SEMI-SUPERVISED TRANSFER LEARNING FOR CONVOLUTIONAL NEURAL NETWORKS FOR GLAUCOMA DETECTION

Manal Al Ghamdi[⋆] *Mingqi Li*[†] *Mohamed Abdel-Mottaleb*[‡] *Mohamed Abou Shousha* [⋆†]

[⋆] Umm Al-Qura University, Department of Computer Science, Saudi Arabia
[†] Harbin Institute of Technology, Department of Computer Science, China
[‡] University of Miami, Department of Electrical and Computer Engineering, USA
[⋆†] Bascom Palmer Eye Institute, Miami, USA

ABSTRACT

Convolutional neural network (CNN) can be applied in glaucoma detection for achieving good performance. However, its performance depends on the availability of a large number of the labelled samples for its training phase. To solve this problem, this paper present a semi-supervised transfer learning CNN model for automatic glaucoma detection based on both labeled and unlabeled data. First, a pre-trained CNN from non-medical data is fine-tuned and trained in a supervised fashion using the labeled data. The self-learning approach is then used to predict the labels for the unlabeled data and utilize it for training. The experimental results on the RIM-ONE database demonstrate the effectiveness of the proposed algorithm despite the lack of initial labeled samples.

***Index Terms*—** Semi-supervised, glaucoma detection, convolutional neural networks, feature learning.

1. INTRODUCTION

Glaucoma is one of the leading causes of permanent blindness in the world. It is a chronic eye disease caused by retinal changes, specifically in the area of the optic nerve head (ONH) [1]. Therefore, an automated screening program for glaucoma detection is urgently required to facilitate the early detection of the disease, which is critical to prevent its progression. Modern fundus cameras with improved technology produce high-quality images with reliable information about diagnostically important retinal structures. Because of that, researchers focus their research on analysing fundus images. One of the important indicators for glaucoma in the fundus images is the size of the optic cup (OC) with respect to the optic disc (OD). Therefore, multiple parameters have been estimated by the previous works to detect glaucoma from images, such as the vertical cup to disc ratio (CDR)[2]. However, the clinical diagnosing by annotating the cup and disc manually from each image is labor-intensive and also time-consuming.

Convolutional Neural Network (CNN) models have been successfully utilized in many computer vision applications such as classification and semantic segmentation. These models have the ability to learn the deep features of the input without pre-processing steps such as segmentation. However, constructing these models to achieve desirable results is restricted by the availability of large and annotated training sets. An alternative solution the intensive and expensive training process, many works have recently presented based on off-the-shelf CNNs to address different problems from those for which they were initially trained [3]. These CNNs were originally trained using extremely large dataset from a different domain and then transferred to extract discriminative features from a different domain. However, adopting a classifier to new domains requires a large dataset to avoid overfitting, a setting that cannot be accomplished with current publicly available sets for glaucoma detection. To address this problem, both supervised and unsupervised learning can be used simultaneously with labeled and unlabeled data. This approach achieved a good performance in previous computer vision tasks.

This paper proposes a simpler method of training neural networks in a semi-supervised fashion. The proposed network consists of two stages. The first stag is the transfer learning that involves using a pre-trained CNN in the context of glaucoma detection with limited-size labeled dataset. The second stage is the self-learning that involves increasing the training set and thus the performance of the pre-trained CNN, using the unlabeled dataset. The self-learning method predicts the labels, which are known as pseudo-Labels, for the unlabeled samples by choosing the class which has the maximum predicted probability and consider them as if they were true labels. In principle, this framework can combine almost all neural network models and training methods. In our experiments, combining the transfer learning for the labeled data and self-learning for the unlabeled data present promising results for glaucoma detection.

2. RELATED WORK

The lack of large dataset of labeled retinal images limits the use of supervised learning for the diagnosis of these images. To deal with this issue, few studies used semi-supervised learning [4] and transfer learning [5]. Mahapatra [4] developed an OD segmentation framework by combining manual annotations with semi-supervised learning. Bechar *et al.*[6] proposed a semi-supervised learning method for OD and OC segmentation based on color features, spatial information and manual annotation of the image's pixels. However, to the best of our knowledge, there are no studies in the literature for glaucoma detection using semi-supervised deep learning.

