Multi-scale Microaneurysms Segmentation Using Embedding Triplet Loss

Mhd Hasan Sarhan^{1,2}, Shadi Albarqouni¹, Nassir Navab^{1,3}, and Abouzar ${\rm Eslami}^2$

¹ Computer Aided Medical Procedures, Technical University of Munich, Germany ² Carl Zeiss Meditec AG, Munich, Germany

³ Computer Aided Medical Procedures, Johns Hopkins University, Baltimore, USA

Abstract. Deep learning techniques are recently being used in fundus image analysis and diabetic retinopathy detection. Microaneurysms are an important indicator of diabetic retinopathy progression. We introduce a two-stage deep learning approach for microaneurysms segmentation using multiple scales of the input with selective sampling and embedding triplet loss. The model first segments on two scales and then the segmentations are refined with a classification model. To enhance the discriminative power of the classification model, we incorporate triplet embedding loss with a selective sampling routine. The model is evaluated quantitatively to assess the segmentation performance and qualitatively to analyze the model predictions. This approach introduces a 30.29% relative improvement over the fully convolutional neural network.

Keywords: Deep Learning · Segmentation · Ophthalmology

1 Introduction

Diabetic retinopathy (DR) is the leading cause of vision impairment and blindness for middle-aged groups [9]. DR early detection is important for the treatment planning. Severity of DR falls into one of five levels (none, mild, moderate, severe, or proliferative) [1] . Microaneurysms are considered as the first signs for detecting early stages of DR. Hence, detecting these lesions is important for Computer Aided Diagnosis systems. Microaneurysms are abnormalities in the microvascular structure and appear as small red dots in color fundus images. Screening programs use colored fundus images of the retina for their rich information and ease of access. Detecting microaneurysms in colored fundus images is a challenging task due to the small size of the lesion which makes up less than 1% of the entire image , and the low contrast between microaneurysms and background.

Microaneurysms are the strongest determinant for DR since they are the first lesion that appears during the early stages. Various approaches for microaneurysms detection using deep learning are proposed [6], [12], [8]. These methods are patch-wise approaches and use deep architectures to extract representative features. These features could be added to a set of hand-crafted features [12]

and passed to a classification model or used solely in an end-to-end network [6], [8]. Deep learning techniques in the literature of microaneurysms detection use random patches selection, hence, they are prone to be biased towards the oversampled class. Moreover, no work in the microaneurysms segmentation context has leveraged the embedding space of the input patches to impose an additional constraint on the learning process.

Contributions: In this work, a multi-scale patch-wise approach for segmenting microaneurysms in retinal fundus images is proposed. The main contributions of this work are 1) fusing segmentation on multiple scales for microaneurysms detection, and 2) using embedding triplet loss [14] with selective sampling [5] to increase the descriptiveness of the feature representation while focusing the training on informative examples. The model is agnostic to other lesions (i.e. the model differentiates between healthy and microaneurysm patches regardless of information about other lesions). Being agnostic to other lesions is important in such cases as it may be difficult to obtain an annotated dataset with all DR lesions annotated

2 Methodology

Our proposed microaneurysms segmentation framework, depicted in Fig. 1, consists of two stages; the *hypothesis generation network* (HGN), where multi-scale fully convolutional networks (FCNs) are employed to propose a region of interest (ROI), and *patch-wise refinement network* (PRN), where extracted patches around ROIs are passed to the classifier. In the next sections we introduce the details of the applied method. First, we go through the fully convolutional hypothesis generation networks, the reasoning behind having multiple scales, and the details of the loss function used for optimization. The second section is dedicated for the PRN. In which, the motive behind this network is explained and the details of triplet loss and selective sampling are presented.

Hypothesis Generation Network (HGN): High-resolution fundus images where a microaneurysm covers a very small part of the image are examined to segment microaneurysm. Using a zoomed-in patch would allow for high spatial accuracy on account of losing semantic information, whilst a zoomed-out patch would have a richer semantic representation on the account of losing spatial resolution [3]. As a trade-off, we use equally sized patches on two scales of the image to build two HGNs, one for each scale.

HGN is a fully convolutional neural network trained on patches of size 256×256 extracted from the fundus images. Two HGNs are trained for two different scales of the fundus image (1x, 0.5x). This allows the extraction of scale-related features while at the same time preserve full resolution image information. The architecture used is the full resolution residual network type A [13] for its good results in segmentation.



