Region of Interest Detection in Fundus Images Using Deep Learning and Blood Vessel Information

Jongwoo Kim, Sema Candemir, and George R. Thoma
Lister Hill National Center for Biomedical Communications
National Library of Medicine
Bethesda, Maryland, USA
jongkim@mail.nih.gov

Emily Y. Chew
Division of Epidemiology & Clinical Applications,
National Eye Institute
National Institutes of Health
Bethesda, Maryland, USA
echew@nei.nih.gov

Abstract— **Ophthalmologists use the optic disc to cup ratio as one of the factors to diagnose glaucoma. The region of interest (ROI) for glaucoma in fundus images is the area that locates optic disc and cup in the center. Therefore, ROI detection is used as a preprocessing step for automatic detection of optic disc and cup areas. This paper proposes an automated method to detect ROI using deep learning. Convolutional Neural Networks (CNNs) are used to classify ROI and non-ROI images. The structure of our CNNs is composed of two convolutional layers, two Max Pooling layers, two fully connected layers, and one output layer. We train two CNNs using fundus images from the MESSIDOR dataset, a public dataset containing 1,200 fundus images. In addition, we estimate blood vessels from the images and use the images embedded with the blood vessels to train two other CNNs. The proposed method moves a window in the horizontal and vertical directions in each fundus image, estimates a probability of each window using the CNNs, and selects the window with the highest probability as ROI. The experimental results are promising. The best-performing CNN from the first CNN group shows over 0.99 accuracy for the MESSIDOR dataset and over 0.93 accuracy for five other public fundus image datasets. The best CNN from the second CNN group shows more robust results: over 0.99 accuracy for the MESSIDOR dataset and over 0.97 accuracy for the five other image datasets.**

Keywords-Region of Interest (ROI); Deep Learning; Convolutional Neural Network (CNN); Optic Disc; Cup; Blood Vessels; Glaucoma.

I. INTRODUCTION

Glaucoma is a serious disease that is caused by damage to the optic nerve, resulting in vision loss. Increased pressure in the eye and/or loss of blood flow to the optic nerve cause nerve fibers to begin to die and this makes the cup become larger in comparison to the optic disc. A cup to disc ratio (the diameter of the cup divided by the diameter of the optic disc) greater than 30% is considered to be suspicious for glaucoma [1]. There are several technical papers that process fundus images to estimate region of interest (ROI), the optic disc, cup, blood vessels, and eye diseases. ROI is the area (in the image) where the optic disc is in the center.

ROI detection is used as a preprocessing step to estimate optic disc and cup areas for glaucoma. In addition, it can be used to distinguish lesions from the optic disc for age-related macular degeneration (AMD) since both lesions and the optic disc have high (bright) pixel values in fundus images.

The sum of high pixel values (e.g., the highest 0.5% in a fundus image) within an area is commonly used to estimate ROI [2, 3] since the optic disc usually has the highest pixel values in fundus images. However, the rim and lesion also can have high pixel values. A circular shape filter is used to remove bright fringes in the rim areas from the optic disc candidates [2]. However, lesions are not removed by the filter. Therefore, blood vessels are used as key information to estimate ROI since the blood vessels are emanating from the optic disc [3, 4, 5, 8].

There are several methods to segment blood vessels in fundus images. Smoothing, low pass filtering, Otsu threshold, and morphology (Opening, Closing, Top-Hat transformation) are commonly used to differentiate blood vessels from surrounding structures [6, 7, 8]. Image enhancement is frequently used as a preprocessing step to normalize images [8, 9] and to improve estimation results. However, it is difficult to find blood vessel detection and image enhancement algorithms that work consistently across images from different sources. The enhancement can degrade the original images. Some fundus images have lower intensity pixels in optic disc areas, making it difficult to recognize blood vessels passing through the optic disc. Therefore, it could be a problem if algorithms depend heavily on the blood vessels as a tool for ROI detection.

Deep learning has been used in fundus image analysis: vessel segmentation [9], disc and cup segmentation [10], etc. Unlike other classifiers, deep learning uses original images as input and extracts features from the images. Therefore, it is not necessary to develop handcrafted features. We propose a deep learning algorithm that uses a Convolutional Neural Network (CNN) to estimate ROI from fundus images. The training datasets are collected by cropping the original images (without preprocessing).

