpensive, broad-based screening programs for DR.

The detection of anatomic structures is fundamental to the subsequent characterization of the normal or disease states that may exist in the retina. In this paper, we compare two methods for detecting a critical structure in images of the human retina, specifically the optic nerve (ON) which is also known as the "optic disk" due to it characteristic circular shape. An example of a retina image with the ON highlighted in black is shown in Figure 1. The literature contains many examples of ON detection in retinal imagery. These methods incorporate techniques such as dynamic contours [4], convergence of the vasculature [5], and geometric models [6]. We compare the model-based method of [7] which uses a principal component-based reconstruction of the ON to perform a Euclidean distance similarity measurement, to the ON localization method of the authors [8] which uses a successful segmentation of the vasculature to generate features that are treated as a Gaussian distribution in a Bayesian discriminator pattern classifier. Both methods have strengths in their approach. The method of [7], which we identify as the PCA method, uses principal component analysis to embody the main structure of the ON and does not require vasculature segmentation, but is more dependent on intensity variations in the image and has lower-throughput. The method of [8], which we identify as the FBLR for feature-based likelihood ratio, uses well-known classical Bayesian pattern recognition techniques. The method requires a successful segmentation of the vasculature in the image, but is less dependent on intensity variations in the image. As a result of our comparison, we develop improvements to the PCA method that use additional information about the ON regions and non-ON regions of the training images to improve accuracy. Finally, we suggest a means of fusing the methods to gain improved accuracy.

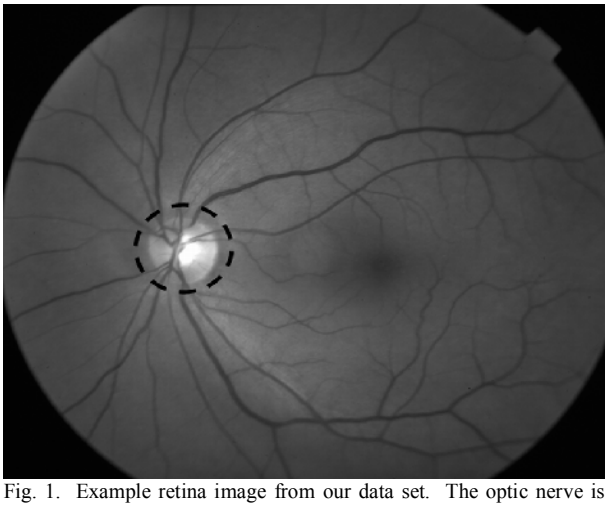

Fig. 1. Example retina image from our data set. The optic nerve is shown circled in black. The optic nerve is a key landmark in identifying the retinal structures.

Manuscript received March 31, 2006. This work was supported in part by the National Eye Institute of the National Institutes of Health (RO1-EY017065). This paper was prepared by the OAK RIDGE NATIONAL LABORATORY, Oak Ridge, Tennessee, 37831-6285, operated by UT-BATTELLE, LLC for the U.S. DEPARTMENT OF ENERGY under contract DE-AC05-00OR22725.

T.P. Karnowski is with the Oak Ridge National Laboratory, Oak Ridge, TN 37831 USA (e-mail: karnowskitp@ ornl.gov).

V.P. Govindasamy is with the Oak Ridge National Laboratory, Oak Ridge, TN 37831 USA (e-mail: muthusamygov@ornl.gov)

K.W. Tobin, Jr., is with the Oak Ridge National Laboratory, Oak Ridge, TN 37831 USA (e-mail: tobinkwjr@ornl.gov)

E. Chaum is the Plough Foundation Professor of Retinal Diseases with the University of Tennessee Health Science Center, 930 Madison Avenue, Suite 731, Memphis TN 38163 (e-mail: echaum@utmem.edu)