CNN and transfer learning have been achieved promising results in various eye diseases, other than glaucoma, such as diabetic retinopathy (DR) [7, 8]. This is caused by the availability of the annotated data for the DR but not for the other diseases. In transfer learning, the pre-trained network is fed with the new images, retrieving the outputs of the fully connected layer as feature vectors, and using them to train a new classifier explicitly devoted to the new task. If the new dataset is different from the original one, then fine-tuning of the weights of the pre-trained network is required. In the glaucoma detection literature, Orlando *et al.*[3] used two different pre-trained CNN models (OverFeat and VGG-S) as feature extraction and combined with regularized logistic regression models. Experiments were performed on two public datasets, which were annotated manually for the glaucoma detection. Al-Bander *et al.*[9] proposed an automated system to detect glaucoma in retinal fundus images. They used a pre-trained CNN model (Alexnet) as a feature extractor fed into support vector machine (SVM) for classification. Cerentini *et al.*[19] used two pre-trained models (GoogLeNet); the first one combined with a sliding-window approach for feature extraction while the second one for image classification. Both Cerentini *et al.*[19] and Al-Bander *et al.*[9] performed their experiments on the publicly available RIM-ONE dataset.

Despite the good performance reported by the above models, none of them fine-tuned the pre-trained models. In addition, no models adopted the large number of unlabeled retina images, which are more available than the labeled ones. Presenting such a model is required and can be generalized to other eye diseases and other medical image analysis tasks since it does not depend on the availability of the large annotated dataset.

3. GLAUCOMA CLASSIFICATION BASED ON SEMI-SUPERVISED CNN

We combined the transfer learning (Section 3.1) with the self-learning (Section 3.2) to perform semi-supervised learning for detecting glaucoma. First, we defined the symbols to be used in this section. The training dataset X is composed of two parts: $X = [L, U] \in R^{d \times N}$, where $L = [x_1, x_2, \ldots, x_l] \in R^{d \times l}$ represents the labeled samples and $U = [x_{l+1}, x_{l+2}, \ldots, x_{l+u}] \in R^{d \times u}$ represents the unlabeled samples. The training process of the proposed method consists of two stages as shown in Figure 1. First, a pre-trained CNN model from a non-medical image domain is transferred and fine-tuned to learn the representative features of retina images. Secondly, the updated CNN model is used to select the most reliable samples from the unlabeled data using the self-learning method. When the self-learning process has finished, the unlabeled samples are included in the training data with the labeled assigned by the updated CNN.

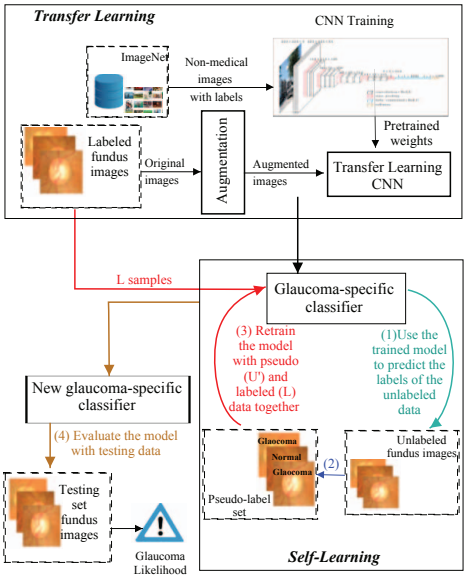

Fig. 1: Flowchart of the semi-supervised transfer learning CNN method for glaucoma detection in fundus images.

3.1. Transfer Learning

In order to label the unlabeled samples and use them later for training, the underlying classifier should have a high accuracy. Instead of training a CNN from scratch with a limited-size dataset, we used a pre-trained CNN and modify it with our small labeled samples. We chose the VGG-16 model for its high performance reported in the literature. It contains convolution layers, max pooling layers, fully connected layers and a softmax layer, which is the output layer. The weights from the model are transferred as the weights for our network and thus the full network (after removing the fully connected layers) is treated as a fixed-feature extractor for the new dataset. To prevent overfitting, data augmentation were employed

including shift, flip and random zoom-in of the image.