All machine learning algorithms, including CNN, need a large dataset to train the algorithm more robustly. Collecting the dataset and training the algorithms are very time consuming. Therefore, we also propose to use a feature (blood vessels) to efficiently train the CNNs using a limited number of training images.

The remainder of this paper is organized as follows. Section II describes overview of the paper. Section III describes vessel segmentation. The details of our CNNs and ROI detection algorithm are presented in Section IV and V. We discuss experimental results in Section VI, and show conclusions in Section VII.

II. OVERVIEW

We use two types of datasets to train the CNNs to detect ROIs: (1) original fundus image dataset, and (2) an additional dataset that contains the original fundus images embedded with blood vessel information.

The procedure is as follows. First, train the CNNs using a dataset collected from original fundus images. This training set will be used to classify ROI and non-ROI images. Second, estimate blood vessels from the original images. Third, train other CNNs using the training dataset collected from the images embedded with the blood vessels. Last, apply each CNN to classify ROI in fundus images and compare the performance of the CNNs.

III. VESSEL SEGMENTATION

ROI in fundus images contains the optic disc in the middle with blood vessels emanating from the optic disc. Therefore, blood vessels are an important feature to estimate ROI in the images. The following method (step sequence) is used to estimate the main blood vessels from fundus images.

1. Let $I(x, y)$ be an input image, $I_g(x, y)$ be the green channel of $I(x, y)$, and $I_r(x, y)$ be the red channel of $I(x, y)$.
2. Estimate a binary image $I_{rb}(x, y)$ from $I_r(x, y)$ using Otsu threshold to distinguish the retina area from background.
3. Estimate an edge image $I_{rbe}(x, y)$ from $I_{rb}(x, y)$ using Canny operator, and estimate an ellipse from $I_{rbe}(x, y)$ using OpenCV [11] to estimate the longest diameter of the ellipse d. d is used to decide the filter sizes and ROI window size later.
4. Apply median operator with window size [3×3] and CLAHE histogram equalization operator to $I(x, y)$ to have a contrast enhanced image $I_c(x, y)$.
5. Apply the average operator to $I_c(x, y)$ with window size = $d/(3.5×4)$ to have $I_{cb}(x, y)$. The window size is close to a third or a half of the optic disc diameter.
6. Subtract $I_c(x, y)$ from $I_{cb}(x, y)$ to have a blood vessel image $I_d(x, y)$.

$$I_d(x, y) = I_{cb}(x, y)\text{-}I_c(x, y), \text{ if } I_{cb}(x, y)>I_c(x, y) \text{ and } I_{rb}(x, y)>0$$
$$0, \qquad\qquad \text{Otherwise}$$

7. Apply morphological closing using [3×3] square to $I_d(x, y)$ to have $I_{dc}(x, y)$. The closing fills missing blood vessel pixels inside the vessels.
8. Normalize $I_{dc}(x, y)$ from 0 to 255 to have $I_{dcn}(x, y)$.
9. Apply Otsu threshold to $I_{dcn}(x, y)$ to have a binary image $I_b(x, y)$. This step removes pixels that belong to minor blood vessels and noise.
10. Select the top five large blood vessel blob pieces and remove other pieces from $I_b(x, y)$. Remove pieces when

number of pixels in the piece is less than 20. Let $I_{vb}(x, y)$ be the binary blood vessel image after removing useless blood vessel pieces in $I_b(x, y)$.
11. Estimate a gray level blood vessel image $I_{vg}(x, y)$.

$$I_{vg}(x, y) = 255 - I_{dcn}(x, y), \text{ if } I_{vb}(x, y) > 0$$
$$0, \qquad\qquad \text{Otherwise}$$

12. Generate an image $I_V(x, y)$ by embedding the blood vessels ($I_{vg}(x, y)$) to original fundus images $I(x, y)$. The blood vessel pixels are added in the channels if $I_g(x, y) > I_{vg}(x, y)$.

$$I_{Vi}(x, y) = I_{vg}(x, y), \text{ if } I_{vg}(x, y) < I_g(x, y)$$
$$I_i(x, y), \text{ Otherwise}$$

where $I_i(x, y)$ is the i channel of $I(x, y)$, $I_{Vi}(x, y)$ means the i channel of $I_V(x, y)$, and i = b, g, and r. I.e., $I_g(x, y)$ means the green channel of $I(x, y)$ and $I_{Vb}(x, y)$ means the blue channel of $I_V(x, y)$.

