

Automatic Exudate Detection with a Naive Bayes Classifier

Akara Sopharak¹, Khine Thet Nwe², Yin Aye Moe³, Matthew N. Dailey⁴, Bunyarit Uyyanonvara⁵

^{1,5}Sirindhorn International Institute of Technology, Thammasat University

^{2,3,4}Computer Science and Information Management, Asian Institute of Technology

¹akara@siit.tu.ac.th, ²khinethetnwe@gmail.com, ³yayemoe@gmail.com,

⁴mdailey@ait.ac.th, ⁵bunyarit@siit.tu.ac.th,

Abstract—Diabetic retinopathy is a major cause of vision loss for diabetic patients, but early detection of its symptoms and treatment can prevent blindness. Exudate is a key indicator of diabetic retinopathy that can potentially be automatically quantified. In this paper, which focuses on automatic exudate detection in images acquired through non-dilated pupils, we present a series of experiments on feature selection and exudate pixel classification using a naive Bayes classifier. We find that the best feature set is a combination of 6 features: the pixel intensity after preprocessing, the standard deviation of the preprocessed intensities in a window around the pixel, the pixel hue, the number of edge pixels in a window around the pixel, the ratio between the size of the pixel's intensity cluster and the optic disc, and the response at the pixel to a derivative of Gaussian filter. The overall sensitivity, specificity, precision, and accuracy are 93.38%, 98.14%, 47.51%, and 98.05%, respectively.

Keywords exudate, diabetic retinopathy, non-dilated retinal images, naive Bayes classifier

I. INTRODUCTION

Diabetic retinopathy is one of the main causes of vision loss, and its prevalence is rising. Screening diabetic patients for developing diabetic retinopathy can potentially reduce the risk of blindness. Diabetes patients need to have their eyes screened each year, but in many countries, there are too few ophthalmologists to meet the demand [1] - [3].

Diabetic retinopathy can lead to several retinal abnormalities, including microaneurysms, haemorrhages, cotton wool spots, and exudate. This paper focuses on exudate because it provides information about early diabetic retinopathy. The main cause of exudate is proteins and lipids leaking from the bloodstream into the retina via damaged blood vessels [1]. In retinal images, exudate manifests as hard white or yellowish localized regions with varying sizes, shapes, and locations.

There has been quite a bit of related research, but most of it is based on imagery acquired after dilating the pupils, e.g. with Tropicamide eye drops, in which exudate and other retinal features are clearly visible. Osareh et al. [1] use fuzzy *c*-means clustering to segment color retinal images into homogeneous regions, then train a neural network to separate exudate and non-exudate areas. Sanchez et al. [2] combine color and sharp edge features to detect exudate.

First they find yellowish objects, then they find objects that have sharp edges using various rotated versions of Kirsch masks on the green component of the original image. Yellowish objects with sharp edges are deemed to be exudate. Sinthanayothin et al. [4] perform automated detection of diabetic retinopathy on digital fundus images using a recursive region growing segmentation algorithm. Ege et al. [5] use a median filter to remove noise, segment bright lesions and dark lesions by thresholding, perform region growing, then identify exudate regions with Bayesian, Mahalanobis, and *k*-nearest neighbor classifiers. Li and Chutatape [6] propose an exudate extraction technique using a combination of region growing and edge detection techniques. Wang et al. [7] use color features as a feature space and a nearest neighbor classifier to identify retinal lesions. Walter et al. [8] use morphological reconstruction techniques to detect contours typical of exudates. Xiaohui and Chutatape [9] use local contrast enhancement and fuzzy *c*-means clustering in Luv color space to segment candidate bright lesion areas. They use a hierarchical support vector machine to classify bright non-lesion areas, exudate, and cotton wool spots.

Since pupil dilation takes time and is uncomfortable for patients, in our work, we investigate methods for automated exudate detection on imagery acquired without pupil dilation. The ultimate goal is a practical system to help ophthalmologists in the screening process.

In previous work, we have proposed and evaluated methods for automatic detection of exudate in non-dilated retinal images using mathematical morphology techniques [10], fuzzy *c*-means [11], and a combination of fuzzy *c*-means and mathematical morphology [12]. In experiments on comparable data sets, the sensitivity and specificity for these methods was 80% and 99.4% [10], 86% and 99% [11], and 92.18% and 91.52% [12]. While these results are encouraging, they are limited by suboptimal feature selection and pixel classification techniques.