II. APPROACH

A. Model-based (PCA Method)

The model-based method of [7] uses principal component analysis (PCA) to capture the information content of a training set of optic nerve images. In this method, "candidate regions" of a retina image that may contain the ON are projected to PCA space, then the PCA coefficients are used to reconstruct the region and the residual error is measured. This process is repeated for all pixels in the candidate and the pixel with the smallest residual error is chosen as the ON location. The method is very elegant and offers advantages over other methods, in particular the absence of any complimentary segmentation aside from simple thresholding to locate the candidate regions. In the training step, each segmented optic nerve is scaled to the same size (N pixels by N pixels) and intensity normalized by subtracting the scalar mean and dividing by the scalar standard deviation. Finally, the principal components are computed. We retain those components that capture 90% of the image information as given by the eigenvalues as proposed in [7]. Test images are processed by thresholding to retain the most intense 1% of pixels. These pixels are labeled with connected component analysis and their area is measured. Blobs with size smaller than 0.04% of the entire image area are rejected as nuisance areas. The threshold step parameter (1% of pixels) and the area rejection threshold (0.04%) were identical to those in [7]. The remaining candidate regions are replaced by a square the size of the average optic nerve centered on each centroid. For each candidate pixel, a square region of size (S*N) pixels centered on pixel (r,c), is extracted, where S is a scaling factor ranging from 0.8, 0.9, 1.0, 1.1, or 1.2 for our implementation. The region is resized to N x N using bicubic interopolation, its intensity is normalized as in the training step, and the resulting image is projected to the PCA space. Finally, the region is reconstructed using the PCA coefficients and the Euclidean distance between the reduced region and the original region is measured. This process is repeated for all candidate pixels at all scales. The selected ON center is that pixel which has the smallest reconstruction error across all scales.

B. Feature-Based Likelihood Ratio (FBLR) method

In the FBLR method of [8], the retinal image is processed by segmenting the vasculature with morphological reconstruction. Next, a set of four features are generated at each pixel. These features are the brightness of the pixel region, the thickness of the vasculature, the orientation of the vasculature, and the density of the vasculature. Optic nerve regions are labeled using the ON center and a surrounding area based on an estimate of the ON radius. A training set of data (with the optic disk hand-segmented by a non-clinical researcher) is analyzed to estimate the parameters of a Gaussian distribution describing the ON regions and the non-ON regions. We also incorporate *a priori* information by using the hand-segmented training set to estimate the ON center probability density function (pdf). The pdf was estimated by summing the ON locations of the training set and convolving them with a window sized N_a x N_a where N_a is the average ON size in pixels. This approximates the pdf of the ON location using Parzen windows [9]. The Gaussian parameters and the pdf are used to compute a likelihood ratio function using maximum *a posteriori* (MAP) estimation.

C. PCA-LDA-LR method

We wanted to include information about the probability distribution of the ON into the estimation, as well as information about non-OD pixels. We thus implemented a likelihood ratio estimator using linear discriminant analysis (LDA). In this case, after generating the PCA coefficients for the ON training set, the reconstruction distances were calculated for the entire image of all training images. Twenty different pixels (the number is somewhat arbitrary) were chosen to comprise a data set in PCA space. We first selected the ON center and then masked it with a region the size of the average ON. We repeated this process on the five smallest reconstruction distances, which comprise a set of training examples which are non-ON but score low on the reconstruction distance. Finally, we chose the remaining 14 vectors by randomly selecting pixels in the candidate region. The vectors corresponding to these twenty pixels were projected back to PCA space. After repeating this process for all images in the training set, LDA was employed to compute a transform to a one-dimensional space. This feature was then used to formulate a MAP likelihood ratio modeling the feature as a Gaussian random process,

$$LR(x,y) = \frac{\frac{1}{\sigma_1} e^{-\frac{(d(x,y)-m_1)^2}{\sigma_1^2}} \cdot P(\omega_1 \mid x,y)}{\frac{1}{\sigma_0} e^{-\frac{(d(x,y)-m_0)^2}{\sigma_0^2}} \cdot [1 - P(\omega_1 \mid x,y)]}, \quad (1)$$

where $d(x,y)$ is the LDA-transformed PCA coefficients, m_0, m_1 and σ_0, σ_1 are the mean and standard deviation of the $d(x,y)$ for the non-ON and ON regions, and $P(\omega_1|x,y)$ is the probability of ON as a function of position. We call this method the PCA-LDA Likelihood Ratio method or PCA-LDA-LR.

