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 principle 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 ways to combine both methods 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. Computer assisted diagnostics using image analysis has the potential to provide a low-cost method using widely distributed systems.

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 red-free 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 principle 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 principle 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. Through 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 both methods to gain improved accuracy with minimal throughput reduction.

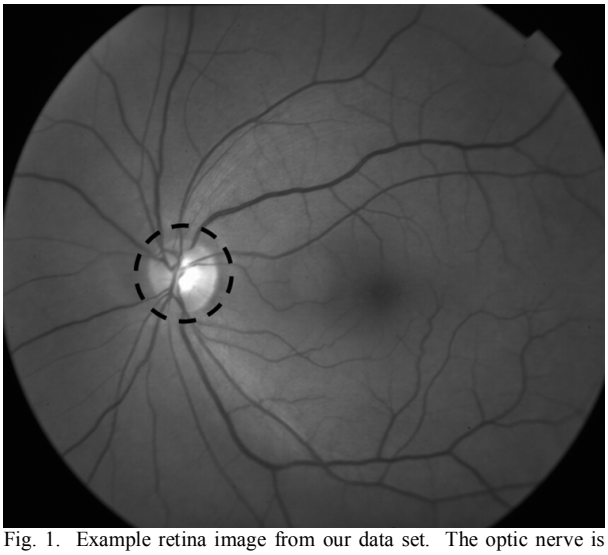

Fig. 1. Example retina image from our data set. The optic nerve is shown circled in black. The optic nerve is a key landmark in identifying the retinal structures.

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 principle 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 are projected to PCA space, then the coefficients are used to reconstruct the region and the residual error is measured. This process is repeated for a number of regions in the image and the pixel with the smallest residual error is chosen as the most likely optic nerve 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 candidate regions. The procedure used to implement the method is based on their work and is summarized as a *training* step and a *testing* step.

In the training stage, all images are aligned so that the optic nerve is on the left side of the image. A training set of manually segmented optic nerve images are selected. Each segmented optic nerve has been scaled to the same size (N pixels by N pixels) and the intensity has been normalized to extend over the full 8-bit range of the image. Next, the images of the training set are rasterized to form a vector of size N^2 x 1. The training set was standardized by computing the mean and standard deviation of each feature or pixel location in the training set. Finally, the principle components were computed. We retain those components that capture 90% of the image information as given by the eigenvalues as proposed in [7].

In the testing stage, a test image is processed by first picking out the most intense 1% of pixels. These pixels are labeled with connected component analysis and their area is measured. Connected components with size smaller than 0.004% of the entire image area are rejected as nuisance areas. Next, the remaining candidate regions are expanded by the size of the average optic nerve, increasing the number of candidate pixels that are potential ON centers. 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, its intensity is normalized, 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.

In our work, we evaluated three different candidate regions, the first given by [7], the second found by thresholding an estimate of the ON locations or *a priori* location model, and the third a combination of the two. The prior location model was formulated by taking 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.

B. Feature-Based Likelihood Ratio (FBLR) method

In the FBLR method, which is summarized here, the red-free 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 identified by the ON center and a surrounding area based on an estimate of the ON radius. A training set of data is analyzed to estimate the parameters of a Gaussian distribution describing the ON regions and the non-ON regions. These parameters are used to compute a likelihood ratio function. We also incorporate *a priori* information about the ON centers by using the training set to estimate the ON center probability density function (pdf).

C. PCA-LDA-LR method

We experimented with incorporating the prior information into the PCA method through expanding the candidate region as described. However, we were also interested in formulating a means of incorporating this information into the estimation process more directly. Therefore, we 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 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 likelihood ratio modeling the feature as a Gaussian random process,

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

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 red-free retinal images representing 19 different retinal pathologies. We manually aligned the images so that the ON was on the left side of the image. The ON was manually located for each image and segmented by hand. 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 each for a 2-fold validation study. 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. For the FBLR method, processing at a reduced resolution is conducted by estimating the vasculature at the highest resolution, resizing, and then generating the remaining features at the smaller resolution. We report our 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 or performance measurements 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 numbers below the average ON radius, indicating that the selected point is actually on the ON as opposed to outside it.

We conducted five different experiments with the data set, one each for the three different candidate regions for the PCA method, and one each for the FBLR and PCA-LDA-LR methods.

A. *Comparison of PCA method with expanded regions*

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 with 43 and 44 principle 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 12%, 28%, and 30% of the total image size, corresponding to roughly 2100, 4900, and 5200 pixels. Searching larger regions hurts the throughput in a linear fashion, but ideally should improve the accuracy of the method. Unfortunately this did not prove to be the case. Figure 2 shows a histogram of the distances from the actual ON center, normalized by the average ON radius. This plot shows that the original candidate region – which is also the smallest – outperforms the larger expanded regions somewhat, and has higher throughput. 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.