Two stages of training were performed. The first training was performed to extract the features from the layer right before the output layer and was used for classification. These higher-level features are available in the deeper layers of the CNN. Therefore, the top fully connected layers and softmax layer were replaced by the following new layers while the rest of the layers were frozen to keep the weights from changing. A flat layer was added to convert the output of the convolutional layers, flattening its whole structure to create a single long feature vector to be used by the dense layer for the final classification. Two batch normalization layers were added between the two fully connected layers and the output layer. This is because the VGG-16 model is deep and our training datasets are small, so it is easy to encounter the vanishing gradient problem. Finally, a softmax layer was added to map the output to two categories. We used the 'Adam' method [10] as our optimizer, which has the ability to work well despite minimal tuning [11]. For the second training, the last convolutional block of the VGG-16 was unfrozen and retrained multiple times to be customized to the new dataset. The optimizer at this stage was changed to Stochastic Gradient Descent (SGD) which empirically performs better than Adam at this stage [11].

Learning the model was achieved by using the cross-entropy loss [12] on the softmax normalization score which is defined as:

$$J = -\frac{1}{L}\left(\sum_{i=1}^{L} y_i \cdot \log\left(\hat{y}_i\right)\right) \tag{1}$$

where y is the ground-truth for x, $\hat{y}$ is the estimated value and L is the dataset size of the labeled samples.

3.2. Self Learning

To improve the robustness of the classifier, especially when we have insufficient labeled samples, it is necessary to extend the initial training set. Thus, the self-learning method [13] is utilized to increase the number of the initial training samples L with samples from unlabeled set that have high confidence when classified with the transfer classifier. The assigned labels of these samples are known as pseudo-labels and generated in an iterative fashion to obtain a stable and accurate classifier. Given a set of labeled data L and a set of unlabeled data U, self-learning proceeds as follows: (1) The pre-trained CNN model from Section 3.1 is used to classify the unlabeled data U. (2) The prediction scores are ordered and the subset $U' \subset U$ for which the classifier has the highest confidence scores is added to L and removed from U. At this stage, the labels for the set U' are called pseudo-labels. (3) The classifier is re-trained using the new $L + U'$ training set. These steps are repeated until the algorithm converges. At the end of the learning process, the remaining unlabeled samples are dropped because they are not reliable. To avoid overfitting, each time the a subset of samples U' is selected, the ratio of the normal to glaucoma should be the same as the distribution of labeled data [14]. The loss function for the self-learning is the same as Eq 1, which was used in the supervised learning. (4) Finally, the new glaucoma-specific classifier is evaluated with the testing data.

4. EXPERIMENTS

RIM-ONE [15], a publicly available dataset of retinal images, was used for training and evaluating the proposed method. It consists of 455 high-resolution images of which 255 images are categorized as normal and 200 images are categorized as belonging to patients with glaucoma. A subset of around $1,500$ images from the RIGA [16], a public dataset used for evaluating the segmentation methods in glaucoma detection tasks, was used as the unlabeled data. This subset of $1,500$ was determined by the self-learning process following the normal to glaucomatous ratio in the labeled samples. The images were cropped to be consisted with the training labeled samples.

The evaluation was based on 20% of the 455 images using accuracy, sensitivity and specificity, defined as:

$$Accuracy = \frac{T_p + T_N}{T_P + F_p + F_N + T_N} \tag{2}$$

$$Sensitivity = \frac{T_P}{T_P + F_N} \tag{3}$$

$$Specificity = \frac{T_N}{T_N + F_P} \tag{4}$$

where T_P, T_N, F_P and F_N are the number of true positives, true negatives, false positives and false negatives, respectively.

4.1. Experimental Setup

We used NVIDIA GTX 1080TI 11GB GPU card with 3584 CUDA parallel-processing core for implementation. For transfer learning (Section 3.1) setup, the VGG-16's weights were downloaded from the Keras GitHub[1] and the model was developed using Keras API [17]. The first training was performed 20 times to train the fully connected layers and the softmax layer. The learning rate for the Adam optimizer was set to 0.001. The second training was performed 50 times to allow the top convolutional layers to extract more detailed features. The fine-tuning requires a stable training process with a lower learning rate relative to the initial rate used in the first stage, otherwise the previously learned features could be destroyed and the optimization could destabilize [11].

[1] https://github.com/fchollet

Therefore, the optimizer at this stage was changed to SGD and the learning rate was decreased to 0.0001.

For self-learning (Section 3.2) setup, the last learning rate and optimizer from the transfer learning were used. The subset U' contains the highest A samples of normal and B samples of glaucoma with ratio of $A : B$ set to $1 : 2$ similar to the labeled data ratio. The initial sizes of A and B were set to 30 and 60 samples, respectively. Since the classifier becomes better with more iterations, these sizes were increased iteratively during the learning process by $30 + 3 \times i$, where number 30 denotes the initial size of A and i denotes the number of iteration. The number of epochs was initially set to 12 to suit the small size of the labeled samples and then increased iteratively by $epoch + 4 \times i$ [18], where i is the number of iteration.