III. COMPARISON EXPERIMENTS

Our data set was composed of 395 retinal images representing 18 different retinal pathologies and normal (non-diseased) retina, with only the green plane of the images available for processing. The images were originally captured at a resolution of 12 microns per pixel. We processed at a smaller scale (roughly 100 microns per pixel) to improve speed of performance which will be essential for automatic screening purposes. We aligned the images so that the manually segmented ON was on the left side of the image. Note that this image set, which represents an actual population from an ophthalmology practice, has large

variability in its intensities and in the physiological structure of the ON. The implications of this variability for automated screening are yet to be addressed, but we should point out that these represent individuals who have sought medical attention and are likely more advanced with respect to their DR variability than real broad-based screening data may encompass. The data set of 395 images was randomly separated into two sets of 198 and 197 images for a 2-fold validation study. We report the combined results in terms of distributions of distances from the manually segmented optic nerve center, normalized to one ON radius. We prefer this metric to specificity / sensitivity because our objective is to locate the ON center as opposed to actually classify pixels as ON or non-ON. Nevertheless, it is also instructive to report results as the number below the average ON radius, indicating that the selected point is actually on the ON as opposed to outside it. We conducted three different experiments with the data set, one each for the PCA, FBLR and PCA-LDA-LR methods.

A. Performance of PCA method

The original training set of the image consisted of 18 x 18 pixel images for a total of 324 features. In the PCA decomposition we found that we could represent 90% of the data (based on the eigenvalues) with 52 and 55 principal components respectively for each step of the 2-fold validation test. Processing the different candidate regions consisted of processing the combined region, then masking out the prior region for the original candidate case and vice-versa. The average size of the regions for processing was 6.3% of the total image size or about 1100 pixels. Figure 2 shows a histogram of the distances from the actual ON center, normalized by the average ON radius. Our results are not as accurate as expected, with 68% located within one ON radius. Inspecting the reconstruction maps for an explanation showed that while the ON location always has a local minimum, there are often other areas that have smaller minima and are therefore selected by the algorithm as more likely ON centers. We suspect our data set is likely more difficult than that used in [7], due to the high degree of DR in the set.

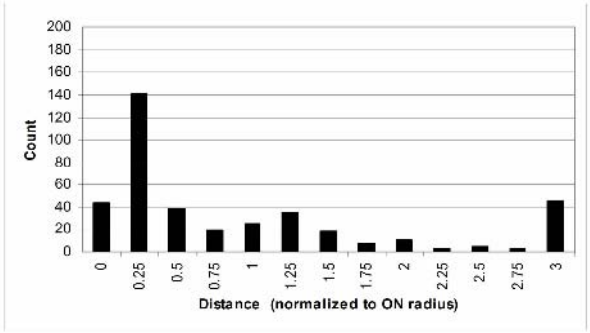

Fig. 2. Histograms of distances from actual ON for the PCA method.

B. Performance of FBLR method

The FBLR method gave similar performance to the reported results in [8], and superior results to the PCA method. Furthermore, the throughput was much greater, processing images in approximately 1/10 the time. While some of the processing time differences may be accounted for by the implementation, in [7] the motivation for selecting candidate regions is reduction of processing time, so we believe this is indeed a true benefit to the FBLR method. Figure 3 shows the histogram of distances from the FBLR method. In this case we see that most (90%) of the ONs are located within one ON radius.

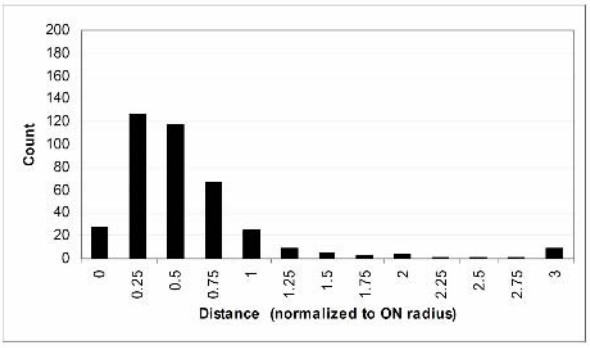

Fig. 3. Histograms of distances from actual ON for the FBLR method. Most (90%) of the sample data has ON located within one ON radius.