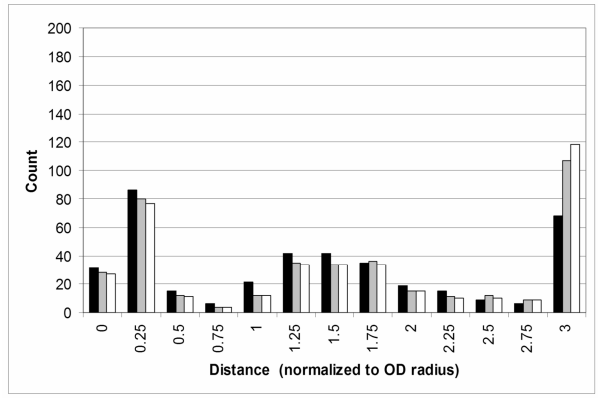

Fig. 2. Histograms of distances from actual ON for the PCA method, with different candidate regions. Black is the initial candidate region; gray, the prior-based region; and white is the combination of the two.

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.

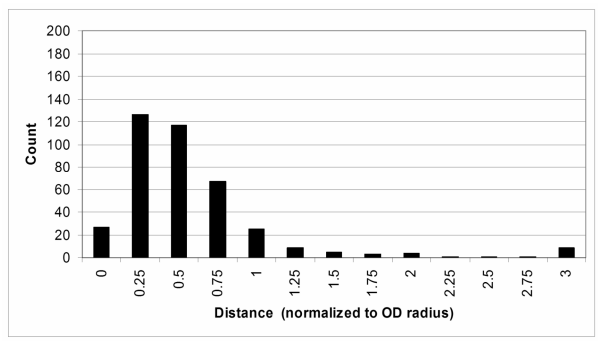

Fig. 3. Histograms of distances from actual ON for the FBLR method. Most (90%) of the sample data has ON located within one ON radius.

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, but the performance is superior to all methods tested 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. 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.

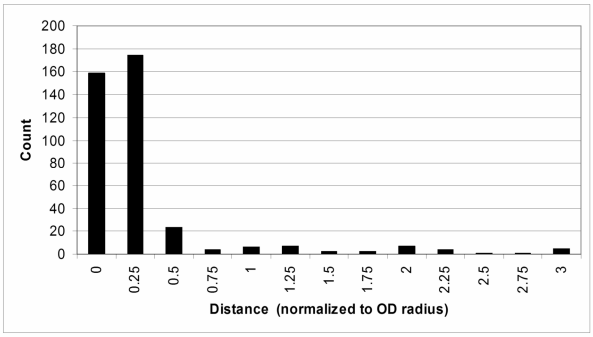

Fig. 4. Histograms of distances from actual ON for the PCA-LDA-LR method. Most (92%) of the sample data has ON located within one ON radius, with 89% within 0.5 ON radius.

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

Ideally we would like to use the high performance and throughput of the FBLR method with the higher accuracy of the PCA-LDA-LR method. We tested one strategy for combining these methods. We used the FBLR method to find the ON and treated this pixel as a candidate region. We then expanded the candidate region slightly and let the PCA-LDA-LR method "fine tune" the location. We tested candidate region sizes of 5, 9, 17 and 33 pixels square. We identify a candidate region as "FBLR 5x5", for exampling, meaning the center of the candidate region is found by performing FBLR, then a 5x5 region around that center is searched with PCA-LDA-LR. These results are shown in Figure 5 with three plots for the median value of the ON distance error, the number found within one ON radius, and a relative time comparison. The relative time comparison was created by taking the time to process an image with the FBLR and adding the time to process each candidate window, then normalizing to the time for the original candidate region as in the initial PCA experiments. Thus 1.0 is the full processing time for conducting only the PCA-LDA-LR method, while the minimum time is that for conducting only the FBLR method. We expect the accuracy increase as we expand the region, but we see that it peaks before the region is expanded to the full candidate region of PCA-LDA-LR. 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 improve throughput but accuracy as well.

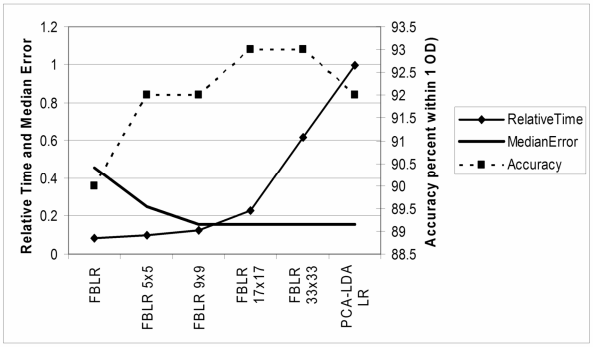

Fig. 5. Relative time, median error, and accuracy (number within one ON) expressed as a fraction of the total.

We believe these results show strong evidence for the benefits of combining the two methods, achieving high accuracy (over 90%) with high throughput. We plan to continue this work by exploring more ways to establish a confidence metric for the ON location, permitting true automatic screening for the ON location with minimal error.

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, p. 33.