Fig. 1 shows fundus images from MESSIDOR. Fig. 2 shows the estimated blood vessel images from Fig 1. Fig. 3 shows fundus images embedded with the blood vessels (Fig. 2). $I_V(x, y)$ has $I_{vg}(x, y)$ in the blue, green, and red channels of the Fig. 1.

Most ROI detection algorithms and optic disc detection algorithms [6, 7, 8] use binary level blood vessels as a feature. Therefore, the threshold for the blood vessel binarization plays a critical role for the performance of the algorithms. We use grey level blood vessels to alleviate the blood vessel detection issue of the binarization process.
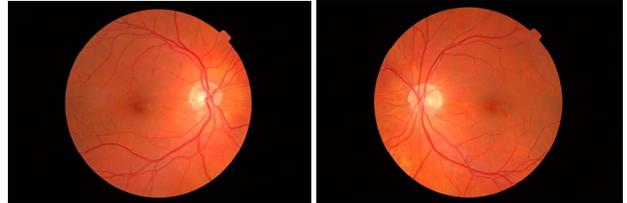
Figure 1. Fundus images ($I(x, y)$).
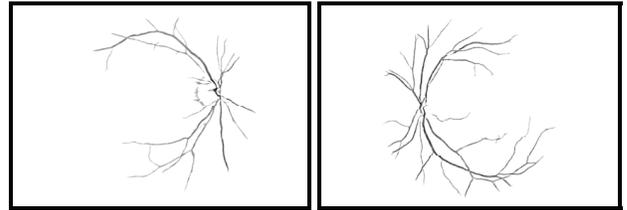
Figure 2. Estimated grey level blood vessels from Fig. 1 ($I_{vg}(x, y)$).
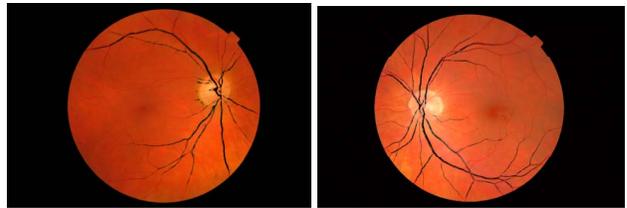
Figure 3. Fundus images embedded with grey level blood vessels (Fig. 2) in the three channels ($I_V(x, y)$).

IV. CONVOLUTIONAL NEURAL NETWORKS

CNN is one of the deep learning algorithms actively used in image classification and segmentation. Therefore, we use the CNN for the classification of ROI images. The architecture of the CNN used in our experiments is shown in Table I. We use two different kernel sizes M (5 and 11) and two different strides N (1 and 2) for the first convolutional layers. Stride is defined as the amount of shift applied to a filter on each step as the filter traverses the image. The kernel filter is used to extract features on each step. We train the CNN using two different training datasets as we discuss in Section II. In the table the output size 80 in layers 1 and 2 means we use 80 filters to extract features. The total feature size depends on M and N. For example, if M=11 and N=1 in layer 1, the total feature size becomes 80×90×90. "Output size 3000" in the Fully Connected layers (5 and 6) means we use a 1×3000 vector for the layers. Figure 4 shows the architecture of the CNN.

TABLE I. CNN ARCHITECTURE USED

Layer	Layer Type	Output Size	Kernel Size	Stride
0	Input (100×100×3)			
1	Convolutional	80	$M×M$	N, N
2	Max Pooling	80	2×2	2, 2
3	Convolutional	160	5×5	1, 1
4	Max Pooling	160	22	2, 2
5	Fully Connected	3000		
	Dropout (0.5)			
6	Fully Connected	3000		
	Dropout (0.5)			
7	Output	2		
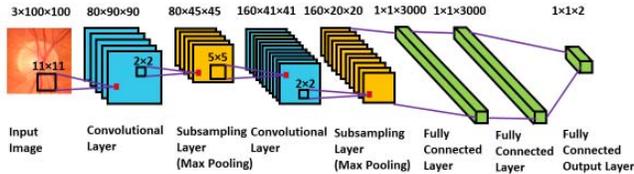
Figure 4. CNN architecture used in the experiments (M=11 and N=1).