Fig. 1: Pipeline for the multi-scale microaneurysms segmentation framework. Top part shows the pipeline in the training mode where each model is trained separately. Bottom part shows the pipeline at inference time where the image is used rather than patches for the hypothesis generation networks.

To select the training patches, we define images that contain no signs of DR as healthy (negative) images and images with microaneurysms as lesion (positive) images. Healthy pixels are extracted only from healthy patients' scans and lesion pixels are extracted from DR patients at the microaneurysms locations. As a loss function, weighted cross entropy loss is used to compensate for the imbalance negative and positive patches. Moreover, dice loss is optimized to enhance the spatial overlap between a segmentation map output and the gold standard segmentation. We use a differentiable approximation of the dice loss as in [10].

Patch-wise Refinement Network PRN is a classification network that is used as on top of the HGN. The input of the network is an image patch and the output is the probability of the patch center pixel being a microaneurysm or healthy. The segmentation maps of the HGN are used as regions of interest for the PRN. The architecture of classification networks allows for receptive fields larger than fully convolutional networks that consume more memory because of the decoder part and skip connections. The larger receptive field allows for feature maps that incorporate more spatial information about the image which enriches the extracted features. The architecture employed for this network is an adopted version of the Resnet-50 [7]. One downsampling step is omitted from the original architecture because the input image size in our case is smaller than what is expected in the Resnet-50 scenario. In the training phase, patches are extracted from images in the same manner of extracting 1x resolution patches for HGN. The only difference is the size of PRN patches is 129×129

To make the refinement network features more descriptive we introduce the To extract discriminative features in PRN we propose the utilization of triplet loss [14]. Triplet loss is applied on the embedding of a patch around pixel x

into a *d*-dimensional feature space. The aim of triplet loss is to make similar patches closer to each other in the embedding space while pushing dissimilar patches away from each other in the embedding space using a predefined distance measure. We found the feature representation of the last convolution layer after the global average pooling (GAP) as a good representation in the embedding space due to its high descriptive power while having a compact representation. The optimization of triplet loss requires three input patches namely the anchor patch x^a , the positive patch x^p and the negative patch x^n . The goal is to make the embedding of the positive patch closer to the anchor patch than the embedding of the negative patch. Patches with Microaneurysms at the center pixels are used as anchor and positive patches, while healthy patches are used as negative patches. The loss is defined as

$$\mathcal{L}_{triplet} = \sum_{i}^{N} \left[d(f(x_{i}^{a}), f(x_{i}^{p})) - d(f(x_{i}^{a}), f(x_{i}^{n})) + a \right]_{+}$$
(1)

where a is a margin to enforce a distance between positive and negative pairs, $d(.,.) \in \mathbb{R}^1$ is the distance measure in the embedding space, and N is the number of all possible triplets. As a distance measure, angular cosine distance is utilized as it shows better performance on high dimensional representations when training deep networks [11]. In addition to triplet loss, cross entropy loss for patches is optimized.

Generating all triplets, in this case, would be computationally prohibitive. Moreover, the imbalance in the dataset is high. To counter these problems, we use selective sampling [5]. This approach of training proved to enhance the results in training scenarios where data from different classes are not balanced. In our use case, the healthy class is over-represented. In selective sampling, patches with higher loss have a higher probability of being picked for the next epoch as they are considered representative samples.

3 Experiments

3.1 Experimental Setup

Dataset For our evaluations of the segmentation pipeline , we use the IDRiD⁴ publicly available dataset. All images are captured with the same device that has 50-degree field of view and have size of 4288×2848 pixels. Before patch extraction, the published train dataset is split into two parts: training, and validation sets. The validation set is used for monitoring the training. Table 1 shows the dataset splits.

Implementation details We employ contrast enhancement following the formula $I_{pre}(x, y) = 4I(x, y) - 4G_{\sigma} * I(x, y) + 1024/30$. To train HGN, we define a minibatch of size 10 and consider each epoch to be 1000 mini-patches. The learning rate for the full-scale network is 1e - 6 and for the half-scale network is 1e - 5.