4.2. The Effectiveness of the Self-learning Method.

The self-learning method based on high confidence was utilized to expand the training set in an iterative manner. During each iteration, the self-learning method ordered the prediction scores for the unlabeled samples and the ones with the highest confidence were selected to join the training samples. To demonstrate the effectiveness of the self-learning method, we compared it with the random selection method. During each iteration, the same number of the unlabeled samples are randomly selected and the highest probability class for each sample was used as the true label. The classification accuracy of the two methods varies with the iterations, as shown in Figure2(a). Obviously, the classification accuracy of the self-learning method was higher than that of the random selection method. Thus, it proves the effectiveness of the self-learning method in utilizing the unlabeled samples.

4.3. The Effectiveness of the Semi-supervised Learning Method.

After the expansion of the initial labeled training set, the semi-supervised learning method was evaluated using the testing set. To demonstrate the effectiveness of expanding the training set with the unlabeled samples, we compared it with the fully supervised model using the same setup. The results are presented in Table 1(a) and Figure 2(b). It is obvious that, the classification accuracy of the semi-supervised method is higher than that of the supervised method. This validates the suggestion that the self-learning method is capable of utilizing the information contained in the unlabeled samples to improve the classification accuracy of the proposed method.

4.4. Comparison with Related Methods.

Extensive comparison with other related methods is not feasible for two reasons: (1) no studies have used the semi-supervised learning concept for glaucoma detection, and (2) most of the existing studies used their own private datasets. As an alternative, we compare our performance with the transfer learning methods applied for glaucoma detection task. To the best of our knowledge, only two studies by Al-Bander *et al.*[9] and Cerentini *et al.*[19] were evaluated on RIM-ONE using the transfer learning CNN. The results are presented in Table 1(b). The results show that the presented method outperformed the other two approaches with a clear margin. This was caused by our fine-tuning of the weights in the pre-trained networks which makes the later layers of the proposed method more specific to the details in our dataset. The results also show that the supervised methods, even with transfer learning, are less accurate than the semi-supervised method due to small number of the labeled samples. Iteratively increasing the training set causes the accuracy of the CNN method to be gradually increased.

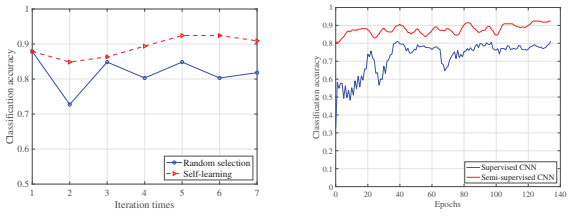

(a) Self-learning evaluation (b) Semi-supervised evaluation

Fig. 2: Evaluating (a) the self-learning against random selection, (b) the supervised against semi-supervised.

(a) Comparison with supervised.

Method	accuracy	sensitivity	specificity
Semi-supervised	**92.4%**	**91.7%**	**93.3%**
Supervised	81.25%	74.2%	86.3%

(b) Comparison with the related works.

Method	accuracy	sensitivity	specificity
Semi-supervised	**92.4%**	**91.7%**	**93.3%**
Al-Bander *et al.*[9]	88.2%	85.0%	90.8%
Cerentini *et al.*[19]	86.2%	n/a	n/a

Table 1: Comparison of the proposed glaucoma detector results with different methods, n/a means not available.

5. CONCLUSIONS

This paper presented a CNN semi-supervised learning framework for detecting glaucoma in fundus images. By employing the self-learning strategy to increase the training samples using the unlabeled data and fine-tuning a pre-trained CNN to define a glaucoma-specific classifier, the presented approach demonstrated promising performance compared to existing methods. Improvements to generalized the framework to detect more eye diseases are part of ongoing research by the authors.

6. REFERENCES

[1] J. Flammer and E. Meier, *Glaucoma: A Guide for Patients : an Introduction for Care-providers : a Quick Reference*, Hogrefe & Huber, 2003.

[2] Gundersen Kjell Gunnar, Heijl Anders, and Bengtsson Bo, "Sensitivity and specificity of structural optic disc parameters in chronic glaucoma," *Acta Ophthalmologica Scandinavica*, vol. 74, no. 2, pp. 120–125, 1996.