V. REGION OF INTEREST (ROI) DETECTION

There are several steps to estimate ROI from fundus images after training CNNs. Fig. 5 shows the pixel values of three channels in the green horizontal line in the fundus image. The red channel shows unique values in the retina and the green channel shows the distinction between background, retina, and disc areas. The blue channel shows similar values as the green channel. Therefore, we use the red channel to classify the retina area from the background. Fig. 6 shows the workflow of the proposed method when original fundus images (Fig. 1) are used as input.

1. Extract the red channel from an input image and apply the average operator with window size [9×9] to remove noise.
2. Binaries the red channel image with Otsu threshold.
3. Estimate an edge image from the binary image using Canny operator.

4. Estimate an ellipse from the edge image using an ellipse detection function in OpenCV.
5. Set the window size W = longest diameter of the ellipse/3.5 and stride $S = W/4$.
6. Move $W×W$ window to the horizontal and vertical directions by S, estimate a CNN result of each window, and choose a window W_1 that has the highest result.
7. Let top-left coordinate of W_1 be (x_1, y_1) and bottom-right coordinate of W_1 be (x_2, y_2).
8. Move $W×W$ window from (x_1-$W/2$, y_1-$W/2$) to (x_2+$W/2$, y_2+$W/2$) by stride $S = W/16$, estimate a CNN result of each $W×W$ window, and choose the window with the highest result as ROI to fine tune the final ROI location. If more than one window has the same highest result, we average the window coordinates to have the final ROI.

Fig. 7 shows an example of a final result. The blue color window is ROI at the sixth step and the green color window is the final ROI estimated at the eighth step. The green color window has the optic disc more in the center than the blue color window.
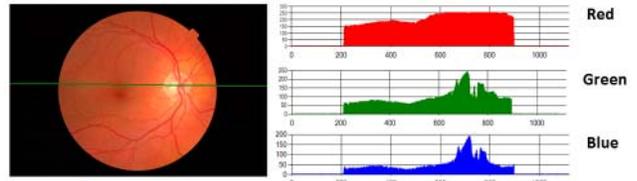
Figure 5. Pixel values of the three channels in the horizontal line in the fundus image.
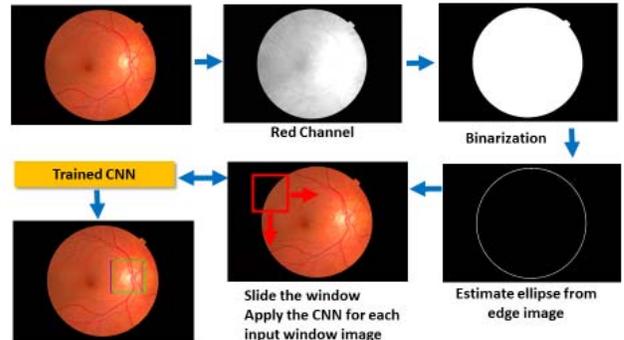
Figure 6. Workflow of the ROI detection method.
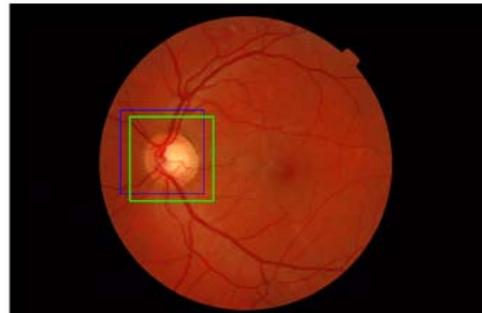
Figure 7. An example of ROI detection result of the proposed method.

VI. EXPERIMENTAL RESULTS

There are two steps in the experiments. First, train four CNNs. Second, estimate ROI from fundus images using the CNNs. The following image datasets are used in these experiments: MESSIDOR [12], Chase_DB1 [13], DIARETDB0 [14], DIARETDB1 [14], DRIONS [15], and DRIVE [16].

