Automated laterality determination in Ultra-Widefield Scanning Laser Ophthalmoscopy

Peter Wakeford
Optos (Nikon)
Queensferry House
Carnegie Campus
Dunfermline
Scotland
KY11 8GR
pwakeford@optos.com

Enrico Pellegrini
Optos (Nikon)
Queensferry House
Carnegie Campus
Dunfermline
Scotland
KY11 8GR

Jano van Hemert
Optos (Nikon)
Queensferry House
Carnegie Campus
Dunfermline
Scotland
KY11 8GR

Ik Siong Heng
Kelvin Building
University of Glasgow
University Avenue
Glasgow
G12 8QQ

Abstract

A deep learning method to determine the laterality of ultra-widefield retinal images is presented. On a test set of 411 images, a 98.3% accuracy was achieved. Saliency maps show that the network activates on the same features which a human reviewer would use to determine laterality.

1 Introduction

Optos plc (a Nikon company) manufactures scanning laser ophthalmoscopes (SLOs) for clinical applications. The current imaging procedure requires an operator to visually establish and manually select the *laterality* of the retinal image (i.e, whether the image is from the right or left eye). This repetitive task is prone to human error, thus an automated solution to this problem is sought.

Convolutional neural networks (CNNs) are a class of machine learning algorithms that have proved to be extremely powerful when applied to medical image classification problems (1). For this reason, we present a CNN-based approach to this problem able to achieve an overall accuracy of 98% on 411 test images.

Other work on laterality determination in fundus camera images has been published with promising results. In (2), a support vector machine classifier trained on the locations of segmented retinal features (vasculature, optic disc, macula) achieved 94% accuracy. In (3), the output of a CNN trained with a transfer learning approach was used in conjunction with extracted anatomical features (such as vasculature density, orientation), also achieving 94% accuracy. To our knowledge, this is the first time that laterality determination has been performed with deep learning on ultra-widefield (UWF) SLO images.

2 Materials

Most images captured on Optos SLOs (4) are central-pole (CP), in which the scan is centred on the fovea region. A single CP image covers $\sim 77\%$ of the retina. Systems also have the capability to capture eyesteered (ES) images, whereby the subject's eye is pointed in one of four directions (superior, inferior, nasal, temporal), allowing a larger portion of the periphery of the retina to be imaged.

In this work, 4974 images in total were used: 3772 for training the model, 841 for validation and 411 images (of which, 266 were CP and 145 ES) were used as a separate test set. Care was taken to ensure that images from the same subject belonged to only one of the three sub-sets.

1st Conference on Medical Imaging with Deep Learning (MIDL 2018), Amsterdam, The Netherlands.

3 Methods

In CP images, the location of the optic disc (OD) with respect to the image axes, despite being difficult to determine precisely in the presence of lesions or imaging artifacts, can often give a good indication of the laterality of the image. In ES images, however, the location of the OD can vary considerably between images of the same laterality (Figure 1). As as result, it is not possible to discriminate between lateralities based only on the location of the OD. A CNN model was chosen for this classification task in order to capture other retinal features from the images, such as the morphology of the retinal vasculature and the orientation of the main vessel arcades.

3.1 Image pre-processing

The images were cropped and downsampled to 400×400 pixels to be passed into the network (Section 3.2). This reduces the computational cost while retaining the image features useful for classification. Examples of pre-processed images are shown in Figure 1.

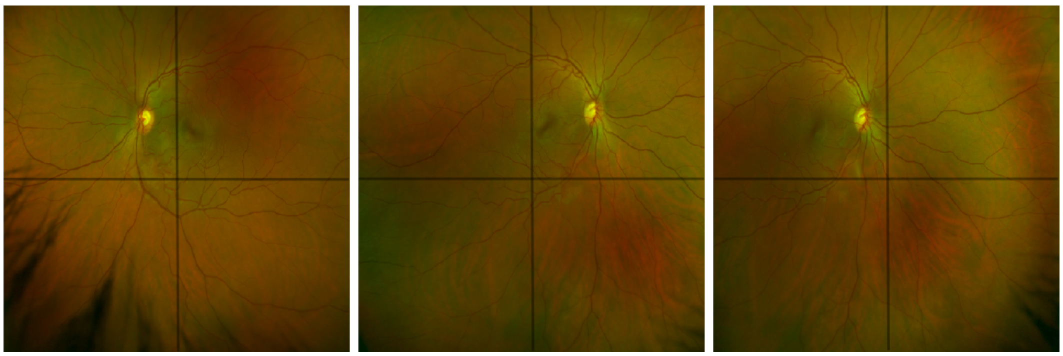

Figure 1: Pre-processed images from one subject. Left to right: Left eye, CP; right eye, CP; right eye, ES. Note that the bright OD is in the same region in the left CP and the right ES image.