⁴ https://idrid.grand-challenge.org/

Table 1: IDRiD dataset splits							
	Healthy		Microaneurysms				
	Images	Patches	Images	Patches			
Train set	80	6M	44	$\sim 500 \mathrm{K}$			
lidation set	9	$^{\sim}6M$	10	$\sim 132 \mathrm{K}$			

Test set

PRN is trained with mini-batches of triplets. The size of a mini-batch is 90×3 patches. We sample 90×2 patches from the pool of lesion patches randomly with uniform distribution, and 90×1 samples from the healthy patches pool with selective sampling. This neural network has a Siamese structure [2], this means that each part of the triplet's three parts is run through identical versions of the network and the gradients are combined at the output to update the weights of the network. In addition to triplet loss, cross entropy loss for pairs is optimized. To this end, we optimize the cross-entropy loss between the anchor and the negative pair. Every 1000 mini-batches is considered as an epoch. We run selective sampling routine every 10 epoch, this is because of the big number of training patches. which takes a significant amount of time to evaluate. Learning rate is set to 1e - 5 and decreased by a factor of 10 after 20 epochs.

3.2 Multi-scale effect

In this evaluation, we study the effect of using multiple scalse. To this end, two HGNs are trained, one for the full resolution image and one for the downsampled image by a factor of two. The evaluation is done on the publicly published test set images. We compare results from each scale with the results of combining the two scales in two different ways . First the output of the half scale is upsampled using linear interpolation, then the prediction maps of the two scales are combined either with pixelwise arithmetic or geometric averaging. The results of this evaluation are presented in Table. 2. FCN 1x, FCN 0.5x represent the evaluation on the prediction map of the full scale and half scale HGNs, respectively. FCN geometric and FCN arithmetic refer to the results of combining the two scales with geometric and arithmetic averaging, respectively. The results show that combining the two scales gives better performance either way. We notice a higher recall from the half scale network but lowest precision, this reflects that the model is very sensitive to microaneurysms and generates a high number of false positives that drops down the overall performance.

3.3 Patch-wise refinement and triplet loss effect

We evaluate the effect of 1) using a classification network to refine the classifications of HGNs and 2) using triplet loss in the classification network to refine HGNs results (i.e PRN). To evaluate the classification network, we utilize patches from the image in a sliding window fashion and use the classification probability of each point to obtain segmentation maps. It is worth noting that sliding

window does not go over all the image, but only the parts higher than a preset probability threshold (0.5 in our case) from HGN. Two segmentation maps will be obtained by sliding over the image masked with two HGNs outputs. Two prediction maps from two HGNs and two prediction maps from refining HGNs results with the classification networks combined as shown in Figure 1.

We first demonstrate the effect of incorporating a classification network to refine the results of HGNs. To this end, we use an edited version of PRN that uses only cross entropy loss without the triplet embedding optimization. This network is denoted as *cls*. Using the classification network on top of the fully convolutional networks enhances the results of the overall segmentation. The larger receptive field allows for more descriptive representations which in turn could suppress false positives that are triggered by HGNs. The effect of utilizing triplet embedding loss is then evaluated by training a PRN using triplets from the training set. We set the margin value *a* from Equation 1 to 0.5. At test time, this network is utilized in a sliding window fashion similar to *cls*.

Using triplet loss in a multi-scale approach with geometric averaging has an overall 30.29% PR AUC improvement over the baseline fully convolutional neural network trained with weighted cross entropy. The improvement when incorporating triplet loss could be attributed to the quality of the learned representations where lesion patches are forced to be close to each other with a certain margin of difference from healthy ones.

Our results come in 4th place in the IDRiD challenge outdated leaderboard based on the metric used on the released test set. The challenge submission is currently closed. iFLYTEK-MIG used Mask-RCNN to segment 3 lesions at the same time. VRT used a U-net to segment four lesions all together. PATech used a patch-wise approach with false positives bootstrapping on lesions simultaneously. We notice that in all these models, information about lesions other than microaneurysms is utilized. This makes the disambiguation between lesion types (e.g. hemorrhages and microaneurysms) learned inherently in the model but has the drawback of requiring full annotation of multiple lesion types. Out model does not require information from other lesions to be trained.

3.4 Visual evaluation

We study the misclassifications of the model by visually examining samples of the results. In Figure. 2 an example of a segmentation is presented. From the example, we notice that false positives mostly lay in the area around hemorrhages or on top of a blood vessel where a higher intensity occur. False negatives are more difficult to be detected because they sometimes appear very close to hemorrhage and blend in or the contrast in the image is low enough to lose the microaneurysm. In the top left example, we see a false negative example where the microaneurysm is misclassified because of very light edges and irregular shape that leans towards hemorrhage. In the other false negative examples (in cyan), the cases are very difficult to be distinguished and variability between raters may occur in such cases. The bottom right example shows a false positive example where a darker area around the bright exudates appears similar