There are 1,200 fundus images in the MESSIDOR dataset. Among them, 600 images are used to train the CNNs for ROI image classification (for the first step) and the remaining 600 are used for testing (for the second step). Fig. 8 shows some ROI images used for training that are cropped from the 600 training images. Top images (first row) are for ROI class, and bottom images (second row) are for non-ROI class. 1,200 images are collected for each class. For ROI class, we collect the images two times from each fundus image manually. For non-ROI class, we manually collect 600 images. We then randomly select 600 additional images, using a random number generator. Fundus images can have different sizes (different widths and heights). Therefore, we use $W \times W$ window to collect images for the training dataset as we discuss in Section V. We use 400×400 window for images with size of 2240×1488 and 250×250 window for images with size of 1440×960 window to collect images for the training dataset and normalize them to 100×100. Fig. 9 shows some training images collected from fundus images embedded with grey level blood vessels (Fig. 3). We use the same sampling window coordinates to collect images in the same locations for the two different datasets as shown in Figs. 8 and 9.
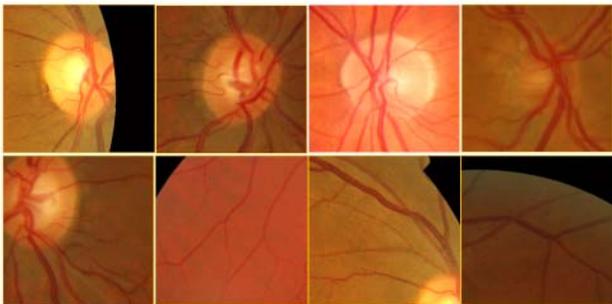
Figure 8. Training Images. (Top) Images for ROI class. (Bottom) Images for non-ROI class.
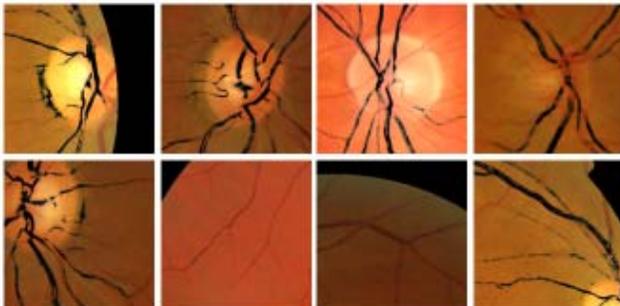
Figure 9. Training Images embedded with grey level blood vessels in three channels. (Top) Images for the ROI class. (Bottom) Images for the non-ROI class.

We train the CNNs without pretrained models and use 5,000 epochs, 0.01 learning rate, Stochastic Gradient Descent, DeepLearning4J [17], Nivida GTX 1080, and a Dell Precision T5810 with Intel® Xeon® CPU E5-1620 v3 @3.50GHz to train and test the CNNs.

We use two different datasets to test the CNNs. The first dataset contains 600 images from the MESSIDOR dataset. The second dataset contains 398 images (called Open Sources) from Chase_DB1, DIARETDB0, DIARETDB1, DRIONS, and DRIVE datasets. The second dataset is not used for training the CNNs. It is used only for testing the CNNs as we mentioned at the beginning of this section.

We first test the CNNs trained using original fundus images (Fig. 8). The results are shown in the first and second rows in Table II. CNN1 in the first row uses the 5×5 kernel in the first Convolutional layer and CNN2 in the second row uses 11×11 kernel. The results are promising. CNN1 and CNN2 show 0.9950 and 0.9883 accuracies for the MESSIDOR dataset as shown in the fourth column. Fig. 10 shows some of CNN1 results. Fig. 10(a) has small lesions. Fig. 10(b) has the optic disc close to the background. In Figs. 10(c) and (d), the optic discs are not as clear as Fig. 10(a), and blood vessels in optic discs are not clear. However, all results show that ROI windows (green window) have the optic discs in the center. CNN1 has errors in three images as shown in Fig. 11. We consider that an ROI location error has occurred when the optic disc is not inside the ROI boundary (Green square). All images have a lesion in or near the optic disc areas. The fifth column shows the test results of the CNNs for the Open Sources dataset. CNN1 has 0.9397 accuracy and CNN2 has 0.9497 accuracy. The accuracies drop 5.79% and 3.9%. CNN1 has errors in 25 images and CNN2 has errors in 20 images. Fig. 12 shows CNN1 errors from images in the Open Sources dataset. In the figure, blue color window is the pre-estimated ROI and green color window is the final ROI as we mentioned in Section V. Since CNN1 and CNN2 are trained on the MESSIDOR dataset and use a limited number of images for training, the CNNs could generate more errors for input images from other sources having different image characteristics. This problem can be resolved if we train the CNNs using fundus images from different sources.