Here we take a machine learning approach to the problem of exudate classification. Our method performs exudate feature selection and classification using a naive Bayes classifier. We filter images for noise, enhance image contrast, detect and remove the optic disc, extract local features describing pixels or regions, then classify those features using a model built from a training set. The naive Bayes classifier, after feature selection, achieves an overall per-pixel sensitivity of 93.38%, specificity of 98.14%,

precision of 47.51%, and an overall accuracy of 98.05% on a test set not used during training. The final classifier is a substantial improvement on previous work.

II. METHODOLOGY

We acquired 39 digital retinal images taken without pupil dilation from the Eye Care Center at Thammasat University Hospital. The images were captured in 24-bit color with a KOWA-7 non-mydratric retinal camera with a 45° field of view. We scaled all images to 752 x 500.

A. Preprocessing

We transformed each original RGB image into the HSI (hue, saturation, and intensity) color space. To reduce noise and improve image contrast, we applied median filtering and Contrast-Limited Adaptive Histogram Equalization (CLAHE) [13] to the I band.

B. Optic Disc Detection

As other researchers have noted [8], [9], [14], after preprocessing, the optic disc has some visual characteristics quite similar to those of hard exudate: bright intensity and sharp boundaries with the background. To prevent the optic disc from interfering with exudate detection, we first detect the optic disc and eliminate it from consideration as exudate during training and testing.

We find that the optic disc is easily distinguished from the rest of the retina by its smooth texture. To determine which regions of the image are smooth or textured, for each point x we obtain a probability mass function $P_{I_x} : \{0,1,\dots,255\} \mapsto [0..1]$ for the intensities I_x of the pixels in a local region around x then compute the entropy

$$H(I_x) = -\sum_{i=0}^{255} P_{I_x}(i) \cdot \log P_{I_x}(i) \quad (1)$$

This local pixel intensity entropy measure is high when the region around a pixel is complex and low when it is smooth. After filtering with the entropy operator we apply Otsu's binarization algorithm [15] to separate the complex regions from the smooth regions.

However, after binarization there is typically more than one candidate region identified. We select the largest connected component R_i whose shape is approximately circular. We measure circularity using the *compactness* measure

$$C(R_i) = 4\pi A(R_i) / P^2(R_i) \quad (2)$$

where $A(R_i)$ is the number of pixels in region R_i and $P(R_i)$ is the length of the boundary of R_i . To ensure that all pixels on the boundary of the optic disc are grouped with the optic disc region, we apply binary dilation to the detected region.



Fig. 1. Optic disc detection result (a) Entropy filtered image, (b) Optic disc area eliminated from the contrast enhanced image.

C. Feature Extraction

As an initial set of candidate per-pixel features, we first propose 15 features potentially able to distinguish exudate pixels from non-exudate pixels:

- 1) *The pixel's intensity value after preprocessing.*
- 2) *The standard deviation of the preprocessed intensity values in a window around the pixel.* We use a window size of 15 x 15.
- 3) *The pixel's hue.*
- 4) *The number of edge pixels in a region around the pixel.* We apply a Sobel edge operator then eliminate the strong edges arising from blood vessels and the optic disc using decorrelation stretch [16] on the red band. We use a 17 x 17 neighborhood.
- 5) *The average intensity of the pixel's cluster.* We cluster the preprocessed image into contiguous regions with similar intensity using agglomerative clustering [17]. This feature is simply the average of the intensities of the pixels assigned to the same cluster as the target pixel.
- 6) *The size (measured in pixels) of the pixel's cluster.*
- 7) *The average intensity of the pixels in the neighborhood of the pixel's cluster.* The pixel's cluster is extracted and dilated using a 3 x 3 structuring matrix. The neighborhood pixels are those obtained by subtracting the original region from the dilated region.
- 8) *The ratio between the size of the pixel's cluster and the size of the optic disc.*
- 9) *The distance between the pixel's cluster and the optic disc.* We use the Euclidean distance between the centroid of the pixel's cluster and the centroid of the optic disc.
- 10) *Six Difference of Gaussian (DoG) filter responses.*

The DoG filter subtracts one blurred version of an original image from another blurred version of the image [18]. We convolve with seven different Gaussian kernels with standard deviations of 0.5, 1, 2, 4, 8, 16, and 32. We use DoG1, DoG2, DoG3, DoG4, DoG5 and DoG6 to refer to the features obtained by subtracting the image at scale $\sigma = 0.5$ from scale $\sigma = 1$, scale $\sigma = 1$ from $\sigma = 2$, scale $\sigma = 2$ from $\sigma = 4$, scale $\sigma = 4$ from $\sigma = 8$, scale $\sigma = 8$ from $\sigma = 16$, and scale $\sigma = 16$ from $\sigma = 32$, respectively.