	1100110	I I DOOLO	1 100101011	rtectan
HGN 1x - baseline	0.3374	0.3618	0.2970	0.4626
HGN $0.5x$	0.3411	0.4001	0.4380	0.3682
HGN geometric	0.3622	0.3866	0.5115	0.3108
HGN arithmetic	0.3701	0.4156	0.4741	0.3701
cls geometric	0.3895	0.4153	0.5402	0.3374
cls arithmetic	0.3905	0.4368	0.4973	0.3895
PRN arithmetic	0.3978	0.4323	0.54051	0.3602
PRN geometric	0.4196	0.38477	0.61128	0.2807
IDRiD iFLYTEK-MIG	0.5017	-	-	-
IDRiD VRT	0.4951	-	-	-
IDRiD PATech	0.4740	-	-	-

Table 2: Ablation test for emphasizing each part of the pipeline



Fig. 2: An example of a segmented microaneurysms in a fundus image. Green is for true positives, red is for false positives, and cvan is for false negatives.

to microaneurysm. The variability in illumination parameters of the capturing device has also a significant effect on the training and may lead to a bias towards a certain image appearance. It is important to note that images in the IDRiD dataset are compressed with a lossy compression which leads to big jumps in intensity values next to each other.

4 Discussion and Conclusion

We hypothesize that using multiple fully convolutional networks for multiple scales of the inputs enhances the segmentation of small objects similar to microaneurysms because it gives a better trade-off between semantic and spatial accuracy. Embedding loss is employed mainly in learning image descriptors [15]. We use the triplet embedding loss in our model to treat deeper layers of the classification network as a local descriptor of the keypoint represented by the healthy or microaneurysm patch. The classification performance increases by adding this additional constraint on the features created by the network.

The segmentation results could be used in report generation for the doctors or in future studies to do big data analysis of populations. Microaneurysms

turnover is also an important factor in the progression analysis of DR [4] and could be studied better with reliable models for Microaneurysms segmentation.

References

- American Academy of Ophthalmology. international clinical diabetic retinopathy disease severity scale detailed table. http://www.icoph.org/downloads/ Diabetic-Retinopathy-Detail.pdf, accessed: 10.09.2018
- Bromley, J., Guyon, I., LeCun, Y., Säckinger, E., Shah, R.: Signature verification using a" siamese" time delay neural network. In: Advances in Neural Information Processing Systems. pp. 737–744 (1994)
- Ghiasi, G., Fowlkes, C.C.: Laplacian pyramid reconstruction and refinement for semantic segmentation. In: European Conference on Computer Vision. pp. 519– 534. Springer (2016)
- Goatman, K.A., Cree, M.J., Olson, J.A., Forrester, J.V., Sharp, P.F.: Automated measurement of microaneurysm turnover. Investigative ophthalmology & visual science 44(12), 5335–5341 (2003)
- van Grinsven, M.J., van Ginneken, B., Hoyng, C.B., Theelen, T., Sánchez, C.I.: Fast convolutional neural network training using selective data sampling: Application to hemorrhage detection in color fundus images. IEEE transactions on medical imaging 35(5), 1273–1284 (2016)
- Haloi, M.: Improved microaneurysm detection using deep neural networks. arXiv preprint arXiv:1505.04424 (2015)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Lam, C., Yu, C., Huang, L., Rubin, D.: Retinal lesion detection with deep learning using image patches. Investigative ophthalmology & visual science 59(1), 590–596 (2018)
- 9. Lee, R., Wong, T.Y., Sabanayagam, C.: Epidemiology of diabetic retinopathy, diabetic macular edema and related vision loss. Eye and vision **2**(1), 17 (2015)
- Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 3D Vision (3DV), 2016 Fourth International Conference on. pp. 565–571. IEEE (2016)
- Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: Proceedings of the 27th international conference on machine learning (ICML-10). pp. 807–814 (2010)
- Orlando, J.I., Prokofyeva, E., del Fresno, M., Blaschko, M.B.: Learning to detect red lesions in fundus photographs: An ensemble approach based on deep learning. arXiv preprint arXiv:1706.03008 (2017)
- 13. Pohlen, T., Hermans, A., Mathias, M., Leibe, B.: Fullresolution residual networks for semantic segmentation in street scenes. arXiv preprint (2017)
- Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 815–823 (2015)
- Wohlhart, P., Lepetit, V.: Learning descriptors for object recognition and 3d pose estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3109–3118 (2015)