The third and fourth rows in Table II show the results of CNNs trained using the images embedded with blood vessels. CNN3 and CNN4 use training images in Fig. 9 that have grey level blood vessels. Each CNN uses different kernel sizes (5×5 and 11×11) for training. The fourth column shows the test results for the MESSIDOR dataset. CNN3 and CNN4 show 0.9933 accuracy. They have slightly lower accuracy than CNN1. CNN1 has three errors, and CNN3 and CNN4 have four errors. The fifth column shows the test results for the Open Sources dataset. CNN4 shows the best accuracy with 0.9774.

TABLE II.   ACCURACY OF ROI DETECTION

CNN	Input Image	CNN Model First conv. layer	MESSIDOR Test set (600)	Open Sources Test set (398)
1	Original Images	Kernel (5,5) Stride (2, 2)	**0.9950** (3 errors)	0.9397 (25 errors)
2	Original Images	Kernel (11,11) Stride (1,1)	0.9883 (7 errors)	0.9497 (20 errors)
3	Images embedded with Grey Blood Vessels	Kernel (5,5) Stride (2, 2)	0.9933 (4 errors)	0.9673 (13 errors)
4	Images embedded with Grey Blood Vessels	Kernel (11,11) Stride (1, 1)	0.9933 (4 errors)	**0.9774** (9 errors)
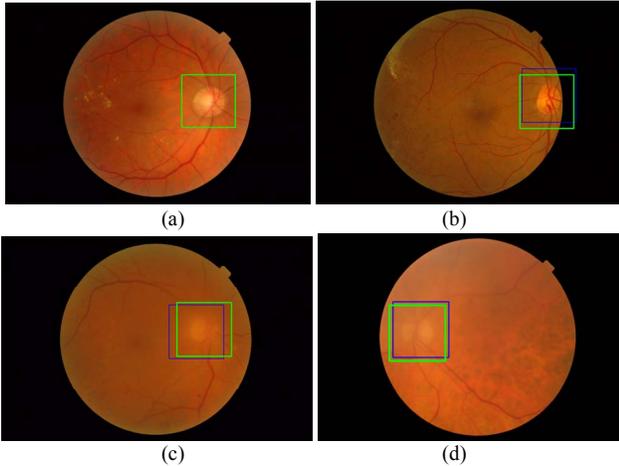
Figure 10. ROI Detection results of CNN1 from images in the MESSIDOR dataset.
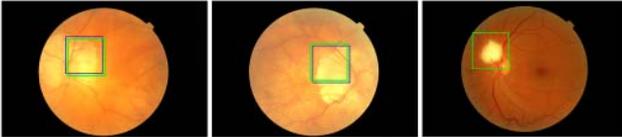
Figure 11. ROI Detection errors of CNN1 from images in the MESSIDOR dataset.
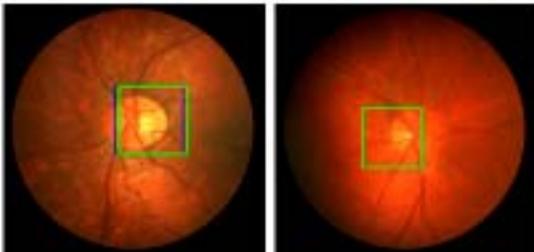
Figure 12. ROI Detection error of CNN1 from images in the Open Sources dataset.

Fig. 13 shows ROI estimation results of CNN4 for images in the Open Sources dataset. CNN4 correctly find ROIs even if there are lesions in the image or the optic disc has lower pixel values. Fig. 14 shows some CNN4 errors from images in the Open Sources dataset. Fig. 14(a) shows that

the center of the ROI window was moved to the right side. The right side of the disc rim area is labeled as blood vessels which may cause the window to move to the right side. Fig. 14(b) shows that some areas of the lesion inside are mislabeled as blood vessels. Therefore, the lesion becomes similar to optic disc. Fig. 14(c) shows that the ROI window is moved upward. Fig. 14(d) shows that the optic disc does not have any common optic disc features. The optic disc area is darker than other retina area. It is also hard to estimate blood vessels passing through the optic disc. In the case of Fig. 14(a) and (b), poor blood vessel estimation affects the errors. We currently use a blood vessel detection algorithm to test the role of blood vessels in training CNNs. Therefore, CNN4 may correctly recognize ROIs of the images if we use more accurate blood vessel classification algorithms.