[3] Jos Ignacio Orlando, Elena Prokofyeva, Mariana del Fresno, and Matthew B. Blaschko, "Convolutional neural network transfer for automated glaucoma identification," in *12th International Symposium on Medical Information Processing and Analysis;*, 2017, vol. 10160, pp. 10160 – 10160 – 10.

[4] Dwarikanath Mahapatra, "Combining multiple expert annotations using semi-supervised learning and graph cuts for medical image segmentation," *Computer Vision and Image Understanding*, vol. 151, pp. 114 – 123, 2016, Probabilistic Models for Biomedical Image Analysis.

[5] Abhijit Guha Roy and Debdoot Sheet, "DASA: domain adaptation in stacked autoencoders using systematic dropout," *CoRR*, 2016.

[6] Mohammed El Amine Bechar, Nesma Settouti, Vincent Barra, and Mohamed Amine Chikh, "Semi-supervised superpixel classification for medical images segmentation: application to detection of glaucoma disease," *Multidimensional Systems and Signal Processing*, vol. 29, pp. 979–998, 2018.

[7] Jeremy Benson, Hector Carrillo, Jeff Wigdahl, Sheila Nemeth, John Maynard, Gilberto Zamora, Simon Barriga, Trilce Estrada, and Peter Soliz, "Transfer learning for diabetic retinopathy," *Proc.SPIE*, vol. 10574, pp. 10574 – 10574 – 9, 2018.

[8] X. Li, T. Pang, B. Xiong, W. Liu, P. Liang, and T. Wang, "Convolutional neural networks based transfer learning for diabetic retinopathy fundus image classification," in *2017 10th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*, Oct 2017, pp. 1–11.

[9] B. Al-Bander, W. Al-Nuaimy, M. A. Al-Taee, and Y. Zheng, "Automated glaucoma diagnosis using deep learning approach," in *2017 14th International Multi-Conference on Systems, Signals Devices (SSD)*, March 2017, pp. 207–210.

[10] Diederik P. Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.

[11] Nitish Shirish Keskar and Richard Socher, "Improving generalization performance by switching from adam to SGD," *CoRR*, vol. abs/1712.07628, 2017.

[12] M. Guillaumin, T. Mensink, J. Verbeek, and C. Schmid, "Tagprop: Discriminative metric learning in nearest neighbor models for image auto-annotation," in *2009 IEEE 12th International Conference on Computer Vision*, Sept 2009, pp. 309–316.

[13] Dong-Hyun Lee, "Pseudo-label : The simple and efficient semi-supervised learning method for deep neural networks," *Workshop on Challenges in Representation Learning, ICML*, 2013.

[14] Y. Zhou, M. Kantarcioglu, and B. Thuraisingham, "Self-training with selection-by-rejection," in *2012 IEEE 12th International Conference on Data Mining*, Dec 2012, pp. 795–803.

[15] F. Fumero, S. Alayon, J. L. Sanchez, J. Sigut, and M. Gonzalez-Hernandez, "Rim-one: An open retinal image database for optic nerve evaluation," in *2011 24th International Symposium on Computer-Based Medical Systems (CBMS)*, June 2011, pp. 1–6.

[16] Ahmed Almazroa, Sami Alodhayb, Essameldin Osman, Eslam Ramadan, Mohammed Hummadi, Mohammed Dlaim, Muhannad Alkatee, Kaamran Raahemifar, and Vasudevan Lakshminarayanan, "Retinal fundus images for glaucoma analysis: the riga dataset," *Proc.SPIE*, vol. 10579, pp. 10579 – 10579 – 8, 2018.

[17] François Chollet et al., "Keras," https://keras.io, 2015.

[18] Fei Gao, Zhenyu Yue, Jun Wang, Jinping Sun, Erfu Yang, and Huiyu Zhou, "A novel active semisupervised convolutional neural network algorithm for sar image recognition," *Computational Intelligence and Neuroscience*, vol. 2017, pp. 1–8, 2017.

[19] Allan Cerentini, Daniel Welfer, Marcos Cordeiro d'Ornellas, Carlos Jesus Pereira Haygert, and Gustavo Nogara Dotto, "Automatic identification of glaucoma using deep learning methods," *Studies in health technology and informatics*, vol. 245, pp. 318–321, 2017.