C. Performance of PCA-LDA-LR method

As our final comparison we generated the PCA-LDA-LR results. The PCA-LDA-LR throughput is similar to the other PCA methods, and the performance is superior to the PCA method and very similar to the FBLR method as shown in Figure 4. Perhaps most striking is the accuracy of the method; not only are most results within one ON, the majority are within 0.5 ON while the FBLR method has 59% within 0.5 ON. Adding the additional information to the method for the position prior and the LDA transformation with special sensitivity to elements with low reconstruction distance enhanced the performance considerably.

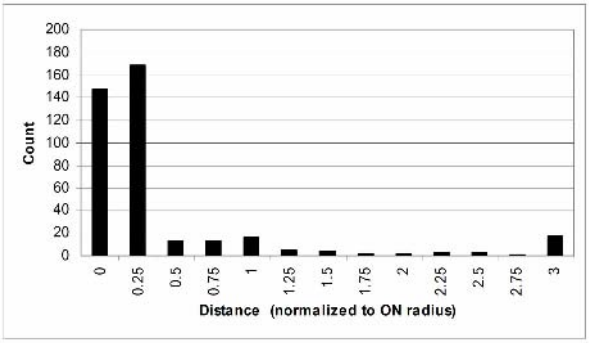

Fig. 4. Histograms of distances from actual ON for the PCA-LDA-LR method. Most (89%) of the sample data has ON located within one ON radius, with 82% within 0.5 ON radius.

D. Combining the FBLR and PCA-LDA-LR method

We tested one strategy for combining methods. We used the FBLR method to find the ON and treated this pixel as the center of a candidate region for PCA-LDA-LR. We tested candidate region sizes of 5, 7, 11, 15, 19 and 25 pixels square. We identify a candidate region as "FBLR 5x5", for example, meaning a 5x5 region around the FBLR selected

center is searched with PCA-LDA-LR. These results are shown in Figure 5 with three plots overlaid for the median value of the ON distance error, the number found within one ON radius, and the number found within 0.5 ON radius. Small searches with the PCA-LDA-LR method decrease accuracy in terms of the number of localizations below one OD radius, but the median error drops consistently to an asymptote likely related to our sampling. Overall performance drops initially then climbs again as we begin to search regions on the order of the size of the OD radius or greater. Clearly, employing this strategy as shown will only benefit if larger areas can be searched, such as FBLR 19x19 or FBLR 25x25. The implications of this effect require further analysis, but it seems to imply that the two methods could be fused more effectively: by using the FBLR results to limit the PCA-LDA-LR method search region, we limit the search space and improve the median error without impacting accuracy.

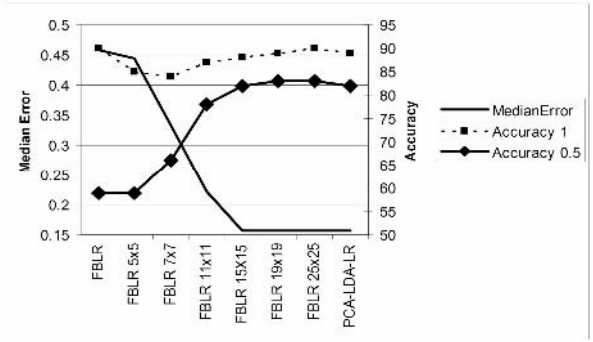

Fig. 5. Median error, and accuracy (number within one ON and 0.5 ON) expressed as a percentage of the total.

E. Example images

We show some results in Figure 6. The manually selected optic nerve location is shown with an "X", the FBLR position with a square, and the PCA-LDA-LR method with a cross. We note that in the example where both methods fail the likely reason is the large number of vertical striations in the region that attract the FBLR method through strong values for the density and orientation and the PCA-LDA-LR method through a high intensity combined with a definitive partition caused by the striations. Similarly, in the PCA-LDA-LR failure case the hard edge of the lesion seems to be a better fit in the likelihood ratio than the actual ON location. The FBLR failure shown can be attributed to the weak vasculature segmentation in the area of the optic nerve, creating disconnected components and weakening the density and orientation features considerably. As a final note, we found that both methods failed in only 2.3% of the cases.