We compare our proposed methods with other algorithms. Zhang *et. al.*'s ROI detection algorithm [2] shows 0.964 accuracy for 1,564 glaucoma and non-glaucoma fundus images. Foracchia *et. al.*'s algorithm [5] for detecting optic disc location shows 0.98 accuracy for the STARE dataset. Marin *et. al.*'s algorithm [8] for detecting optical disc location uses optic disc radius (*R*) to evaluate their results using the MESSIDOR dataset. They have 0.9975 accuracy for $1 \times R$ threshold. This means the distance between the real optic disc center and the estimated optic disc center should be equal or less than $1 \times R$ to be labeled as correct. In our case, we consider an error when the optic disc is not inside ROI boundary (green square) as shown in Figs. 11 to 14. Therefore, our performance in Table II will be improved if we use the $1 \times R$ threshold. The proposed method shows a slightly lower performance than Marin *et. al.*'s algorithm. However, we have a limited number of images for training the CNNs, and the training datasets for the CNNs do not have many images similar to Figs. 11 and 14. Therefore, our performance will improve if we use more images similar to Figs. 11 and 14 to train the CNNs.

As shown in Fig. 8, blood vessels usually are located on the left or right side of the optic discs and the brightest areas in the optic discs are located in the opposite side of the blood vessels (not in the center of the optic discs). Therefore, when blood vessels and/or the brightest optic disc areas are used to estimate ROI, the ROI will inaccurately appear to the left or right of the real ROI. In addition, the method will heavily rely on blood vessel estimation and the brightest area (in optic disc) estimation algorithms. However, the proposed method uses original images and rely less on the feature (blood vessels) than other algorithms. Therefore, the proposed method has more chance to have real ROI with the optic discs in the middle especially when the optic disc and/or blood vessels are not clear in fundus images.
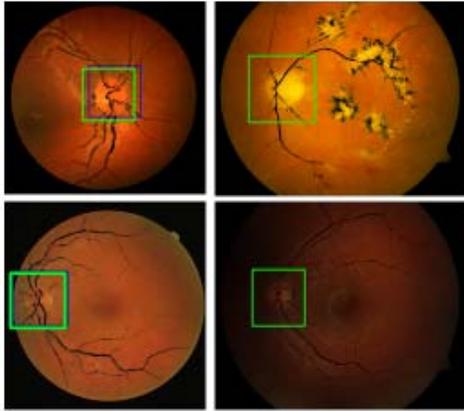
Figure 13. ROI estimation results of CNN4 from images in the Open Sources dataset.

(a)　　　　　　　　(b)

(c)　　　　　　　　(d)
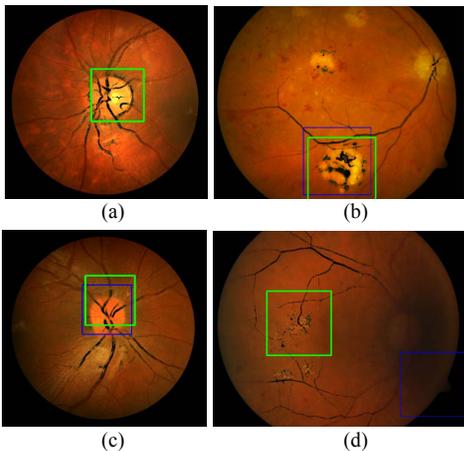
Figure 14. ROI Detection errors of CNN4 from images in the Open Sources dataset.