3.1.1 Image augmentation

Before training, the size of the dataset was doubled by flipping the images horizontally and labeling them as the opposite laterality. Left and right images from a subject are not perfectly symmetrical (Figure 1), so the network was exposed to more of the retinal feature-space during training. Image augmentation was also performed online during training by applying random horizontal and vertical shift, shear, zoom, vertical flip and rotation.

3.2 Convolutional neural network

The CNN used in this work is made of three convolutional blocks (convolutional, rectified linear unit, max-pooling layers), followed by a flattening layer and a fully-connected layer. A dropout of 50% is applied to prevent overfitting before the final sigmoid activation layer. In total, the network has 9,465,953 trainable parameters.

The network was coded in Python 3.6.2, making use of the Keras library (5) (version 2.1.3) using Tensorflow (6) (version 1.3.0) as the backend. The training was done on a NVIDIA GeForce GTX 1070 GPU. The network was trained for 300 epochs, and the model state was saved at the epoch with the lowest validation loss. In this setup, the model took 4 hours and 46 minutes to train.

4 Results

The network was tested on a total of 822 test images (411 original and 411 flipped—see section 3.1.1). Results are shown in Table 1.

Saliency maps (Figure 2) generated with the keras-vis package (7) show that the network is paying attention to the orientation of the main retinal blood vessels radiating from the OD to determine laterality. A human reviewer would also use these features in laterality discrimination.

Table 1: Test set accuracy

Image mode	Left accuracy	Right accuracy	Combined accuracy
Central-pole	98.9%	98.9%	98.9%
Eyesteered	98.6%	96.0%	97.3%
		Total	**98.3%**

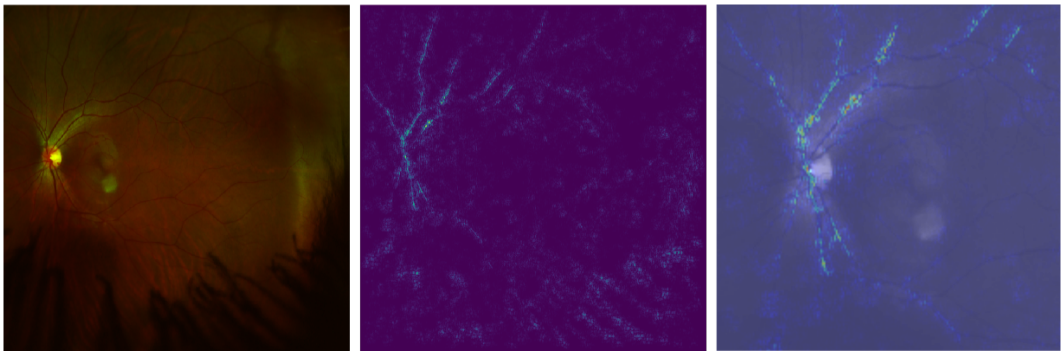

Figure 2: Left to right: Input image (left eye ES); saliency map of input image; saliency map superimposed onto input image, cropped around the OD.

5 Conclusion

In this paper we have presented a CNN-based method to determine the laterality of UWF retinal images. To our best knowledge, this is the first time that a deep learning approach to address this task on CP and ES UWF SLO retinal images has been proposed. Our model achieved a total accuracy of 98.3%; 98.9% on CP, and 97.3% on ES images.

By generating saliency maps, it can be seen that the network is activating on the vasculature around the OD, the same as a human reviewer would use to determine the image laterality.

References

[1] Geert J. S. Litjens et al. A survey on deep learning in medical image analysis. *CoRR*, abs/1702.05747, 2017. URL http://arxiv.org/abs/1702.05747.

[2] Ngan Meng et al. Classification of left and right eye retinal images, 2010. URL https://doi.org/10.1117/12.844638.

[3] Pallab Kanti Roy et al. Automatic eye type detection in retinal fundus image using fusion of transfer learning and anatomical features. In *2016 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–7, Nov 2016. doi: 10.1109/DICTA.2016.7797012.

[4] Robert H Webb et al. Scanning laser ophthalmoscope. *IEEE transactions on bio-medical engineering*, 28(7):488–92, 1981. ISSN 0018-9294. doi: 10.1109/TBME.1981.324734. URL http://www.ncbi.nlm.nih.gov/pubmed/16965154{%}5Cnhttp://www.ncbi.nlm.nih.gov/pubmed/7275128.

[5] François Chollet et al. Keras. https://github.com/fchollet/keras, 2015.

[6] Martín Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL https://www.tensorflow.org/. Software available from tensorflow.org.

[7] Raghavendra Kotikalapudi et al. keras-vis. https://github.com/raghakot/keras-vis, 2017.