Before feature selection or classification, we z-scale (transform to a mean of 0 and a standard deviation of 1) all 15 features using the statistics of each feature over the training set. Examples of some of the features are shown in Fig. 2.

D. Feature Selection and Classification Using Naive Bayes

The naive Bayes classifier [19] - [21] uses the principle of Bayesian maximum a posteriori (MAP) classification: measure a finite set of features $\mathbf{x} = (x_1, \dots, x_n)$ then select the class

$$\hat{y} = \arg \max_y P(y|\mathbf{x})$$

where

$$P(y|\mathbf{x}) \propto P(\mathbf{x}|y)P(y) \quad (3)$$




Fig. 2. Input features. (a) Standard deviation of intensity. (b) Hue. (c) Number of edge pixels. (d) Cluster intensity. (e) DoG1. (f) DoG2. (g) DoG3. (h) DoG4. (i) DoG5. (j) DoG6.

$P(\mathbf{x}|y)$ is the likelihood of feature vector \mathbf{x} given class y , and $P(y)$ is the priori probability of class y . Naive Bayes assumes that the features are conditionally independent given the class:

$$P(\mathbf{x}|y) = \prod_i P(x_i|y)$$

We estimate the parameters $P(x_i|y)$ and $P(y)$ from training data.

After z-scaling, all of our features x_i are continuous, but the simple version of naive Bayes just described requires discrete features, so we perform unsupervised proportional k -interval discretization as implemented in Weka [22]. The technique uses equal-frequency binning, where the number of bins is the square root of the number of values.

Feature selection proceeds as follows. We first estimate the model of Equation 3 from a training set using all features then evaluate the resulting classifier's performance on a separate test set. Then we iteratively delete features until the average of the precision and recall (PR, see next section) stops improving. On each step, for each feature, we delete that feature from the model, train a new classifier, and evaluate its performance on the test set. The PR of the best such classifier is compared to the PR of the classifier without deleted features. If PR improves, we permanently delete that feature then repeat the process. Finally, the best feature set and classifier are retained.

E. Performance Measurement

We evaluate performance on the test set quantitatively by comparing the classifier's result to ground truth. To obtain ground truth, for each image, we used image processing software to hand label candidate exudate regions, then we asked an ophthalmologist to verify or reject each candidate region. We split the 39-image data set into a training set containing 29 images and a test set containing 10 images. To evaluate classifier performance, we use sensitivity (recall),

specificity, precision, PR and accuracy. Sensitivity (recall) is the percentage of the actual exudate pixels that are detected; specificity is the percentage of non-exudate pixels that are correctly classified as non-exudate pixels. Precision is the percentage of detected pixels that are actually exudate; PR is the average of the precision and recall. Accuracy is the overall per-pixel success rate of the classifier.

III. RESULTS

As a first experiment, we fit the naive Bayes model (Equation 3) to the training set using all 15 features. The resulting classifier had overall per-pixel sensitivity, specificity, precision, PR and accuracy of 95.84%, 96.56%, 33.49%, 64.67% and 96.55%, respectively. On the next step, we removed features from the classifier one at a time and compared each resulting PR value to the previous feature set's performance. We obtained the best PR value (65.78%) by deleting cluster intensity, presumably due to its redundancy with the pixel intensity feature. We continued this process until the PR stopped improving. We obtained the final best-performing classifier by deleting cluster intensity, DoG3, neighborhood intensity, cluster size, DoG5, DoG2, distance between the cluster and optic disc, DoG6, and DoG1. Finally, the best-performing classifier contained six features: the pixel's intensity after preprocessing, the standard deviation of the preprocessed intensities in a window around the pixel, the pixel hue, the number of edge pixels in a window around the pixel, the ratio between the size of the pixel's intensity cluster and the optic disc, and DoG4. The complete test results are listed in Table 1.

Fig. 3 shows an example of the exudate pixels detected by the system in comparison with ground truth.