VII.　CONCLUSIONS

This paper proposes an automatic method to classify ROI from fundus images using deep learning. Two types of image datasets are used to train the CNNs: original fundus images, and original fundus images embedded with grey level blood vessels. We implement a blood vessel detection algorithm using smoothing, image enhancement, and morphological algorithms. We train multiple CNNs for each dataset using different sizes of kernel and stride in the first convolutional layer. The proposed method moves a window in the horizontal and vertical directions in each image, and applies CNNs to each window to find a ROI with the highest CNN result. The test results show promise. All CNNs in the two groups show over 98% accuracy, in the case of the MESSIDOR dataset. In the case of the Open Sources dataset, the CNN using the blood vessels shows the best performance overall: over 99% accuracy for the MESSIDOR dataset and over 97% accuracy for the Open Sources dataset. Therefore, we conclude that it is efficient to train CNNs using features (blood vessels) to make the CNNs more robust.

As a future task, we plan to improve the blood vessel detection algorithm to increase the accuracy of the proposed method. In addition, we will extend the method to segment optic disc and cup areas in fundus images.

REFERENCES

[1] Glaucoma Research Foundation, https://www.glaucoma.org/.

[2] Z. Zhang, B. H. Lee, J. Liu, D. W. K. Wong, N. M. Tan, J. H. Lim, F. Yin, W. Huang, H. Li, T, and Y. Wong, "Optic Disc Region of Interest Localization in Fundus Image for Glaucoma Detection in ARGALI," IEEE Conference on Industrial Electronics and Application, pp. 1686-1689, June 2010.

[3] S. Roychowdhury, D. D. Koozekanani, S. N. Kuchinka, and K. K. Parhi, "Optic Disc Boundary and Vessel Origin Segmentation of Fundus Images," IEEE Journal of Biomedical and Health Informatics, vol. 20, no 6, pp. 1562-1574, November 2016.

[4] A. Hoover and M. Goldbaum., "Locating the Optic Nerve in a Retinal Image Using the Fuzzy Convergence of the Blood Vessels," IEEE Transactions on Medical Imaging, vol. 22, no 8, pp. 951-958, August 2003.

[5] M. Foracchia, E. Grisan and A. Ruggeri., "Detection of Optic Disc in Retinal Images by Means of a Geometrical Model of Vessel Structure," IEEE Transactions on Medical Imaging, vol. 23, no 10, pp. 1189-1195, October 2004.

[6] OpenCV, More Morphology Transformations, https://docs.opencv.org/3.3.0/d3/dbe/tutorial_opening_closing_hats.html.

[7] S. Roychowdhury, D. D. Koozekanani, and K. K. Parhi, "Blood Vessel Segmentation of Fundus Image by Major Vessel Extraction and Subimage Classification," IEEE Journal of Biomedical and Health Informatics, vol. 19, no 3, pp. 1118-1128, May 2015.

[8] D. Marin, M.E. Gegundez-Arias, A. Suero, and J. M. Bravo, "Obtaining optic disc center and pixel region by automatic thresholding methods on morphologically processed fundus images", Computer Methods and Programs in Biomedicine, pp. 173-185, 2015.

[9] B. Savelli, A. Bria, A. Galdran, C. Marrocco, and M. Molinara, "Illumination correction by dehazing for retinal vessel segmentation" IEEE 30th International Symposium on Computer-Based Medical Systems, pp. 219-224, June 2017.

[10] C. Lim, Y Cheng, W. Hsu, and M. L. Lee, "Integrated Optic Disc and Cup Segmentation with Deep Learning," IEEE 27th International Conference on Tools with Artificial Intelligence, vol. 19, no 3, pp. 162-169, November 2015.

[11] OpenCV document, https://docs.opencv.org/2.4/doc/tutorials/introduction/desktop_java/java_dev_intro.html

[12] E. Decencière et al., "Feedback on a publicly distributed database: the Messidor database", Image Analysis & Stereology, v. 33, n. 3, p. 231-234, August. 2014, http://messidor.crihan.fr/ download-en.php.

[13] Retina Image Analysis, http://blogs.kingston.ac.uk/retinal/chasedb1/

[14] IMAGERET, http://www.it.lut.fi/project/imageret/

[15] DRIONS-DB: Digital Retinal Images for Optic Nerve Segmentation Database, http://www.ia.uned.es/~ejcarmona/DRIONS-DB.html

[16] DRIVE: Digital Retinal Images for Vessel Extraction, https://www.isi.uu.nl/Research/Databases/DRIVE/

[17] Deep Learning for Java, https://deeplearning4j.org