IV. CONCLUSIONS

We believe our results show evidence for the benefits of including additional prior knowledge into the model-based approach with the LDA-PCA-LR results. Both FBLR and LDA-PCA-LR methods are more robust to the high degree of DR in our set because they incorporate many ON and non-ON pixel regions in training as well as the use of the *a priori* ON center distribution. Also, we showed that by combining the two methods we are able to achieve high performance (90% within 1 ON radius) and accuracy with smaller search spaces. One problem with fusing the results using the method suggested is there is still a dependence on the segmentation of the vasculature. We plan to continue this work by exploring more ways to fuse the localization methods, including establishing a confidence metric for the ON location which will permit more automatic screening for the ON location with minimal error and manual intervention.

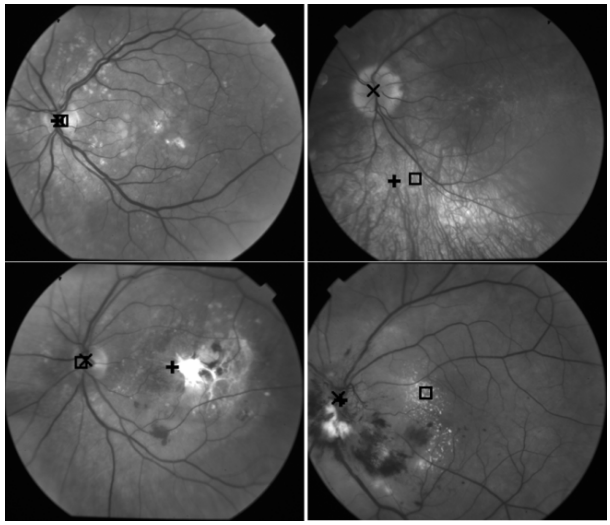

Fig. 6. Result examples. Upper left: both methods succeed. Upper right: both methods fail. Lower left: FBLR succeeds. Lower right: PCA-LDA-LR succeeds.

REFERENCES

[1] A.F. Amos, D.J. McCarty, P. Zimmet, "The rising global burden of diabetes and its complications: estimates and projections to the year 2010," Diabetic Med 1997;14: S57-85.

[2] Centers for Disease Control and Prevention (2003) National Diabetes Fact Sheet. (http://www.cdc.gov).

[3] C. Stellingwerf, P. Hardus, J. Hooymans, "Two-field photography can identify patients with vision-threatening diabetic retinopathy: a screening approach in the primary care setting," Diabetes Care. 2001, 24:2086-90.

[4] D.T. Morris, C. Donnison, "Identifying the neuroretinal rim boundary using dynamic contours," Image and Vision Computing, Vol. 17, 1999, p. 169-174.

[5] A. Hoover, M. Goldbaum, "Locating the optic nerve in a retinal image using the fuzzy convergence of the blood vessels," IEEE Tran. on Medical Imaging, Vol. 22, No. 8, Aug. 2003, p. 951-958.

[6] M. Foracchia, E. Grisan, A. Ruggeri, "Detection of the optic disc in retinal images by means of a geometrical model of vessel structure," IEEE Trans. on Medical Imaging, Vol. 23, No. 10, Oct. 2004, p. 1189-1195.

[7] H. Li, O. Chutatape, "Automated feature extraction in color retinal images by a model based approach," IEEE Trans Biomed Eng. 2004 Feb;51(2):246-54.

[8] Tobin, K.W., Chaum, E., Govindasamy, V.P., Karnowski, T.P., Sezer, O., "Characterization of the Optic Disk in Retinal Imagery using a Probabilistic Approach", SPIE International Symposium on Medical Imaging, San Diego, California, Proceedings of SPIE, Vol. 6144, February 2006.

[9] R.O. Duda, P.E. Hart, D.G. Stork, Pattern Classification, 2nd Ed., John Wiley and Sons, Inc, New York, 2001.