Fig. 3. Comparison of exudates detection (a) Detected image, (b) Ground truth image.

IV. DISCUSSION AND CONCLUSION

In this paper, we proposed 15 per-pixel features as potentially indicative of exudate and used a naive Bayes classifier for feature selection and pixel classification.

The final classifier contained six features: the pixel's intensity after preprocessing, the standard deviation of the preprocessed intensities in a window around the pixel, the pixel hue, the number of edge pixels in a window around the pixel, the ratio between the size of the pixel's intensity cluster and the optic disc, and the response at the pixel to a derivative of Gaussian filter. The final classifier's overall sensitivity, specificity, precision, PR, and accuracy, on a test set not used for training, were 93.38%, 98.14%, 47.51%, 70.45%, and 98.05%, respectively.

Our results show that careful preprocessing, good

features, and an appropriate classifier together provide excellent exudate detection performance even on retinal images acquired from patients whose pupils have not been dilated. In future work we plan to expand the data set and explore using the system as a practical aid to help ophthalmologists screen patients for diabetic retinopathy symptoms quickly and easily.

ACKNOWLEDGMENTS

We would like to thank the Eye Care Center at Thammasat University Hospital for supplying the images used in this work. This work was partly supported by Thailand Research Fund grant #MRG4780209 to MND.

REFERENCES

- [1] A. Osareh, M. Mirmehdi, B. Thomas, and R. Markham, "Automated Identification of Diabetic Retinal Exudates in Digital Colour Images," *Br J Ophthalmol*. Vol. 87, pp. 1220-1223, 2003.
- [2] C.I. Sanchez, R. Hornero, M.I. Lopez, and J. Poza, "Retinal Image Analysis to Detect and Quantify Lesions Associated with Diabetic Retinopathy," In : *Internat. Conf. on Engineering in Medicine and Biology Society (EMBC)*, pp. 1624 – 1627, 2004.
- [3] J.A. Olson, F.M Strachana, and J.H. Hipwell, "A comparative evaluation of digital imaging, retinal photography and optometrist examination in screening for diabetic retinopathy," *Diabet Med*. Vol. 20, pp. 528-534, 2003.
- [4] C. Sinthanayothin, J.F. Boyce, T.H. Williamson, and H.L. Cook, Automated Detection of Diabetic Retinopathy on Digital Fundus Image, *Diabet Med*. Vol. 19, pp. 105-112, 2002.
- [5] B.M. Ege, O.K. Hejlese, O.V. Larsen, K. Moller, B. Jennings, D. Kerr et al, "Screening for diabetic retinopathy using computer based image analysis and statistical classification," *Comput Meth Programs Biomed*. Vol. 62, pp. 165-175, 2000.
- [6] Huiqi Li and Opas Chutatape, "A model-based approach for automated feature extraction in fundus images," *Internat. Conf. on Computer Vision (ICCV)*, pp. 394-399, 2003..
- [7] H. Wang, W. Hsu, K.G. Goh, and M.L. Lee, "An effective approach to detect lesions in color retinal image," *IEEE Conf. on Computer Vision and Pattern Recognition*, Vol. 2, pp. 181-186, 2000.
- [8] T. Walter, J.C. Klevin, P. Massin, 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*. Vol. 21, pp. 1236-1243, 2002.
- [9] Zhang Xiaohui and Opas Chutatape, "Detection and Classification of Bright Lesions in Colour Fundus Images," *Proceeding of IEEE International conference on Image Processing (ICIP)*, Vol.1, pp.139-142, 2004.
- [10] A. Sopharak and B. Uyyanonvara, "Automatic Exudates Detection on Thai Diabetic Retinopathy Patients' Retinal Images," *Proceedings of the 2006 ECTI International Conference*, pp.709-712, 2006.
- [11] A. Sopharak and B. Uyyanonvara, "Automatic Exudates Detection from Non-dilated Diabetic Retinopathy Retinal Image Using Fuzzy C-Means Clustering," *Proceedings of the Third WACBE World Congress on Bioengineering*, 2007.
- [12] A. Sopharak and B. Uyyanonvara, "Automatic Exudates Detection from Diabetic Retinopathy Retinal Image Using Fuzzy C-Means and Morphological Methods," *3rd IASTED International Conference on Advances in Computer Science and Technology*, pp. 359-364, 2007.
- [13] R.C. Gonzales and R.E. Woods, *Digital Image Processing*, Addison-Wesley publishing Co., 1993, pp. 88-103, 519-634.
- [14] C. Sinthanayothin, J.F. Boyce, H.L. Cook, and T.H. Williamson, "Automated Localization of the Optic Disc, Fovea, and Retinal Blood Vessels from Digital Colour Fundus Images," *Br J Ophthalmol*. Vol. 83, pp. 231-238, 1999.
- [15] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Trans. on Syst Man and Cybern*. SMC-9 (1979) 62–66.
- [16] Son Lam Phung, Abdesselam Bouzerdoum, and Douglas Chai, "Skin Segmentation Using Color Pixel Classification: Analysis and Comparison," *Proc. IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 27, pp. 148 – 154, 2005.
- [17] D.A. Forsyth, and J. Ponce, *Computer Vision: A Modern Approach*, Prentice Hall, 2003, pp.433-462.
- [18] R.S. Kenneth, C.R. John, J. Matthew et al. (2006, June 15). Difference of Gaussians Edge Enhancement [Online]. Available: http://micro.magnet.fsu.edu/primer/java/digitalimaging/processing/diff_gaussians/index.html
- [19] N. Friedman, D. Geiger, and M. Goldszmidt, "Bayesian network classifiers," *Machine Learning*. Vol. 29, pp.131-163, 1997.
- [20] O.D. Richard, E.H. Peter, and G.S. David, *Pattern Classification 2nd edition*, A Wiley-Interscience Publication, 2000, pp. 20-83.
- [21] X.Y. Wang, J. Garibaldi, and T. Ozen, "Application of The Fuzzy C-Means clustering Method on the Analysis of non Pre-processed FTIR Data for Cancer Diagnosis," *Internat. Conf. on Australian and New Zealand Intelligent Information Systems (ANZIS)*, pp. 233-238, 2003.
- [22] I.H. Witten and E. Frank, *Data Mining: Practical machine learning tools and techniques*. 2nd edition. Morgan Kaufmann, 2005.

Table1. Naive Bayes performance results.

| Features | SE | SP | Precision | PR | Accuracy |
|---|--------|--------|-----------|--------|----------|
| All features | 95.84% | 96.56% | 33.49% | 64.67% | 96.55% |
| Without CI | 95.68% | 96.91% | 35.89% | 65.78% | 96.89% |
| Without CI and DoG 3 | 95.58% | 97.12% | 37.51% | 66.55% | 97.09% |
| Without CI, DoG 3 and NI | 95.20% | 97.36% | 39.49% | 67.35% | 97.33% |
| Without CI, DoG 3, NI and CS | 96.10% | 97.41% | 40.17% | 68.14% | 97.39% |
| Without CI, DoG 3, NI, CS and DoG 5 | 95.76% | 97.62% | 42.12% | 68.94% | 97.59% |
| Without CI, DoG 3, NI, CS, DoG and DoG 2 | 95.41% | 97.80% | 43.98% | 69.70% | 97.76% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2 and distance | 94.82% | 97.95% | 45.56% | 70.19% | 97.90% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2, distance and DoG 6 | 93.82% | 98.06% | 46.62% | 70.22% | 97.98% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2, distance, DoG 6 and DoG 1 | 93.38% | 98.14% | 47.51% | 70.45% | 98.05% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2, distance, DoG 6, DoG 1 and NE | 92.12% | 96.46% | 31.98% | 62.05% | 96.38% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2, distance, DoG 6, DoG 1 and SD | 90.03% | 97.70% | 41.48% | 65.76% | 97.57% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2, distance, DoG 6, DoG 1 and SR | 93.70% | 97.52% | 40.61% | 67.16% | 97.46% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2, distance, DoG 6, DoG 1 and intensity | 92.08% | 97.81% | 43.18% | 67.63% | 97.71% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2, distance, DoG 6, DoG 1 and hue | 94.14% | 97.77% | 43.27% | 68.70% | 97.70% |
| Without CI, DoG 3, NI, CS, DoG 5, DoG 2, distance, DoG 6, DoG 1 and DoG 4 | 91.42% | 98.11% | 91.42% | 69.01% | 97.99% |

*SE= sensitivity, SP = specificity, CI = cluster intensity, NI = neighborhood intensity, CS = cluster size, NE = number of edge pixels, SD = standard deviation and SR = size ratio